

Feynman's Preface

These are the lectures in physics that I gave last year and the year before to the freshman and sophomore classes at Caltech. The lectures are, of course, not verbatim—they have been edited, sometimes extensively and sometimes less so. The lectures form only part of the complete course. The whole group of 180 students gathered in a big lecture room twice a week to hear these lectures and then they broke up into small groups of 15 to 20 students in recitation sections under the guidance of a teaching assistant. In addition, there was a laboratory session once a week.

The special problem we tried to get at with these lectures was to maintain the interest of the very enthusiastic and rather smart students coming out of the high schools and into Caltech. They have heard a lot about how interesting and exciting physics is—the theory of relativity, quantum mechanics, and other modern ideas. By the end of two years of our previous course, many would be very discouraged because there were really very few grand, new, modern ideas presented to them. They were made to study inclined planes, electrostatics, and so forth, and after two years it was quite stultifying. The problem was whether or not we could make a course which would save the more advanced and excited student by maintaining his enthusiasm.

The lectures here are not in any way meant to be a survey course, but are very serious. I thought to address them to the most intelligent in the class and to make sure, if possible, that even the most intelligent student was unable to completely encompass everything that was in the lectures—by putting in suggestions of applications of the ideas and concepts in various directions outside the main line of attack. For this reason, though, I tried very hard to make all the statements as accurate as possible, to point out in every case where the equations and ideas fitted into the body of physics, and how—when they learned more—things would be modified. I also felt that for such students it is important to indicate what it is that they should—if they are sufficiently clever—be able to understand by deduction from what has been said before, and what is being put in as something new. When new ideas came in, I would try either to deduce them if they were deducible, or to explain that it *was* a new idea which hadn't any basis in terms of things they had already learned and which was not supposed to be provable—but was just added in.

At the start of these lectures, I assumed that the students knew something when they came out of high school—such things as geometrical optics, simple chemistry ideas, and so on. I also didn't see that there was any reason to make the lectures

in a definite order, in the sense that I would not be allowed to mention something until I was ready to discuss it in detail. There was a great deal of mention of things **to** come, without complete discussions. These more complete discussions would come later when the preparation became more advanced. Examples are the discussions of inductance, and of energy levels, which are at first brought in in a very qualitative way and are later developed more completely.

At the same time that I was aiming at the more active student, I also wanted to take care of the fellow for whom the extra fireworks and side applications are merely disquieting and who cannot be expected to learn most of the material in the lecture at all. For such students I wanted there to be at least a central core or backbone of material which he *could* get. Even if he didn't understand everything in a lecture, I hoped he wouldn't get nervous. I didn't expect him to understand everything, but only the central and most direct features. It takes, of course, a certain intelligence on his part to see which are the central theorems and central ideas, and which are the more advanced side issues and applications which he may understand only in later years.

In giving these lectures there was one serious difficulty: in the way the course was given, there wasn't any feedback from the students to the lecturer to indicate how well the lectures were going over. This is indeed a very serious difficulty, and I don't know how good the lectures really are. The whole thing was essentially an experiment. And if I did it again I wouldn't do it the same way—I hope I *don't* have to do it again! I think, though, that things worked out—so far as the physics is concerned—quite satisfactorily in the first year.

In the second year I was not so satisfied. In the first part of the course, dealing with electricity and magnetism, I couldn't think of any really unique or different way of doing it—of any way that would be particularly more exciting than the usual way of presenting it. So I don't think I did very much in the lectures on electricity and magnetism. At the end of the second year I had originally intended to go on, after the electricity and magnetism, by giving some more lectures on the properties of materials, but mainly to take up things like fundamental modes, solutions of the diffusion equation, vibrating systems, orthogonal functions,... developing the first stages of what are usually called "the mathematical methods of physics." In retrospect, I think that if I were doing it again I would go back to that original idea. But since it was not planned that I would be giving these lectures again, it was suggested that it might be a good idea to try to give an introduction to the quantum mechanics—what you will find in Volume III.

It is perfectly clear that students who will major in physics can wait until their third year for quantum mechanics. On the other hand, the argument was made that many of the students in our course study physics as a background for their primary interest in other fields. And the usual way of dealing with quantum mechanics makes that subject almost unavailable for the great majority of students because they have to take so long to learn it. Yet, in its real applications—especially in its more complex applications, such as in electrical engineering and chemistry—the full machinery of the differential equation approach is not actually used. So I tried to describe the principles of quantum mechanics in a way which wouldn't require that one first know the mathematics of partial differential equations. Even for a physicist I think that is an interesting thing to try to do—to present quantum mechanics in this reverse fashion—for several reasons which may be apparent in the lectures themselves. However, I think that the experiment in the quantum mechanics part was not completely successful—in large part because I really did not have enough time at the end (I should, for instance, have had three or four more lectures in order to deal more completely with such matters as energy bands and the spatial dependence of amplitudes). Also, I had never presented the subject this way before, so the lack of feedback was particularly serious. I now believe the quantum mechanics should be given at a later time. Maybe I'll have a chance to do it again someday. Then I'll do it right.

The reason there are no lectures on how to solve problems is because there **were** recitation sections. Although I did put in three lectures in the first year on how to solve problems, they are not included here. Also there was a lecture on inertial

guidance which certainly belongs after the lecture on rotating systems, but which was, unfortunately, omitted. The fifth and sixth lectures are actually due to Matthew Sands, as I was out of town.

The question, of course, is how well this experiment has succeeded. My own point of view—which, however, does not seem to be shared by most of the people who worked with the students—is pessimistic. I don't think I did very well by the students. When I look at the way the majority of the students handled the problems on the examinations, I think that the system is a failure. Of course, my friends point out to me that there were one or two dozen students who—very surprisingly—understood almost everything in all of the lectures, and who were quite active in working with the material and worrying about the many points in an excited and interested way. These people have now, I believe, a first-rate background in physics—and they are, after all, the ones I was trying to get at. But then, "The power of instruction is seldom of much efficacy except in those happy dispositions where it is almost superfluous." (Gibbon)

Still, I didn't want to leave any student completely behind, as perhaps I did. I think one way we could help the students more would be by putting more hard work into developing a set of problems which would elucidate some of the ideas in the lectures. Problems give a good opportunity to fill out the material of the lectures and make more realistic, more complete, and more settled in the mind the ideas that have been exposed.

I think, however, that there isn't any solution to this problem of education other than to realize that the best teaching can be done only when there is a direct individual relationship between a student and a good teacher—a situation in which the student discusses the ideas, thinks about the things, and talks about the things. It's impossible to learn very much by simply sitting in a lecture, or even by simply doing problems that are assigned. But in our modern times we have so many students to teach that we have to try to find some substitute for the ideal. Perhaps my lectures can make some contribution. Perhaps in some small place where there are individual teachers and students, they may get some inspiration or some ideas from the lectures. Perhaps they will have fun thinking them through—or going on to develop some of the ideas further.

RICHARD P. FEYNMAN

June, 1963

Foreword

This book is based upon a course of lectures in introductory physics given by Prof. R. P. Feynman at the California Institute of Technology during the academic year 1961-62; it covers the first year of the two-year introductory course taken by all Caltech freshmen and sophomores, and was followed in 1962-63 by a similar series covering the second year. The lectures constitute a major part of a fundamental revision of the introductory course, carried out over a four-year period.

The need for a basic revision arose both from the rapid development of physics in recent decades and from the fact that entering freshmen have shown a steady increase in mathematical ability as a result of improvements in high school mathematics course content. We hoped to take advantage of this improved mathematical background, and also to introduce enough modern subject matter to make the course challenging, interesting, and more representative of present-day physics.

In order to generate a variety of ideas on what material to include and how to present it, a substantial number of the physics faculty were encouraged to offer their ideas in the form of topical outlines for a revised course. Several of these were presented and were thoroughly and critically discussed. It was agreed almost at once that a basic revision of the course could not be accomplished either by merely adopting a different textbook, or even by writing one *ab initio*, but that the new course should be centered about a set of lectures, to be presented at the rate of two or three per week; the appropriate text material would then be produced as a secondary operation as the course developed, and suitable laboratory experiments would also be arranged to fit the lecture material. Accordingly, a rough outline of the course was established, but this was recognized as being incomplete, tentative, and subject to considerable modification by whoever was to bear the responsibility for actually preparing the lectures.

Concerning the mechanism by which the course would finally be brought **to** life, several plans were considered. These plans were mostly rather similar, involving a cooperative effort by N staff members who would share the total burden symmetrically and equally: each man would take responsibility for $1/N$ of the material, deliver the lectures, and write text material for his part. However, the unavailability of sufficient staff, and the difficulty of maintaining a uniform point of view because of differences in personality and philosophy of individual participants, made such plans seem unworkable.

The realization that we actually possessed the means to create not just a new and different physics course, but possibly a unique one, came as a happy inspiration to Professor Sands. He suggested that Professor R. P. Feynman prepare and deliver the lectures, and that these be tape-recorded. When transcribed and edited, they would then become the textbook for the new course. This is essentially the plan that was adopted.

It was expected that the necessary editing would be minor, mainly consisting of supplying figures, and checking punctuation and grammar; it was to be done by one or two graduate students on a part-time basis. Unfortunately, this expectation was short-lived. It was, in fact, a major editorial operation to transform the verbatim transcript into readable form, even without the reorganization or revision of The subject matter that was sometimes required. Furthermore, it was not a job for a technical editor or for a graduate student, but one that required the close attention of a professional physicist for from ten to twenty hours per lecture!

The difficulty of the editorial task, together with the need to place the material in the hands of the students as soon as possible, set a strict limit upon the amount of "polishing" of the material that could be accomplished, and thus we were forced to aim toward a preliminary but technically correct product that could be used immediately, rather than one that might be considered final or finished. Because of an urgent need for more copies for our students, and a heartening interest on the part of instructors and students at several other institutions, we decided to publish the material in its preliminary form rather than wait for a further major revision which might never occur. We have no illusions as to the completeness, smoothness, or logical organization of the material; in fact, we plan several minor modifications in the course in the immediate future, and we hope that it will not become static in form or content.

In addition to the lectures, which constitute a centrally important part of the course, it was necessary also to provide suitable exercises to develop the students' experience and ability, and suitable experiments to provide first-hand contact with the lecture material in the laboratory. Neither of these aspects is in as advanced a state as the lecture material, but considerable progress has been made. Some exercises were made up as the lectures progressed, and these were expanded and amplified for use in the following year. However, because we are not yet satisfied that the exercises provide sufficient variety and depth of application of the lecture material to make the student fully aware of the tremendous power being placed at his disposal, the exercises are published separately in a less permanent form in order to encourage frequent revision.

A number of new experiments for the new course have been devised by Professor H. V. Neher. Among these are several which utilize the extremely low friction exhibited by a gas bearing: a novel linear air trough, with which quantitative measurements of one-dimensional motion, impacts, and harmonic motion can be made, and an air-supported, air-driven Maxwell top, with which accelerated rotational motion and gyroscopic precession and nutation can be studied. The development of new laboratory experiments is expected to continue **for a** considerable period of time.

The revision program was under the direction of Professors R. B. Leighton, H. V. Neher, and M. Sands. Officially participating in the program were Professors R. P. Feynman, G. Neugebauer, R. M. Sutton, H. P. Stabler,* F. Strong, and R. Vogt, from the division of Physics, Mathematics and Astronomy, and Professors T. Caughey, M. Plesset, and C. H. Wilts from the division of Engineering Science. The valuable assistance of all those contributing to the revision program is gratefully acknowledged. We are particularly indebted to the Ford Foundation, without whose financial assistance this program could not have been carried out.

ROBERT B. LEIGHTON

July, 1963

* 1961-62, while on leave from Williams College, Williamstown, Mass.

Contents

CHAPTER 1. ATOMS IN MOTION

1-1 Introduction 1-1
1-2 Matter is made of atoms 1-2
1-3 Atomic processes 1-5
1-4 Chemical reactions 1-6

CHAPTER 2. BASIC PHYSICS

2-1 Introduction 2-1
2-2 Physics before 1920 2-3
2-3 Quantum physics 2-6
2-4 Nuclei and particles 2-8

CHAPTER 3. THE RELATION OF PHYSICS TO OTHER SCIENCES

3-1 Introduction 3-1
3-2 Chemistry 3-1
3-3 Biology 3-2
3-4 Astronomy 3-6
3-5 Geology 3-7
3-6 Psychology 3-8
3-7 How did it get that way? 3-9

CHAPTER 4. CONSERVATION OF ENERGY

4-1 What is energy? 4-1
4-2 Gravitational potential energy 4-2
4-3 Kinetic energy 4-5
4–4 Other forms of energy 4-6

CHAPTER 5. TIME AND DISTANCE

5-1 Motion 5-1
5-2 Time 5-1
5-3 Short times 5-2
5-4 Long times 5-3
5-5 Units and standards of time 5-5
5-6 Large distances 5-5
5-7 Short distances 5-8

CHAPTER 6. PROBABILITY

6-1 Chance and likelihood 6-1
6-2 Fluctuations 6-3
6-3 The random walk 6-5
6-4 A probability distribution 6-7
6-5 The uncertainty principle 6-10

CHAPTER 7. THE THEORY OF GRAVITATION

7-1 Planetary motions 7-1
7-2 Kepler's laws 7-1
7-3 Development of dynamics 7-2
7–4 Newton's law of gravitation 7-3
7-5 Universal gravitation 7-5
7-6 Cavendish's experiment 7-9
7-7 What is gravity? 7-9
7-8 Gravity and relativity 7-11

CHAPTER 8. MOTION

8-1 Description of motion 8-1
8-2 Speed 8-2
8-3 Speed as a derivative 8-5
8-4 Distance as an integral 8-7
8-5 Acceleration 8-8

CHAPTER 9. NEWTON'S LAWS OF **DYNAMICS**

9-1 Momentum and force 9-1
9-2 Speed and velocity 9-2
9-3 Components of velocity, **acceleration, and force 9-3**
9-4 What is the force? 9-3
9-5 Meaning of the dynamical equations 9-4
9-6 Numerical solution of the equations 9-5
9-7 Planetary motions 9-6

CHAPTER 10. CONSERVATION OF MOMENTUM

10-1 Newton's Third Law 10-1
10-2 Conservation of momentum 10-2
10-3 Momentum *is* conserved! 10-5
10-4 Momentum and energy 10-7
10-5 Relativistic momentum 10-8

CHAPTER 11. VECTORS

11-1 Symmetry in physics 1**1-1**
11-2 Translations 11-1
11-3 Rotations 11-3
11-4 Vectors 11-5
11-5 Vector algebra **11-6**
11-6 Newton's laws in vector notation **11-7**
11-7 Scalar product of vectors 11-8

CHAPTER 12. CHARACTERISTICS OF FORCE

12-1 What is a force? 12-1
12-2 Friction 12-3
12-3 Molecular forces 12-6
12-4 Fundamental forces. Fields 12-7
12-5 Pseudo forces 12-10
12-6 Nuclear forces 12-12

CHAPTER 13. WORK AND POTENTIAL ENERGY (A)

13-1 Energy of a falling body 13-1
13-2 Work done by gravity 13-3
13-3 Summation of energy 13-6
13-4 Gravitational field of large objects 13-8

CHAPTER 14. WORK AND POTENTIAL ENERGY (conclusion)

14-1 Work 14-1
14-2 Constrained motion 14-3
14-3 Conservative forces 14-3
14-4 Nonconservative forces 14-6
14-5 Potentials and fields 14-7

CHAPTER 15. THE SPECIAL THEORY OF RELATIVITY

15-1 The principle of relativity 15-1
15-2 The Lorentz transformation 15-3
15-3 The Michelson-Morley experiment 15-3
15-4 Transformation of time 15-5
15-5 The Lorentz contraction 15-7
15-6 Simultaneity 15-7
15-7 Four-vectors 15-8
15-8 Relativistic dynamics 15-9
15-9 Equivalence of mass and energy 15-10

CHAPTER 16. RELATIVISTIC ENERGY AND MOMENTUM

16-1 Relativity and the philosophers **16-1**
16-2 The twin paradox 16-3
16-3 Transformation of velocities 16-4
16-4 Relativistic mass 16-6
16-5 Relativistic energy 16-8

CHAPTER 17. SPACE-TIME

17-1 The geometry of space-time 17-1
17-2 Space-time intervals 17-2
17-3 Past, present, and future 17-4
17-4 More about four-vectors 17-5
17-5 Four-vector algebra 17-7.

CHAPTER 18. ROTATION IN Two DIMENSIONS

18-1 The center of mass 18-1
18-2 Rotation of a rigid body 18-2
18-3 Angular momentum 18-5
18-4 Conservation of angular momentum 18-6

CHAPTER 19. CENTER OF MASS; MOMENT OF **INERTIA**

19-1 Properties of the center of mass 19-1
19-2 Locating the center of mass 19-4
19-3 Finding the moment of inertia 19-5
19–4 Rotational kinetic energy 19-7

CHAPTER 20. ROTATION IN SPACE

20-1 Torques in three dimensions 20-1
20-2 The rotation equations using cross **products 20–4**
20-3 The gyroscope 20-5
20-4 Angular momentum of a solid body 20-8

CHAPTER 21. THE HARMONIC OSCILLATOR

21-1 Linear differential equations 21-1
21-2 The harmonic oscillator 21-1
21-3 Harmonic motion and circular motion 21-4
21-4 Initial conditions 21-4
21-5 Forced oscillations 21-5

CHAPTER 22. ALGEBRA

22-1 Addition and multiplication 22-1
22-2 The inverse operations 22-2
22-3 Abstraction and generalization 22-3
22-4 Approximating irrational numbers 22-4
22-5 Complex numbers 22-7
22-6 Imaginary exponents 22-9

CHAPTER 23. RESONANCE

23-1 Complex numbers and harmonic motion 23-1
23-2 The forced oscillator with damping 23-3
23-3 Electrical resonance 23-5
23-4 Resonance in nature 23-7

CHAPTER 24. TRANSIENTS

24-1 The energy of an oscillator 24-1
24-2 Damped oscillations 24-2
24-3 Electrical transients 24-5

CHAPTER 25. LINEAR SYSTEMS AND **REVIEW**

25-1 Linear differential equations 25-1
25-2 Superposition of solutions 25-2
25-3 Oscillations in linear systems 25-5
25-4 Analogs in physics 25-6
25-5 Series and parallel impedances 25-8

CHAPTER 26. OPTICS: THE PRINCIPLE OF LEAST TIME

26-1 Light 26-1
26-2 Reflection and refraction 26-2
26-3 Fermat's principle of least time 26-3
26-4 Applications of Fermat's principle 26-5
26-5 A more precise statement of Fermat's principle 26-7
26-6 How it works 26-8

CHAPTER 27. GEOMETRICAL OPTICS

27-1 Introduction 27-1
27-2 The focal length of a spherical surface 27-1
27-3 The focal length of a lens 27-4
27-4 Magnification 27-5
27-5 Compound lenses **27-6**
27-6 Aberrations 27-7
27-7 Resolving power 27-7

CHAPTER 28. ELECTROMAGNETIC RADIATION

28-1 Electromagnetism 28-1
28-2 Radiation 28-3
28-3 **The** dipole radiator 28-5
28-4 Interference 28-6

CHAPTER 29. INTERFERENCE

29-1 Electromagnetic waves 29-1
29-2 Energy of radiation 29-2
29-3 Sinusoidal waves 29-2
29-4 Two dipole radiators 29-3
29-5 The mathematics of interference 29-5

CHAPTER 30. DIFFRACTION

30-1 The resultant amplitude due to *n* equal oscillators 30-1
30-2 The diffraction grating 30-3
30-3 Resolving power of a grating **30-5**
30-4 The parabolic antenna 30-6
30-5 Colored films; crystals 30-7
30-6 Diffraction by opaque screens 30-8
30-7 The field of a plane of oscillating charges 30-10

CHAPTER 31. THE ORIGIN OF THE REFRACTIVE INDEX

31-1 The index of refraction 31-1
31-2 The field due to the material 31-4
31-3 Dispersion 31-6
31-4 Absorption 31-8
31-5 The energy carried by an electric wave 31-9
31-6 Diffraction of light by a screen 31-10

CHAPTER 32. RADIATION DAMPING. **LIGHT** SCATTERING

32-1 Radiation resistance 32-1
32-2 The rate of radiation of energy 3.2-2
32-3 Radiation damping 32-3
32-4 Independent sources 32-5
32-5 Scattering of light 32-6

CHAPTER 33. POLARIZATION

33-1 The electric vector of light 33-1
33-2 Polarization of scattered light 33-3
33-3 Birefringence 33-3
33-4 Polarizers 33-5
33-5 Optical activity 33-6
33-6 The intensity of reflected light 33-7
33-7 Anomalous refraction 33-9

CHAPTER 34. RELATIVISTIC EFFECTS IN RADIATION

34-1 Moving sources 34-1
34-2 Finding the "apparent" motion 34-2
34-3 Synchrotron radiation 34-3
34-4 Cosmic synchrotron radiation 34-6
34-5 Bremsstrahlung 34-6
34-6 The Doppler effect 34-7
34-7 The ω, k four-vector 34-9
34-8 Aberration 34-10
34-9 The momentum of light 34-10

CHAPTER 35. COLOR VISION

35-1 The human eye 35-1
35-2 Color depends on intensity 35-2
35-3 Measuring the color sensation 35-3
35-4 The chromaticity diagram 35-6
35-5 The mechanism of color vision 35-7
35-6 Physiochemistry of color vision 35-9

CHAPTER 36. MECHANISMS OF SEEING

36-1 The sensation of color 36-1
36-2 The physiology of the eye 36-3
36-3 The rod cells 36-6
36-4 The compound (insect) eye 36-6
36-5 Other eyes 36-9
36-6 Neurology of vision 36-9

CHAPTER 37. QUANTUM **BEHAVIOR**

37-1 Atomic mechanics 37-1
37-2 An experiment with bullets 37-2
37-3 An experiment with waves 37-3
37-4 An experiment with electrons 37-4
37-5 The interference of electron waves 37-5
37-6 Watching the electrons 37-7
37-7 First principles of quantum mechanics 37-10
37-8 The uncertainty principle 37-11

CHAPTER 38. THE RELATION OF WAVE AND PARTICLE VIEWPOINTS

38-1 Probability wave amplitudes 38-1
38-2 Measurement of position and momentum 38-2
38-3 Crystal diffraction 38-4
38-4 The size of an atom 38-5
38-5 Energy levels 38-7
38-6 Philosophical implications 38-8

CHAPTER 39. THE KINETIC THEORY OF GASES

39-1 Properties of matter 39-1
39-2 The pressure of a gas 39-2
39-3 Compressibility of radiation 39-6
39-4 Temperature and kinetic energy 39-6
39-5 The ideal gas law 39-10

CHAPTER 40. THE PRINCIPLES OF STATISTICAL MECHANICS

40-1 The exponential atmosphere 40-1
40-2 The Boltzmann law 40-2
40-3 Evaporation of a liquid 40-3
40-4 The distribution of molecular speeds 40-4
40-5 The specific heats of gases 40-7
40-6 The failure of classical physics 40-8

CHAPTER 41. THE BROWNIAN MOVEMENT

41-1 Equipartition of energy 41-1
41-2 Thermal equilibrium of radiation 41-3
41-3 Equipartition and the quantum oscillator 41-6
41–4 The random walk 41-8

CHAPTER 42. APPLICATIONS OP KINETIC THEORY

42-1 Evaporation 42-1
42-2 Thermionic emission 42-4
42-3 Thermal ionization 42-5
42–4 Chemical kinetics 42-7
42-5 Einstein's laws of radiation 42-8

CHAPTER 43. DIFFUSION

43-1 Collisions between molecules 43-1
43-2 The mean free path 43-3
43-3 The drift speed 43-4
43-4 Ionic conductivity 43-6
43-5 Molecular diffusion 43-7
43-6 Thermal conductivity 43-9

CHAPTER 44. THE LAWS OF THERMODYNAMICS

44-1 Heat engines; the first law 44-1
44-2 The second law 44-3
44-3 Reversible engines 44-4
44-4 The efficiency of an ideal engine 44-7
44-5 The thermodynamic temperature 44-9
44-6 Entropy 44-10

CHAPTER 45. ILLUSTRATIONS OF THERMODYNAMICS

45-1 Internal energy 45-1
45-2 Applications 45-4
45-3 The Clausius-Clapeyron equation 45-6

CHAPTER 46. RATCHET AND PAWL

46-1 How a ratchet works 46-1
46-2 The ratchet as an engine 46-2
46-3 Reversibility in mechanics 46-4
46-4 Irreversibility 46-5
46-5 Order and entropy 46-7

CHAPTER 47. SOUND. THE WAVE EQUATION

47-1 Waves 47-1
47-2 The propagation of sound 47-3
47-3 The wave equation 47-4
47-4 Solutions of the wave equation 47-6
47-5 **The speed of sound** 47-7

CHAPTER 48. BEATS

48-1 Adding two waves 48-1
48-2 Beat notes and modulation 48-3
48-3 Side bands 48-4
48-4 Localized wave trains 48-5
48-5 Probability amplitudes for particles 48-7
48-6 Waves in three dimensions 48-9
48-7 Normal modes 48-10

CHAPTER 49. MODES

49-1 The reflection of waves 49-1
49-2 Confined waves, with natural frequencies 49-2
49-3 Modes in two dimensions 49-3
49-4 Coupled pendulums 49-6
49-5 Linear systems 49-7

INDEX

CHAPTER 50. HARMONICS

50-1 Musical tones 50-1
50-2 The Fourier series 50-2
50-3 Quality and consonance 50-3
50-4 The Fourier coefficients 50-5
50-5 The energy theorem 50-7
50-6 Nonlinear responses 50-8

CHAPTER 51. WAVES

51-1 Bow waves **51-1**
51-2 Shock waves 51-2
51-3 Waves in solids 51-4
51-4 Surface waves 51-7

CHAPTER 52. SYMMETRY **IN** PHYSICAL LAWS

52-1 Symmetry operations 52-1
52-2 Symmetry **in** space **and time 52-1**
52-3 Symmetry and conservation laws 52-3
52-4 Mirror reflections 52-4
52-5 Polar and axial vectors 52-6
52-6 Which hand is right? 52-8
52-7 Parity is not conserved! 52-8
52-8 Antimatter 52-10
52-9 Broken symmetries 52-11

Atoms in Motion

1-1 Introduction

1-1 Introduction

1-2 Matter is made of atoms

1-3 Atomic processes

1-4 Chemical reactions

This two-year course in physics is presented from the point of view that you, the reader, are going to be a physicist. This is not necessarily the case of course, but that is what every professor in every subject assumes! If you are going to be a physicist, you will have a lot to study: two hundred years of the most rapidly developing field of knowledge that there is. So much knowledge, in fact, that you might think that you cannot learn all of it in four years, and truly you cannot; you will have to go to graduate school too!

Surprisingly enough, in spite of the tremendous amount of work that has been done for all this time it is possible to condense the enormous mass of results to a large extent—that is, to find *laws* which summarize all our knowledge. Even so, the laws are so hard to grasp that it is unfair to you to start exploring this tremendous subject without some kind of map or outline of the relationship of one part of the subject of science to another. Following these preliminary remarks, the first three chapters will therefore outline the relation of physics to the rest of the sciences, the relations of the sciences to each other, and the meaning of science, to help us develop a "feel" for the subject.

You might ask why we cannot teach physics by just giving the basic laws on page one and then showing how they work in all possible circumstances, as we do in Euclidean geometry, where we state the axioms and then make all sorts of deductions. (So, not satisfied to learn physics in four years, you want to learn it in four minutes?) We cannot do it in this way for two reasons. First, we do not yet *know* all the basic laws: there is an expanding frontier of ignorance. Second, the correct statement of the laws of physics involves some very unfamiliar ideas which require advanced mathematics for their description. Therefore, one needs a considerable amount of preparatory training even to learn what the *words* mean. No, it is not possible to do it that way. We can only do it piece by piece.

Each piece, or part, of the whole of nature is always merely an *approximation to* the complete truth, or the complete truth so far as we know it. In fact, everything we know is only some kind of approximation, because *we know that we do not know all the laws* as yet. Therefore, things must be learned only to be unlearned again or, more likely, to be corrected.

The principle of science, the definition, almost, is the following: *The test of all knowledge is experiment.* Experiment is the *sole judge* of scientific "truth." But what is the source of knowledge? Where do the laws that are to be tested come from? Experiment, itself, helps to produce these laws, in the sense that it gives us hints. But also needed is *imagination* to create from these hints the great generalizations—to guess at the wonderful, simple, but very strange patterns beneath them all, and then to experiment to check again whether we have made the right guess. This imagining process is so difficult that there is a division of labor in physics: there are *theoretical* physicists who imagine, deduce, and guess at new laws, but do not experiment; and then there are *experimental* physicists who experiment, imagine, deduce, and guess.

We said that the laws of nature are approximate: that we first find the "wrong" ones, and then we find the "right" ones. Now, how can an experiment be "wrong"? First, in a trivial way: if something is wrong with the apparatus that you did not notice. But these things are easily fixed, and checked back and forth. So without snatching at such minor things, how *can* the results of an experiment be wrong? Only by being inaccurate. For example, the mass of an object never seems to

change; a spinning top has the same weight as a still one. So a "law" was invented: mass is constant, independent of speed. That "law" is now found to be incorrect. Mass is found to increase with velocity, but appreciable increases require velocities near that of light. A *true* law is: if an object moves with a speed of less than one hundred miles a second the mass is constant to within one part in a million. In some such approximate form this is a correct law. So in practice one might think that the new law makes no significant difference. Well, yes and no. For ordinary speeds we can certainly forget it and use the simple constant-mass law as a good approximation. But for high speeds we are wrong, and the higher the speed, the more wrong we are.

Finally, and most interesting, *philosophically we are completely wrong* with the approximate law. Our entire picture of the world has to be altered even though the mass changes only by a little bit. This is a very peculiar thing about the philosophy, or the ideas, behind the laws. Even a very small effect sometimes requires profound changes in our ideas.

Now, what should we teach first? Should we teach the *correct* but unfamiliar law with its strange and difficult conceptual ideas, for example the theory of relativity, four-dimensional space-time, and so on? Or should we first teach the simple "constant-mass" law, which is only approximate, but does not involve such difficult ideas? The first is more exciting, more wonderful, and more fun, but the second is easier to get at first, and is a first step to a real understanding of the second idea. This point arises again and again in teaching physics. At different times we shall have to resolve it in different ways, but at each stage it is worth learning what is now known, how accurate it is, how it fits into everything else, and how it may be changed when we learn more.

Let us now proceed with our outline, or general map, of our understanding of science today (in particular, physics, but also of other sciences on the periphery), so that when we later concentrate on some particular point we will have some idea of the background, why that particular point is interesting, and how it fits into the big structure. So, what *is* our over-all picture of the world?

1-2 Matter is made of atoms

If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms—little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world, if just a little imagination and thinking are applied.

To illustrate the power of the atomic idea, suppose that we have a drop of water a quarter of an inch on the side. If we look at it very closely we see nothing but water—smooth, continuous water. Even if we magnify it with the best optical microscope available—roughly two thousand times—then the water drop will be roughly forty feet across, about as big as a large room, and if we looked rather closely, we would *still* see relatively smooth water—but here and there small football-shaped things swimming back and forth. Very interesting. These are paramecia. You may stop at this point and get so curious about the paramecia with their wiggling cilia and twisting bodies that you go no further, except perhaps to magnify the paramecia still more and see inside. This, of course, is a subject for biology, but for the present we pass on and look still more closely at the water material itself, magnifying it two thousand times again. Now the drop of water extends about fifteen miles across, and if we look very closely at it we see a kind of teeming, something which no longer has a smooth appearance—it looks something like a crowd at a football game as seen from a very great distance. In order to see what this teeming is about, we will magnify it another two hundred and fifty times and we will see something similar to what is shown in Fig. 1-1. This is a picture of water magnified a billion times, but idealized in several ways.

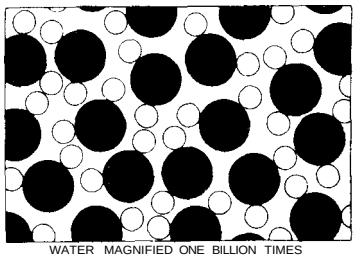


Figure 1-1

In the first place, the particles are drawn in a simple manner with sharp edges, which is inaccurate. Secondly, for simplicity, they are sketched almost schematically in a two-dimensional arrangement, but of course they are moving around in three dimensions. Notice that there are two kinds of "blobs" or circles to represent the atoms of oxygen (black) and hydrogen (white), and that each oxygen has two hydrogens tied to it. (Each little group of an oxygen with its two hydrogens is called a molecule.) The picture is idealized further in that the real particles in nature are continually jiggling and bouncing, turning and twisting around one another. You will have to imagine this as a dynamic rather than a static picture. Another thing that cannot be illustrated in a drawing is the fact that the particles are "stuck together"—that they attract each other, this one pulled by that one, etc. The whole group is "glued together," so to speak. On the other hand, the particles do not squeeze through each other. If you try to squeeze two of them too close together, they repel.

The atoms are 1 or 2 X 10^{-8} cm in radius. Now 10^{-8} cm is called an *angstrom* (just as another name), so we say they are 1 or 2 angstroms (Å) in radius. Another way to remember their size is this: if an apple is magnified to the size of the earth, then the atoms in the apple are approximately the size of the original apple.

Now imagine this great drop of water with all of these jiggling particles stuck together and tagging along with each other. The water keeps its volume; it does not fall apart, because of the attraction of the molecules for each other. If the drop is on a slope, where it can move from one place to another, the water will flow, but it does not just disappear—things do not just fly apart—because of the molecular attraction. Now the jiggling motion is what we represent as *heat:* when we increase the temperature, we increase the motion. If we heat the water, the jiggling increases and the volume between the atoms increases, and if the heating continues there comes a time when the pull between the molecules is not enough to hold them together and they *do* fly apart and become separated from one another. Of course, this is how we manufacture steam out of water—by increasing the temperature; the particles fly apart because of the increased motion.

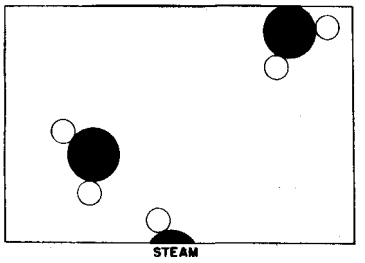


Figure 1-2

In Fig. 1-2 we have a picture of steam. This picture of steam fails in one respect: at ordinary atmospheric pressure there might be only a few molecules in a whole room, and there certainly would not be as many as three in this figure. Most squares this size would contain none—but we accidentally have two and a half or three in the picture (just so it would not be completely blank). Now in the case of steam we see the characteristic molecules more clearly than in the case of water. For simplicity, the molecules are drawn so that there is a 120° angle between them. In actual fact the angle is 105°3', and the distance between the center of a hydrogen and the center of the oxygen is 0.957 Å, so we know this molecule very well.

Let us see what some of the properties of steam vapor or any other gas are. The molecules, being separated from one another, will bounce against the walls. Imagine a room with a number of tennis balls (a hundred or so) bouncing around in perpetual motion. When they bombard the wall, this pushes the wall away. (Of course we would have to push the wall back.) This means that the gas exerts a jittery force which our coarse senses (not being ourselves magnified a billion times) feels only as an *average push.* In order to confine a gas we must apply a pressure. Figure 1-3 shows a standard vessel for holding gases (used in all textbooks), a cylinder with a piston in it. Now, it makes no difference what the shapes of water molecules are, so for simplicity we shall draw them as tennis balls or little dots. These things are in perpetual motion in all directions. So many of them are hitting the top piston all the time that to keep it from being patiently knocked out of the tank by this continuous banging, we shall have to hold the piston down by a certain force, which we call the *pressure* (really, the pressure times the area is the force). Clearly, the force is proportional to the area, for if we increase the area but keep the number of molecules per cubic centimeter the same, we increase the number of collisions with the piston in the same proportion as the area was increased.

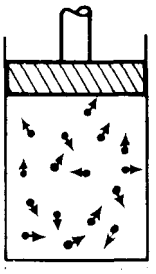


Figure 1-3

Now let us put twice as many molecules in this tank, so as to double the density, and let them have the same speed, i.e., the same temperature. Then, to a close approximation, the number of collisions will be doubled, and since each will be just as "energetic" as before, the pressure is proportional to the density. If we consider the true nature of the forces between the atoms, we would expect a slight decrease in pressure because of the attraction between the atoms, and a slight increase because of the finite volume they occupy. Nevertheless, to an excellent approximation, if the density is low enough that there are not many atoms, *the pressure is proportional to the density.*

We can also see something else: If we increase the temperature without changing the density of the gas, i.e., if we increase the speed of the atoms, what is going to happen to the pressure? Well, the atoms hit harder because they are moving faster, and in addition they hit more often, so the pressure increases. You see how simple the ideas of atomic theory are.

Let us consider another situation. Suppose that the piston moves inward, so that the atoms are slowly compressed into a smaller space. What happens when an atom hits the moving piston? Evidently it picks up speed from the collision. You can try it by bouncing a ping-pong ball from a forward-moving paddle, for example, and you will find that it comes off with more speed than that with which it struck. (Special example: if an atom happens to be standing still and the piston hits it, it will certainly move.) So the atoms are "hotter" when they come away from the piston than they were before they struck it. Therefore all the atoms which are in the vessel will have picked up speed. This means that *when we compress a gas slowly, the temperature of the gas increases. So,* under slow *compression,* a gas will *increase* in temperature, and under slow *expansion it* will *decrease* in temperature.

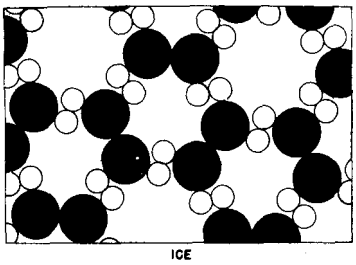


Figure 1-4

We now return to our drop of water and look in another direction. Suppose that we decrease the temperature of our drop of water. Suppose that the jiggling of the molecules of the atoms in the water is steadily decreasing. We know that there are forces of attraction between the atoms, so that after a while they will not be able to jiggle so well. What will happen at very low temperatures is indicated in Fig. 1-4: the molecules lock into a new pattern which is *ice.* This particular schematic diagram of ice is wrong because it is in two dimensions, but it is right qualitatively. The interesting point is that the material has a *definite place for every atom,* and you can easily appreciate that if somehow or other we were to hold all the atoms at one end of the drop in a certain arrangement, each atom in a certain place, then because of the structure of interconnections, which is rigid, the other end miles away (at our magnified scale) will have a definite location. So if we hold a needle of ice at one end, the other end resists our pushing it aside, unlike the case of water, in which the structure is broken down because of the increased jiggling so that the atoms all move around in different ways. The difference between solids and liquids is, then, that in a solid the atoms are arranged in some kind of an array, called a *crystalline array,* and they do not have a random position at long distances; the position of the atoms on one side of the crystal is determined by that of other atoms millions of atoms away on the other side of the crystal. Figure 1-4 is an invented arrangement for ice, and although it contains many of the correct features of ice, it is not the true arrangement. One of the correct features is that there is a part of the symmetry that is hexagonal. You can see that if we turn the picture around an axis by 120°, the picture returns to itself. So there is a *symmetry* in the ice which accounts for the six-sided appearance of snowflakes. Another thing we can see from Fig. 1-4 is why ice shrinks when it melts. The particular crystal pattern of ice shown here has many "holes" in it, as does the true ice structure. When the organization breaks down, these holes can be occupied by molecules. Most simple substances, with the exception of water and type metal, *expand* upon melting, because the atoms are closely packed in the solid crystal and upon melting need more room to jiggle around, but an open structure collapses, as in the case of water.

Now although ice has a "rigid" crystalline form, its temperature can change—ice has heat. If we wish, we can change the amount of heat. What is the heat in

the case of ice? The atoms are not standing still. They are jiggling and vibrating. So even though there is a definite order to the crystal—a definite structure—all of the atoms are vibrating "in place." As we increase the temperature, they vibrate with greater and greater amplitude, until they shake themselves out of place. We call this *melting.* As we decrease the temperature, the vibration decreases and decreases until, at absolute zero, there is a minimum amount of vibration that the atoms can have, but *not zero.* This minimum amount of motion that atoms can have is not enough to melt a substance, with one exception: helium. Helium merely decreases the atomic motions as much as it can, but even at absolute zero there is still enough motion to keep it from freezing. Helium, even at absolute zero, does not freeze, unless the pressure is made so great as to make the atoms squash together. If we increase the pressure, we *can* make it solidify.

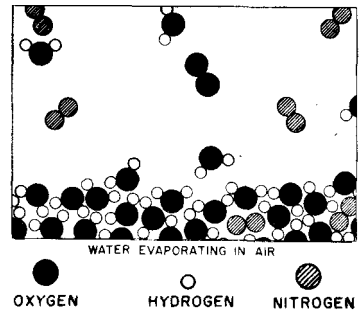


Figure 1-5

1-3 Atomic processes

So much for the description of solids, liquids, and gases from the atomic point of view. However, the atomic hypothesis also describes *processes,* and so we shall now look at a number of processes from an atomic standpoint. The first process that we shall look at is associated with the surface of the water. What happens at the surface of the water? We shall now make the picture more complicated—and more realistic—by imagining that the surface is in air. Figure 1-5 shows the surface of water in air. We see the water molecules as before, forming a body of liquid water, but now we also see the surface of the water. Above the surface we find a number of things: First of all there are water molecules, as in steam. This is *water vapor,* which is always found above liquid water. (There is an equilibrium between the steam vapor and the water which will be described later.) In addition we find some other molecules—here two oxygen atoms stuck together by themselves, forming an *oxygen molecule*, there two nitrogen atoms also stuck together to make a nitrogen molecule. Air consists almost entirely of nitrogen, oxygen, some water vapor, and lesser amounts of carbon dioxide, argon, and other things. So above the water surface is the air, a gas, containing some water vapor. Now what is happening in this picture? The molecules in the water are always jiggling around. From time to time, one on the surface happens to be hit a little harder than usual, and gets knocked away. It is hard to see that happening in the picture because it is a *still* picture. But we can imagine that one molecule near the surface has just been hit and is flying out, or perhaps another one has been hit and is flying out. Thus, molecule by molecule, the water disappears—it evaporates. But if we *close* the vessel above, after a while we shall find a large number of molecules of water amongst the air molecules. From time to time, one of these vapor molecules comes flying down to the water and gets stuck again. So we see that what looks like a dead, uninteresting thing—a glass of water with a cover, that has been sitting there for perhaps twenty years—really contains a dynamic and interesting phenomenon which is going on all the time. To our eyes, our crude eyes, nothing is changing, but if we could see it a billion times magnified, we would see that from its own point of view it is always changing: molecules are leaving the surface, molecules are coming back.

Why do *we* see *no change?* Because just as many molecules are leaving as are coming back! In the long run "nothing happens." If we then take the top of the vessel off and blow the moist air away, replacing it with dry air, then the number of molecules leaving is just the same as it was before, because this depends on the jiggling of the water, but the number coming back is greatly reduced because there are so many fewer water molecules above the water. Therefore there are more going out than coming in, and the water evaporates. Hence, if you wish to evaporate water turn on the fan!

Here is something else: Which molecules leave? When a molecule leaves it is due to an accidental, extra accumulation of a little bit more than ordinary energy, which it needs if it is to break away from the attractions of its neighbors. Therefore, since those that leave have more energy than the average, the ones that are left have *less* average motion than they had before. So the liquid gradually

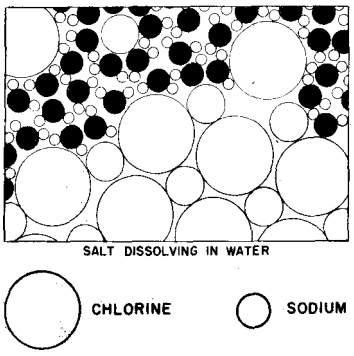
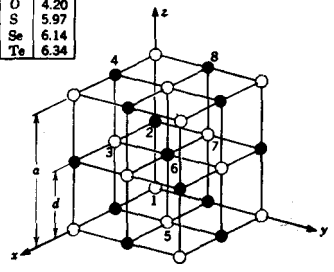


Figure 1-6

Crystal	●	○	a(Å)
Rocksalt	Na	Cl	5.64
Sylvine	K	Cl	6.28
	Ag	Cl	5.54
	Mg	O	4.20
Galena	Pb	S	5.97
	Pb	Se	6.14
	Pb	Te	6.34



Nearest neighbor distance d=a/2

Figure 1-7

cools if it evaporates. Of course, when a molecule of vapor comes from the air to the water below there is a sudden great attraction as the molecule approaches the surface. This speeds up the incoming molecule and results in generation of heat. So when they leave they take away heat; when they come back they generate heat. Of course when there is no net evaporation the result is nothing—the water is not changing temperature. If we blow on the water so as to maintain a continuous preponderance in the number evaporating, then the water is cooled. Hence, blow on soup to cool it!

Of course you should realize that the processes just described are more complicated than we have indicated. Not only does the water go into the air, but also, from time to time, one of the oxygen or nitrogen molecules will come in and "get lost" in the mass of water molecules, and work its way into the water. Thus the air dissolves in the water; oxygen and nitrogen molecules will work their way into the water and the water will contain air. If we suddenly take the air away from the vessel, then the air molecules will leave more rapidly than they come in, and in doing so will make bubbles. This is very bad for divers, as you may know.

Now we go on to another process. In Fig. 1-6 we see, from an atomic point of view, a solid dissolving in water. If we put a crystal of salt in the water, what will happen? Salt is a solid, a crystal, an organized arrangement of "salt atoms." Figure 1-7 is an illustration of the three-dimensional structure of common salt, sodium chloride. Strictly speaking, the crystal is not made of atoms, but of what we call *ions.* An ion is an atom which either has a few extra electrons or has lost a few electrons. In a salt crystal we find chlorine ions (chlorine atoms with an extra electron) and sodium ions (sodium atoms with one electron missing). The ions all stick together by electrical attraction in the solid salt, but when we put them in the water we find, because of the attractions of the negative oxygen and positive hydrogen for the ions, that some of the ions jiggle loose. In Fig. 1-6 we see a chlorine ion getting loose, and other atoms floating in the water in the form of ions. This picture was made with some care. Notice, for example, that the hydrogen ends of the water molecules are more likely to be near the chlorine ion, while near the sodium ion we are more likely to find the oxygen end, because the sodium is positive and the oxygen end of the water is negative, and they attract electrically. Can we tell from this picture whether the salt is *dissolving in* water or *crystallizing out* of water? Of course we *cannot* tell, because while some of the atoms are leaving the crystal other atoms are rejoining it. The process is a *dynamic* one, just as in the case of evaporation, and it depends on whether there is more or less salt in the water than the amount needed for equilibrium. By equilibrium we mean that situation in which the rate at which atoms are leaving just matches the rate at which they are coming back. If there is almost no salt in the water, more atoms leave than return, and the salt dissolves. If, on the other hand, there are too many "salt atoms," more return than leave, and the salt is crystallizing.

In passing, we mention that the concept of a *molecule* of a substance is only approximate and exists only for a certain class of substances. It is clear in the case of water that the three atoms are actually stuck together. It is not so clear in the case of sodium chloride in the solid. There is just an arrangement of sodium and chlorine ions in a cubic pattern. There is no natural way to group them as "molecules of salt."

Returning to our discussion of solution and precipitation, if we increase the temperature of the salt solution, then the rate at which atoms are taken away is increased, and so is the rate at which atoms are brought back. It turns out to be very difficult, in general, to predict which way it is going to go, whether more or less of the solid will dissolve. Most substances dissolve more, but some substances dissolve less, as the temperature increases.

1-4 Chemical reactions

In all of the processes which have been described so far, the atoms and the ions have not changed partners, but of course there are circumstances in which the atoms do change combinations, forming new molecules. This is illustrated in

Fig. 1-8. A process in which the rearrangement of the atomic partners occurs is what we call a *chemical reaction.* The other processes so far described are called physical processes, but there is no sharp distinction between the two. (Nature does not care what we call it, she just keeps on doing it.) This figure is supposed to represent carbon burning in oxygen. In the case of oxygen, *two* oxygen atoms stick together very strongly. (Why do not *three* or *even four* stick together? That is one of the very peculiar characteristics of such atomic processes. Atoms are very special: they like certain particular partners, certain particular directions, and so on. It is the job of physics to analyze why each one wants what it wants. At any rate, two oxygen atoms form, saturated and happy, a molecule.)

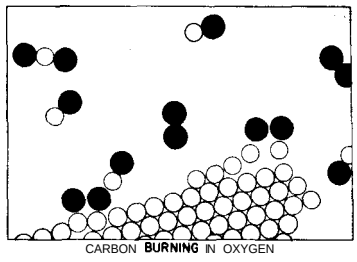


Figure 1-8

The carbon atoms are supposed to be in a solid crystal (which could be graphite or diamond*). Now, for example, one of the oxygen molecules can come over to the carbon, and each atom can pick up a carbon atom and go flying off in a new combination—"carbon-oxygen"—which is a molecule of the gas called carbon monoxide. It is given the chemical name CO. It is very simple: the letters "CO" are practically a picture of that molecule. But carbon attracts oxygen much more than oxygen attracts oxygen or carbon attracts carbon. Therefore in this process the oxygen may arrive with only a little energy, but the oxygen and carbon will snap together with a tremendous vengeance and commotion, and everything near them will pick up the energy. A large amount of motion energy, kinetic energy, is thus generated. This of course is *burning;* we are getting *heat* from the combination of oxygen and carbon. The heat is ordinarily in the form of the molecular motion of the hot gas, but in certain circumstances it can be so enormous that it generates *light.* That is how one *gets flames.*

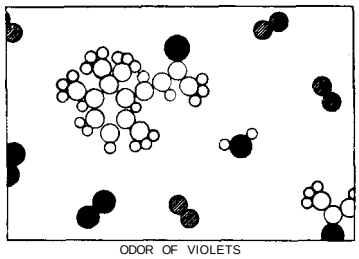


Figure 1-9

In addition, the carbon monoxide is not quite satisfied. It is possible for it to attach another oxygen, so that we might have a much more complicated reaction in which the oxygen is combining with the carbon, while at the same time there happens to be a collision with a carbon monoxide molecule. One oxygen atom could attach itself to the CO and ultimately form a molecule, composed of one carbon and two oxygens, which is designated CO_2 and called carbon dioxide. If we burn the carbon with very little oxygen in a very rapid reaction (for example, in an automobile engine, where the explosion is so fast that there is not time for it to make carbon dioxide) a considerable amount of carbon monoxide is formed. In many such rearrangements, a very large amount of energy is released, forming explosions, flames, etc., depending on the reactions. Chemists have studied these arrangements of the atoms, and found that every substance is some type of *arrangement of atoms.*

To illustrate this idea, let us consider another example. If we go into a field of small violets, we know what "that smell" is. It is some kind of *molecule,* or arrangement of atoms, that has worked its way into our noses. First of all, *how* did it work its way in? That is rather easy. If the smell is some kind of molecule in the air, jiggling around and being knocked every which way, it might have *accidentally* worked its way into the nose. Certainly it has no particular desire to get into our nose. It is merely one helpless part of a jostling crowd of molecules, and in its aimless wanderings this particular chunk of matter happens to find itself in the nose.

Now chemists can take special molecules like the odor of violets, and analyze them and tell us the *exact arrangement* of the atoms in space. We know that the carbon dioxide molecule is straight and symmetrical: O—C—O. (That can be determined easily, too, by physical methods.) However, even for the vastly more complicated arrangements of atoms that there are in chemistry, one can, by a long, remarkable process of detective work, find the arrangements of the atoms. Figure 1-9 is a picture of the air in the neighborhood of a violet; again we find nitrogen and oxygen in the air, and water vapor. (Why is there water vapor? Because the violet is *wet.* All plants transpire.) However, we also see a "monster" composed of carbon atoms, hydrogen atoms, and oxygen atoms, which have picked a certain particular pattern in which to be arranged. It is a much more complicated arrange-

*One *can* burn a diamond in air.

ment than that of carbon dioxide; in fact, it is an enormously complicated arrangement. Unfortunately, we cannot picture all that is really known about it chemically, because the precise arrangement of all the atoms is actually known in three dimensions, while our picture is in only two dimensions. The six carbons which form a ring do not form a flat ring, but a kind of "puckered" ring. All of the angles and distances are known. So a chemical *formula* is merely a picture of such a molecule. When the chemist writes such a thing on the blackboard, he is trying to "draw," roughly speaking, in two dimensions. For example, we see a "ring" of six carbons, and a "chain" of carbons hanging on the end, with an oxygen second from the end, three hydrogens tied to that carbon, two carbons and three hydrogens sticking up here, etc.

How does the chemist find what the arrangement is? He mixes bottles full of stuff together, and if it turns red, it tells him that it consists of one hydrogen and two carbons tied on here; if it turns blue, on the other hand, that is not the way it is at all. This is one of the most fantastic pieces of detective work that has ever been done—organic chemistry. To discover the arrangement of the atoms in these enormously complicated arrays the chemist looks at what happens when he mixes two different substances together. The physicist could never quite believe that the chemist knew what he was talking about when he described the arrangement of the atoms. For about twenty years it has been possible, in some cases, to look at such molecules (not quite as complicated as this one, but some which contain parts of it) by a physical method, and it has been possible to locate every atom, not by looking at colors, but by *measuring where they are.* And lo and behold!, the chemists are almost always correct.

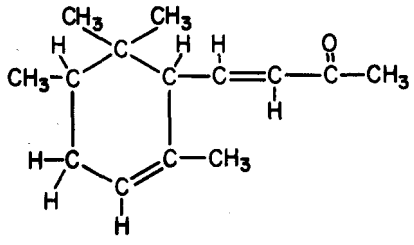


Fig. **1-10.** The substance pictured is α-irone.

It turns out, in fact, that in the odor of violets there are three slightly different molecules, which differ only in the arrangement of the hydrogen atoms.

One problem of chemistry is to name a substance, so that we will know what it is. Find a name for this shape! Not only must the name tell the shape, but it must also tell that here is an oxygen atom, there a hydrogen—exactly what and where each atom is. So we can appreciate that the chemical names must be complex in order to be complete. You see that the name of this thing in the more complete form that will tell you the structure of it is 4-(2, 2, 3, 6 tetramethyl-5-cyclohexanyl)-3-buten-2-one, and that tells you that this is the arrangement. We can appreciate the difficulties that the chemists have, and also appreciate the reason for such long names. It is not that they wish to be obscure, but they have an extremely difficult problem in trying to describe the molecules in words!

How do we *know* that there are atoms? By one of the tricks mentioned earlier: we make the *hypothesis* that there are atoms, and one after the other results come out the way we predict, as they ought to if things *are* made of atoms. There is also somewhat more direct evidence, a good example of which is the following: The atoms are so small that you cannot see them with a light microscope—in fact, not even with an *electron* microscope. (With a light microscope you can only see things which are much bigger.) Now if the atoms are always in motion, say in water, and we put a big ball of something in the water, a ball much bigger than the atoms, the ball will jiggle around—much as in a push ball game, where a great big ball is pushed around by a lot of people. The people are pushing in various directions, and the ball moves around the field in an irregular fashion. So, in the same way, the "large ball" will move because of the inequalities of the collisions on one side to the other, from one moment to the next. Therefore, if we look at very tiny particles (colloids) in water through an excellent microscope, we see a perpetual jiggling of the particles, which is the result of the bombardment of the atoms. This is called the *Brownian motion.*

We can see further evidence for atoms in the structure of crystals. In many cases the structures deduced by x-ray analysis agree in their spatial "shapes" with the forms actually exhibited by crystals as they occur in nature. The angles between the various "faces" of a crystal agree, within seconds of arc, with angles deduced on the assumption that a crystal is made of many "layers" of atoms.

Everything is made of atoms. That is the key hypothesis. The most important hypothesis in all of biology, for example, is that *everything that animals do, atoms*

do. In other words, *there is nothing that living things do that cannot be understood from the point of view that they are made of atoms acting according to the laws of physics.* This was not known from the beginning: it took some experimenting and theorizing to suggest this hypothesis, but now it is accepted, and it is the most useful theory for producing new ideas in the field of biology.

If a piece of steel or a piece of salt, consisting of atoms one next to the other, can have such interesting properties; if water—which is nothing but these little blobs, mile upon mile of the same thing over the earth—can form waves and foam, and make rushing noises and strange patterns as it runs over cement; if all of this, all the life of a stream of water, can be nothing but a pile of atoms, *how much more is possible?* If instead of arranging the atoms in some definite pattern, again and again repeated, on and on, or even forming little lumps of complexity like the odor of violets, we make an arrangement which is *always different* from place to place, with different kinds of atoms arranged in many ways, continually changing, not repeating, how much more marvelously is it possible that this thing might behave? Is it possible that that "thing" walking back and forth in front of you, talking to you, is a great glob of these atoms in a very complex arrangement, such that the sheer complexity of it staggers the imagination as to what it can do? When we say we are a pile of atoms, we do not mean we are *merely* a pile of atoms, because a pile of atoms which is not repeated from one to the other might well have the possibilities which you see before you in the mirror.

2

Basic Physics

2-1 Introduction

2-1 Introduction
2-2 Physics before 1920
2-3 Quantum physics
2-4 Nuclei and particles

In this chapter, we shall examine the most fundamental ideas that we have about physics—the nature of things as we see them at the present time. We shall not discuss the history of how we know that all these ideas are true; you will learn these details in due time.

The things with which we concern ourselves in science appear in myriad forms, and with a multitude of attributes. For example, if we stand on the shore and look at the sea, we see the water, the waves breaking, the foam, the sloshing motion of the water, the sound, the air, the winds and the clouds, the sun and the blue sky, and light; there is sand and there are rocks of various hardness and permanence, color and texture. There are animals and seaweed, hunger and disease, and the observer on the beach; there may be even happiness and thought. Any other spot in nature has a similar variety of things and influences. It is always as complicated as that, no matter where it is. Curiosity demands that we ask questions, that we try to put things together and try to understand this multitude of aspects as perhaps resulting from the action of a relatively small number of elemental things and forces acting in an infinite variety of combinations.

For example: Is the sand other than the rocks? That is, is the sand perhaps nothing but a great number of very tiny stones? Is the moon a great rock? If we understood rocks, would we also understand the sand and the moon? Is the wind a sloshing of the air analogous to the sloshing motion of the water in the sea? What common features do different movements have? What is common to different kinds of sound? How many different colors are there? And so on. In this way we try gradually to analyze all things, to put together things which at first sight look different, with the hope that we may be able to *reduce* the number of *different* things and thereby understand them better.

A few hundred years ago, a method was devised to find partial answers to such questions. *Observation, reason,* and *experiment* make up what we call the *scientific method.* We shall have to limit ourselves to a bare description of **our** basic view of what is sometimes called *fundamental physics,* or fundamental ideas which have arisen from the application of the scientific method.

What do we mean by "understanding" something? We can imagine that this complicated array of moving things which constitutes "the world" is something like a great chess game being played by the gods, and we are observers of the game. We do not know what the rules of the game are; all we are allowed to do is to *watch* the playing. Of course, if we watch long enough, we may eventually catch on to a few of the rules. *The rules of the game* are what we mean *by fundamental physics.* Even if we knew every rule, however, we might not be able to understand why a particular move is made in the game, merely because it is too complicated and our minds are limited. If you play chess you must know that it is easy to learn all the rules, and yet it is often very hard to select the best move or to understand why a player moves as he does. So it is in nature, only much more so; but we may be able at least to find all the rules. Actually, we do not have all the rules now. (Every once in a while something like castling is going on that we still do not understand.) Aside from not knowing all of the rules, what we really can explain in terms of those rules is very limited, because almost all situations are so enormously complicated that we cannot follow the plays of the game using the rules, much less tell what is going to happen next. We must, therefore, limit ourselves to the more basic question of the rules of the game. If we know the rules, we consider that we "understand" the world.

How can we tell whether the rules which we "guess" at are really right if **we** cannot analyze the game very well? There are, roughly speaking, three ways. First, there may be situations where nature has arranged, or we arrange nature, to be simple and to have so few parts that we can predict exactly what will happen, and thus we can check how our rules work. (In one corner of the board there may be only a few chess pieces at work, and that we can figure out exactly.)

A second good way to check rules is in terms of less specific rules derived from them. For example, the rule on the move of a bishop on a chessboard is that it moves only on the diagonal. One can deduce, no matter how many moves may be made, that a certain bishop will always be on a red square. So, without being able to follow the details, we can always check our idea about the bishop's motion by finding out whether it is always on a red square. Of course it will be, for a long time, until all of a sudden we find that it is on a *black* square (what happened of course, is that in the meantime it was captured, another pawn crossed for queening, and it turned into a bishop on a black square). That is the way it is in physics. For a long time we will have a rule that works excellently in an over-all way, even when we cannot follow the details, and then some time we may discover a *new rule.* From the point of view of basic physics, the most interesting phenomena are of course in the *new* places, the places where the rules do not work—not the places where they *do* work! That is the way in which we discover new rules.

The third way to tell whether our ideas are right is relatively crude but probably the most powerful of them all. That is, by rough *approximation.* While we may not be able to tell why Alekhine moves *this particular piece,* perhaps we can *roughly* understand that he is gathering his pieces around the king to protect it, more or less, since that is the sensible thing to do in the circumstances. In the same way, we can often understand nature, more or less, without being able to see what *every little piece* is doing, in terms of our understanding of the game.

At first the phenomena of nature were roughly divided into classes, like heat, electricity, mechanics, magnetism, properties of substances, chemical phenomena, light or optics, x-rays, nuclear physics, gravitation, meson phenomena, etc. However, the aim is to see *complete nature* as different aspects of *one set* of phenomena. That is the problem in basic theoretical physics, today—to *find the laws behind experiment;* to *amalgamate these classes.* Historically, we have always been able to amalgamate them, but as time goes on new things are found. We were amalgamating very well, when all of a sudden x-rays were found. Then we amalgamated some more, and mesons were found. Therefore, at any stage of the game, it always looks rather messy. A great deal is amalgamated, but there are always many wires or threads hanging out in all directions. That is the situation today, which we shall try to describe.

Some historic examples of amalgamation are the following. First, take *heat* and *mechanics.* When atoms are in motion, the more motion, the more heat the system contains, and so *heat and all temperature effects can be represented by the laws of mechanics.* Another tremendous amalgamation was the discovery of the relation between electricity, magnetism, and light, which were found to be different aspects of the same thing, which we call today the *electromagnetic field.* Another amalgamation is the unification of chemical phenomena, the various properties of various substances, and the behavior of atomic particles, which is in the *quantum mechanics of chemistry.*

The question is, of course, is it going to be possible to amalgamate *everything,* and merely discover that this world represents different aspects of *one* thing? Nobody knows. All we know is that as we go along, we find that we can amalgamate pieces, and then we find some pieces that do not fit, and we keep trying to put the jigsaw puzzle together. Whether there are a finite number of pieces, and whether there is even a border to the puzzle, is of course unknown. It will never be known until we finish the picture, if ever. What we wish to do here is to see to what extent this amalgamation process has gone on, and what the situation is at present, in understanding basic phenomena in terms of the smallest set of principles. To express it in a simple manner, *what are things made of and how few elements are there ?*

2-2 Physics before 1920

It is a little difficult to begin at once with the present view, so we shall first see how things looked in about 1920 and then take a few things out of that picture. Before 1920, our world picture was something like this: The "stage" on which the universe goes is the three-dimensional *space* of geometry, as described by Euclid, and things change in a medium called *time.* The elements on the stage are *particles,* for example the atoms, which have some *properties.* First, the property of inertia: if a particle is moving it keeps on going in the same direction unless *forces* act upon it. The second element, then, is *forces,* which were then thought to be of two varieties: First, an enormously complicated, detailed kind of interaction force which held the various atoms in different combinations in a complicated way, which determined whether salt would dissolve faster or slower when we raise the temperature. The other force that was known was a long-range interaction—a smooth and quiet attraction—which varied inversely as the square of the distance, and was called *gravitation.* This law was known and was very simple. *Why* things remain in motion when they are moving, or *why* there is a law of gravitation was, of course, not known.

A description of nature is what we are concerned with here. From this point of view, then, a gas, and indeed *all* matter, is a myriad of moving particles. Thus many of the things we saw while standing at the seashore can immediately be connected. First the pressure: this comes from the collisions of the atoms with the walls or whatever; the drift of the atoms, if they are all moving in one direction on the average, is wind; the *random* internal motions are the *heat.* There are waves of excess density, where too many particles have collected, and so as they Tush off they push up piles of particles farther out, and so on. This wave of excess density is *sound.* It is a tremendous achievement to be able to understand so much. Some of these things were described in the previous chapter.

What *kinds* of particles are there? There were considered to be 92 at that time: 92 different kinds of atoms were ultimately discovered. They had different names associated with their chemical properties.

The next part of the problem was, *what are the short-range forces?* Why does carbon attract one oxygen or perhaps two oxygens, but not three oxygens? What is the machinery of interaction between atoms? Is it gravitation? The answer is no. Gravity is entirely too weak. But imagine a force analogous to gravity, varying inversely with the square of the distance, but enormously more powerful and having one difference. In gravity everything attracts everything else, but now imagine that there are *two kinds* of "things," and that this new force (which is the electrical force, of course) has the property that likes *repel* but unlikes *attract.* The "thing" that carries this strong interaction is called *charge.*

Then what do we have? Suppose that we have two unlikes that attract each other, a plus and a minus, and that they stick very close together. Suppose we have another charge some distance away. Would it feel any attraction? It would feel *practically none,* because if the first two are equal in size, the attraction for the one and the repulsion for the other balance out. Therefore there is very little force at any appreciable distance. On the other hand, if we get *very close* with the extra charge, *attraction* arises, because the repulsion of likes and attraction of unlikes will tend to bring unlikes closer together and push likes farther apart. Then the repulsion will be *less* than the attraction. This is the reason why the atoms, which are constituted out of plus and minus electric charges, feel very little force when they are separated by appreciable distance (aside from gravity). When they come close together, they can "see inside" each other and rearrange their charges, with the result that they have a very strong interaction. The ultimate basis of an interaction between the atoms is *electrical.* Since this force is so enormous, all the plusses and all minuses will normally come together in as intimate a combination as they can. All things, even ourselves, are made of fine-grained, enormously strongly interacting plus and minus parts, all neatly balanced out. Once in a while, by accident, we may rub off a few minuses or a few plusses (usually it is easier to rub off minuses), and in those circumstances we find the force of electricity *unbalanced,* and we can then see the effects of these electrical attractions.

To give an idea of how much stronger electricity is than gravitation, consider two grains of sand, a millimeter across, thirty meters apart. If the force between them were not balanced, if everything attracted everything else instead of likes repelling, so that there were no cancellation, how much force would there be? There would be a force of *three million tons* between the two! You see, there is very, *very* little excess or deficit of the number of negative or positive charges necessary to produce appreciable electrical effects. This is, of course, the reason why you cannot see the difference between an electrically charged or uncharged thing—so few particles are involved that they hardly make a difference in the weight or size of an object.

With this picture the atoms were easier to understand. They were thought to have a "nucleus" at the center, which is positively electrically charged and very massive, and the nucleus is surrounded by a certain number of "electrons" which are very light and negatively charged. Now we go a little ahead in our story to remark that in the nucleus itself there were found two kinds of particles, protons and neutrons, almost of the same weight and very heavy. The protons are electrically charged and the neutrons are neutral. If we have an atom with six protons inside its nucleus, and this is surrounded by six electrons (the negative particles in the ordinary world of matter are all electrons, and these are very light compared with the protons and neutrons which make nuclei), this would be atom number six in the chemical table, and it is called carbon. Atom number eight is called oxygen, etc., because the chemical properties depend upon the electrons on the *outside,* and in fact only upon *how many* electrons there are. So the *chemical* properties of a substance depend only on a number, the number of electrons. (The whole list of elements of the chemists really could have been called 1, 2, 3, 4, 5, etc. Instead of saying "carbon," we could say "element six," meaning six electrons, but of course, when the elements were first discovered, it was not known that they could be numbered that way, and secondly, it would make everything look rather complicated. It is better to have names and symbols for these things, rather than to call everything by number.)

More was discovered about the electrical force. The natural interpretation of electrical interaction is that two objects simply attract each other: plus against minus. However, this was discovered to be an inadequate idea to represent it. A more adequate representation of the situation is to say that the existence of the positive charge, in some sense, distorts, or creates a "condition" in space, so that when we put the negative charge in, it feels a force. This potentiality for producing a force is called an *electric field.* When we put an electron in an electric field, we say it is "pulled." We then have two rules: (a) charges make a field, and (b) charges in fields have forces on them and move. The reason for this will become clear when we discuss the following phenomena: If we were to charge a body, say a comb, electrically, and then place a charged piece of paper at a distance and move the comb back and forth, the paper will respond by always pointing to the comb. If we shake it faster, it will be discovered that the paper is a little behind, *there is a delay* in the action. (At the first stage, when we move the comb rather slowly, we find a complication which is *magnetism.* Magnetic influences have to do with *charges in relative motion,* so magnetic forces and electric forces can really be attributed to one field, as two different aspects of exactly the same thing. A changing electric field cannot exist without magnetism.) If we move the charged paper farther out, the delay is greater. Then an interesting thing is observed. Although the forces between two charged objects should go inversely as the *square* of the distance, it is found, when we shake a charge, that the influence extends *very much farther out* than we would guess at first sight. That is, the effect falls off more slowly than the inverse square.

Here is an analogy: If we are in a pool of water and there is a floating cork very close by, we can move it "directly" by pushing the water with another cork. If you looked only at the two *corks,* all you would see would be that one moved immediately in response to the motion of the other—there is some kind of *"interaction"* between them. Of course, what we really do is to disturb the *water;* the *water* then disturbs the other cork. We could make up a "law" that if you pushed

the water a little bit, an object close by in the water would move. If it were farther away, of course, the second cork would scarcely move, for we move the water *locally.* On the other hand, if we jiggle the cork a new phenomenon is involved, in which the motion of the water moves the water there, etc., and *waves* travel away, so that by jiggling, there is an influence *wry much farther out,* an oscillatory influence, that cannot be understood from the direct interaction. Therefore the-idea of direct interaction must be replaced with the existence of the water, or-in-the electrical case, with what we call the *electromagnetic field.*

The electromagnetic field can carry waves; some of these waves are *light,* others are used in *radio broadcasts,* but the general name is *electromagnetic waves.* These oscillatory waves can have various *frequencies.* The only thing that is really different from one wave to another is the *frequency of oscillation.* If we shake a charge back and forth more and more rapidly, and look at the effects, we get a whole series of different kinds of effects, which are all unified by specifying but one number, the number of oscillations per second. The usual "pickup" that we get from electric currents in the circuits in the walls of a building have a frequency of about one hundred cycles per second. If we increase the frequency to 500 or 1000 kilocycles (1 kilocycle = 1000 cycles) per second, we are "on the air," for this is the frequency range which is used for radio broadcasts. (Of course it has nothing to do with the *air!* We can have radio broadcasts without any air.) If we again increase the frequency, we come into the range that is used for FM and TV. Going still further, we use certain short waves, for example for *radar.* Still higher, and we do not need an instrument to "see" the stuff, we can see it with the human eye. In the range of frequency from 5×10^{14} to 5×10^{15} cycles per second our eyes would see the oscillation of the charged comb, if we could shake it that fast, as red, blue, or violet light, depending on the frequency. Frequencies below this range are called infrared, and above it, ultraviolet. The fact that we can see in a particular frequency range makes that part of the electromagnetic spectrum no more impressive than the other parts from a physicist's standpoint, but from a human standpoint, of course, it *is* more interesting. If we go up even higher in frequency, we get x-rays. X-rays are nothing but very high-frequency light. If we go still higher, we get gamma rays. These two terms, x-rays and gamma rays, are used almost synonymously. Usually electromagnetic rays coming from nuclei are called gamma rays, while those of high energy from atoms are called x-rays, but at the same frequency they are indistinguishable physically, no matter what their source. If we go to still higher frequencies, say to 10^{24} cycles per second, we find that we can make those waves artificially, for example with the synchrotron here at Caltech. We can find electromagnetic waves with stupendously high frequencies—with even a thousand times more rapid oscillation—in the waves found in *cosmic rays.* These waves cannot be controlled by us.

Table 2-1

The Electromagnetic Spectrum

Frequency in oscillations/sec	Name	Rough behavior
10^2	Electrical disturbance	Field
$5 \times 10^5 - 10^6$	Radio broadcast	Waves
10^8	FM — TV	
10^{10}	Radar	
$5 \times 10^{14} - 10^{15}$	Light	
10^{18}	X-rays	Particle
10^{21}	γ-rays, nuclear	
10^{24}	γ-rays, "artificial"	
10^{27}	γ-rays, in cosmic rays	

2-3 Quantum physics

Having described the idea of the electromagnetic field, and that this field can carry waves, we soon learn that these waves actually behave in a strange way which seems very unwavelike. At higher frequencies they behave much more like *particles!* It is *quantum mechanics*, discovered just after 1920, which explains this strange behavior. In the years before 1920, the picture of space as a three-dimensional space, and of time as a separate thing, was changed by Einstein, first into a combination which we call space-time, and then still further into a *curved* space-time to represent gravitation. So the "stage" is changed into space-time, and gravitation is presumably a modification of space-time. Then it was also found that the rules for the motions of particles were incorrect. The mechanical rules of "inertia" and "forces" are *wrong*—Newton's laws are *wrong*—in the world of atoms. Instead, it was discovered that things on a small scale behave *nothing like* things on a large scale. That is what makes physics difficult—and very interesting. It is hard because the way things behave on a small scale is so "unnatural"; we have no direct experience with it. Here things behave like nothing we know of, so that it is impossible to describe this behavior in any other than analytic ways. It is difficult, and takes a lot of imagination.

Quantum mechanics has many aspects. In the first place, the idea that a particle has a definite location and a definite speed is no longer allowed; that is wrong. To give an example of how wrong classical physics is, there is a rule in quantum mechanics that says that one cannot know both where something is and how fast it is moving. The uncertainty of the momentum and the uncertainty of the position are complementary, and the product of the two is constant. We can write the law like this: $\Delta x \, \Delta p \geq \hbar/2$, but we shall explain it in more detail later. This rule is the explanation of a very mysterious paradox: if the atoms are made out of plus and minus charges, why don't the minus charges simply sit on top of the plus charges (they attract each other) and get so close as to completely cancel them out? *Why are atoms so big?* Why is the nucleus at the center with the electrons around it? It was first thought that this was because the nucleus was so big; but no, the nucleus is *very small.* An atom has a diameter of about 10^{-8} cm. The nucleus has a diameter of about 10^{-13} cm. If we had an atom and wished to see the nucleus, we would have to magnify it until the whole atom was the size of a large room, and then the nucleus would be a bare speck which you could just about make out with the eye, but very nearly *all the weight* of the atom is in that infinitesimal *nucleus.* What keeps the electrons from simply falling in? This principle: If they were in the nucleus, we would know their position precisely, and the uncertainty principle would then require that they have a very *large* (but uncertain) momentum, i.e., a very large *kinetic energy.* With this energy they would break away from the nucleus. They make a compromise: they leave themselves a little room for this uncertainty and then jiggle with a certain amount of minimum motion in accordance with this rule. (Remember that when a crystal is cooled to absolute zero, we said that the atoms do not stop moving, they still jiggle. Why? If they stopped moving, we would know where they were and that they had zero motion, and that is against the uncertainty principle. We cannot know where they are and how fast they are moving, so they must be continually wiggling in there!)

Another most interesting change in the ideas and philosophy of science brought about by quantum mechanics is this: it is not possible to predict *exactly* what will happen in any circumstance. For example, it is possible to arrange an atom which is ready to emit light, and we can measure when it has emitted light by picking up a photon particle, which we shall describe shortly. We cannot, however, predict *when* it is going to emit the light or, with several atoms, *which one* is going to. You may say that this is because there are some internal "wheels" which we have not looked at closely enough. No, there *are* no internal wheels; nature, as we understand it today, behaves in such a way that *it is fundamentally impossible* to make a precise prediction of *exactly what will happen* in a given experiment. This is a horrible thing; in fact, philosophers have said before that one of the fundamental requisites of science is that whenever you set **up** the same

conditions, the same thing must happen. This is simply *not true*, it is *not* a fundamental condition of science. The fact is that the same thing does not happen, that we can find only an average, statistically, **as to** what happens. Nevertheless, science has not completely collapsed. Philosophers, incidentally, say a great deal about what is *absolutely necessary* for science, and it is always, so far as one can see, rather naive, and probably wrong. For example, some philosopher or other said it is fundamental to the scientific effort that if an experiment is performed in, say, Stockholm, and then the same experiment is done in, say, Quito, the *same results* must occur. That is quite false. It is not necessary that *science* do that; it may be *a fact of experience*, but it is not necessary. For example, if one of the experiments is to look out at the sky and see the aurora borealis in Stockholm, you do not see it in Quito; that is a different phenomenon. "But," you say, "that is something that has to do with the outside; can you close yourself up in a box in Stockholm and pull down the shade and get any difference?" Surely. If we take a pendulum on a universal joint, and pull it out and let go, then the pendulum will swing almost in a plane, but not quite. Slowly the plane keeps changing in Stockholm, but not in Quito. The blinds are down, too. The fact that this happened does not bring on the destruction of science. What *is* the fundamental hypothesis of science, the fundamental philosophy? We stated it in the first chapter: *the sole test of the validity of any idea is experiment.* If it turns out that most experiments work out the same in Quito as they do in Stockholm, then those "most experiments" will be used to formulate some general law, and those experiments which do not come out the same we will say were a result of the environment near Stockholm. We will invent some way to summarize the results of the experiment, and we do not have to be told ahead of time what this way will look like. If we are told that the same experiment will always produce the same result, that is all very well, but if when we try it, it does *not*, then it does *not.* We just have to take what we see, and then formulate all the rest of our ideas in terms of our actual experience.

Returning again to quantum mechanics and fundamental physics, we cannot go into details of the quantum-mechanical principles at this time, of course, because these are rather difficult to understand. We shall assume that they are there, and go on to describe what some of the consequences are. One of the consequences is that things which we used to consider as waves also behave like particles, and particles behave like waves; in fact everything behaves the same way. There is no distinction between a wave and a particle. So quantum mechanics *unifies* the idea of the field and its waves, and the particles, all into one. Now it is true that when the frequency is low, the field aspect of the phenomenon is more evident, or more useful as an approximate description in terms of everyday experiences. But as the frequency increases, the particle aspects of the phenomenon become more evident with the equipment with which we usually make the measurements. In fact, although we mentioned many frequencies, no phenomenon directly involving a frequency has yet been detected above approximately 10^{12} cycles per second. We only *deduce* the higher frequencies from the energy of the particles, by a rule which assumes that the particle-wave idea of quantum mechanics is valid.

Thus we have a new view of electromagnetic interaction. We have a new kind of *particle* to add to the electron, the proton, and the neutron. That new particle is called a *photon.* The new view of the interaction of electrons and protons that is electromagnetic theory, but with everything quantum-mechanically correct, is called *quantum electrodynamics.* This fundamental theory of the interaction of light and matter, or electric field and charges, is our greatest success so far in physics. In this one theory we have the basic rules for all ordinary phenomena except for gravitation and nuclear processes. For example, out of quantum electrodynamics come all known electrical, mechanical, and chemical laws: the laws for the collision of billiard balls, the motions of wires in magnetic fields, the specific heat of carbon monoxide, the color of neon signs, the density of salt, and the reactions of hydrogen and oxygen to make water are all consequences of this one law. All these details can be worked out if the situation is simple enough for us to make an approximation, which is almost never, but often we can understand more

or less what is happening. At the present time no exceptions are found to the quantum-electrodynamic laws outside the nucleus, and there we do not know whether there is an exception because we simply do not know what is going on in the nucleus.

In principle, then, quantum electrodynamics is the theory of all chemistry, and of life, if life is ultimately reduced to chemistry and therefore just to physics because chemistry is already reduced (the part of physics which is involved in chemistry being already known). Furthermore, the same quantum electrodynamics, this great thing, predicts a lot of new things. In the first place, it tells the properties of very high-energy photons, gamma rays, etc. It predicted another very remarkable thing: besides the electron, there should be another particle of the same mass, but of opposite charge, called a *positron*, and these two, coming together, could annihilate each other with the emission of light or gamma rays. (After all, light and gamma rays are all the same, they are just different points on a frequency scale.) The generalization of this, that for each particle there is an antiparticle, turns out to be true. In the case of electrons, the antiparticle has another name—it is called a positron, but for most other particles, it is called anti-so-and-so, like antiproton or antineutron. In quantum electrodynamics, *two numbers* are put in and most of the other numbers in the world are supposed to come out. The two numbers that are put in are called the mass of the electron and the charge of the electron. Actually, that is not quite true, for we have a whole set of numbers for chemistry which tells how heavy the nuclei are. That leads us to the next part.

2-4 Nuclei and particles

What are the nuclei made of, and how are they held together? It is found that the nuclei are held together by enormous forces. When these are released, the energy released is tremendous compared with chemical energy, in the same ratio as the atomic bomb explosion is to a TNT explosion, because, of course, the atomic bomb has to do with changes inside the nucleus, while the explosion of TNT has to do with the changes of the electrons on the outside of the atoms. The question is, what are the forces which hold the protons and neutrons together in the nucleus? Just as the electrical interaction can be connected to a particle, a photon, Yukawa suggested that the forces between neutrons and protons also have a field of some kind, and that when this field jiggles it behaves like a particle. Thus there could be some other particles in the world besides protons and neutrons, and he was able to deduce the properties of these particles from the already known characteristics of nuclear forces. For example, he predicted they should have a mass of two or three hundred times that of an electron; and lo and behold, in cosmic rays there was discovered a particle of the right mass! But it later turned out to be the wrong particle. It was called a m-meson, or muon.

However, a little while later, in 1947 or 1948, another particle was found, the p-meson, or pion, which satisfied Yukawa's criterion. Besides the proton and the neutron, then, in order to get nuclear forces we must add the pion. Now, you say, "Oh great!, with this theory we make quantum nucleodynamics using the pions just like Yukawa wanted to do, and see if it works, and everything will be explained." Bad luck. It turns out that the calculations that are involved in this theory are so difficult that no one has ever been able to figure out what the consequences of the theory are, or to check it against experiment, and this has been going on now for almost twenty years!

So we are stuck with a theory, and we do not know whether it is right or wrong, but we do know that it is a *little* wrong, or at least incomplete. While we have been dawdling around theoretically, trying to calculate the consequences of this theory, the experimentalists have been discovering some things. For example, they had already discovered this m-meson or muon, and we do not yet know where it fits. Also, in cosmic rays, a large number of other "extra" particles were found. It turns out that today we have approximately thirty particles, and it is very difficult to understand the relationships of all these particles, and what nature

, wants them for, or what the connections are from one to another. We do not today understand these various particles as different aspects of the same thing, and the fact that we have so many unconnected particles is a representation of the fact that we have so much unconnected information without a good theory. After the great successes of quantum electrodynamics, there is a certain amount of knowledge of nuclear physics which is rough knowledge, sort of half experience and half theory, assuming a type of force between protons and neutrons and seeing what will happen, but not really understanding where the force comes from. Aside from that, we have made very little progress. We have collected an enormous number of chemical elements. In the chemical case, there suddenly appeared a relationship among these elements which was unexpected, and which is embodied in the periodic table of Mendeleev. For example, sodium and potassium are about the same in their chemical properties and are found in the same column in the Mendeleev chart. We have been seeking a Mendeleev-type chart for the new particles. One such chart of the new particles was made independently by Gell-Mann in the U.S.A. and Nishijima in Japan. The basis of their classification is a new number, like the electric charge, which can be assigned to each particle, called its "strangeness," *S.* This number is conserved, like the electric charge, in reactions which take place by nuclear forces.

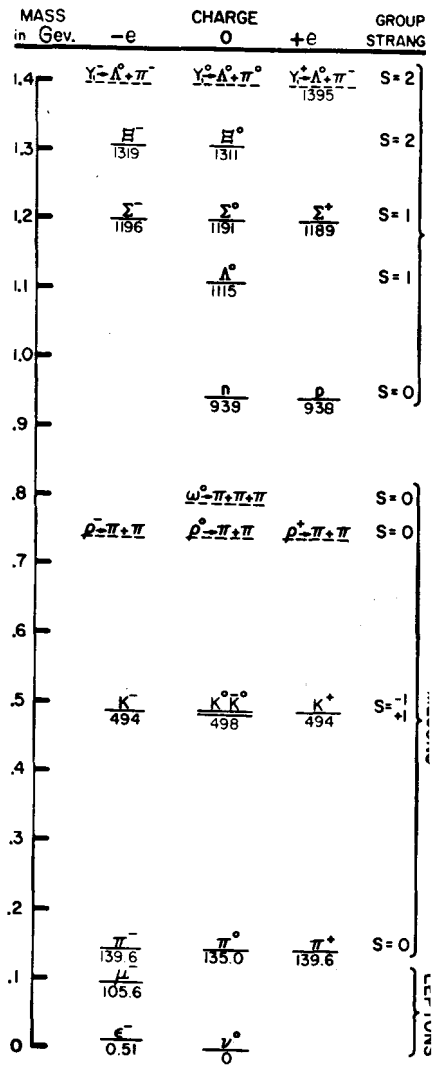
In Table 2-2 are listed all the particles. We cannot discuss them much at this stage, but the table will at least show you how much we do not know. Underneath each particle its mass is given in a certain unit, called the Mev. One Mev is equal to 1.782 X 10^{-27} gram. The reason this unit was chosen is historical, and we shall not go into it now. More massive particles are put higher up on the chart; we see that a neutron and a proton have almost the same mass. In vertical columns we have put the particles with the same electrical charge, all neutral objects in one column, all positively charged ones to the right of this one, and all negatively charged objects to the left.

Particles are shown with a solid line and "resonances" with a dashed one. Several particles have been omitted from the table. These include the important zero-mass, zero-charge particles, the photon and the graviton, which do not fall into the baryon-meson-lepton classification scheme, and also some of the newer resonances *(K*,* <p, *ri).* The antiparticles of the mesons are listed in the table, but the antiparticles of the leptons and baryons would have to be listed in another table which would look exactly like this one reflected on the zero-charge column. Although all of the particles except the electron, neutrino, photon, graviton, and proton are unstable, decay products have been shown only for the resonances. Strangeness assignments are not applicable for leptons, since they do not interact strongly with nuclei.

All particles which are together with the neutrons and protons are called *baryons,* and the following ones exist: There is a "lambda," with a mass of 1154 Mev, and three others, called sigmas, minus, neutral, and plus, with several masses almost the same. There are groups or multiplets with almost the same mass, within one or two percent. Each particle in a multiplet has the same strangeness. The first multiplet is the proton-neutron doublet, and then there is a singlet (the lambda) then the sigma triplet, and finally the xi doublet. Very recently, in 1961, even a few more particles were found. Or *are* they particles? They live so short a time, they disintegrate almost instantaneously, as soon as they are formed, that we do not know whether they should be considered as new particles, or some kind of "resonance" interaction of a certain definite energy between the A and T products into which they disintegrate.

In addition to the baryons the other particles which are involved in the nuclear interaction are called *mesons.* There are first the pions, which come in three varieties, positive, negative, and neutral; they form another multiplet. We have also found some new things called A'-mesons, and they occur as a doublet, K^+ and $K°$. Also, every particle has its antiparticle, unless a particle *is its own* antiparticle. For example, the π^- and the π^+ are antiparticles, but the $\pi°$ is its own antiparticle. The K~ and K^+ are antiparticles, and the $K°$ and $K°$. In addition, in 1961 we also found some more mesons or *maybe* mesons which disintegrate almost immediately.

Table 2-2. Elementary Particles



A thing called *w* which goes into three pions has a mass 780 on this scale, and somewhat less certain is an object which disintegrates into two pions. These particles, called mesons and baryons, and the antiparticles of the mesons are on the same chart, but the antiparticles of the baryons must be put on another chart, "reflected" through the charge-zero column.

Just as Mendeleev's chart was very good, except for the fact that there were a number of rare earth elements which were hanging out loose from it, so we have a number of things hanging out loose from this chart—particles which do not interact strongly in nuclei, have nothing to do with a nuclear interaction, and do not have a strong interaction (I mean the powerful kind of interaction of nuclear energy). These are called leptons, and they are the following: there is the electron, which has a very small mass on this scale, only 0.510 Mev. Then there is that other, the μ-meson, the muon, which has a mass much higher, 206 times as heavy as an electron. So far as we can tell, by all experiments so far, the difference between the electron and the muon is nothing but the mass. Everything works exactly the same for the muon as for the electron, except that one is heavier than the other. Why is there another one heavier; what is the use for it? We do not know. In addition, there is a lepton which is neutral, called a neutrino, and this particle has zero mass. In fact, it is now known that there are *two* different kinds of neutrinos, one related to electrons and the other related to muons.

Finally, we have two other particles which do not interact strongly with the nuclear ones: one is a photon, and perhaps, if the field of gravity also has a quantum-mechanical analog (a quantum theory of gravitation has not yet been worked out), then there will be a particle, a graviton, which will have zero mass.

What is this "zero mass"? The masses given here are the masses of the particles *at rest.* The fact that a particle has zero mass means, in a way, that it cannot *be* at *rest.* A photon is never at rest, it is always moving at 186,000 miles a second. We will understand more what mass means when we understand the theory of relativity, which will come in due time.

Thus we are confronted with a large number of particles, which together seem to be the fundamental constituents of matter. Fortunately, these particles are not *all* different in their *interactions* with one another. In fact, there seem to be just *four kinds* of interaction between particles which, in the order of decreasing strength, are the nuclear force, electrical interactions, the beta-decay interaction, and gravity. The photon is coupled to all charged particles and the strength of the interaction is measured by some number, which is 1/137. The detailed law of this coupling is known, that is quantum electrodynamics. • Gravity is coupled to all *energy,* but its coupling is extremely weak, much weaker than that of electricity. This law is also known. Then there are the so-called weak decays—beta decay, which causes the neutron to disintegrate into proton, electron, and neutrino, relatively slowly. This law is only partly known. The so-called strong interaction, the meson-baryon interaction, has a strength of 1 in this scale, and the law is completely unknown, although there are a number of known rules, such as that the number of baryons does not change in any reaction.

Table 2-3. Elementary Interactions

Coupling	Strength*	Law
Photon to charged particles	$\sim 10^{-2}$	Law known
Gravity to all energy	$\sim 10^{-40}$	Law known
Weak decays	$\sim 10^{-5}$	Law partly known
Mesons to baryons	~ 1	Law unknown (some rules known)

* The "strength" is a dimensionless measure of the coupling constant involved in each interaction (~ means "approximately").

This then, is the horrible condition of our physics today. To summarize it, I would say this: outside the nucleus, we seem to know all; inside it, quantum mechanics is valid—the principles of quantum mechanics have not been found to fail. The stage on which we put all of our knowledge, we would say, is relativistic space-time; perhaps gravity is involved in space-time. We do not know how the universe got started, and we have never made experiments which check our ideas of space and time accurately, below some tiny distance, so we only *know* that our ideas work above that distance. We should also add that the rules of the game are the quantum-mechanical principles, and those principles apply, so far as we can tell, to the new particles as well as to the old. The origin of the forces in nuclei leads us to new particles, but unfortunately they appear in great profusion and we lack a complete understanding of their interrelationship, although we already know that there are some very surprising relationships among them. We seem gradually to be groping toward an understanding of the world of sub-atomic particles, but we really do not know how far we have yet to go in this task.

3

The Relation of Physics to Other Sciences

3-1 Introduction

3-2 Chemistry

3-3 Biology

3-4 Astronomy

3-5 Geology

3-6 Psychology

3-7 How did it get that way?

3-1 Introduction

Physics is the most fundamental and all-inclusive of the sciences, and has had a profound effect on all scientific development. In fact, physics is the present-day equivalent of what used to be called *natural philosophy*, from which most of our modern sciences arose. Students of many fields find themselves studying physics because of the basic role it plays in all phenomena. In this chapter we shall try to explain what the fundamental problems in the other sciences are, but of course it is impossible in so small a space really to deal with the complex, subtle, beautiful matters in these other fields. Lack of space also prevents our discussing the relation of physics to engineering, industry, society, and war, or even the most remarkable relationship between mathematics and physics. (Mathematics is not a science from our point of view, in the sense that it is not a *natural* science. The test of its validity is not experiment.) We must, incidentally, make it clear from the beginning that if a thing is not a science, it is not necessarily bad. For example, love is not a science. So, if something is said not to be a science, it does not mean that there is something wrong with it; it just means that it is not a science.

3-2 Chemistry

The science which is perhaps the most deeply affected by physics is chemistry. Historically, the early days of chemistry dealt almost entirely with what we now call inorganic chemistry, the chemistry of substances which are not associated with living things. Considerable analysis was required to discover the existence of the many elements and their relationships—how they make the various relatively simple compounds found in rocks, earth, etc. This early chemistry was very important for physics. The interaction between the two sciences was very great because the theory of atoms was substantiated to a large extent by experiments in chemistry. The theory of chemistry, i.e., of the reactions themselves, was summarized to a large extent in the periodic chart of Mendeleev, which brings out many strange relationships among the various elements, and it was the collection of rules as to which substance is combined with which, and how, that constituted inorganic chemistry. All these rules were ultimately explained in principle by quantum mechanics, so that theoretical chemistry is in fact physics. On the other hand, it must be emphasized that this explanation is *in principle.* We have already discussed the difference between knowing the rules of the game of chess, and being able to play. So it is that we may know the rules, but we cannot play very well. It turns out to be very difficult to predict precisely what will happen in a given chemical reaction; nevertheless, the deepest part of theoretical chemistry must end up in quantum mechanics.

There is also a branch of physics and chemistry which was developed by both sciences together, and which is extremely important. This is the method of statistics applied in a situation in which there are mechanical laws, which is aptly called *statistical mechanics.* In any chemical situation a large number of atoms are involved, and we have seen that the atoms are all jiggling around in a very random and complicated way. If we could analyze each collision, and be able to follow in detail the motion of each molecule, we might hope to figure out what would happen, but the many numbers needed to keep track of all these molecules exceeds so enormously the capacity of any computer, and certainly the capacity of

the mind, that it was important to develop a method for dealing with such complicated situations. Statistical mechanics, then, is the science of the phenomena of heat, or thermodynamics. Inorganic chemistry is, as a science, now reduced essentially to what are called physical chemistry and quantum chemistry; physical chemistry to study the rates at which reactions occur and what is happening in detail (How do the molecules hit? Which pieces fly off first?, etc.), and quantum chemistry to help us understand what happens in terms of the physical laws.

The other branch of chemistry is *organic chemistry,* the chemistry of the substances which are associated with living things. For a time it was believed that the substances which are associated with living things were so marvelous that they could not be made by hand, from inorganic materials. This is not at all true—they are just the same as the substances made in inorganic chemistry, but more complicated arrangements of atoms are involved. Organic chemistry obviously has a very close relationship to the biology which supplies its substances, and to industry, and furthermore, much physical chemistry and quantum mechanics can be applied to organic as well as to inorganic compounds. However, the main problems of organic chemistry are not in these aspects, but rather in the analysis and synthesis of the substances which are formed in biological systems, in living things. This leads imperceptibly, in steps, toward biochemistry, and then into biology itself, or molecular biology.

3-3 Biology

Thus we come to the science of *biology,* which is the study of living things. In the early days of biology, the biologists had to deal with the purely descriptive problem of finding out *what* living things there were, and so they just had to count such things as the hairs of the limbs of fleas. After these matters were worked out with a great deal of interest, the biologists went into the *machinery* inside the living bodies, first from a gross standpoint, naturally, because it takes some effort to get into the finer details.

There was an interesting early relationship between physics and biology in which biology helped physics in the discovery of the *conservation of energy,* which was first demonstrated by Mayer in connection with the amount of heat taken in and given out by a living creature.

If we look at the processes of biology of living animals more closely, we see *many* physical phenomena: the circulation of blood, pumps, pressure, etc. There are nerves: we know what is happening when we step on a sharp stone, and that somehow or other the information goes from the leg up. It is interesting how that happens. In their study of nerves, the biologists have come to the conclusion that nerves are very fine tubes with a complex wall which is very thin; through this wall the cell pumps ions, so that there are positive ions on the outside and negative ions on the inside, like a capacitor. Now this membrane has an interesting property; if it "discharges" in one place, i.e., if some of the ions were able to move through one place, so that the electric voltage is reduced there, that electrical influence makes itself felt on the ions in the neighborhood, and it affects the membrane in such a way that it lets the ions through at neighboring points also. This in turn affects it farther along, etc., and so there is a wave of "penetrability" of the membrane which runs down the fiber when it is "excited" at one end by stepping on the sharp stone. This wave is somewhat analogous to a long sequence of vertical dominoes; if the end one is pushed over, that one pushes the next, etc. Of course this will transmit only one message unless the dominoes are set up again; and similarly in the nerve cell, there are processes which pump the ions slowly out again, to get the nerve ready for the next impulse. So it is that we know what we are doing (or at least where we are). Of course the electrical effects associated with this nerve impulse can be picked up with electrical instruments, and because there *are* electrical effects, obviously the physics of electrical effects has had a great deal of influence on understanding the phenomenon.

The opposite effect is that, from somewhere in the brain, a message is sent out along a nerve. What happens at the end of the nerve? There the nerve branches

out into fine little things, connected to a structure near a muscle, called an end-plate. For reasons which are not exactly understood, when the impulse reaches the end of the nerve, little packets of a chemical called acetylcholine are shot off (five or ten molecules at a time) and they affect the muscle fiber and make it contract—how simple! What makes a muscle contract? A muscle is a very large number of fibers close together, containing two different substances, myosin and actomyosin, but the machinery by which the chemical reaction induced by acetylcholine can modify the dimensions of the molecule is not yet known. Thus the fundamental processes in the muscle that make mechanical motions are not known.

Biology is such an enormously wide field that there are hosts of other problems that we cannot mention at all—problems on how vision works (what the light does in the eye), how hearing works, etc. (The way in which *thinking* works we shall discuss later under psychology.) Now, these things concerning biology which we have just discussed are, from a biological standpoint, really not fundamental, at the bottom of life, in the sense that even if we understood them we still would not understand life itself. To illustrate: the men who study nerves feel their work is very important, because after all you cannot have animals without nerves. But you *can* have *life* without nerves. Plants have neither nerves nor muscles, but they are working, they are alive, just the same. So for the fundamental problems of biology we must look deeper; when we do, we discover that all living things have a great many characteristics in common. The most common feature is that they are made of *cells,* within each of which is complex machinery for doing things chemically. In plant cells, for example, there is machinery for picking up light and generating sucrose, which is consumed in the dark to keep the plant alive. When the plant is eaten the sucrose itself generates in the animal a series of chemical reactions very closely related to photosynthesis (and its opposite effect in the dark) in plants.

In the cells of living systems there are many elaborate chemical reactions, in which one compound is changed into another and another. To give some impression of the enormous efforts that have gone into the study of biochemistry, the chart in Fig. 3-1 summarizes our knowledge to date on just one small part of the many series of reactions which occur in cells, perhaps a percent or so of it.

Here we see a whole series of molecules which change from one to another in a sequence or cycle of rather small steps. It is called the Krebs cycle, the respiratory cycle. Each of the chemicals and each of the steps is fairly simple, in terms of what change is made in the molecule, but—and this is a centrally important discovery in biochemistry—these changes are *relatively difficult to accomplish in a laboratory.* If we have one substance and another very similar substance, the one does not just turn into the other, because the two forms are usually separated by

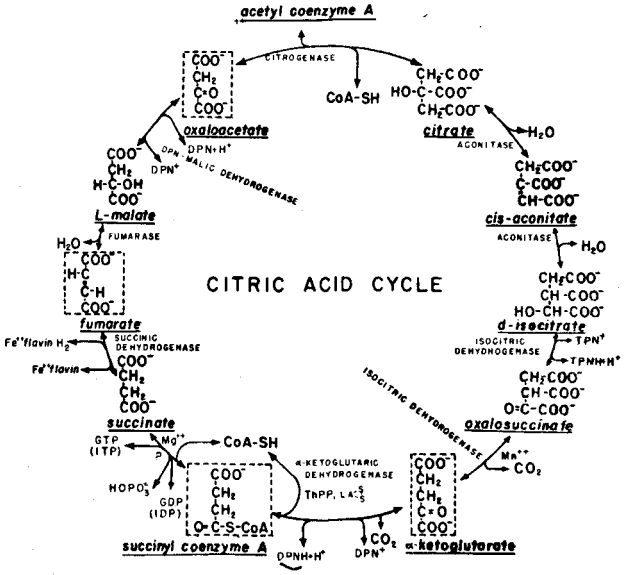


Fig. 3–1. The Krebs cycle.

an energy barrier or "hill." Consider this analogy: If we wanted to take an object from one place to another, at the same level but on the other side of a hill, we could push it over the top, but to do so requires the addition of some energy. Thus most chemical reactions do not occur, because there is what is called an *activation energy* in the way. In order to add an extra atom to our chemical requires that we get it *close* enough that some rearrangement can occur; then it will stick. But if we cannot give it enough energy to get it close enough, it will not go to completion, it will just go part way up the "hill" and back down again. However, if we could literally take the molecules in our hands and push and pull the atoms around in such a way as to open a hole to let the new atom in, and then let it snap back, we would have found another way, *around* the hill, which would not require extra energy, and the reaction would go easily. Now there actually *are*, in the cells, *very* large molecules, much larger than the ones whose changes we have been describing, which in some complicated way hold the smaller molecules just right, so that the reaction can occur easily. These very large and complicated things are called *enzymes.* (They were first called ferments, because they were originally discovered in the fermentation of sugar. In fact, some of the first reactions in the cycle were discovered there.) In the presence of an enzyme the reaction will go.

An enzyme is made of another substance called *protein.* Enzymes are very big and complicated, and each one is different, each being built to control a certain special reaction. The names of the enzymes are written in Fig. 3-1 at each reaction. (Sometimes the same enzyme may control two reactions.) We emphasize that the enzymes themselves are not involved in the reaction directly. They do not change; they merely let an atom go from one place to another. Having done so, the enzyme is ready to do it to the next molecule, like a machine in a factory. Of course, there must be a supply of certain atoms and a way of disposing of other atoms. Take hydrogen, for example: there are enzymes which have special units on them which carry the hydrogen for all chemical reactions. For example, there are three or four hydrogen-reducing enzymes which are used all over our cycle in different places. It is interesting that the machinery which liberates some hydrogen at one place will take that hydrogen and use it somewhere else.

The most important feature of the cycle of Fig. 3-1 is the transformation from GDP to GTP (guanadine-di-phosphate to guanadine-tri-phosphate) because the one substance has much more energy in it than the other. Just as there is a "box" in certain enzymes for carrying hydrogen atoms around, there are special *energy-carrying* "boxes" which involve the triphosphate group. So, GTP has more energy than GDP and if the cycle is going one way, we are producing molecules which have extra energy and which can go drive some other cycle which *requires* energy, for example the contraction of muscle. The muscle will not contract unless there is GTP. We can take muscle fiber, put it in water, and add GTP, and the fibers contract, changing GTP to GDP if the right enzymes are present. So the real system is in the GDP-GTP transformation; in the dark the GTP which has been stored up during the day is used to run the whole cycle around the other way. An enzyme you see, does not care in which direction the reaction goes, for if it did it would violate one of the laws of physics.

Physics is of great importance in biology and other sciences for still another reason, that has to do with *experimental techniques.* In fact, if it were not for the great development of experimental physics, these biochemistry charts would not be known today. The reason is that the most useful tool of all for analyzing this fantastically complex system is to *label* the atoms which are used in the reactions. Thus, if we could introduce into the cycle some carbon dioxide which has a "green mark" on it, and then measure after three seconds where the green mark is, and again measure after ten seconds, etc., we could trace out the course of the reactions. What are the "green marks"? They are different *isotopes.* We recall that the chemical properties of atoms are determined by the number of *electrons,* not by the mass of the nucleus. But there can be, for example in carbon, six neutrons or seven neutrons, together with the six protons which all carbon nuclei have. Chemically, the two atoms C^{12} and C^{13} are the same, but they differ in weight and they have different nuclear properties, and so they are distinguishable.

By using these isotopes of different weights, or even radioactive isotopes like C_{14}, which provide a more sensitive means for tracing very small quantities, it is possible to trace the reactions.

Now, we return to the description of enzymes and proteins. All proteins are not enzymes, but all enzymes are proteins. There are many proteins, such as the proteins in muscle, the structural proteins which are, for example, in cartilage and hair, skin, etc., that are not themselves enzymes. However, proteins are a very characteristic substance of life: first of all they make up all the enzymes, and second, they make up much of the rest of living material. Proteins have a very interesting and simple structure. They are a series, or chain, of different *ammo acids*. There are twenty different amino acids, and they all can combine with each other to form chains in which the backbone is CO-NH, etc. Proteins are nothing but chains of various ones of these twenty amino acids. Each of the amino acids probably serves some special purpose. Some, for example, have a sulphur atom at a certain place; when two sulphur atoms are in the same protein, they form a bond, that is, they tie the chain together at two points and form a loop. Another has extra oxygen atoms which make it an acidic substance, another has a basic characteristic. Some of them have big groups hanging out to one side, so that they take up a lot of space. One of the amino acids, called prolene, is not really an amino acid, but imino acid. There is a slight difference, with the result that when prolene is in the chain, there is a kink in the chain. If we wished to manufacture a particular protein, we would give these instructions: put one of those sulphur hooks here; next, add something to take up space; then attach something to put a kink in the chain. In this way, we will get a complicated-looking chain, hooked together and having some complex structure; this is presumably just the manner in which all the various enzymes are made. One of the great triumphs in recent times (since 1960), was at last to discover the exact spatial atomic arrangement of certain proteins, which involve some fifty-six or sixty amino acids in a row. Over a thousand atoms (more nearly two thousand, if we count the hydrogen atoms) have been located in a complex pattern in two proteins. The first was hemoglobin. One of the sad aspects of this discovery is that we cannot see anything from the pattern; we do not understand why it works the way it does. Of course, that is the next problem to be attacked.

Another problem is how do the enzymes know what to be? A red-eyed fly makes a red-eyed fly baby, and so the information for the whole pattern of enzymes to make red pigment must be passed from one fly to the next. This is done by a substance in the nucleus of the cell, not a protein, called DNA (short for desoxyribose nucleic acid). This is the key substance which is passed from one cell to another (for instance, sperm cells consist mostly of DNA) and carries the information as to how to make the enzymes. DNA is the "blueprint." What does the blueprint look like and how does it work? First, the blueprint must be able to reproduce itself. Secondly, it must be able to instruct the protein. Concerning the reproduction, we might think that this proceeds like cell reproduction. Cells simply grow bigger and then divide in half. Must it be thus with DNA molecules, then, that they too grow bigger and divide in half? Every *atom* certainly does not grow bigger and divide in half! No, it is impossible to reproduce a molecule except by some more clever way.

The structure of the substance DNA was studied for a long time, first chemically to find the composition, and then with x-rays to find the pattern in space. The result was the following remarkable discovery: The DNA molecule is a pair of chains, twisted upon each other. The backbone of each of these chains, which are analogous to the chains of proteins but chemically quite different, is a series of sugar and phosphate groups, as shown in Fig. 3-2. Now we see how the chain can contain instructions, for if we could split this chain down the middle, we would have a series *BAADC* . . . and every living thing could have a different series. Thus perhaps, in some way, the specific *instructions* for the manufacture of proteins are contained in the specific *series* of the DNA.

Attached to each sugar along the line, and linking the two chains together, are certain pairs of cross-links. However, they are not all of the same kind; there are

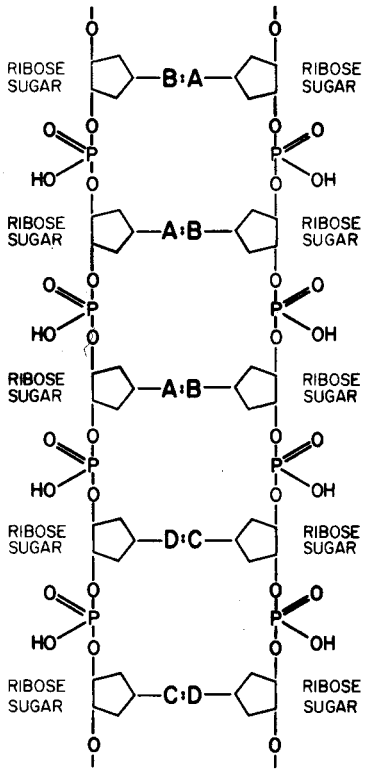


Fig. 3–2. Schematic diagram of DNA.

four kinds, called adenine, thymine, cytosine, and guanine, but let us call them *A, B, C,* and *D.* The interesting thing is that only certain pairs can sit opposite each other, for example *A* with *B* and C with *D.* These pairs are put on the two chains in such a way that they "fit together," and have a strong energy of interaction. However, C will not fit with *A,* and *B* will not fit with C; they will only fit in pairs, *A* against *B* and C against *D.* Therefore if one is C, the other must be *D,* etc. Whatever the letters may be in one chain, each one must have its specific complementary letter on the other chain.

What then about reproduction? Suppose we split this chain in two. How can we make another one just like it? If, in the substances of the cells, there is a manufacturing department which brings up phosphate, sugar, and *A, B, C, D* units not connected in a chain, the only ones which will attach to our split chain will be the correct ones, the complements of *BAADC . . .,* namely, *ABBCD ...* Thus what happens is that the chain splits down the middle during cell division, one half ultimately to go with one cell, the other half to end up in the other cell; when separated, a new complementary chain is made by each half-chain.

Next comes the question, precisely how does the order of the *A, B, C, D* units determine the arrangement of the amino acids in the protein? This is the central unsolved problem in biology today. The first clues, or pieces of information, however, are these: There are in the cell tiny particles called microsomes, and it is now known that that is the place where proteins are made. But the microsomes are not in the nucleus, where the DNA and its instructions are. Something seems to be the matter. However, it is also known that little molecule pieces come off the DNA—not as long as the big DNA molecule that carries all the information itself, but like a small section of it. This is called RNA, but that is not essential. It is a kind of copy of the DNA, a short copy. The RNA, which somehow carries a message as to what kind of protein to make goes over to the microsome; that is known. When it gets there, protein is synthesized at the microsome. That is also known. However, the details of how the amino acids come in and are arranged in accordance with a code that is on the RNA are, as yet, still unknown. We do not know how to read it. If we knew, for example, the "lineup" *A, B, C, C, A,* we could not tell you what protein is to be made.

Certainly no subject or field is making more progress on so many fronts at the present moment, than biology, and if we were to name the most powerful assumption of all, which leads one on and on in an attempt to understand life, it is that *all things are made of atoms,* and that everything that living things do can be understood in terms of the jigglings and wigglings of atoms.

3-4 Astronomy

In this rapid-fire explanation of the whole world, we must now turn to astronomy. Astronomy is older than physics. In fact, it got physics started by showing the beautiful simplicity of the motion of the stars and planets, the understanding of which was the *beginning* of physics. But the most remarkable discovery in all of astronomy is that *the stars are made of atoms of the same kind as those on the earth** How was this done? Atoms liberate light which has definite fre-

* How I'm rushing through this! How much each sentence in this brief story contains. "The stars are made of the same atoms as the earth." I usually pick one small topic like this to give a lecture on. Poets say science takes away from the beauty of the stars—mere globs of gas atoms. Nothing is "mere." I too can see the stars on a desert night, and feel them. But do I see less or more? The vastness of the heavens stretches my imagination—stuck on this carousel my little eye can catch one-million-year-old light. A vast pattern—of which I am a part—perhaps my stuff was belched from some forgotten star, as one is belching there. Or see them with the greater eye of Palomar, rushing all apart from some common starting point when they were perhaps all together. What is the pattern, or the meaning, or the *why*? It does not do harm to the mystery to know a little about it. For far more marvelous is the truth than any artists of the past imagined! Why do the poets of the present not speak of it? What men are poets who can speak of Jupiter if he were like a man, but if he is an immense spinning sphere of methane and ammonia must be silent?

quencies, something like the timbre of a musical instrument, which has definite pitches or frequencies of sound. When we are listening to several different tones we can tell them apart, but when we look with our eyes at a mixture of colors we cannot tell the parts from which it was made, because the eye is nowhere near as discerning as the ear in this connection. However, with a spectroscope we *can* analyze the frequencies of the light waves and in this way we can see the very tunes of the atoms that are in the different stars. As a matter of fact, two of the chemical elements were discovered on a star before they were discovered on the earth. Helium was discovered on the sun, whence its name, and technetium was discovered in certain cool stars. This, of course, permits us to make headway in understanding the stars, because they are made of the same kinds of atoms which are on the earth. Now we know a great deal about the atoms, especially concerning their behavior under conditions of high temperature but not very great density, so that we can analyze by statistical mechanics the behavior of the stellar substance. Even though we cannot reproduce the conditions on the earth, using the basic physical laws we often can tell precisely, or very closely, what will happen. So it is that physics aids astronomy. Strange as it may seem, we understand the distribution of matter in the interior of the sun far better than we understand the interior of the earth. What goes on *inside* a star is better understood than one might guess from the difficulty of having to look at a little dot of light through a telescope, because we can *calculate* what the atoms in the stars should do in most circumstances.

One of the most impressive discoveries was the origin of the energy of the stars, that makes them continue to burn. One of the men who discovered this was out with his girl friend the night after he realized that *nuclear reactions* must be going on in the stars in order to make them shine. She said "Look at how pretty the stars shine!" He said "Yes, and right now I am the only man in the world who knows *why* they shine." She merely laughed at him. She was not impressed with being out with the only man who, at that moment, knew why stars shine. Well, it is sad to be alone, but that is the way it is in this world.

It is the nuclear "burning" of hydrogen which supplies the energy of the sun; the hydrogen is converted into helium. Furthermore, ultimately, the manufacture of various chemical elements proceeds in the centers of the stars, from hydrogen. The stuff of which *we* are made, was "cooked" once, in a star, and spit out. How do we know? Because there is a clue. The proportion of the different isotopes—how much C^{12}, how much C^{13}, etc., is something which is never changed by *chemical* reactions, because the chemical reactions are so much the same for the two. The proportions are purely the result of *nuclear* reactions. By looking at the proportions of the isotopes in the cold, dead ember which we are, we can discover what the *furnace* was like in which the stuff of which we are made was formed. That furnace was like the stars, and so it is very likely that our elements were "made" in the stars and spit out in the explosions which we call novae and supernovae. Astronomy is so close to physics that we shall study many astronomical things as we go along.

3-5 Geology

We turn now to what are called *earth sciences,* or *geology.* First, meteorology and the weather. Of course the *instruments* of meteorology are physical instruments, and the development of experimental physics made these instruments possible, as was explained before. However, the theory of meteorology has never been satisfactorily worked out by the physicist. "Well," you say, "there is nothing but air, and we know the equations of the motions of air." Yes we do. "So if we know the condition of air today, why can't we figure out the condition of the air tomorrow?" First, we do not *really* know what the condition is today, because the air is swirling and twisting everywhere. It turns out to be very sensitive, and even unstable. If you have ever seen water run smoothly over a dam, and then turn into a large number of blobs and drops as it falls, you will understand what I mean by unstable. You know the condition of the water before it goes over the

spillway; it is perfectly smooth; but the moment it begins to fall, where do the drops begin? What determines how big the lumps are going to be and where they will be? That is not known, because the water is unstable. Even a smooth moving mass of air, in going over a mountain turns into complex whirlpools and eddies. In many fields we find this situation of *turbulent flow* that we cannot analyze today. Quickly we leave the subject of weather, and discuss geology!

The question basic to geology is, what makes the earth the way it is? The most obvious processes are in front of your very eyes, the erosion processes of the rivers, the winds, etc. It is easy enough to understand these, but for every bit of erosion there is an equal amount of something else going on. Mountains are no lower today, on the average, than they were in the past. There must be moun-*tsim-forming* processes. You will find, if you study geology, that there *are* mountain-forming processes and vulcanism, which nobody understands but which is half of geology. The phenomenon of volcanoes is really not understood. What makes an earthquake is, ultimately, not understood. It is understood that if something is pushing something else, it snaps and will slide—that is all right. But what pushes, and why? The theory is that there are currents inside the earth—circulating currents, due to the difference in temperature inside and outside—which, in their motion, push the surface slightly. Thus if there are two opposite circulations next to each other, the matter will collect in the region where they meet and make belts of mountains which are in unhappy stressed conditions, and so produce volcanoes and earthquakes.

What about the inside of the earth? A great deal is known about the speed of earthquake waves through the earth and the density of distribution of the earth. However, physicists have been unable to get a good theory as to how dense a substance should be at the pressures that would be expected at the center of the earth. In other words, we cannot figure out the properties of matter very well in these circumstances. We do much less well with the earth than we do with the conditions of matter in the stars. The mathematics involved seems a little too difficult, so far, but perhaps it will not be too long before someone realizes that it is an important problem, and really work it out. The other aspect, of course, is that even if we did know the density, we cannot figure out the circulating currents. Nor can we really work out the properties of rocks at high pressure. We cannot tell how fast the rocks should "give"; that must all be worked out by experiment.

3-6 Psychology

Next, we consider the science of *psychology*. Incidentally, psychoanalysis is not a science: it is at best a medical process, and perhaps even more like witch-doctoring. It has a theory as to what causes disease—lots of different "spirits," etc. The witch doctor has a theory that a disease like malaria is caused by a spirit which comes into the air; it is not cured by shaking a snake over it, but quinine does help malaria. So, if you are sick, I would advise that you go to the witch doctor because he is the man in the tribe who knows the most about the disease; on the other hand, his knowledge is not science. Psychoanalysis has not been checked carefully by experiment, and there is no way to find a list of the number of cases in which it works, the number of cases in which it does not work, etc.

The other branches of psychology, which involve things like the physiology of sensation—what happens in the eye, and what happens in the brain—are, if you wish, less interesting. But some small but real progress has been made in studying them. One of the most interesting technical problems may or may not be called psychology. The central problem of the mind, if you will, or the nervous system, is this: when an animal learns something, it can do something different than it could before, and its brain cell must have changed too, if it is made out of atoms. *In what way is it different ?* We do not know where to look, or what to look for, when something is memorized. We do not know what it means, or what change there is in the nervous system, when a fact is learned. This is a very important problem which has not been solved at all. Assuming, however, that there is some kind of memory thing, the brain is such an enormous mass of interconnect-

ing wires and nerves that it probably cannot be analyzed in a straightforward manner. There is an analog of this to computing machines and computing elements, in that they also have a lot of lines, and they have some kind of element, analogous, perhaps, to the synapse, or connection of one nerve to another. This is a very interesting subject which we have not the time to discuss further—the relationship between thinking and computing machines. It must be appreciated, of course, that this subject will tell us very little about the real complexities of ordinary human behavior. All human beings are so different. It will be a long time before we get there. We must start much further back. If we could even figure out how a *dog* works, we would have gone pretty far. Dogs are easier to understand, but nobody yet knows how dogs work.

3-7 How did it get that way?

In order for physics to be useful to other sciences in a *theoretical* way, other than in the invention of instruments, the science in question must supply to the physicist a description of the object in a physicist's language. They can say "why does a frog jump?," and the physicist cannot answer. If they tell him what a frog is, that there are so many molecules, there is a nerve here, etc., that is different. If they will tell us, more or less, what the earth or the stars are like, then we can figure it out. In order for physical theory to be of any use, we must know where the atoms are located. In order to understand the chemistry, we must know exactly what atoms are present, for otherwise we cannot analyze it. That is but one limitation, of course.

There is another *kind* of problem in the sister sciences which does not exist in physics; we might call it, for lack of a better term, the historical question. How did it get that way? If we understand all about biology, we will want **to** know how all the things which are on the earth got there. There is the theory of evolution, an important part of biology. In geology, we not only want to know how the mountains are forming, but how the entire earth was formed in the beginning, the origin of the solar system, etc. That, of course, leads us to want to know what kind of matter there was in the world. How did the stars evolve? What were the initial conditions? That is the problem of astronomical history. A great deal has been found out about the formation of stars, the formation of elements from which we were made, and even a little about the origin of the universe.

There is no historical question being studied in physics at the present time. We do not have a question, "Here are the laws of physics, how did they get that way?" We do not imagine, at the moment, that the laws of physics are somehow changing with time, that they were different in the past than they are at present. Of course they *may* be, and the moment we find they *are*, the historical question of physics will be wrapped up with the rest of the history of the universe, and then the physicist will be talking about the same problems as astronomers, geologists, and biologists.

Finally, there is a physical problem that is common to many fields, that is very old, and that has not been solved. It is not the problem of finding new fundamental particles, but something left over from a long time ago—over a hundred years. Nobody in physics has really been able to analyze it mathematically satisfactorily in spite of its importance to the sister sciences. It is the analysis of *circulating or turbulent fluids.* If we watch the evolution of a star, there comes a point where we can deduce that it is going to start convection, and thereafter we can no longer deduce what should happen. A few million years later the star explodes, but we cannot figure out the reason. We cannot analyze the weather. We do not know the patterns of motions that there should be inside the earth. The simplest form of the problem is to take a pipe that is very long and push water through it at high speed. We ask: to push a given amount of water through that pipe, how much pressure is needed? No one can analyze it from first principles and the properties of water. If the water flows very slowly, or if we use a thick goo like honey, then we can do it nicely. You will find that in your textbook.

What we really cannot do is deal with actual, wet water running through a pipe. That is the central problem which we ought to solve some day, and we have not.

A poet once said, "The whole universe is in a glass of wine." We will probably never know in what sense he meant that, for poets do not write to be understood. But it is true that if we look at a glass of wine closely enough we see the entire universe. There are the things of physics: the twisting liquid which evaporates depending on the wind and weather, the reflections in the glass, and our imagination adds the atoms. The glass is a distillation of the earth's rocks, and in its composition we see the secrets of the universe's age, and the evolution of stars. What strange array of chemicals are in the wine? How did they come to be? There are the ferments, the enzymes, the substrates, and the products. There in wine is found the great generalization: all life is fermentation. Nobody can discover the chemistry of wine without discovering, as did Louis Pasteur, the cause of much disease. How vivid is the claret, pressing its existence into the consciousness that watches it! If our small minds, for some convenience, divide this glass of wine, this universe, into parts—physics, biology, geology, astronomy, psychology, and so on—remember that nature does not know it! So let us put it all back together, not forgetting ultimately what it is for. Let it give us one more final pleasure: drink it and forget it all!

4

Conservation of Energy

4-1 What is energy?

4-2 Gravitational potential energy

4-3 Kinetic energy

4-4 Other forms of energy

4-1 What is energy?

In this chapter, we begin our more detailed study of the different aspects of physics, having finished our description of things in general. To illustrate the ideas and the kind of reasoning that might be used in theoretical physics, we shall now examine one of the most basic laws of physics, the conservation of energy.

There is a fact, or if you wish, a *law*, governing all natural phenomena that are known to date. There is no known exception to this law—it is exact so far as we know. The law is called the *conservation of energy*. It states that there is a certain quantity, which we call energy, that does not change in the manifold changes which nature undergoes. That is a most abstract idea, because it is a mathematical principle; it says that there is a numerical quantity which does not change when something happens. It is not a description of a mechanism, or anything concrete; it is just a strange fact that we can calculate some number and when we finish watching nature go through her tricks and calculate the number again, it is the same. (Something like the bishop on a red square, and after a number of moves—details unknown—it is still on some red square. It is a law of this nature.) Since it is an abstract idea, we shall illustrate the meaning of it by an analogy.

Imagine a child, perhaps "Dennis the Menace," who has blocks which are absolutely indestructible, and cannot be divided into pieces. Each is the same as the other. Let us suppose that he has 28 blocks. His mother puts him with his 28 blocks into a room at the beginning of the day. At the end of the day, being curious, she counts the blocks very carefully, and discovers a phenomenal law—no matter what he does with the blocks, there are always 28 remaining! This continues for a number of days, until one day there are only 27 blocks, but a little investigating shows that there is one under the rug—she must look everywhere to be sure that the number of blocks has not changed. One day, however, the number appears to change—there are only 26 blocks. Careful investigation indicates that the window was open, and upon looking outside, the other two blocks are found. Another day, careful count indicates that there are 30 blocks! This causes considerable consternation, until it is realized that Bruce came to visit, bringing his blocks with him, and he left a few at Dennis' house. After she has disposed of the extra blocks, she closes the window, does not let Bruce in, and then everything is going along all right, until one time she counts and finds only 25 blocks. However, there is a box in the room, a toy box, and the mother goes to open the toy box, but the boy says "No, do not open my toy box," and screams. Mother is not allowed to open the toy box. Being extremely curious, and somewhat ingenious, she invents a scheme! She knows that a block weighs three ounces, so she weighs the box at a time when she sees 28 blocks, and it weighs 16 ounces. The next time she wishes to check, she weighs the box again, subtracts sixteen ounces and divides by three. She discovers the following:

$$\begin{pmatrix}\text{number of}\\ \text{blocks seen}\end{pmatrix} + \frac{(\text{weight of box}) - 16 \text{ ounces}}{3 \text{ ounces}} = \text{constant}. \qquad (4.1)$$

There then appear to be some new deviations, but careful study indicates that the dirty water in the bathtub is changing its level. The child is throwing blocks into the water, and she cannot see them because it is so dirty, but she can find out how many blocks are in the water by adding another term to her formula. Since the original height of the water was 6 inches and each block raises the water a quarter

of an inch, this new formula would be:

$$\begin{pmatrix}\text{number of}\\ \text{blocks seen}\end{pmatrix} + \frac{(\text{weight of box}) - 16\ \text{ounces}}{3\ \text{ounces}} + \frac{(\text{height of water}) - 6\ \text{inches}}{1/4\ \text{inch}} = \text{constant.} \qquad (4.2)$$

In the gradual increase in the complexity of her world, she finds a whole series of terms representing ways of calculating how many blocks are in places where she is not allowed to look. As a result, she finds a complex formula, a quantity which *has to be computed*, which always stays the same in her situation.

What is the analogy of this to the conservation of energy? The most remarkable aspect that must be abstracted from this picture is that *there are no blocks.* Take away the first terms in (4.1) and (4.2) and we find ourselves calculating more or less abstract things. The analogy has the following points. First, when we are calculating the energy, sometimes some of it leaves the system and goes away, or sometimes some comes in. In order to verify the conservation of energy, we must be careful that we have not put any in or taken any out. Second, the energy has a large number of *different forms,* and there is a formula for each one. These are: gravitational energy, kinetic energy, heat energy, elastic energy, electrical energy, chemical energy, radiant energy, nuclear energy, mass energy. If we total up the formulas for each of these contributions, it will not change except for energy going in and out.

It is important to realize that in physics today, we have no knowledge of what energy *is.* We do not have a picture that energy comes in little blobs of a definite amount. It is not that way. However, there are formulas for calculating some numerical quantity, and when we add it all together it gives "28"—always the same number. It is an abstract thing in that it does not tell us the mechanism or the *reasons* for the various formulas.

4-2 Gravitational potential energy

Conservation of energy can be understood only if we have the formula for all of its forms. I wish to discuss the formula for gravitational energy near the surface of the Earth, and I wish to derive this formula in a way which has nothing to do with history but is simply a line of reasoning invented for this particular lecture to give you an illustration of the remarkable fact that a great deal about nature can be extracted from a few facts and close reasoning. It is an illustration of the kind of work theoretical physicists become involved in. It is patterned after a most excellent argument by Mr. Carnot on the efficiency of steam engines.*

Consider weight-lifting machines—machines which have the property that they lift one weight by lowering another. Let us also make a hypothesis: that *there is no such thing as perpetual motion* with these weight-lifting machines. (In fact, that there is no perpetual motion at all is a general statement of the law of conservation of energy.) We must be careful to define perpetual motion. First, let us do it for weight-lifting machines. If, when we have lifted and lowered a lot of weights and restored the machine to the original condition, we find that the net result is to have *lifted a weight,* then we have a perpetual motion machine because we can use that lifted weight to run something else. That is, *provided* the machine which lifted the weight is brought back to its exact *original condition,* and furthermore that it is completely *self-contained*—that it has not received the energy to lift that weight from some external source—like Bruce's blocks.



Fig. 4–1. Simple weight-lifting machine.

A very simple weight-lifting machine is shown in Fig. 4-1. This machine lifts weights three units "strong." We place three units on one balance pan, and one unit on the other. However, in order to get it actually to work, we must lift a little weight off the left pan. On the other hand, we could lift a one-unit weight

* Our point here is not so much the result, (4.3), which in fact you may already know, as the possibility of arriving at it by theoretical reasoning.

by lowering the three-unit weight, if we cheat a little by lifting a little weight off the other pan. Of course, we realize that with any *actual* lifting machine, we must add a little extra to get it to run. This we disregard, *temporarily.* Ideal machines, although they do not exist, do not require anything extra. A machine that we actually use can be, in a sense, *almost* reversible: that is, if it will lift the weight of three by lowering a weight of one, then it will also lift nearly the weight of one the same amount by lowering the weight of three.

We imagine that there are two classes of machines, those that are *not* reversible, which includes all real machines, and those that *are* reversible, which of course are actually not attainable no matter how careful we may be in our design of bearings, levers, etc. We suppose, however, that there is such a thing—a reversible machine—which lowers one unit of weight (a pound or any other unit) by one unit of distance, and at the same time lifts a three-unit weight. Call this reversible machine, Machine *A.* Suppose this particular reversible machine lifts the three-unit weight a distance *X.* Then suppose we have another machine, Machine *B,* which is not necessarily reversible, which also lowers a unit weight a unit distance, but which lifts three units a distance *Y.* We can now prove that *Y* is not higher than *X;* that is, it is impossible to build a machine that will lift a weight *any higher* than it will be lifted by a reversible machine. Let us see why. Let us suppose that *Y* were higher than *X.* We take a one-unit weight and lower it one unit height with Machine *B,* and that lifts the three-unit weight up a distance V. Then we could lower the weight from *Y* to *X, obtaining free power,* and use the reversible Machine *A,* running backwards, to lower the three-unit weight a distance *X* and lift the one-unit weight by one unit height. This will put the one-unit weight back where it was before, and leave both machines ready to be used again! We would therefore have perpetual motion if *Y* were higher than *X,* which we assumed was impossible. With those assumptions, we thus deduce that ***Y*** *is not higher than X,* so that of all machines that can be designed, the reversible machine is the best.

We can also see that all reversible machines must lift to *exactly the same height.* Suppose that *B* were really reversible also. The argument that *Y* is not higher than *X* is, of course, just as good as it was before, but we can also make our argument the other way around, using the machines in the opposite order, and prove that *X is not higher than Y.* This, then, is a very remarkable observation because it permits us to analyze the height to which different machines are going to lift something *without looking at the interior mechanism.* We know at once that if somebody makes an enormously elaborate series of levers that lift three units a certain distance by lowering one unit by one unit distance, and we compare it with a simple lever which does the same thing and is fundamentally reversible, his machine will lift it no higher, but perhaps less high. If his machine is reversible, we also know exactly *how* high it will lift. To summarize: every reversible machine, no matter how it operates, which drops one pound one foot and lifts a three-pound weight always lifts it the same distance, *X.* This is clearly a universal law of great utility. The next question is, of course, what is *X*?

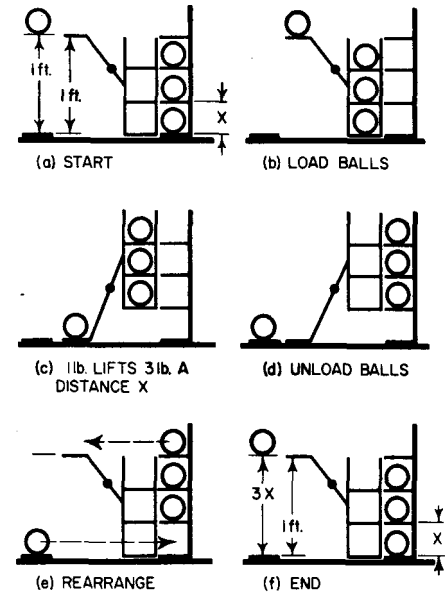


Fig. 4–2. A reversible machine.

Suppose we have a reversible machine which is going to lift this distance *X,* three for one. We set up three balls in a rack which does not move, as shown in Fig. 4-2. One ball is held on a stage at a distance one foot above the ground. The machine can lift three balls, lowering one by a distance 1. Now, we have arranged that the platform which holds three balls has a floor and two shelves, exactly spaced at distance *X,* and further, that the rack which holds the balls is spaced at distance *X,* (a). First we roll the balls horizontally from the rack to the shelves, (b), and we suppose that this takes no energy because we do not change the height. The reversible machine then operates: it lowers the single ball to the floor, and it lifts the rack a distance *X,* (c). Now we have ingeniously arranged the rack so that these balls are again even with the platforms. Thus we unload the balls onto the rack, (d); having unloaded the balls, we can restore the machine to its original condition. Now we have three balls on the upper three shelves and one at the bottom. But the strange thing is that, in a certain way of speaking, we have not lifted *two* of them at all because, after all, there were balls on shelves 2 and 3

before. The resulting effect has been to lift *one ball* a distance *3X*. Now, if *3X* exceeds one foot, then we can *lower* the ball to return the machine to the initial condition, (f), and we can run the apparatus again. Therefore 3 *X cannot* exceed one foot, for if 3 *X* exceeds one foot we can make perpetual motion. Likewise, we can prove that *one foot cannot exceed 3X*, by making the whole machine run the opposite way, since it is a reversible machine. Therefore *3X* is neither *greater nor less than a foot*, and we discover then, by argument alone, the law that $X = \frac{1}{3}$ foot. The generalization is clear: one pound falls a certain distance in operating a reversible machine; then the machine can lift *p* pounds this distance divided by *p*. Another way of putting the result is that three pounds times the height lifted, which in our problem was *X*, is equal to one pound times the distance lowered, which is one foot in this case. If we take all the weights and multiply them by the heights at which they are now, above the floor, let the machine operate, and then multiply all the weights by all the heights again, *there will be no change.* (We have to generalize the example where we moved only one weight to the case where when we lower one we lift several different ones—but that is easy.)

We call the sum of the weights times the heights *gravitational potential energy*—the energy which an object has because of its relationship in space, relative to the earth. The formula for gravitational energy, then, so long as we are not too far from the earth (the force weakens as we go higher) is

$$\begin{pmatrix}\text{gravitational}\\ \text{potential energy}\\ \text{for one object}\end{pmatrix} = (\text{weight}) \times (\text{height}). \tag{4.3}$$

It is a very beautiful line of reasoning. The only problem is that perhaps it is not true. (After all, nature does not *have* to go along with our reasoning.) ,For example, perhaps perpetual motion is, in fact, possible. Some of the assumptions may be wrong, or we may have made a mistake in reasoning, so it is always necessary to check. *It turns out experimentally,* in fact, to be true.

The general name of energy which has to do with location relative to something else is called *potential* energy. In this particular case, of course, we call it *gravitational potential energy.* If it is a question of electrical forces against which we are working, instead of gravitational forces, if we are "lifting" charges away from other charges with a lot of levers, then the energy content is called *electrical potential energy.* The general principle is that the change in the energy is the force times the distance that the force is pushed, and that this is a change in energy in general:

$$\begin{pmatrix}\text{change in}\\ \text{energy}\end{pmatrix} = (\text{force}) \times \begin{pmatrix}\text{distance force}\\ \text{acts through}\end{pmatrix}. \tag{4.4}$$

We will return to many of these other kinds of energy as we continue the course.

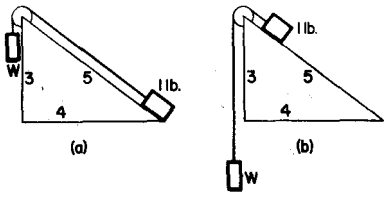


Fig. 4–3. Inclined plane.

The principle of the conservation of energy is very useful for deducing what will happen in a number of circumstances. In high school we learned a lot of laws about pulleys and levers used in different ways. We can now see that these "laws" are *all the same thing,* and that we did not have to memorize 75 rules to figure it out. A simple example is a smooth inclined plane which is, happily, a three-four-five triangle (Fig. 4-3). We hang a one-pound weight on the inclined plane with a pulley, and on the other side of the pulley, a weight *W.* We want to know how heavy *W* must be to balance the one pound on the plane. How can we figure that out? If we say it is just balanced, it is reversible and so can move up and down, and we can consider the following situation. In the initial circumstance, (a), the one pound weight is at the bottom and weight *W* is at the top. When *W* has slipped down in a reversible way, we have a one-pound weight at the top and the weight *W* the slant distance, (b), or five feet, from the plane in which it was before. We *lifted* the one-pound weight only *three* feet and we lowered *W* pounds by *five* feet. Therefore $W = \frac{3}{5}$ of a pound. Note that we deduced this from the *conservation of energy,* and not from force components. Cleverness, however, is relative. It can be deduced in a way which is even more brilliant, discovered by

Stevinus and inscribed on his tombstone. Figure 4-4 explains that it has to be $\frac{3}{5}$ of a pound, because the chain does not go around. It is evident that the lower part of the chain is balanced by itself, so that the pull of the five weights on one side must balance the pull of three weights on the other, or whatever the ratio of the legs. You see, by looking at this diagram, that W must be $\frac{3}{5}$ of a pound. (If you get an epitaph like that on your gravestone, you are doing fine.)

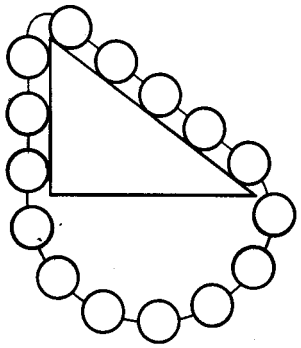


Fig. 4–4. The epitaph of Stevinus.

Let us now illustrate the energy principle with a more complicated problem, the screw jack shown in Fig. 4-5. A handle 20 inches long is used to turn the screw, which has 10 threads to the inch. We would like to know how much force would be needed at the handle to lift one ton (2000 pounds). If we want to lift the ton one inch, say, then we must turn the handle around ten times. When it goes around once it goes approximately 126 inches. The handle must thus travel 1260 inches, and if we used various pulleys, etc., we would be lifting our one ton with an unknown smaller weight W applied to the end of the handle. So we find out that W is about 1.6 pounds. This is a result of the conservation of energy.

Take now the somewhat more complicated example shown in Fig. 4-6. A rod or bar, 8 feet long, is supported at one end. In the middle of the bar is a weight of 60 pounds, and at a distance of two feet from the support there is a weight of 100 pounds. How hard do we have to lift the end of the bar in order to keep it balanced, disregarding the weight of the bar? Suppose we put a pulley at one end and hang a weight on the pulley. How big would the weight W have to be in order for it to balance? We imagine that the weight falls any arbitrary distance—to make it easy for ourselves suppose it goes down 4 inches—how high would-the two load weights rise? The center rises 2 inches, and the point a quarter of the way from the fixed end lifts 1 inch. Therefore, the principle that the sum of the heights times the weights does not change tells us that the weight W times 4 inches down, plus 60 pounds times 2 inches up, plus 100 pounds times 1 inch has to add up to nothing:

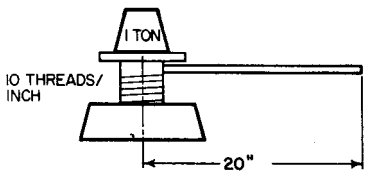


Fig. 4–5. A screw jack.

$$-4W + (2)(60) + (1)(100) = 0, \qquad W = 55 \text{ lb.} \tag{4.5}$$

Thus we must have a 55-pound weight to balance the bar. In this way we can work out the laws of "balance"—the statics of complicated bridge arrangements, and so on. This approach is called *the principle of virtual work*, because in order to apply this argument we had to *imagine* that the structure moves a little—even though it is not *really* moving or even *movable.* We use the very small imagined motion to apply the principle of conservation of energy.

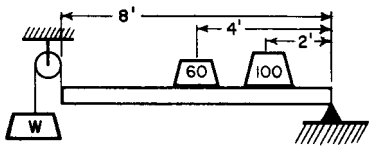


Fig. 4–6. Weighted rod supported on one end.

4-3 Kinetic energy

To illustrate another type of energy we consider a pendulum (Fig. 4-7). If we pull the mass aside and release it, it swings back and forth. In its motion, it loses height in going from either end to the center. Where does the potential energy go? Gravitational energy disappears when it is down at the bottom; nevertheless, it will climb up again. The gravitational energy must have gone into another form. Evidently it is by virtue of its *motion* that it is able to climb up again, so we have the conversion of gravitational energy into some other form when **it** reaches the bottom.

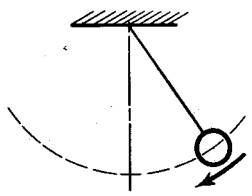


Fig. 4–7. Pendulum.

We must get a formula for the energy of motion. Now, recalling our arguments about reversible machines, we can easily see that in the motion at the bottom must be a quantity of energy which permits it to rise a certain height, and which has nothing to do with the *machinery* by which it comes up or the *path* by which it comes up. So we have an equivalence formula something like the one we wrote for the child's blocks. We have another form to represent the energy. It is easy to say what it is. The kinetic energy at the bottom equals the weight times the height that it could go, corresponding to its velocity: K.E. = WH. What we need is the formula which tells us the height by some rule that has to do with the motion of objects. If we start something out with a certain velocity, say straight up, it will reach a certain height; we do not know what it is yet, but it depends on the velocity—there is a formula for that. Then to find the formula for kinetic energy

for an object moving with velocity V, we must calculate the height that it could reach, and multiply by the weight. We shall soon find that we can write it this way:

$$\text{K.E.} = WV^2/2g. \tag{4.6}$$

Of course, the fact that motion has energy has nothing to do with the fact that we are in a gravitational field. It makes no difference *where* the motion came from. This is a general formula for various velocities. Both (4.3) and (4.6) are approximate formulas, the first because it is incorrect when the heights are great, i.e., when the heights are so high that gravity is weakening; the second, because of the relativistic correction at high speeds. However, when we do finally get the exact formula for the energy, then the law of conservation of energy is correct.

4-4 Other forms of energy

We can continue in this way to illustrate the existence of energy in other forms. First, consider elastic energy. If we pull down on a spring, we must do some work, for when we have it down, we can lift weights with it. Therefore in its stretched condition it has a possibility of doing some work. If we were to evaluate the sums of weights times heights, it would not check out—we must add something else to account for the fact that the spring is under tension. Elastic energy is the formula for a spring when it is stretched. How much energy is it? If we let go, the elastic energy, as the spring passes through the equilibrium point, is converted to kinetic energy and it goes back and forth between compressing or stretching the spring and kinetic energy of motion. (There is also some gravitational energy going in and out, but we can do this experiment "sideways" if we like.) It keeps going until the losses—Aha! We have cheated all the way through by putting on little weights to move things or saying that the machines are reversible, or that they go on forever, but we can see that things do stop, eventually. Where is the energy when the spring has finished moving up and down? This brings in *another* form of energy: *heat energy*.

Inside a spring or a lever there are crystals which are made up of lots of atoms, and with great care and delicacy in the arrangement of the parts one can try to adjust things so that as something rolls on something else, none of the atoms do any jiggling at all. But one must be very careful. Ordinarily when things roll, there is bumping and jiggling because of the irregularities of the material, and the atoms start to wiggle inside. So we lose track of that energy; we find the atoms are wiggling inside in a random and confused manner after the motion slows down. There is still kinetic energy, all right, but it is not associated with visible motion. What a dream! How do we *know* there is still kinetic energy? It turns out that with thermometers you can find out that, in fact, the spring or the lever is *warmer*, and that there is really an increase of kinetic energy by a definite amount. We call this form of energy *heat energy*, but we know that it is not really a new form, it is just kinetic energy—internal motion. (One of the difficulties with all these experiments with matter that we do on a large scale is that we cannot really demonstrate the conservation of energy and we cannot really make our reversible machines, because every time we move a large clump of stuff, the atoms do not remain absolutely undisturbed, and so a certain amount of random motion goes into the atomic system. We cannot see it, but we can measure it with thermometers, etc.)

There are many other forms of energy, and of course we cannot describe them in any more detail just now. There is electrical energy, which has to do with pushing and pulling by electric charges. There is radiant energy, the energy of light, which we know is a form of electrical energy because light can be represented as wigglings in the electromagnetic field. There is chemical energy, the energy which is released in chemical reactions. Actually, elastic energy is, to a certain extent, like chemical energy, because chemical energy is the energy of the attraction of the atoms, one for the other, and so is elastic energy. Our modern understanding is the following: chemical energy has two parts, kinetic energy of the electrons inside the atoms, so part of it is kinetic, and electrical energy of interaction of the

electrons and the protons—the rest of it, therefore, is electrical. Next we come to nuclear energy, the energy which is involved with the arrangement of particles inside the nucleus, and we have formulas for that, but we do not have the fundamental laws. We know that it is not electrical, not gravitational, and not purely chemical, but we do not know what it is. It seems to be an additional form of energy. Finally, associated with the relativity theory, there is a modification of the laws of kinetic energy, or whatever you wish to call it, so that kinetic energy is combined with another thing called *mass energy*. An object has energy from its sheer *existence*. If I have a positron and an electron, standing still doing nothing —never mind gravity, never mind anything—and they come together and disappear, radiant energy will be liberated, in a definite amount, and the amount can be calculated. All we need know is the mass of the object. It does not depend on what it is—we make two things disappear, and we get a certain amount of energy. The formula was first found by Einstein; it is $E = mc^2$.

It is obvious from our discussion that the law of conservation of energy is enormously useful in making analyses, as we have illustrated in a few examples without knowing all the formulas. If we had all the formulas for all kinds of energy, we could analyze how many processes should work without having to go into the details. Therefore conservation laws are very interesting. The question naturally arises as to what other conservation laws there are in physics. There are two other conservation laws which are analogous to the conservation of energy. One is called the conservation of linear momentum. The other is called the conservation of angular momentum. We will find out more about these later. In the last analysis, we do not understand the conservation laws deeply. We do not understand the conservation of energy. We do not understand energy as a certain number of little blobs. You may have heard that photons come out in blobs and that the energy of a photon is Planck's constant times the frequency. That is true, but since the frequency of light can be anything, there is no law that says that energy has to be a certain definite amount. Unlike Dennis' blocks, there can be any amount of energy, at least as presently understood. So we do not understand this energy as counting something at the moment, but just as a mathematical quantity, which is an abstract and rather peculiar circumstance. In quantum mechanics it turns out that the conservation of energy is very closely related to another important property of the world, *things do not depend on the absolute time.* We can set up an experiment at a given moment and try it out, and then do the same experiment at a later moment, and it will behave in exactly the same way. Whether this is strictly true or not, we do not know. If we assume that it *is* true, and add the principles of quantum mechanics, then we can deduce Jhe principle of the conservation of energy. It is a rather subtle and interesting thing, and it is not easy to explain. The other conservation laws are also linked together. The conservation of momentum is associated in quantum mechanics with the proposition that it makes no difference *where* you do the experiment, the results will always be the same. As independence in space has to do with the conservation of momentum, independence of time has to do with the conservation of energy, and finally, if we *turn* our apparatus, this too makes no difference, and so the invariance of the world to angular orientation is related to the conservation of *angular momentum.* Besides these, there are three other conservation laws, that are exact so far as we can tell today, which *are* much simpler to understand because they are in the nature of counting blocks.

The first of the three is the *conservation of charge,* and that merely means that you count how many positive, minus how many negative electrical charges you have, and the number is never changed. You may get rid of a positive with a negative, but you do not create any net excess of positives over negatives. Two other laws are analogous to this one—one is called the *conservation of baryons.* There are a number of strange particles, a neutron and a proton are examples, which are called baryons. In any reaction whatever in nature, if we count how many baryons are coming into a process, the number of baryons* which come out

*Counting antibaryons as —1 baryon.

will be exactly the same. There is another law, the *conservation of leptons.* We can say that the group of particles called leptons are: electron, mu meson, and neutrino. There is an antielectron which is a positron, that is, a —1 lepton. Counting the total number of leptons in a reaction reveals that the number in and out never changes, at least so far as we know at present.

These are the six conservation laws, three of them subtle, involving space and time, and three of them simple, in the sense of counting something.

With regard to the conservation of energy, we should note that *available* energy is another matter—there is a lot of jiggling around in the atoms of the water of the sea, because the sea has a certain temperature, but it is impossible to get them herded into a definite motion without taking energy from somewhere else. That is, although we know for a fact that energy is conserved, the energy available for human utility is not conserved so easily. The laws which govern how much energy is available are called the *laws of thermodynamics* and involve a concept called entropy for irreversible thermodynamic processes.

Finally, we remark on the question of where we can get our supplies of energy today. Our supplies of energy are from the sun, rain, coal, uranium, and hydrogen. The sun makes the rain, and the coal also, so that all these are from the sun. Although energy is conserved, nature does not seem to be interested in it; she liberates a lot of energy from the sun, but only one part in two billion falls on the earth. Nature has conservation of energy, but does not really care; she spends a lot of it in all directions. We have already obtained energy from uranium; we can also get energy from hydrogen, but at present only in an explosive and dangerous condition. If it can be controlled in thermonuclear reactions, it turns out that the energy that can be obtained from 10 quarts of water per second is equal to all of the electrical power generated in the United States. With 150 gallons of running water a minute, you have enough fuel to supply all the energy which is used in the United States today! Therefore it is up to the physicist to figure out how to liberate us from the need for having energy. It can be done.

5

Time and Distance

5-1 Motion

5-2 **Time**

5-3 Short times

5-4 Long times

5-5 Units and standards of time

5-6 Large distances

5-7 Short distances

5-1 Motion

In this chapter we shall consider some aspects of the concepts of *time* and *distance.* It has been emphasized earlier that physics, as do all the sciences, depends on *observation.* One might also say that the development of the physical sciences to their present form has depended to a large extent on the emphasis which has been placed on the making of *quantitative* observations. Only with quantitative observations can one arrive at quantitative relationships, which are the heart of physics.

Many people would like to place the beginnings of physics with the work done 350 years ago by Galileo, and to call him the first physicist. Until that time, the study of motion had been a philosophical one based on arguments that could be thought up in one's head. Most of the arguments had been presented by Aristotle and other Greek philosophers, and were taken as "proven." Galileo was skeptical, and did an experiment on motion which was essentially this: **He** allowed a ball to roll down an inclined trough and observed the motion. He did not, however, just look; he measured *how far* the ball went in *how long a time.*

The way to measure a distance was well known long before Galileo, but there were no accurate ways of measuring time, particularly short times. Although he later devised more satisfactory clocks (though not like the ones we know), Galileo's first experiments on motion were done by using his pulse to count off equal intervals of time. Let us do the same.

We may count off beats of a pulse as the ball rolls down the track: "one .. . two ... three .. . four .. . five ... six ... seven . . . eight..." We ask a friend to make a small mark at the location of the ball at each count; we can then measure the *distance* the ball travelled from the point of release in one, or two, or three, etc., equal intervals of time. Galileo expressed the result of *his* observations in this way: if the location of the ball is marked at 1, 2, 3, 4,... units of time from the instant of Its release, those marks are distant from the starting point in proportion to the numbers 1, 4, 9, 16, ... Today we would say the distance is proportional to the square of the time:

$$D \propto t^2.$$

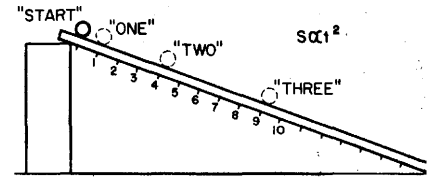


Fig. 5-1. A ball rolls down an dined track.

The study of *motion,* which is basic to all of physics, treats with the questions: where? and when?

5-2 Time

Let us consider first what we mean by *time.* What *is* time? It would be nice if we could find a good definition of time. Webster defines "a time" as "a period," and the latter as "a time," which doesn't seem to be very useful. Perhaps we should say: "Time is what happens when nothing else happens." Which also doesn't get us very far. Maybe it is just as well if we face the fact that time is one of the things we probably cannot define (in the dictionary sense), and just say that it is what we already know it to be: it is how long we wait!

What really matters anyway is not how we *define* time, but how we measure it. One way of measuring time is to utilize something which happens over and over again in a regular fashion—something which is *periodic.* For example, a day. A day seems to happen over and over again. But when you begin to think

about it, you might well ask: "Are days periodic; are they regular? Are all days the same length?" One certainly has the impression that days in summer are longer than days in winter. Of course, some of the days in winter seem to get awfully long if one is very bored. You have certainly heard someone say, "My, but this has been a long day!"

It does seem, however, that days are about the same length *on the average*. Is there any way we can test whether the days are the same length—either from one day to the next, or at least on the average? One way is to make a comparison with some other periodic phenomenon. Let us see how such a comparison might be made with an hour glass. With an hour glass, we can "create" a periodic occurrence if we have someone standing by it day and night to turn it over whenever the last grain of sand runs out.

We could then count the turnings of the glass from each morning to the next. We would find, this time, that the number of "hours" (i.e., turnings of the glass) was not the same each "day." We should distrust the sun, or the glass, or both. After some thought, it might occur to us to count the "hours" from noon to noon. (Noon is here defined *not* as 12:00 o'clock, but that instant when the sun is at its highest point.) We would find, this time, that the number of "hours" each day is the same.

We now have some confidence that both the "hour" and the "day" have a regular periodicity, i.e., mark off successive equal intervals of time, although we have not *proved* that either one is "really" periodic. Someone might question whether there might not be some omnipotent being who would slow down the flow of sand every night and speed it up during the day. Our experiment does not, of course, give us an answer to this sort of question. All we can say is that we find that a regularity of one kind fits together with a regularity of another kind. We can just say that we base our *definition* of time on the repetition of some apparently periodic event.

5-3 Short times

We should now notice that in the process of checking on the reproducibility of the day, we have received an important by-product. We have found a way of measuring, more accurately, *fractions* of a day. We have found a way of counting time in smaller pieces. Can we carry the process further, and learn to measure even smaller intervals of time?

Galileo decided that a given pendulum always swings back and forth in equal intervals of time so long as the size of the swing is kept small. A test comparing the number of swings of a pendulum in one "hour" shows that such is indeed the case. We can in this way mark fractions of an hour. If we use a mechanical device to count the swings—and to keep them going—we have the pendulum clock of our grandfathers.

Let us agree that if our pendulum oscillates 3600 times in one hour (and if there are 24 such hours in a day), we shall call each period of the pendulum one "second." We have then divided our original unit of time into approximately 10^5 parts. We can apply the same principles to divide the second into smaller and smaller intervals. It is, you will realize, not practical to make mechanical pendulums which go arbitrarily fast, but we can now make *electrical* pendulums, called oscillators, which can provide a periodic occurrence with a very short period of swing. In these electronic oscillators it is an electrical current which swings to and fro, in a manner analogous to the swinging of the bob of the pendulum.

We can make a series of such electronic oscillators, each with a period 10 times shorter than the previous one. We may "calibrate" each oscillator against the next slower one by counting the number of swings it makes for one swing of the slower oscillator. When the period of oscillation of our clock is shorter than a fraction of a second, we cannot count the oscillations without the help of some device which extends our powers of observation. One such device is the electron-beam oscilloscope, which acts as a sort of microscope for short times. This device plots on a fluorescent screen a graph of electrical current (or voltage) versus time.

By connecting the oscilloscope to two of our oscillators in sequence, so that it plots a graph first of the current in one of our oscillators and then of the current in the other, we get two graphs like those shown in Fig. 5-2. We can readily determine the number of periods of the faster oscillator in one period of the slower oscillator.

With modern electronic techniques, oscillators have been built with periods as short as about 10^{-12} second, and they have been calibrated (by comparison methods such as we have described) in terms of our standard unit of time, the second. With the invention and perfection of the "laser," or light amplifier, in the past few years, it has become possible to make oscillators with even shorter periods than 10^{-12} second, but it has not yet been possible to calibrate them by the methods which have been described, although it will no doubt soon be possible.

Times shorter than 10^{-12} second have been measured, but by a different technique. In effect, a different *definition* of "time" has been used. One way has been to observe the *distance* between two happenings on a moving object. If, for example, the headlights of a moving automobile are turned on and then off, we can figure out *how long* the lights were on if we know *where* they were turned on and off and how fast the car was moving. The time is the distance over which the lights were on divided by the speed.

Within the past few years, just such a technique was used to measure the lifetime of the π^0meson. By observing in a microscope the minute tracks left in a photographic emulsion in which π^0-mesons had been created one saw that a π^0-meson (known to be travelling at a certain speed nearly that of light) went a distance of about 10^{-7} meter, on the average, before disintegrating. It lived for only about 10^{-16} sec. It should be emphasized that we have here used a somewhat different definition of "time" than before. So long as there are no inconsistencies in our understanding, however, we feel fairly confident that our definitions are sufficiently equivalent.

By extending our techniques—and if necessary our definitions—still further we can infer the time duration of still faster physical events. We can speak of the period of a nuclear vibration. We can speak of the lifetime of the newly discovered strange resonances (particles) mentioned in Chapter 2. Their complete life occupies a time span of only 10^{-24} second, approximately the time it would take light (which moves at the fastest known speed) to cross the nucleus of hydrogen (the smallest known object).

What about still smaller times? Does "time" exist on a still smaller scale? Does it make any sense to speak of smaller times if we cannot measure—or perhaps even think sensibly about—something which happens in a shorter time? Perhaps not. These are some of the open questions which you will be asking and perhaps answering in the next twenty or thirty years.

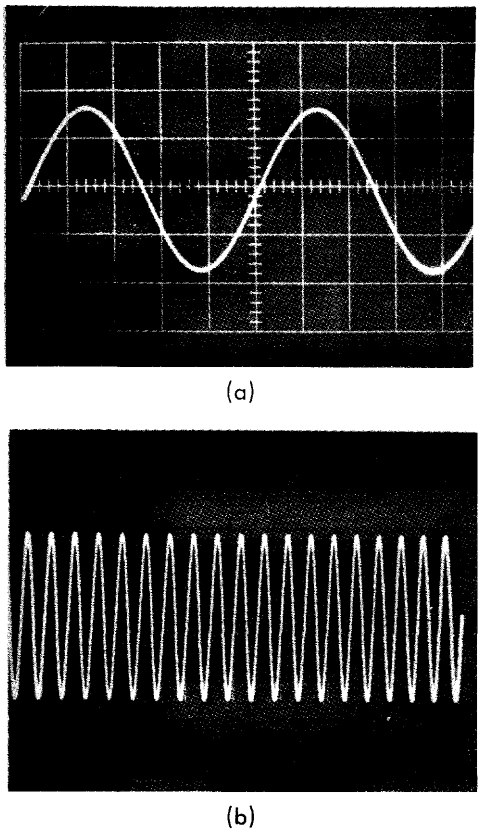


Fig. 5-2. Two views of an oscilloscope screen. In (a) the oscilloscope is connected to one oscillator, in (b) it is connected to an oscillator with a period one-tenth as long.

5-4 Long times

Let us now consider times longer than one day. Measurement of longer times is easy; we just count the days—so long as there is someone around to do the counting. First we find that there is another natural periodicity: the year, about 365 days. We have also discovered that nature has sometimes provided a counter for the years, in the form of tree rings or river-bottom sediments. In some cases we can use these natural time markers to determine the time which has passed since some early event.

When we cannot count the years for the measurement of long times, we must look for other ways to measure. One of the most successful is the use of radioactive material as a "clock." In this case we do not have a periodic occurrence, as for the day or the pendulum, but a new kind of "regularity." We find that the radioactivity of a particular sample of material decreases by the same *fraction* for successive equal increases in its age. If we plot a graph of the radioactivity observed as a function of time (say in days), we obtain a curve like that shown in Fig. 5-3. We observe that if the radioactivity decreases to one-half in T days (called the "half-life"), then it decreases to one-quarter in another T days, and so

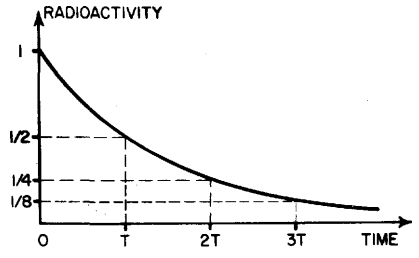


Fig. 5-3. The decrease with time of radioactivity. The activity decreases by one-half in each "half-life," *T*.

TIMES			MEAN LIFE OF
YEARS	SECONDS		
		? ? ? ? ? ? ? ?	
	10^{18}	Age of universe	
10^{9}		Age of earth	U^{238}
	10^{15}		
10^{6}		Earliest men	
	10^{12}	Age of pyramids	
10^{3}			Ra^{226}
		Age of U.S.	
	10^{9}	Life of a man	H^{3}
1			
	10^{6}		
		One day	
	10^{3}	Light goes from sun to earth	Neutron
	1	One heart beat	
	10^{-3}	Period of sound wave	
	10^{-6}	Period of radiowave	Muon
			$\pi^{\pm}$-meson
	10^{-9}	Light travels one foot	
	10^{-12}	Period of molecular rotation	
	10^{-15}	Period of atomic vibration	
			π°-meson
	10^{-18}	Light crosses an atom	
	10^{-21}		
		Period of nuclear vibration	
	10^{-24}	Light crosses a nucleus	Strange particle
		? ? ? ? ? ? ? ?	

on. In an arbitrary time interval t there are t/T "half-lives," and the fraction left after this time t is $(\frac{1}{2})^{t/T}$.

If we knew that a piece of material, say a piece of wood, had contained an amount A of radioactive material when it was formed, and we found out by a direct measurement that it now contains the amount B, we could compute the age of the object, t, by solving the equation

$$(\tfrac{1}{2})^{t/T} = B/A.$$

There are, fortunately, cases in which we can know the amount of radioactivity that was in an object when it was formed. We know, for example, that the carbon dioxide in the air contains a certain small fraction of the radioactive carbon isotope C^{14} (replenished continuously by the action of cosmic rays). If we measure the *total* carbon content of an object, we know that a certain fraction of that amount was originally the radioactive C^{14}; we know, therefore, the starting amount A to use in the formula above. Carbon-14 has a half-life of 5000 years. By careful measurements we can measure the amount left after 20 half-lives or so and can therefore "date" organic objects which grew as long as 100,000 years ago.

We would like to know, and we think we do know, the life of still older things. Much of our knowledge is based on the measurements of other radioactive isotopes which have different half-lives. If we make measurements with an isotope with a longer half-life, then we are able to measure longer times. Uranium, for example, has an isotope whose half-life is about 10^9 years, so that if some material was formed with uranium in it 10^9 years ago, only half the uranium would remain today. When the uranium disintegrates, it changes into lead. Consider a piece of rock which was formed a long time ago in some chemical process. Lead, being of a chemical nature different from uranium, would appear in one part of the rock and uranium would appear in another part of the rock. The uranium and lead

would be separate. If we look at that piece of rock today, where there should only be uranium we will how find a certain fraction of uranium and a certain fraction of lead. By comparing these fractions, we can tell what percent of the uranium disappeared and changed into lead. By this method, the age of certain rocks has been determined to be several billion years. An extension of this method, not using particular rocks but looking at the uranium and lead in the oceans and using averages over the earth, has been used to determine (within the past few years) that the age of the earth itself is approximately 5.5 billion years.

It is encouraging that the age of the earth is found to be the same as the age of the meteorites which land on the earth, as determined by the uranium method. It appears that the earth was formed out of rocks floating in space, and that the meteorites are, quite likely, some of that material left over. At some time more than five billion years ago, the universe started. It is now believed that at least our part of the universe had its beginning about ten or twelve billion years ago. We do not know what happened before then. In fact, we may well ask again: Does the question make any sense? Does an earlier time have any meaning?

5-5 Units and standards of time

We have implied that it is convenient if we start with some standard unit of time, say a day or a second, and refer all other times to some multiple or fraction of this unit. What shall we take as our basic standard of time? Shall we take the human pulse? If we compare pulses, we find that they seem to vary a lot. On comparing two clocks, one finds they do not vary so much. You might then say, well, let us take a clock. But whose clock? There is a story of a Swiss boy who wanted all of the clocks in his town to ring noon at the same time. So he went around trying to convince everyone of the value of this. Everyone thought it was a marvelous idea so long as all of the other clocks rang noon when his did! It is rather difficult to decide whose clock we should take as a standard. Fortunately, we all share one clock—the earth. For a long time the rotational period of the earth has been taken as the basic standard of time. As measurements have been made more and more precise, however, it has been found that the rotation of the earth is not exactly periodic, when measured in terms of the best clocks. These "best" clocks are those which we have reason to believe are accurate because they agree with each other. We now believe that, for various reasons, some days are longer than others, some days are shorter, and on the average the period of the earth becomes a little longer as the centuries pass.

Until very recently we had found nothing much better than the earth's period, so all clocks have been related to the length of the day, and the second has been defined as 1/86400 of an average day. Recently we have been gaining experience with some natural oscillators which we now believe would provide a more constant time reference than the earth, and which are also based on a natural phenomenon available to everyone. These are the so-called "atomic clocks." Their basic internal period is that of an atomic vibration which is very insensitive to the temperature or any other external effects. These clocks keep time to an accuracy of one part in 10^9 or better. Within the past two years an improved atomic clock which operates on the vibration of the hydrogen atom has been designed and built by Professor Norman Ramsey at Harvard University. He believes that this clock might be 100 times more accurate still. Measurements now in progress will show whether this is true or not.

We may expect that since it has been possible to build clocks much more accurate than astronomical time, there will soon be an agreement among scientists to define the unit of time in terms of one of the atomic clock standards.

5-6 Large distances

Let us now turn to the question of *distance*. How far, or how big, are things? Everybody knows that the way you measure distance is to start with a stick and count. Or start with a thumb and count. You begin with a unit and count. How

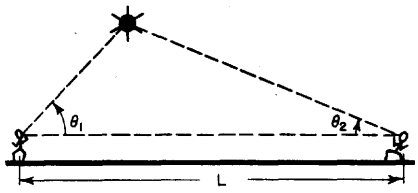


Fig. 5-4. The height of a Sputnik is determined by triangulation

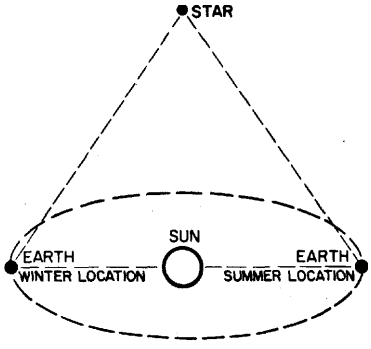


Fig. 5-5. The distance of nearby stars can be measured by triangulation, using the diameter of the earth's orbit as a baseline.

does one measure smaller things? How does one subdivide distance? In the same way that we subdivided time: we take a smaller unit and count the number of such units it takes to make up the longer unit. So we can measure smaller and smaller lengths.

But we do not always mean by distance what one gets by counting off with a meter stick. It would be difficult to measure the horizontal distance between two mountain tops using only a meter stick. We have found by experience that distance can be measured in another fashion: by triangulation. Although this means that we are really using a different definition of distance, when they can both be used they agree with each other. Space is more or less what Euclid thought it was, so the two types of definitions of distance agree. Since they do agree on the earth it gives us some confidence in using triangulation for still larger distances. For example, we were able to use triangulation to measure the height of the first Sputnik. We found that it was roughly 5×10^5 meters high. By more careful measurements the distance to the moon can be measured in the same way. Two telescopes at different places on the earth can give us the two angles we need. It has been found in this way that the moon is 4×10^8 meters away.

We cannot do the same with the sun, or at least no one has been able to yet. The accuracy with which one can focus on a given point on the sun and with which one can measure angles is not good enough to permit us to measure the distance to the sun. Then how can we measure the distance to the sun? We must invent an extension of the idea of triangulation. We measure the relative distances of all the planets by astronomical observations of where the planets appear to be, and we get a picture of the solar system with the proper *relative* distances of everything, but with no *absolute* distance. One absolute measurement is then required, which has been obtained in a number of ways. One of the ways, which was believed until recently to be the most accurate, was to measure the distance from the earth to Eros, one of the small planetoids which passes near the earth every now and then. By triangulation on this little object, one could get the one required scale measurement. Knowing the relative distances of the rest, we can then tell the distance, for example, from the earth to the sun, or from the earth to Pluto.

Within the past year there has been a big improvement in our knowledge of the scale of the solar system. At the Jet Propulsion Laboratory the distance from the earth to Venus was measured quite accurately by a direct radar observation. This, of course, is a still different type of inferred distance. We say we know the speed at which light travels (and therefore, at which radar waves travel), and we assume that it is the same speed everywhere between the earth and Venus. We send the radio wave out, and count the time until the reflected wave comes back. From the *time* we infer a *distance,* assuming we know the speed. We have really another definition of a measurement of distance.

How do we measure the distance to a star, which is much farther away? Fortunately, we can go back to our triangulation method, because the earth moving around the sun gives us a large baseline for measurements of objects outside the solar system. If we focus a telescope on a star in summer and in winter, we might hope to determine these two angles accurately enough to be able to measure the distance to a star.

What if the stars are too far away for us to use triangulation? Astronomers are always inventing new ways of measuring distance. They find, for example, that they can estimate the size and brightness of a star by its color. The color and brightness of many nearby stars—whose distances are known by triangulation—have been measured, and it is found that there is a smooth relationship between the color and the intrinsic brightness of stars (in most cases). If one now measures the color of a distant star, one may use the color-brightness relationship to determine the intrinsic brightness of the star. By measuring how bright the star *appears* to us at the earth (or perhaps we should say how *dim* it appears), we can compute how far away it is. (For a given intrinsic brightness, the apparent brightness decreases with the square of the distance.) A nice confirmation of the correctness of this method of measuring stellar distances is given by the results obtained for groups of stars known as globular clusters. A photograph of such a group is

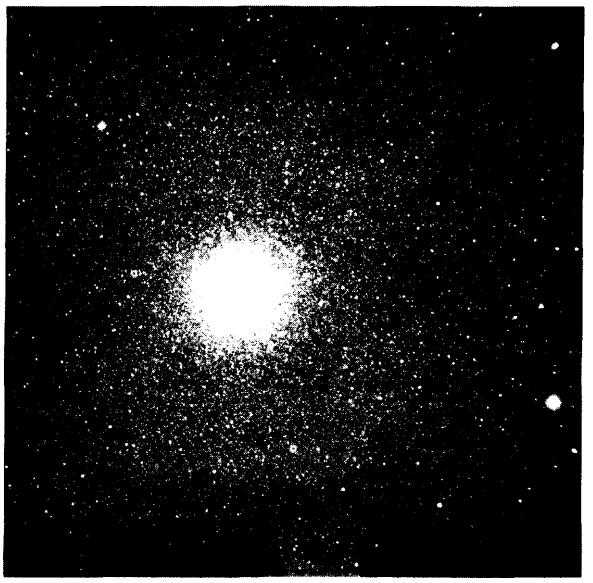


Fig. 5-6. A cluster of stars near the center of our galaxy. Their distance from the earth is 30,000 light-years, or about 3 X 10^{20} meters.

shown in Fig. 5-6. Just from looking at the photograph one is convinced that these stars are all together. The same result is obtained from distance measurements by the color-brightness method.

A study of many globular clusters gives another important bit of information. It is found that there is a high concentration of such clusters in a certain part of the sky and that most of them are about the same distance from us. Coupling this information with other evidence, we conclude that this concentration of clusters marks the center of our galaxy. We then know the distance to the center of the galaxy—about 10^{20} meters.

Knowing the size of our own galaxy, we have a key to the measurement of still larger distances—the distances to other galaxies. Figure 5-7 is a photograph of a galaxy, which has much the same shape as our own. Probably it is the same size, too. (Other evidence supports the idea that galaxies are all about the same size.) If it is the same size as ours, we can tell its distance. We measure the angle it subtends in the sky; we know its diameter, and we compute its distance—triangulation again!



Fig. 5-7. A spiral galaxy like our own. Presuming that its diameter is similar to that of our own galaxy, we may compute its distance from its apparent size. It is 30 million light-years (3 X 10^{23} meters) from the earth.



Fig. 5-8. The most distant object, 3C295 in BOOTES (indicated by the arrow), measured by the 200-inch telescope to date (1960).

Photographs of exceedingly distant galaxies have recently been obtained with the giant Palomar telescope. One is shown in Fig. 5-8. It is now believed that some of these galaxies are about halfway to the limit of the universe—10^{26} meters away—the largest distance we can contemplate!

5-7 Short distances

Now let's think about smaller distances. Subdividing the meter is easy. Without much difficulty we can mark off one thousand equal spaces which add up to one meter. With somewhat more difficulty, but in a similar way (using a good microscope), we can mark off a thousand equal subdivisions of the millimeter to make a scale of microns (millionths of a meter). It is difficult to continue to smaller scales, because we cannot "see" objects smaller than the wavelength of visible light (about 5×10^{-7} meter).

We need not stop, however, at what we can see. With an electron microscope, we can continue the process by making photographs on a still smaller scale, say down to 10^{-8} meter (Fig. 5-9). By indirect measurements—by a kind of triangulation on a microscopic scale— we can continue to measure to smaller and smaller scales. First, from an observation of the way light of short wavelength (x-radiation) is reflected from a pattern of marks of known separation, we determine the wave-

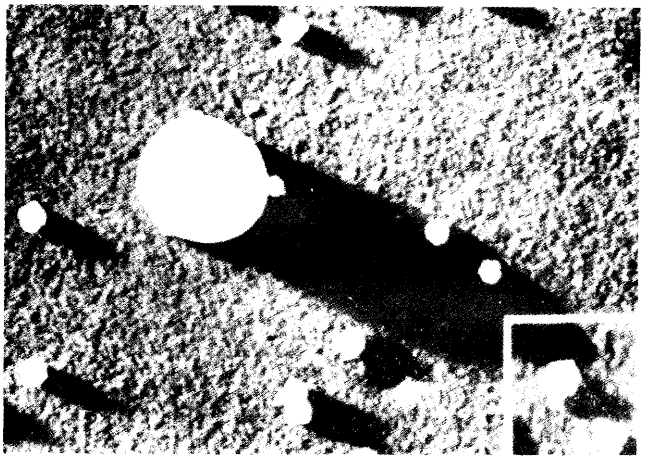


Fig. 5-9. Electron micrograph of some virus molecules. The "large" sphere is for calibration and is known to have a diameter of 2×10^{-7} meter (2000 A).

DISTANCES

LIGHT-YEARS	METERS	
		? ? ? ? ? ? ? ?
	10^{27}	
		Edge of universe
10^9		
	10^{24}	
10^6		To nearest neighbor galaxy
	10^{21}	
		To center of our galaxy
10^3		
	10^{18}	
		To nearest star
1		
	10^{15}	
		Radius of orbit of Pluto
	10^{12}	
		To the sun
	10^9	
		To the moon
	10^6	
		Height of a Sputnik
	10^3	
		Height of a TV antenna tower
	1	Height of a child
	10^{-3}	
		A grain of salt
	10^{-6}	
		A virus
	10^{-9}	
		Radius of an atom
	10^{-12}	
	10^{-15}	Radius of a nucleus
		? ? ? ? ? ? ? ?

length of the light vibrations. Then, from the pattern of the scattering of the same light from a crystal, we can determine the relative location of the atoms in the crystal, obtaining results which agree with the atomic spacings also determined by chemical means. We find in this way that atoms have a diameter of about 10^{-10} meter.

There is a large "gap" in physical sizes between the typical atomic dimension of about 10^{-10} meter and the nuclear dimensions 10^{-15} meter, 10^{-5} times smaller. For nuclear sizes, a different way of measuring size becomes convenient. We measure the *apparent area*, σ, called the effective *cross section.* If we wish the radius, we can obtain it from $\sigma = \pi r^2$, since nuclei are nearly spherical.

Measurement of a nuclear cross section can be made by passing a beam of high-energy particles through a thin slab of material and observing the number of particles which do not get through. These high-energy particles will plow right through the thin cloud of electrons and will be stopped or deflected only if they hit the concentrated weight of a nucleus. Suppose we have a piece of material 1 centimeter thick. There will be about 10^8 atomic layers. But the nuclei are so small that there is little chance that any nucleus will lie behind another. We might *imagine* that a highly magnified view of the situation—looking along the particle beam—would look like Fig. 5-10.

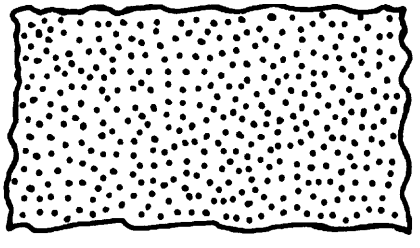


Fig. 5-10. Imagined view through a block of carbon 1 cm thick if only the nuclei were observed.

The chance that a very small particle will hit a nucleus on the trip through is just the total area covered by the profiles of the nuclei divided by the total area in the picture. Suppose that we know that in an area A of our slab of material there are N atoms (each with one nucleus, of course). Then the total area "covered" by the nuclei is $N\sigma/A$. Now let the number of particles of our beam which arrive at the slab be n_1 and the number which come out the other side be n_2. The fraction which do *not* get through is $(n_1 - n_2)/n_1$, which should just equal the

fraction of the area covered. We can obtain the radius of the nucleus from the equation*

$$\pi r^2 = \sigma = \frac{A}{N}\,\frac{n_1 - n_2}{n_1}.$$

From such an experiment we find that the radii of the nuclei are from about 1 to 6 times 10^{-15} meter. The length unit 10^{-15} meter is called *fhe fermi,* in honor of Enrico Fermi (1901-1958).

What do we find if we go to smaller distances? Can we measure smaller distances? Such questions are not yet answerable. It has been suggested that the still unsolved mystery of nuclear forces may be unravelled only by some modification of our idea of space, or measurement, at such small distances.

It might be thought that it would be a good idea to use some natural length as our unit of length—say the radius of the earth or some fraction of it. The meter was originally intended to be such a unit and was defined to be $(\pi/2) \times 10^{-7}$ times the earth's radius. It is neither convenient nor very accurate to determine the unit of length in this way. For a long time it has been agreed internationally that the meter would be defined as the distance between two scratches on a bar kept in a special laboratory in France. More recently, it has been realized that this definition is neither as precise as would be useful, nor as permanent or universal as one would like. It is currently being considered that a new definition be adopted, an agreed-upon (arbitrary) number of wavelengths of a chosen spectral line.

Measurements of distance and of time give results which depend on the observer. Two observers moving with respect to each other will not measure the same distances and times when measuring what appear to be the same things. Distances and time intervals have different magnitudes, depending on the coordinate system (or "frame of reference") used for making the measurements. We shall study this subject in more detail in a later chapter.

Perfectly precise measurements of distances or times are not permitted by the laws of nature. We have mentioned earlier that the errors in a measurement of the position of an object must be at least as large as

$$\Delta x = h/\Delta p,$$

where h is a small quantity called "Planck's constant" and Δp is the error in our knowledge of the momentum (mass times velocity) of the object whose position we are measuring. It was also mentioned that the uncertainty in position measurements is related to the wave nature of particles.

The relativity of space and time implies that time measurements have also a minimum error, given in fact by

$$\Delta t = h/\Delta E,$$

where ΔE is the error in our knowledge of the energy of the process whose time period we are measuring. If we wish to know *more* precisely *when* something happened we must know less about *what* happened, because our knowledge of the energy involved will be less. The time uncertainty is also related to the wave nature of matter.

* This equation is right only if the area covered by the nuclei is a small fraction of the total, i.e., if $(n_1 - n_2)/n_1$ is much less than 1. Otherwise we must make a correction for the fact that some nuclei will be partly obscured by the nuclei in front of them.

6

Probability

"The true logic of this world is in the calculus of probabilities."
—James Clerk Maxwell

6–1 Chance and likelihood
6–2 Fluctuations
6–3 The random walk
6–4 A probability distribution
6–5 The uncertainty principle

6–1 Chance and likelihood

"Chance" is a word which is in common use in everyday living. The radio reports speaking of tomorrow's weather may say: "There is a sixty percent chance of rain." You might say: "There is a small chance that I shall live to be one hundred years old." Scientists also use the word chance. A seismologist may be interested in the question: "What is the chance that there will be an earthquake of a certain size in Southern California next year?" A physicist might ask the question: "What is the chance that a particular geiger counter will register twenty counts in the next ten seconds?" A politician or statesman might be interested in the question: "What is the chance that there will be a nuclear war within the next ten years?" You may be interested in the chance that you will learn something from this chapter.

By *chance*, we mean something like a guess. Why do we make guesses? We make guesses when we wish to make a judgment but have incomplete information or uncertain knowledge. We want to make a guess as to what things are, or what things are likely to happen. Often we wish to make a guess because we have to make a decision. For example: Shall I take my raincoat with me tomorrow? For what earth movement should I design a new building? Shall I build myself a fallout shelter? Shall I change my stand in international negotiations? Shall I go to class today?

Sometimes we make guesses because we wish, with our limited knowledge, to say as much as we *can* about some situation. Really, any generalization is in the nature of a guess. Any physical theory is a kind of guesswork. There are good guesses and there are bad guesses. The theory of probability is a system for making better guesses. The language of probability allows us to speak quantitatively about some situation which may be highly variable, but which does have some consistent average behavior.

Let us consider the flipping of a coin. If the toss—and the coin—are "honest," we have no way of knowing what to expect for the outcome of any particular toss. Yet we would feel that in a large number of tosses there should be about equal numbers of heads and tails. We say: "The probability that a toss will land heads is 0.5."

We speak of probability only for observations that we contemplate being made in the future. *By the "probability" of a particular outcome of an observation we mean our estimate for the most likely fraction of a number of repeated observations that will yield that particular outcome.* If we imagine repeating an observation—such as looking at a freshly tossed coin—N times, and if we call N_A *our estimate* of the most likely number of our observations that will give some specified result A, say the result "heads," then by $P(A)$, the probability of observing A, we mean

$$P(A) = N_A/N. \tag{6.1}$$

Our definition requires several comments. First of all, we may speak of a probability of something happening only if the occurrence is a possible outcome of some *repeatable* observation. It is not clear that it would make any sense to ask: "What is the probability that there is a ghost in that house?"

You may object that no situation is *exactly* repeatable. That is right. Every different observation must at least be at a different time or place. All we can say is that the "repeated" observations should, for our intended purposes, *appear to be equivalent.* We should assume, at least, that each observation was made from an equivalently prepared situation, and especially with the same degree of ignorance at the start. (If we sneak a look at an opponent's hand in a card game, our estimate of our chances of winning are different than if we do not!)

We should emphasize that N and N_A in Eq. (6.1) are *not* intended to represent numbers based on actual observations. N_A is our best *estimate* of what *would* occur in N *imagined* observations. Probability depends, therefore, on our knowledge and on our ability to make estimates. In effect, on our common sense! Fortunately, there is a certain amount of agreement in the common sense of many things, so that different people will make the same estimate. Probabilities need not, however, be "absolute" numbers. Since they depend on our ignorance, they may become different if our knowledge changes.

You may have noticed another rather "subjective" aspect of our definition of probability. We have referred to N_A as "our estimate of the most likely number . . ." We do not mean that we expect to observe *exactly* N_A, but that we expect a number *near* N_A, and that the number N_A is *more likely* than any other number in the vicinity. If we toss a coin, say, 30 times, we should expect that the number of heads would not be very likely to be exactly 15, but rather only some number near to 15, say 12, 13, 14, 15, 16, or 17. However, if we *must* choose, we would decide that 15 heads is *more likely* than any other number. We would write $P(\text{heads}) = 0.5$.

Why did we choose 15 as more likely than any other number? We must have argued with ourselves in the following manner: If the most likely number of heads is N_H in a total number of tosses N, then the most likely number of tails N_T is $(N - N_H)$. (We are assuming that every toss gives *either* heads *or* tails, and no "other" result!) But if the coin is "honest," there is no preference for heads or tails. Until we have some reason to think the coin (or toss) is dishonest, we must give equal likelihoods for heads and tails. So we must set $N_T = N_H$. It follows that $N_T = N_H = N/2$, or $P(H) = p(T) = 0.5$.

We can generalize our reasoning to *any* situation in which there are m different but "equivalent" (that is, equally likely) possible results of an observation. If an observation can yield m different results, and we have reason to believe that any one of them is as likely as any other, then the probability of a *particular* outcome A is $P(A) = 1/m$.

If there are seven different-colored balls in an opaque box and we pick one out "at random" (that is, without looking), the probability of getting a ball of a particular color is $\frac{1}{7}$. The probability that a "blind draw" from a shuffled deck of 52 cards will show the ten of hearts is $\frac{1}{52}$. The probability of throwing a double-one with dice is $\frac{1}{36}$.

In Chapter 5 we described the size of a nucleus in terms of its apparent area, or "cross section." When we did so we were really talking about probabilities. When we shoot a high-energy particle at a thin slab of material, there is some chance that it will pass right through and some chance that it will hit a nucleus. (Since the nucleus is so small that we cannot *see* it, we cannot aim right at a nucleus. We must "shoot blind.") If there are n atoms in our slab and the nucleus of each atom has a cross-sectional area σ, then the total area "shadowed" by the nuclei is $n\sigma$. In a large number N of random shots, we expect that the number of hits N_C of *some* nucleus will be in the ratio to N as the shadowed area is to the total area of the slab:

$$N_C/N = n\sigma/A. \tag{6.2}$$

We may say, therefore, that the *probability* that any one projectile particle will suffer a collision in passing through the slab is

$$P_C = \frac{n}{A}\,\sigma, \tag{6.3}$$

where n/A is the number of atoms per unit area in our slab.

6–2 Fluctuations

We would like now to use our ideas about probability to consider in some greater detail the question: "How many heads do I really *expect* to get if I toss a coin N times?" Before answering the question, however, let us look at what does happen in such an "experiment." Figure 6–1 shows the results obtained in the first three "runs" of such an experiment in which $N = 30$. The sequences of "heads" and "tails" are shown just as they were obtained. The first game gave 11 heads; the second also 11; the third 16. In three trials we did not once get 15 heads. Should we begin to suspect the coin? Or were we wrong in thinking that the most likely number of "heads" in such a game is 15? Ninety-seven more runs were made to obtain a total of 100 experiments of 30 tosses each. The results of the experiment are given in Table 6–1.*

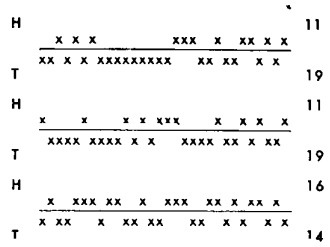


Fig. 6–1. Observed sequences of heads and tails in three games of 30 tosses each.

Table 6–1

Number of heads in successive trials of 30 tosses of a coin.

11	16	17	15	17	16	19	18	15	13
11	17	17	12	20	23	11	16	17	14
16	12	15	10	18	17	13	15	14	15
16	12	11	22	12	20	12	15	16	12
16	10	15	13	14	16	15	16	13	18
14	14	13	16	15	19	21	14	12	15
16	11	16	14	17	14	11	16	17	16
19	15	14	12	18	15	14	21	11	16
17	17	12	13	14	17	9	13	16	13

} 90 trials

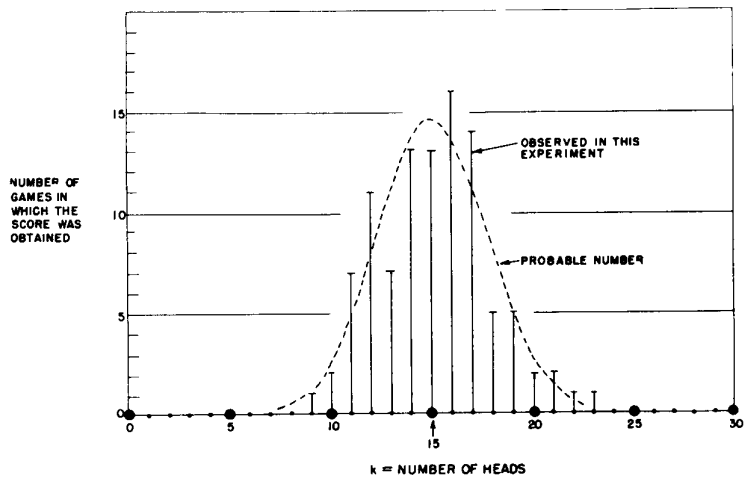


Fig. 6–2. Summary of the results of 100 games of 30 tosses each. The vertical bars show the number of games in which a score of k heads was obtained. The dashed curve shows the expected numbers of games with the score k obtained by a probability computation.

Looking at the numbers in Table 6–1, we see that most of the results are "near" 15, in that they are between 12 and 18. We can get a better feeling for the details of these results if we plot a graph of the *distribution* of the results. We count the number of games in which a score of k was obtained, and plot this number for each k. Such a graph is shown in Fig. 6–2. A score of 15 heads was obtained in 13 games. A score of 14 heads was also obtained 13 times. Scores of 16 and 17 were each obtained *more* than 13 times. Are we to conclude that there is some bias toward heads? Was our "best estimate" not good enough? Should

* After the first three games, the experiment was actually done by shaking 30 pennies violently in a box and then counting the number of heads that showed.

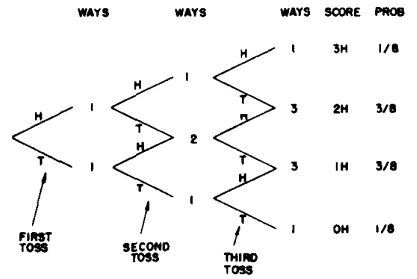


Fig. 6–3. A diagram for showing the number of ways a score of 0, 1, 2, or 3 heads can be obtained in a game of 3 tosses.

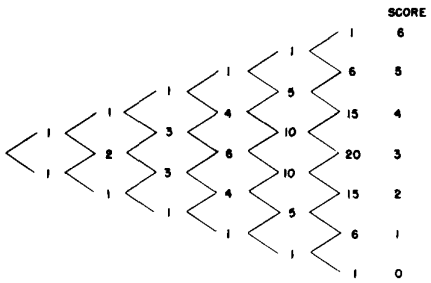


Fig. 6–4. A diagram like that of Fig. 6–3, for a game of 6 tosses.

we conclude now that the "most likely" score for a run of 30 tosses is really 16 heads? But wait! In all the games taken together, there were 3000 tosses. And the total number of heads obtained was 1492. The fraction of tosses that gave heads is 0.497, very nearly, but slightly *less* than half. We should certainly *not* assume that the probability of throwing heads is greater than 0.5! The fact that one *particular* set of observations gave 16 heads most often, is a *fluctuation*. We still expect that the *most likely* number of heads is 15.

We may ask the question: "What *is* the probability that a game of 30 tosses will yield 15 heads—or 16, or any other number?" We have said that in a game of one toss, the probability of obtaining *one* head is 0.5, and the probability of obtaining no head is 0.5. In a game of two tosses there are *four* possible outcomes: *HH*, *HT*, *TH*, *TT*. Since each of these sequences is equally likely, we conclude that (a) the probability of a score of two heads is $\frac{1}{4}$, (b) the probability of a score of one head is $\frac{2}{4}$, (c) the probability of a zero score is $\frac{1}{4}$. There are *two* ways of obtaining one head, but only one of obtaining either zero or two heads.

Consider now a game of 3 tosses. The third toss is equally likely to be heads or tails. There is only one way to obtain 3 heads: we *must* have obtained 2 heads on the first two tosses, and then heads on the last. There are, however, *three* ways of obtaining 2 heads. We could throw tails after having thrown two heads (one way) or we could throw heads after throwing only one head in the first two tosses (two ways). So for scores of 3-*H*, 2-*H*, 1-*H*, 0-*H* we have that the number of equally likely ways is 1, 3, 3, 1, with a total of 8 different possible sequences. The probabilities are $\frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8}$.

The argument we have been making can be summarized by a diagram like that in Fig. 6–3. It is clear how the diagram should be continued for games with a larger number of tosses. Figure 6–4 shows such a diagram for a game of 6 tosses. The number of "ways" to any point on the diagram is just the number of different "paths" (sequences of heads and tails) which can be taken from the starting point. The vertical position gives us the total number of heads thrown. The set of numbers which appears in such a diagram is known as *Pascal's triangle*. The numbers are also known as the *binomial coefficients*, because they also appear in the expansion of $(a + b)^n$. If we call n the number of tosses and k the number of heads thrown, then the numbers in the diagram are usually designated by the symbol $\binom{n}{k}$. We may remark in passing that the binomial coefficients can also be computed from

$$\binom{n}{k} = \frac{n!}{k!(n-k)!}, \tag{6.4}$$

where $n!$, called "n-factorial," represents the product $(n)(n-1)(n-2)\ldots(3)(2)(1)$.

We are now ready to compute the probability $P(k, n)$ of throwing k heads in n tosses, using our definition Eq. (6.1). The total number of possible sequences is 2^n (since there are 2 outcomes for each toss), and the number of ways of obtaining k heads is $\binom{n}{k}$, all equally likely, so we have

$$P(k, n) = \frac{\binom{n}{k}}{2^n}. \tag{6.5}$$

Since $P(k, n)$ is the fraction of games which we expect to yield k heads, then in 100 games we should expect to find k heads $100 \cdot P(k, n)$ times. The dashed curve in Fig. 6–2 passes through the points computed from $100 \cdot P(k, 30)$. We see that we *expect* to obtain a score of 15 heads in 14 or 15 games, whereas this score was observed in 13 games. We *expect* a score of 16 in 13 or 14 games, but we obtained that score in 16 games. Such fluctuations are "part of the game."

The method we have just used can be applied to the most general situation in which there are only two possible outcomes of a single observation. Let us designate the two outcomes by W (for "win") and L (for "lose"). In the general case, the probability of W or L in a single event need not be equal. Let p be the probability of obtaining the result W. Then q, the probability of L, is necessarily

$(1 - p)$. In a set of n trials, the probability $P(k, n)$ that W will be obtained k times is

$$P(k, n) = \binom{n}{k} p^k q^{n-k}. \tag{6.6}$$

This probability function is called the *Bernoulli* or, also, the *binomial* probability.

6–3 The random walk

There is another interesting problem in which the idea of probability is required. It is the problem of the "random walk." In its simplest version, we imagine a "game" in which a "player" starts at the point $x = 0$ and at each "move" is required to take a step *either* forward (toward $+x$) *or* backward (toward $-x$). The choice is to be made *randomly*, determined, for example, by the toss of a coin. How shall we describe the resulting motion? In its general form the problem is related to the motion of atoms (or other particles) in a gas—called Brownian motion—and also to the combination of errors in measurements. You will see that the random-walk problem is closely related to the coin-tossing problem we have already discussed.

First, let us look at a few examples of a random walk. We may characterize the walker's progress by the net distance D_N traveled in N steps. We show in the graph of Fig. 6–5 three examples of the path of a random walker. (We have used for the random sequence of choices the results of the coin tosses shown in Fig. 6–1.)

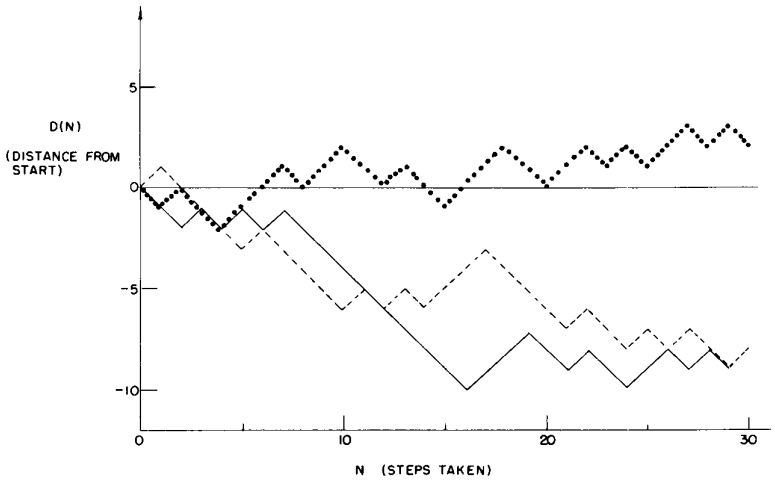


Fig. 6–5. The progress made in a random walk. The horizontal coordinate N is the total number of steps taken; the vertical coordinate $D(N)$ is the net distance moved from the starting position.

What can we say about such a motion? We might first ask: "How far does he get on the average?" We must *expect* that his average progress will be zero, since he is equally likely to go either forward or backward. But we have the feeling that as N increases, he is more likely to have strayed farther from the starting point. We might, therefore, ask what is his average distance travelled in *absolute value*, that is, what is the average of $|D|$. It is, however, more convenient to deal with another measure of "progress," the square of the distance: D^2 is positive for either positive or negative motion, and is therefore a reasonable *measure* of such random wandering.

We can show that the expected value of D_N^2 is just N, the number of steps taken. By "expected value" we mean the probable value (our best guess), which we can think of as the *expected* average behavior in *many repeated* sequences. We represent such an expected value by $\langle D_N^2 \rangle$, and may refer to it also as the "mean square distance." After one step, D^2 is always $+1$, so we have certainly $\langle D_1^2 \rangle = 1$. (All distances will be measured in terms of a unit of one step. We shall not continue to write the units of distance.)

The expected value of D_N^2 for $N > 1$ can be obtained from D_{N-1}. If, after $(N - 1)$ steps, we have D_{N-1}, then after N steps we have $D_N = D_{N-1} + 1$ *or* $D_N = D_{N-1} - 1$. For the squares,

$$D_N^2 = \begin{cases} D_{N-1}^2 + 2D_{N-1} + 1, \\ \quad or \\ D_{N-1}^2 - 2D_{N-1} + 1. \end{cases} \tag{6.7}$$

In a number of independent sequences, we expect to obtain each value one-half of the time, so our average expectation is just the average of the two possible values. The expected value of D_N is then $D_{N-1}^2 + 1$. *In general,* we should *expect* for D_{N-1}^2 its "expected value" $\langle D_{N-1}^2 \rangle$ (by definition!). So

$$\langle D_N^2 \rangle = \langle D_{N-1}^2 \rangle + 1. \tag{6.8}$$

We have already shown that $\langle D_1^2 \rangle = 1$; it follows then that

$$D_N^2 = N, \tag{6.9}$$

a particularly simple result!

If we wish a number like a distance, rather than a distance squared, to represent the "progress made away from the origin" in a random walk, we can use the "root-mean-square distance" D_{rms}:

$$D_{\text{rms}} = \sqrt{\langle D^2 \rangle} = \sqrt{N}. \tag{6.10}$$

We have pointed out that the random walk is closely similar in its mathematics to the coin-tossing game we considered at the beginning of the chapter. If we imagine the direction of each step to be in correspondence with the appearance of heads or tails in a coin toss, then D is just $N_H - N_T$, the difference in the number of heads and tails. Since $N_H + N_T = N$, the total number of steps (and tosses), we have $D = 2N_H - N$. We have derived earlier an expression for the expected distribution of N_H (also called k) and obtained the result of Eq. (6.5). Since N is just a constant, we have the corresponding distribution for D. (Since for every head more than $N/2$ there is a tail "missing," we have the factor of 2 between N_H and D.) The graphs of Fig. 6–2 represent the distribution of distances we might get in 30 random steps (where $k = 15$ is to be read $D = 0$; $k = 16$, $D = 2$; etc.).

The variation of N_H from its expected value $N/2$ is

$$N_H - \frac{N}{2} = \frac{D}{2}. \tag{6.11}$$

The rms deviation is

$$\left(N_H - \frac{N}{2}\right)_{\text{rms}} = \tfrac{1}{2}\sqrt{N}. \tag{6.12}$$

According to our result for D_{rms}, we expect that the "typical" distance in 30 steps ought to be $\sqrt{30} = 5.5$, or a typical k should be about $5.5/2 = 2.8$ units from 15. We see that the "width" of the curve in Fig. 6–2, measured from the center, is just about 3 units, in agreement with this result.

We are now in a position to consider a question we have avoided until now. How shall we tell whether a coin is "honest" or "loaded"? We can give now at least a partial answer. For an honest coin, we expect the fraction of the times heads appears to be 0.5, that is,

$$\frac{\langle N_H \rangle}{N} = 0.5. \tag{6.13}$$

We *also* expect an actual N_H to deviate from $N/2$ by about $\sqrt{N}/2$, or the *fraction* to deviate by

$$\frac{1}{N}\frac{\sqrt{N}}{2} = \frac{1}{2\sqrt{N}}.$$

The larger N is, the closer we *expect* the fraction N_H/N to be to one-half.

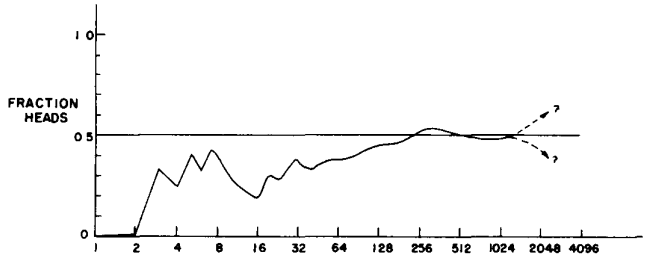


Fig. 6–6. The fraction of the tosses that gave heads in a particular sequence of N tosses of a penny.

In Fig. 6–6 we have plotted the fraction N_H/N for the coin tosses reported earlier in this chapter. We see the tendency for the fraction of heads to approach 0.5 for large N. Unfortunately, for any given run or combination of runs there is no *guarantee* that the observed deviation will be even *near* the *expected* deviation. There is always the finite chance that a large fluctuation—a long string of heads or tails—will give an arbitrarily large deviation. All we can say is that *if* the deviation is near the expected $1/2\sqrt{N}$ (say within a factor of 2 or 3), we have no reason to suspect the honesty of the coin. If it is much larger, we may be suspicious, but cannot prove, that the coin is loaded (or that the tosser is clever!).

We have also not considered how we should treat the case of a "coin" or some similar "chancy" object (say a stone that always lands in either of two positions) that we have good reason to believe *should* have a different probability for heads and tails. We have defined $P(H) = \langle N_H \rangle/N$. How shall we know what to *expect* for N_H? In some cases, the best we can do is to observe the number of heads obtained in large numbers of tosses. For want of anything better, we must set $\langle N_H \rangle = N_H$(observed). (How could we expect anything else?) We must understand, however, that in such a case a different experiment, or a different observer, might conclude that $P(H)$ was different. We would *expect*, however, that the various answers should agree within the deviation $1/2\sqrt{N}$ [if $P(H)$ is near one-half]. An experimental physicist usually says that an "experimentally determined" probability has an "error," and writes

$$P(H) = \frac{N_H}{N} \pm \frac{1}{2\sqrt{N}}. \tag{6.14}$$

There is an implication in such an expression that there *is* a "true" or "correct" probability which *could* be computed if we knew enough, and that the observation may be in "error" due to a fluctuation. There is, however, no way to make such thinking logically consistent. It is probably better to realize that the probability concept is in a sense subjective, that it is always based on uncertain knowledge, and that its quantitative evaluation is subject to change as we obtain more information.

6–4 A probability distribution

Let us return now to the random walk and consider a modification of it. Suppose that in addition to a random choice of the *direction* (+ or −) of each step, the *length* of each step also varied in some unpredictable way, the only condition being that *on the average* the step length was one unit. This case is more representative of something like the thermal motion of a molecule in a gas. If we call the length of a step S, then S may have any value at all, but most often will be "near" 1. To be specific, we shall let $\langle S^2 \rangle = 1$ or, equivalently, $S_{\text{rms}} = 1$. Our derivation for $\langle D^2 \rangle$ would proceed as before except that Eq. (6.8) would be changed now to read

$$\langle D_N^2 \rangle = \langle D_{N-1}^2 \rangle + \langle S^2 \rangle = \langle D_{N-1}^2 \rangle + 1. \tag{6.15}$$

We have, as before, that

$$\langle D_N^2 \rangle = N. \tag{6.16}$$

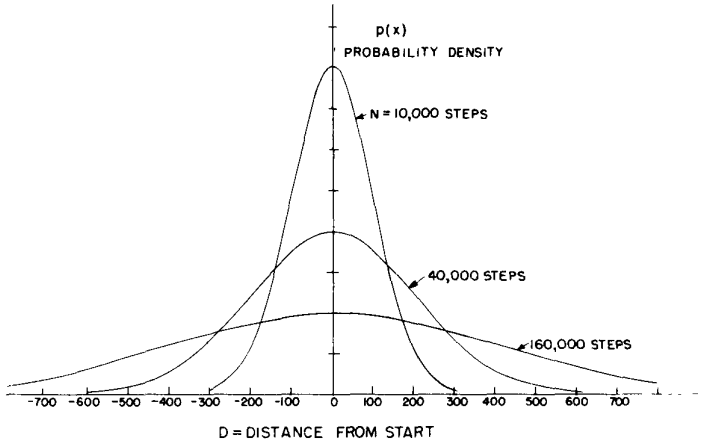


Fig. 6–7. The probability density for ending up at the distance D from the starting place in a random walk of N steps. (D is measured in units of the rms step length.)

What would we expect now for the distribution of distances D? What is, for example, the probability that $D = 0$ after 30 steps? The answer is zero! The probability is zero that D will be *any particular* value, since there is no chance at all that the sum of the backward steps (of varying lengths) would exactly equal the sum of forward steps. We cannot plot a graph like that of Fig. 6–2.

We can, however, obtain a representation similar to that of Fig. 6–2, if we ask, not what is the probability of obtaining D exactly equal to 0, 1, or 2, but instead what is the probability of obtaining D *near* 0, 1, or 2. Let us define $P(x, \Delta x)$ as the probability that D will lie in the interval Δx located at x (say from x to $x + \Delta x$). We expect that for small Δx the chance of D landing in the interval is proportional to Δx, the width of the interval. So we can write

$$P(x, \Delta x) = p(x)\,\Delta x. \tag{6.17}$$

The function $p(x)$ is called the *probability density*.

The form of $p(x)$ will depend on N, the number of steps taken, and also on the distribution of individual step lengths. We cannot demonstrate the proofs here, but for large N, $p(x)$ is the *same* for all reasonable distributions in individual step lengths, and depends only on N. We plot $p(x)$ for three values of N in Fig. 6–7. You will notice that the "half-widths" (typical spread from $x = 0$) of these curves is $\sqrt{N}$, as we have shown it should be.

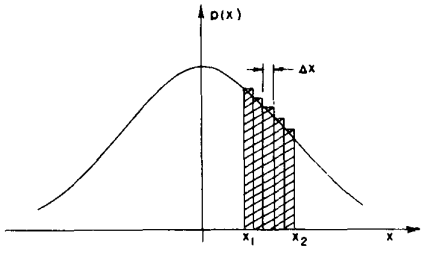


Fig. 6–8. The probability that the distance D traveled in a random walk is between x_1 and x_2 is the area under the curve of $p(x)$ from x_1 to x_2.

You may notice also that the value of $p(x)$ near zero is inversely proportional to $\sqrt{N}$. This comes about because the curves are all of a similar shape and their areas under the curves must all be equal. Since $p(x)\,\Delta x$ is the probability of finding D in Δx when Δx is small, we can determine the chance of finding D *somewhere* inside an arbitrary interval from x_1 to x_2, by cutting the interval in a number of small increments Δx and evaluating the sum of the terms $p(x)\,\Delta x$ for each increment. The probability that D lands somewhere between x_1 and x_2, which we may write $P(x_1 < D < x_2)$, is equal to the shaded area in Fig. 6–8. The smaller we take the increments Δx, the more correct is our result. We can write, therefore,

$$P(x_1 < D < x_2) = \sum p(x)\,\Delta x = \int_{x_1}^{x_2} p(x)\,dx. \tag{6.18}$$

The area under the whole curve is the probability that D lands somewhere (that is, has *some* value between $x = -\infty$ and $x = +\infty$). That probability is surely 1. We must have that

$$\int_{-\infty}^{+\infty} p(x)\,dx = 1. \tag{6.19}$$

Since the curves in Fig. 6–7 get wider in proportion to $\sqrt{N}$, their heights must be proportional to $1/\sqrt{N}$ to maintain the total area equal to 1.

The probability density function we have been describing is one that is encountered most commonly. It is known as the *normal* or *gaussian* probability density. It has the mathematical form

$$p(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-x^2/2\sigma^2}, \tag{6.20}$$

where σ is called the *standard deviation* and is given, in our case, by $\sigma = \sqrt{N}$ or, if the rms step size is different from 1, by $\sigma = \sqrt{N} S_{\text{rms}}$.

We remarked earlier that the motion of a molecule, or of any particle, in a gas is like a random walk. Suppose we open a bottle of an organic compound and let some of its vapor escape into the air. If there are air currents, so that the air is circulating, the currents will also carry the vapor with them. But even in *perfectly still air*, the vapor will gradually spread out—will diffuse—until it has penetrated throughout the room. We might detect it by its color or odor. The individual molecules of the organic vapor spread out in still air because of the molecular motions caused by collisions with other molecules. If we know the average "step" size, and the number of steps taken per second, we can find the probability that one, or several, molecules will be found at some distance from their starting point after any particular passage of time. As time passes, more steps are taken and the gas spreads out as in the successive curves of Fig. 6–7. In a later chapter, we shall find out how the step sizes and step frequencies are related to the temperature and pressure of a gas.

Earlier, we said that the pressure of a gas is due to the molecules bouncing against the walls of the container. When we come later to make a more quantitative description, we will wish to know how fast the molecules are going when they bounce, since the impact they make will depend on that speed. We cannot, however, speak of *the* speed of the molecules. It is necessary to use a probability description. A molecule may have any speed, but some speeds are more likely than others. We describe what is going on by saying that the probability that any particular molecule will have a speed between v and $v + \Delta v$ is $p(v)\,\Delta v$, where $p(v)$, a probability density, is a given function of the speed v. We shall see later how Maxwell, using common sense and the ideas of probability, was able to find a mathematical expression for $p(v)$. The form* of the function $p(v)$ is shown in Fig. 6–9. Velocities may have any value, but are most likely to be near the most probable or expected value $\langle v \rangle$.

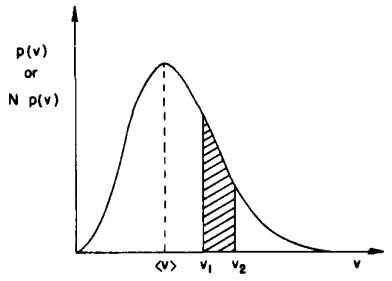


Fig. 6–9. The distribution of velocities of the molecules in a gas.

We often think of the curve of Fig. 6–9 in a somewhat different way. If we consider the molecules in a typical container (with a volume of, say, one liter), then there are a very large number N of molecules present ($N \approx 10^{22}$). Since $p(v)\,\Delta v$ is the probability that *one* molecule will have its velocity in Δv, by our definition of probability we mean that the *expected* number $\langle \Delta N \rangle$ to be found with a velocity in the interval Δv is given by

$$\langle \Delta N \rangle = N p(v)\,\Delta v. \tag{6.21}$$

We call $Np(v)$ the "distribution in velocity." The area under the curve between two velocities v_1 and v_2, for example the shaded area in Fig. 6–9, represents [for the curve $Np(v)$] the expected number of molecules with velocities between v_1 and v_2. Since with a gas we are usually dealing with large numbers of molecules, we expect the deviations from the expected numbers to be small (like $1/\sqrt{N}$), so we often neglect to say the "expected" number, and say instead: "The number of molecules with velocities between v_1 and v_2 *is* the area under the curve." We should remember, however, that such statements are always about *probable* numbers.

---

* Maxwell's expression is $p(v) = Cv^2 e^{-av^2}$, where a is a constant related to the temperature and C is chosen so that the total probability is one.

6–5 The uncertainty principle

The ideas of probability are certainly useful in describing the behavior of the 10^{22} or so molecules in a sample of a gas, for it is clearly impractical even to attempt to write down the position or velocity of each molecule. When probability was first applied to such problems, it was considered to be a *convenience*—a way of dealing with very complex situations. We now believe that the ideas of probability are *essential* to a description of atomic happenings. According to quantum mechanics, the mathematical theory of particles, there is always some uncertainty in the *specification* of positions and velocities. We can, at best, say that there is a certain probability that any particle will have a position near some coordinate x.

We can give a probability density $p_1(x)$, such that $p_1(x)\,\Delta x$ is the probability that the particle will be found between x and $x + \Delta x$. If the particle is reasonably well localized, say near x_1, the function $p_1(x)$ might be given by the graph of Fig. 6–10(a). Similarly, we must specify the velocity of the particle by means of a probability density $p_2(v)$, with $p_2(v)\,\Delta v$ the probability that the velocity will be found between v and $v + \Delta v$.

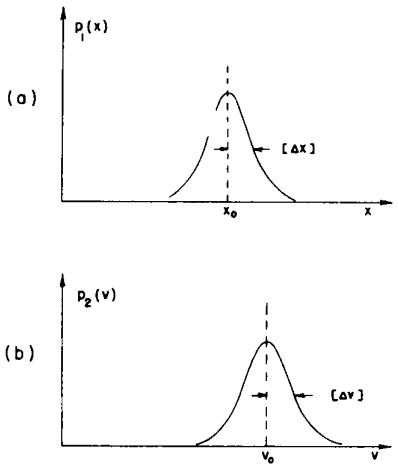


Fig. 6–10. Probability densities for observation of the position and velocity of a particle.

It is one of the fundamental results of quantum mechanics that the two functions $p_1(x)$ and $p_2(v)$ cannot be chosen independently and, in particular, cannot both be made arbitrarily narrow. If we call the typical "width" of the $p_1(x)$ curve $[\Delta x]$, and that of the $p_2(v)$ curve $[\Delta v]$ (as shown in the figure), nature demands that the *product* of the two widths be at least as big as the number h/m, where m is the mass of the particle and h is a fundamental physical constant called *Planck's constant*. We may write this basic relationship as

$$[\Delta x] \cdot [\Delta v] \geq h/m. \tag{6.22}$$

This equation is a statement of the *Heisenberg uncertainty principle* that we mentioned earlier.

Since the right-hand side of Eq. (6.22) is a constant, this equation says that if we try to "pin down" a particle by forcing it to be at a particular place, it ends up by having a high speed. Or if we try to force it to go very slowly, or at a precise velocity, it "spreads out" so that we do not know very well just where it is. Particles behave in a funny way!

The uncertainty principle describes an inherent fuzziness that must exist in any attempt to describe nature. Our most precise description of nature *must* be in terms of *probabilities*. There are some people who do not like this way of describing nature. They feel somehow that if they could only tell what is *really* going on with a particle, they could know its speed and position simultaneously. In the early days of the development of quantum mechanics, Einstein was quite worried about this problem. He used to shake his head and say, "But, surely God does not throw dice in determining how electrons should go!" He worried about that problem for a long time and he probably never really reconciled himself to the fact that this is the best description of nature that one can give. There are still one or two physicists who are working on the problem who have an intuitive conviction that it is possible somehow to describe the world in a different way and that all of this uncertainty about the way things are can be removed. No one has yet been successful.

The necessary uncertainty in our specification of the position of a particle becomes most important when we wish to describe the structure of atoms. In the hydrogen atom, which has a nucleus of one proton with one electron outside of the nucleus, the uncertainty in the position of the electron is as large as the atom itself! We cannot, therefore, properly speak of the electron moving in some "orbit" around the proton. The most we can say is that there is a certain *chance* $p(r)\,\Delta V$, of observing the electron in an element of volume ΔV at the distance r from the proton. The probability density $p(r)$ is given by quantum mechanics. For an undisturbed hydrogen atom $p(r) = Ae^{-r^2/a^2}$, which is a bell-shaped function like that in Fig. 6–8. The number a is the "typical" radius, where the function is decreasing rapidly. Since there is a small probability of finding the electron at distances



Fig. 6–11. A way of visualizing a hydrogen atom. The density (whiteness) of the cloud represents the probability density for observing the electron.

from the nucleus much greater than a, we may think of a as "the radius of the atom," about 10^{-10} meter.

We can form an image of the hydrogen atom by imagining a "cloud" whose density is proportional to the probability density for observing the electron. A sample of such a cloud is shown in Fig. 6–11. Thus our best "picture" of a hydrogen atom is a nucleus surrounded by an "electron cloud" (although we *really* mean a "probability cloud"). The electron is there somewhere, but nature permits us to know only the *chance* of finding it at any particular place.

In its efforts to learn as much as possible about nature, modern physics has found that certain things can never be "known" with certainty. Much of our knowledge must always remain uncertain. The *most* we can know is in terms of probabilities.

7

The Theory of Gravitation

7–1 Planetary motions

7–2 Kepler's laws

7–3 Development of dynamics

7–4 Newton's law of gravitation

7–5 Universal gravitation

7–6 Cavendish's experiment

7–7 What is gravity?

7–8 Gravity and relativity

7–1 Planetary motions

In this chapter we shall discuss one of the most far-reaching generalizations of the human mind. While we are admiring the human mind, we should take some time off to stand in awe of a *nature* that could follow with such completeness and generality such an elegantly simple principle as the law of gravitation. What is this law of gravitation? It is that every object in the universe attracts every other object with a force which for any two bodies is proportional to the mass of each and varies inversely as the square of the distance between them. This statement can be expressed mathematically by the equation

$$F = G\frac{mm'}{r^2}.$$

If to this we add the fact that an object responds to a force by accelerating in the direction of the force by an amount that is inversely proportional to the mass of the object, we shall have said everything required, for a sufficiently talented mathematician could then deduce all the consequences of these two principles. However, since you are not assumed to be sufficiently talented yet, we shall discuss the consequences in more detail, and not just leave you with only these two bare principles. We shall briefly relate the story of the discovery of the law of gravitation and discuss some of its consequences, its effects on history, the mysteries that such a law entails, and some refinements of the law made by Einstein; we shall also discuss the relationships of the law to the other laws of physics. All this cannot be done in one chapter, but these subjects will be treated in due time in subsequent chapters.

The story begins with the ancients observing the motions of planets among the stars, and finally deducing that they went around the sun, a fact that was rediscovered later by Copernicus. Exactly *how* the planets went around the sun, with exactly *what motion*, took a little more work to discover. In the beginning of the fifteenth century there were great debates as to whether they really went around the sun or not. Tycho Brahe had an idea that was different from anything proposed by the ancients: his idea was that these debates about the nature of the motions of the planets would best be resolved if the actual positions of the planets in the sky were measured sufficiently accurately. If measurement showed exactly how the planets moved, then perhaps it would be possible to establish one or another viewpoint. This was a tremendous idea—that to find something out, it is better to perform some careful experiments than to carry on deep philosophical arguments. Pursuing this idea, Tycho Brahe studied the positions of the planets for many years in his observatory on the island of Hven, near Copenhagen. He made voluminous tables, which were then studied by the mathematician Kepler, after Tycho's death. Kepler discovered from the data some very beautiful and remarkable, but simple, laws regarding planetary motion.

7–2 Kepler's laws

First of all, Kepler found that each planet goes around the sun in a curve called an *ellipse*, with the sun at a focus of the ellipse. An ellipse is not just an oval, but is a very specific and precise curve that can be obtained by using two tacks, one at each focus, a loop of string, and a pencil; more mathematically, it

is the locus of all points the sum of whose distances from two fixed points (the foci) is a constant. Or, if you will, it is a foreshortened circle (Fig. 7–1).

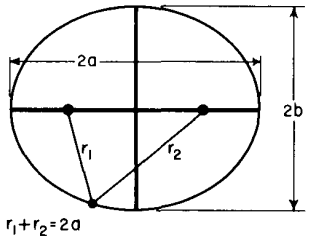


Fig. 7–1. An ellipse.

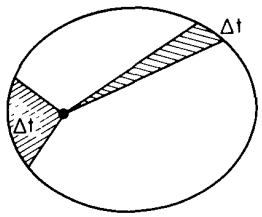


Fig. 7–2. Kepler's law of areas.

Kepler's second observation was that the planets do not go around the sun at a uniform speed, but move faster when they are nearer the sun and more slowly when they are farther from the sun, in precisely this way: Suppose a planet is observed at any two successive times, let us say a week apart, and that the radius vector* is drawn to the planet for each observed position. The orbital arc traversed by the planet during the week, and the two radius vectors, bound a certain plane area, the shaded area shown in Fig. 7–2. If two similar observations are made a week apart, at a part of the orbit farther from the sun (where the planet moves more slowly), the similarly bounded area is exactly the same as in the first case. So, in accordance with the second law, the orbital speed of each planet is such that the radius "sweeps out" equal areas in equal times.

Finally, a third law was discovered by Kepler much later; this law is of a different category from the other two, because it deals not with only a single planet, but relates one planet to another. This law says that when the orbital period and orbit size of any two planets are compared, the periods are proportional to the 3/2 power of the orbit size. In this statement the period is the time interval it takes a planet to go completely around its orbit, and the size is measured by the length of the greatest diameter of the elliptical orbit, technically known as the major axis. More simply, if the planets went in circles, as they nearly do, the time required to go around the circle would be proportional to the 3/2 power of the diameter (or radius). Thus Kepler's three laws are:

I. Each planet moves around the sun in an ellipse, with the sun at one focus.

II. The radius vector from the sun to the planet sweeps out equal areas in equal intervals of time.

III. The squares of the periods of any two planets are proportional to the cubes of the semimajor axes of their respective orbits: $T \sim a^{3/2}$.

7–3 Development of dynamics

While Kepler was discovering these laws, Galileo was studying the laws of motion. The problem was, what makes the planets go around? (In those days, one of the theories proposed was that the planets went around because behind them were invisible angels, beating their wings and driving the planets forward. You will see that this theory is now modified! It turns out that in order to keep the planets going around, the invisible angels must fly in a different direction and they have no wings. Otherwise, it is a somewhat similar theory!) Galileo discovered a very remarkable fact about motion, which was essential for understanding these laws. That is the principle of *inertia*—if something is moving, with nothing touching it and completely undisturbed, it will go on forever, coasting at a uniform speed in a straight line. (*Why* does it keep on coasting? We do not know, but that is the way it is.)

Newton modified this idea, saying that the only way to change the motion of a body is to use *force*. If the body speeds up, a force has been applied *in the direction of motion.* On the other hand, if its motion is changed to a new *direction*, a force has been applied *sideways.* Newton thus added the idea that a force is needed to change the speed *or the direction* of motion of a body. For example, if a stone is attached to a string and is whirling around in a circle, it takes a force to keep it in the circle. We have to *pull* on the string. In fact, the law is that the acceleration produced by the force is inversely proportional to the mass, or the force is proportional to the mass times the acceleration. The more massive a thing is, the stronger the force required to produce a given acceleration. (The mass can be measured by putting other stones on the end of the same string and making them go around the same circle at the same speed. In this way it is found that more or less force is required, the more massive object requiring more force.)

* A radius vector is a line drawn from the sun to any point in a planet's orbit.

The brilliant idea resulting from these considerations is that no *tangential* force is needed to keep a planet in its orbit (the angels do not have to fly tangentially) because the planet would coast in that direction anyway. If there were nothing at all to disturb it, the planet would go off in a *straight line*. But the actual motion deviates from the line on which the body would have gone if there were no force, the deviation being essentially *at right angles* to the motion, not in the direction of the motion. In other words, because of the principle of inertia, the force needed to control the motion of a planet around the sun is not a force *around* the sun but *toward* the sun. (If there is a force toward the sun, the sun might be the angel, of course!)

7–4 Newton's law of gravitation

From his better understanding of the theory of motion, Newton appreciated that *the sun* could be the seat or organization of forces that govern the motion of the planets. Newton proved to himself (and perhaps we shall be able to prove it soon) that the very fact that equal areas are swept out in equal times is a precise sign post of the proposition that all deviations are precisely *radial*—that the law of areas is a direct consequence of the idea that all of the forces are directed exactly *toward the sun.*

Next, by analyzing Kepler's third law it is possible to show that the farther away the planet, the weaker the forces. If two planets at different distances from the sun are compared, the analysis shows that the forces are inversely proportional to the squares of the respective distances. With the combination of the two laws, Newton concluded that there must be a force, inversely as the square of the distance, directed in a line between the two objects.

Being a man of considerable feeling for generalities, Newton supposed, of course, that this relationship applied more generally than just to the sun holding the planets. It was already known, for example, that the planet Jupiter had moons going around it as the moon of the earth goes around the earth, and Newton felt certain that each planet held its moons with a force. He already knew of the force holding *us* on the earth, so he proposed that this was a *universal* force—*that everything pulls everything else.*

The next problem was whether the pull of the earth on its people was the "same" as its pull on the moon, i.e., inversely as the square of the distance. If an object on the surface of the earth falls 16 feet in the first second after it is released from rest, how far does the moon fall in the same time? We might say that the moon does not fall at all. But if there were no force on the moon, it would go off in a straight line, whereas it goes in a circle instead, so it really *falls in* from where it would have been if there were no force at all. We can calculate from the radius of the moon's orbit (which is about 240,000 miles) and how long it takes to go around the earth (approximately 29 days), how far the moon moves in its orbit in 1 second, and can then calculate how far it falls in one second.* This distance turns out to be roughly 1/20 of an inch in a second. That fits very well with the inverse square law, because the earth's radius is 4000 miles, and if something which is 4000 miles from the center of the earth falls 16 feet in a second, something 240,000 miles, or 60 times as far away, should fall only 1/3600 of 16 feet, which also is roughly 1/20 of an inch. Wishing to put this theory of gravitation to a test by similar calculations, Newton made his calculations very carefully and found a discrepancy so large that he regarded the theory as contradicted by facts, and did not publish his results. Six years later a new measurement of the size of the earth showed that the astronomers had been using an incorrect distance to the moon. When Newton heard of this, he made the calculation again, with the corrected figures, and obtained beautiful agreement.

This idea that the moon "falls" is somewhat confusing, because, as you see, it does not come any *closer.* The idea is sufficiently interesting to merit further

---

* That is, how far the circle of the moon's orbit falls below the straight line tangent to it at the point where the moon was one second before.

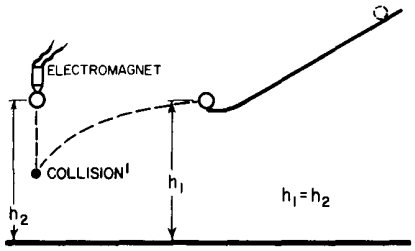


Fig. 7–3. Apparatus for showing the independence of vertical and horizontal motions.

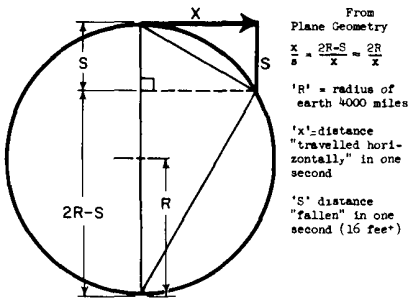


Fig. 7–4. Acceleration toward the center of a circular path. From plane geometry, $x/s = (2R - S)/x \approx 2R/x$, where R is the radius of the earth, 4000 miles; x is the distance "travelled horizontally" in one second; and S is the distance "fallen" in one second (16 feet).

explanation: the moon falls in the sense that *it falls away from the straight line that it would pursue if there were no forces.* Let us take an example on the surface of the earth. An object released near the earth's surface will fall 16 feet in the first second. An object shot out *horizontally* will also fall 16 feet; even though it is moving horizontally, it still falls the same 16 feet in the same time. Figure 7–3 shows an apparatus which demonstrates this. On the horizontal track is a ball which is going to be driven forward a little distance away. At the same height is a ball which is going to fall vertically, and there is an electrical switch arranged so that at the moment the first ball leaves the track, the second ball is released. That they come to the same depth at the same time is witnessed by the fact that they collide in midair. An object like a bullet, shot horizontally, might go a long way in one second—perhaps 2000 feet—but it will still fall 16 feet if it is aimed horizontally. What happens if we shoot a bullet faster and faster? Do not forget that the earth's surface is curved. If we shoot it fast enough, then when it falls 16 feet it may be at just the same height above the ground as it was before. How can that be? It still falls, but the earth curves away, so it falls "around" the earth. The question is, how far does it have to go in one second so that the earth is 16 feet below the horizon? In Fig. 7–4 we see the earth with its 4000-mile radius, and the tangential, straightline path that the bullet would take if there were no force. Now, if we use one of those wonderful theorems in geometry, which says that our tangent is the mean proportional between the two parts of the diameter cut by an equal chord, we see that the horizontal distance travelled is the mean proportional between the 16 feet fallen and the 8000-mile diameter of the earth. The square root of (16/5280) $\times$ 8000 comes out very close to 5 miles. Thus we see that if the bullet moves at 5 miles a second, it then will continue to fall toward the earth at the same rate of 16 feet each second, but will never get any closer because the earth keeps curving away from it. Thus it was that Mr. Gagarin maintained himself in space while going 25,000 miles around the earth at approximately 5 miles per second. (He took a little longer because he was a little higher.)

Any great discovery of a new law is useful only if we can take more out than we put in. Now, Newton *used* the second and third of Kepler's laws to deduce his law of gravitation. What did he *predict?* First, his analysis of the moon's motion was a prediction because it connected the falling of objects on the earth's surface with that of the moon. Second, the question is, *is the orbit an ellipse?* We shall see in a later chapter how it is possible to calculate the motion exactly, and indeed one can prove that it should be an ellipse,* so no extra fact is needed to explain Kepler's *first* law. Thus Newton made his first powerful prediction.

The law of gravitation explains many phenomena not previously understood. For example, the pull of the moon on the earth causes the tides, hitherto mysterious. The moon pulls the water up under it and makes the tides—people had thought of that before, but they were not as clever as Newton, and so they thought there ought to be only one tide during the day. The reasoning was that the moon pulls the water up under it, making a high tide and a low tide, and since the earth spins underneath, that makes the tide at one station go up and down every 24 hours. Actually the tide goes up and down in 12 hours. Another school of thought claimed that the high tide should be on the other side of the earth because, so they argued, the moon pulls the earth away from the water! Both of these theories are wrong. It actually works like this: the pull of the moon for the earth and for the water is "balanced" at the center. But the water which is closer to the moon is pulled *more* than the average and the water which is farther away from it is pulled *less* than the average. Furthermore, the water can flow while the more rigid earth cannot. The true picture is a combination of these two things.

What do we mean by "balanced"? What balances? If the moon pulls the whole earth toward it, why doesn't the earth fall right "up" to the moon? Because the earth does the same trick as the moon, it goes in a circle around a point which is inside the earth but not at its center. The moon does not just go around the

* The proof is not given in this course.

earth, the earth and the moon both go around a central position, each falling toward this common position, as shown in Fig. 7–5. This motion around the common center is what balances the fall of each. So the earth is not going in a straight line either; it travels in a circle. The water on the far side is "unbalanced" because the moon's attraction there is weaker than it is at the center of the earth, where it just balances the "centrifugal force." The result of this imbalance is that the water rises up, away from the center of the earth. On the near side, the attraction from the moon is stronger, and the imbalance is in the opposite direction in space, but again *away* from the center of the earth. The net result is that we get *two* tidal bulges.

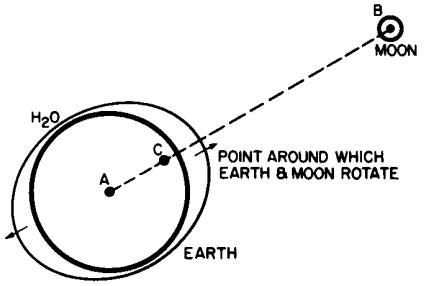


Fig. 7–5. The earth-moon system, with tides.

7–5 Universal gravitation

What else can we understand when we understand gravity? Everyone knows the earth is round. Why is the earth round? That is easy; it is due to gravitation. The earth can be understood to be round merely because everything attracts everything else and so it has attracted itself together as far as it can! If we go even further, the earth is not *exactly* a sphere because it is rotating, and this brings in centrifugal effects which tend to oppose gravity near the equator. It turns out that the earth should be elliptical, and we even get the right shape for the ellipse. We can thus deduce that the sun, the moon, and the earth should be (nearly) spheres, just from the law of gravitation.

What else can you do with the law of gravitation? If we look at the moons of Jupiter we can understand everything about the way they move around that planet. Incidentally, there was once a certain difficulty with the moons of Jupiter that is worth remarking on. These satellites were studied very carefully by Roemer, who noticed that the moons sometimes seemed to be ahead of schedule, and sometimes behind. (One can find their schedules by waiting a very long time and finding out how long it takes on the average for the moons to go around.) Now they were *ahead* when Jupiter was particularly *close* to the earth and they were *behind* when Jupiter was *farther* from the earth. This would have been a very difficult thing to explain according to the law of gravitation—it would have been, in fact, the death of this wonderful theory if there were no other explanation. If a law does not work even in *one place* where it ought to, it is just wrong. But the reason for this discrepancy was very simple and beautiful: it takes a little while to *see* the moons of Jupiter because of the time it takes light to travel from Jupiter to the earth. When Jupiter is closer to the earth the time is a little less, and when it is farther from the earth, the time is more. This is why moons appear to be, on the average, a little ahead or a little behind, depending on whether they are closer to or farther from the earth. This phenomenon showed that light does not travel instantaneously, and furnished the first estimate of the speed of light. This was done in 1656.

If all of the planets push and pull on each other, the force which controls, let us say, Jupiter in going around the sun is not just the force from the sun; there is also a pull from, say, Saturn. This force is not really strong, since the sun is much more massive than Saturn, but there is *some* pull, so the orbit of Jupiter should not be a perfect ellipse, and it is not; it is slightly off, and "wobbles" around the correct elliptical orbit. Such a motion is a little more complicated. Attempts were made to analyze the motions of Jupiter, Saturn, and Uranus on the basis of the law of gravitation. The effects of each of these planets on each other were calculated to see whether or not the tiny deviations and irregularities in these motions could be completely understood from this one law. Lo and behold, for Jupiter and Saturn, all was well, but Uranus was "weird." It behaved in a very peculiar manner. It was not travelling in an exact ellipse, but that was understandable, because of the attractions of Jupiter and Saturn. But even if allowance were made for these attractions, Uranus *still* was not going right, so the laws of gravitation were in danger of being overturned, a possibility that could not be ruled out. Two men, Adams and Leverrier, in England and France, independently,

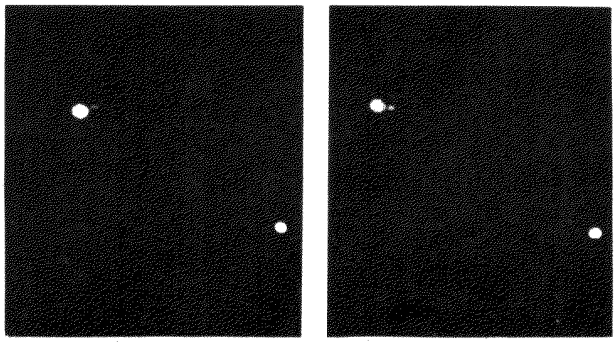


Fig. 7–6. A double-star system.

arrived at another possibility: perhaps there is *another* planet, dark and invisible, which men had not seen. This planet, *N*, could pull on Uranus. They calculated where such a planet would have to be in order to cause the observed perturbations. They sent messages to the respective observatories, saying, "Gentlemen, point your telescope to such and such a place, and you will see a new planet." It often depends on with whom you are working as to whether they pay any attention to you or not. They did pay attention to Leverrier; they looked, and there planet *N* was! The other observatory then also looked very quickly in the next few days and saw it too.

This discovery shows that Newton's laws are absolutely right in the solar system; but do they extend beyond the relatively small distances of the nearest planets? The first test lies in the question, do *stars* attract *each other* as well as planets? We have definite evidence that they do in the *double stars*. Figure 7–6 shows a double star—two stars very close together (there is also a third star in the picture so that we will know that the photograph was not turned). The stars are also shown as they appeared several years later. We see that, relative to the "fixed" star, the axis of the pair has rotated, i.e., the two stars are going around each other. Do they rotate according to Newton's laws? Careful measurements of the relative positions of one such double star system are shown in Fig. 7–7. There we see a beautiful ellipse, the measures starting in 1862 and going all the way around to 1904 (by now it must have gone around once more). Everything coincides with Newton's laws, except that the star Sirius A is *not at the focus*. Why should that be? Because the plane of the ellipse is not in the "plane of the sky." We are not looking at right angles to the orbit plane, and when an ellipse is viewed at a tilt, it remains an ellipse but the focus is no longer at the same place. Thus we can analyze double stars, moving about each other, according to the requirements of the gravitational law.

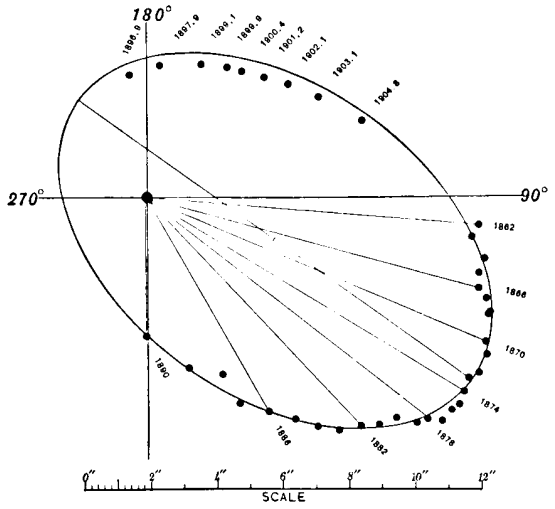


Fig. 7–7. Orbit of Sirius B with respect to Sirius A.

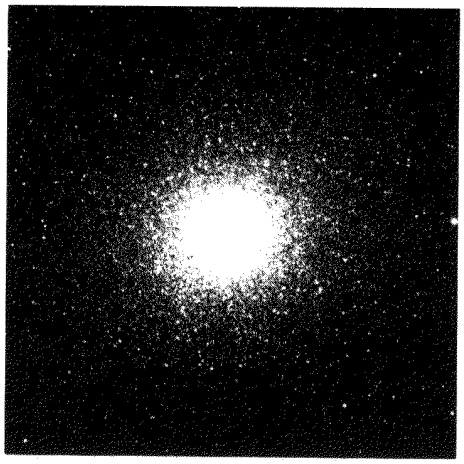


Fig. 7–8. A globular star cluster.

That the law of gravitation is true at even bigger distances is indicated in Fig. 7–8. If one cannot see gravitation acting here, he has no soul. This figure shows one of the most beautiful things in the sky—a globular star cluster. All of the dots are stars. Although they look as if they are packed solid toward the center, that is due to the fallibility of our instruments. Actually, the distances between even the centermost stars are very great and they very rarely collide. There are more stars in the interior than farther out, and as we move outward there are fewer and fewer. It is obvious that there is an attraction among these stars. It is clear that gravitation exists at these enormous dimensions, perhaps 100,000 times the size of the solar system. Let us now go further, and look at an *entire galaxy*, shown in Fig. 7–9. The shape of this galaxy indicates an obvious tendency for its matter to agglomerate. Of course we cannot prove that the law here is precisely inverse square, only that there is still an attraction, at this enormous dimension, that holds the whole thing together. One may say, "Well, that is all very clever but why is it not just a ball?" Because it is *spinning* and has *angular momentum* which it cannot give up as it contracts; it must contract mostly in a plane. (Incidentally, if you are looking for a good problem, the exact details of how the arms are formed and what determines the shapes of these galaxies has not been worked out.) It is, however, clear that the shape of the galaxy is due to gravitation even though the complexities of its structure have not yet allowed



Fig. 7–9. A galaxy.

us to analyze it completely. In a galaxy we have a scale of perhaps 50,000 to 100,000 light years. The earth's distance from the sun is $8\frac{1}{3}$ light *minutes*, so you can see how large these dimensions are.

Gravity appears to exist at even bigger dimensions, as indicated by Fig. 7–10, which shows many "little" things clustered together. This is a *cluster of galaxies*, just like a star cluster. Thus galaxies attract each other at such distances that they too are agglomerated into clusters. Perhaps gravitation exists even over distances of *tens of millions* of light years; so far as we now know, gravity seems to go out forever inversely as the square of the distance.

Not only can we understand the nebulae, but from the law of gravitation we can even get some ideas about the origin of the stars. If we have a big cloud of dust and gas, as indicated in Fig. 7–11, the gravitational attractions of the pieces of dust for one another might make them form little lumps. Barely visible in the figure are "little" black spots which may be the beginning of the accumulations of dust and gases which, due to their gravitation, begin to form stars. Whether we have ever seen a star form or not is still debatable. Figure 7–12 shows the one piece of evidence which suggests that we have. At the left is a picture of a region of gas with some stars in it taken in 1947, and at the right is another picture, taken only 7 years later, which shows two new bright spots. Has gas accumulated, has gravity acted hard enough and collected it into a ball big enough that the stellar nuclear reaction starts in the interior and turns it into a star? Perhaps, and perhaps not. It is unreasonable that in only seven years we should be so lucky as to see a star change itself into visible form; it is much less probable that we should see *two!*



Fig. 7–10. A cluster of galaxies.



Fig. 7–11. An interstellar dust cloud.

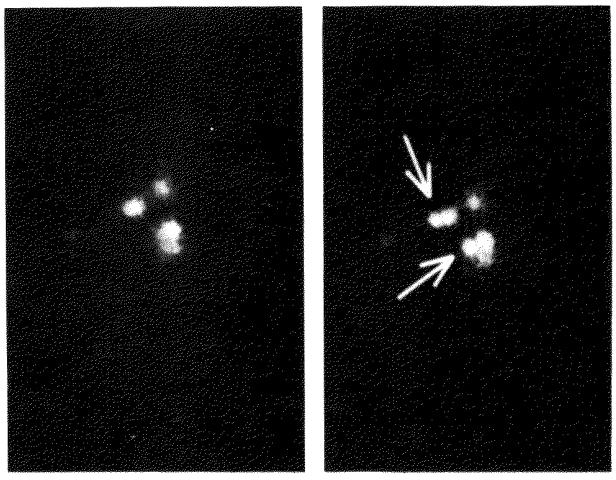


Fig. 7–12. The formation of new stars?

7–6 Cavendish's experiment

Gravitation, therefore, extends over enormous distances. But if there is a force between *any* pair of objects, we ought to be able to measure the force between our own objects. Instead of having to watch the stars go around each other, why can we not take a ball of lead and a marble and watch the marble go toward the ball of lead? The difficulty of this experiment when done in such a simple manner is the very weakness or delicacy of the force. It must be done with extreme care, which means covering the apparatus to keep the air out, making sure it is not electrically charged, and so on; then the force can be measured. It was first measured by Cavendish with an apparatus which is schematically indicated in Fig. 7–13. This first demonstrated the direct force between two large, fixed balls of lead and two smaller balls of lead on the ends of an arm supported by a very fine fiber, called a torsion fiber. By measuring how much the fiber gets twisted, one can measure the strength of the force, verify that it is inversely proportional to the square of the distance, and determine how strong it is. Thus, one may accurately determine the coefficient G in the formula

$$F = G\frac{mm'}{r^2}.$$

All the masses and distances are known. You say, "We knew it already for the earth." Yes, but we did not know the *mass* of the earth. By knowing G from this experiment and by knowing how strongly the earth attracts, we can indirectly learn how great is the mass of the earth! This experiment has been called "weighing the earth." Cavendish claimed he was weighing the earth, but what he was measuring was the coefficient G of the gravity law. This is the only way in which the mass of the earth can be determined. G turns out to be

$$6.670 \times 10^{-11}\ \text{newton} \cdot \text{m}^2/\text{kg}^2.$$

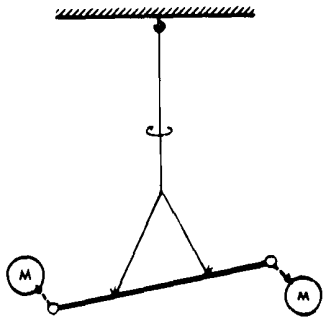


Fig. 7–13. A simplified diagram of the apparatus used by Cavendish to verify the law of universal gravitation for small objects and to measure the gravitational constant G.

It is hard to exaggerate the importance of the effect on the history of science produced by this great success of the theory of gravitation. Compare the confusion, the lack of confidence, the incomplete knowledge that prevailed in the earlier ages, when there were endless debates and paradoxes, with the clarity and simplicity of this law—this fact that all the moons and planets and stars have such a *simple rule* to govern them, and further that man could *understand* it and deduce how the planets should move! This is the reason for the success of the sciences in following years, for it gave hope that the other phenomena of the world might also have such beautifully simple laws.

7–7 What is gravity?

But is this such a simple law? What about the machinery of it? All we have done is to describe *how* the earth moves around the sun, but we have not said *what makes it go*. Newton made no hypotheses about this; he was satisfied to find *what* it did without getting into the machinery of it. *No one has since given any machinery*. It is characteristic of the physical laws that they have this abstract character. The law of conservation of energy is a theorem concerning quantities that have to be calculated and added together, with no mention of the machinery, and likewise the great laws of mechanics are quantitative mathematical laws for which no machinery is available. Why can we use mathematics to describe nature without a mechanism behind it? No one knows. We have to keep going because we find out more that way.

Many mechanisms for gravitation have been suggested. It is interesting to consider one of these, which many people have thought of from time to time. At first, one is quite excited and happy when he "discovers" it, but he soon finds that it is not correct. It was first discovered about 1750. Suppose there were many particles moving in space at a very high speed in all directions and being only slightly absorbed in going through matter. When they *are* absorbed, they give an impulse to the earth. However, since there are as many going one way as another, the

impulses all balance. But when the sun is nearby, the particles coming toward the earth through the sun are partially absorbed, so fewer of them are coming from the sun than are coming from the other side. Therefore, the earth feels a net impulse toward the sun and it does not take one long to see that it is inversely as the square of the distance—because of the variation of the solid angle that the sun subtends as we vary the distance. What is wrong with that machinery? It involves some new consequences which are *not true*. This particular idea has the following trouble: the earth, in moving around the sun, would impinge on more particles which are coming from its forward side than from its hind side (when you run in the rain, the rain in your face is stronger than that on the back of your head!). Therefore there would be more impulse given the earth from the front, and the earth would feel a *resistance to motion* and would be slowing up in its orbit. One can calculate how long it would take for the earth to stop as a result of this resistance, and it would not take long enough for the earth to still be in its orbit, so this mechanism does not work. No machinery has ever been invented that "explains" gravity without also predicting some other phenomenon that does *not* exist.

Next we shall discuss the possible relation of gravitation to other forces. There is no explanation of gravitation in terms of other forces at the present time. It is not an aspect of electricity or anything like that, so we have no explanation. However, gravitation and other forces are very similar, and it is interesting to note analogies. For example, the force of electricity between two charged objects looks just like the law of gravitation: the force of electricity is a constant, with a minus sign, times the product of the charges, and varies inversely as the square of the distance. It is in the opposite direction—likes repel. But is it still not very remarkable that the two laws involve the same function of distance? Perhaps gravitation and electricity are much more closely related than we think. Many attempts have been made to unify them; the so-called unified field theory is only a very elegant attempt to combine electricity and gravitation; but, in comparing gravitation and electricity, the most interesting thing is the *relative strengths* of the forces. Any theory that contains them both must also deduce how strong the gravity is.

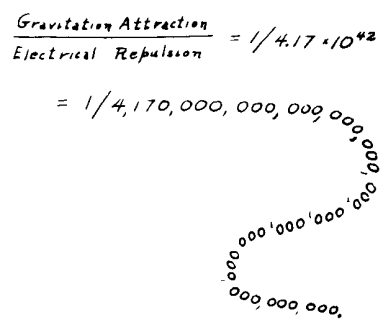


Fig. 7–14. The relative strengths of electrical and gravitational interactions between two electrons.

If we take, in some natural units, the repulsion of two electrons (nature's universal charge) due to electricity, and the attraction of two electrons due to their masses, we can measure the ratio of electrical repulsion to the gravitational attraction. The ratio is independent of the distance and is a fundamental constant of nature. The ratio is shown in Fig. 7–14. The gravitational attraction relative to the electrical repulsion between two electrons is 1 divided by 4.17×10^{42}! The question is, where does such a large number come from? It is not accidental, like the ratio of the volume of the earth to the volume of a flea. We have considered two natural aspects of the same thing, an electron. This fantastic number is a natural constant, so it involves something deep in nature. Where could such a tremendous number come from? Some say that we shall one day find the "universal equation," and in it, one of the roots will be this number. It is very difficult to find an equation for which such a fantastic number is a natural root. Other possibilities have been thought of; one is to relate it to the age of the universe. Clearly, we have to find *another* large number somewhere. But do we mean the age of the universe in *years?* No, because years are not "natural"; they were devised by men. As an example of something natural, let us consider the time it takes light to go across a proton, 10^{-24} second. If we compare this time with the *age of the universe*, 2×10^{10} years, the answer is 10^{-42}. It has about the same number of zeros going off it, so it has been proposed that the gravitational constant is related to the age of the universe. If that were the case, the gravitational constant would change with time, because as the universe got older the ratio of the age of the universe to the time which it takes for light to go across a proton would be gradually increasing. Is it possible that the gravitational constant *is* changing with time? Of course the changes would be so small that it is quite difficult to be sure.

One test which we can think of is to determine what would have been the effect of the change during the past 10^9 years, which is approximately the age from the earliest life on the earth to now, and one-tenth of the age of the universe. In this time, the gravity constant would have increased by about 10 percent. It

turns out that if we consider the structure of the sun—the balance between the weight of its material and the rate at which radiant energy is generated inside it—we can deduce that if the gravity were 10 percent stronger, the sun would be much more than 10 percent brighter—by the *sixth power* of the gravity constant! If we calculate what happens to the orbit of the earth when the gravity is changing, we find that the earth was then *closer in.* Altogether, the earth would be about 100 degrees centigrade hotter, and all of the water would not have been in the sea, but vapor in the air, so life would not have started in the sea. So we do *not* now believe that the gravity constant is changing with the age of the universe. But such arguments as the one we have just given are not very convincing, and the subject is not completely closed.

It is a fact that the force of gravitation is proportional to the *mass*, the quantity which is fundamentally a measure of *inertia*—of how hard it is to hold something which is going around in a circle. Therefore two objects, one heavy and one light, going around a larger object in the same circle at the same speed because of gravity, will stay together because to go in a circle *requires* a force which is stronger for a bigger mass. That is, the gravity is stronger for a given mass in *just the right proportion* so that the two objects will go around together. If one object were inside the other it would *stay* inside; it is a perfect balance. Therefore, Gagarin or Titov would find things "weightless" inside a space ship; if they happened to let go of a piece of chalk, for example, it would go around the earth in exactly the same way as the whole space ship, and so it would appear to remain suspended before them in space. It is very interesting that this force is *exactly* proportional to the mass with great precision, because if it were not exactly proportional there would be some effect by which inertia and weight would differ. The absence of such an effect has been checked with great accuracy by an experiment done first by Eotvos in 1909 and more recently by Dicke. For all substances tried, the masses and weights are exactly proportional within 1 part in 1,000,000,000, or less. This is a remarkable experiment.

7–8 Gravity and relativity

Another topic deserving discussion is Einstein's modification of Newton's law of gravitation. In spite of all the excitement it created, Newton's law of gravitation is not correct! It was modified by Einstein to take into account the theory of relativity. According to Newton, the gravitational effect is instantaneous, that is, if we were to move a mass, we would at once feel a new force because of the new position of that mass; by such means we could send signals at infinite speed. Einstein advanced arguments which suggest that we *cannot send signals faster than the speed of light*, so the law of gravitation must be wrong. By correcting it to take the delays into account, we have a new law, called Einstein's law of gravitation. One feature of this new law which is quite easy to understand is this: In the Einstein relativity theory, anything which has *energy* has mass—mass in the sense that it is attracted gravitationally. Even light, which has an energy, has a "mass." When a light beam, which has energy in it, comes past the sun there is an attraction on it by the sun. Thus the light does not go straight, but is deflected. During the eclipse of the sun, for example, the stars which are around the sun should appear displaced from where they would be if the sun were not there, and this has been observed.

Finally, let us compare gravitation with other theories. In recent years we have discovered that all mass is made of tiny particles and that there are several kinds of interactions, such as nuclear forces, etc. None of these nuclear or electrical forces has yet been found to explain gravitation. The quantum-mechanical aspects of nature have not yet been carried over to gravitation. When the scale is so small that we need the quantum effects, the gravitational effects are so weak that the need for a quantum theory of gravitation has not yet developed. On the other hand, for consistency in our physical theories it would be important to see whether Newton's law modified to Einstein's law can be further modified to be consistent with the uncertainty principle. This last modification has not yet been completed.

8

Motion

8–1 Description of motion

8–2 Speed

8–3 Speed as a derivative

8–4 Distance as an integral

8–5 Acceleration

8–1 Description of motion

In order to find the laws governing the various changes that take place in bodies as time goes on, we must be able to *describe* the changes and have some way to record them. The simplest change to observe in a body is the apparent change in its position with time, which we call motion. Let us consider some solid object with a permanent mark, which we shall call a point, that we can observe. We shall discuss the motion of the little marker, which might be the radiator cap of an automobile or the center of a falling ball, and shall try to describe the fact that it moves and how it moves.

These examples may sound trivial, but many subtleties enter into the description of change. Some changes are more difficult to describe than the motion of a point on a solid object, for example the speed of drift of a cloud that is drifting very slowly, but rapidly forming or evaporating, or the change of a woman's mind. We do not know a simple way to analyze a change of mind, but since the cloud can be represented or described by many molecules, perhaps we can describe the motion of the cloud in principle by describing the motion of all its individual molecules. Likewise, perhaps even the changes in the mind may have a parallel in changes of the atoms inside the brain, but we have no such knowledge yet.

At any rate, that is why we begin with the motion of points; perhaps we should think of them as atoms, but it is probably better to be more rough in the beginning and simply to think of some kind of small objects—small, that is, compared with the distance moved. For instance, in describing the motion of a car that is going a hundred miles, we do not have to distinguish between the front and the back of the car. To be sure, there are slight differences, but for rough purposes we say "the car," and likewise it does not matter that our points are not absolute points; for our present purposes it is not necessary to be extremely precise. Also, while we take a first look at this subject we are going to forget about the three dimensions of the world. We shall just concentrate on moving in one direction, as in a car on one road. We shall return to three dimensions after we see how to describe motion in one dimension. Now, you may say, "This is all some kind of trivia," and indeed it is. How can we describe such a one-dimensional motion—let us say, of a car? Nothing could be simpler. Among many possible ways, one would be the following. To determine the position of the car at different times, we measure its distance from the starting point and record all the observations. In Table 8–1, s represents the distance of the car, in feet, from the starting point, and t represents the time in minutes. The first line in the table represents zero distance and zero time—the car has not started yet. After one minute it has started and has gone 1200 feet. Then in two minutes, it goes farther—notice that it picked up more distance in the second minute—it has accelerated; but something happened between 3 and 4 and even more so at 5—it stopped at a light perhaps? Then it speeds up again and goes 13,000 feet by the end of 6 minutes, 18,000 feet at the end of 7 minutes, and 23,500 feet in 8 minutes; at 9 minutes it has advanced to only 24,000 feet, because in the last minute it was stopped by a cop.

Table 8–1

t (min)	s (ft)
0	0
1	1200
2	4000
3	9000
4	9500
5	9600
6	13000
7	18000
8	23500
9	24000

That is one way to describe the motion. Another way is by means of a graph. If we plot the time horizontally and the distance vertically, we obtain a curve something like that shown in Fig. 8–1. As the time increases, the distance increases, at first very slowly and then more rapidly, and very slowly again for a little while at 4 minutes; then it increases again for a few minutes and finally, at 9 minutes, appears to have stopped increasing. These observations can be made from the

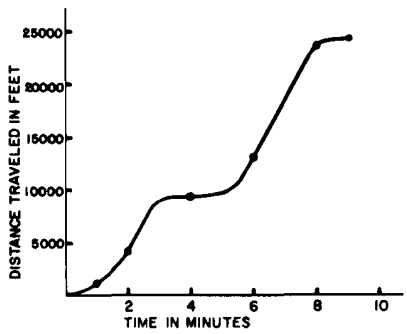


Fig. 8–1. Graph of distance versus time for the car.

Table 8–2

t (sec)	s (ft)
0	0
1	16
2	64
3	144
4	256
5	400
6	576

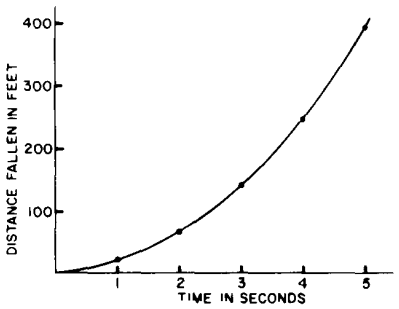


Fig. 8–2. Graph of distance versus time for a falling body.

graph, without a table. Obviously, for a complete description one would have to know where the car is at the half-minute marks, too, but we suppose that the graph means something, that the car has some position at all the intermediate times.

The motion of a car is complicated. For another example we take something that moves in a simpler manner, following more simple laws: a falling ball. Table 8–2 gives the time in seconds and the distance in feet for a falling body. At zero seconds the ball starts out at zero feet, and at the end of 1 second it has fallen 16 feet. At the end of 2 seconds, it has fallen 64 feet, at the end of 3 seconds, 144 feet, and so on; if the tabulated numbers are plotted, we get the nice parabolic curve shown in Fig. 8–2. The formula for this curve can be written as

$$s = 16t^2. \tag{8.1}$$

This formula enables us to calculate the distances at any time. You might say there ought to be a formula for the first graph too. Actually, one may write such a formula abstractly, as

$$s = f(t), \tag{8.2}$$

meaning that s is some quantity depending on t or, in mathematical phraseology, s is a function of t. Since we do not know what the function is, there is no way we can write it in definite algebraic form.

We have now seen two examples of motion, adequately described with very simple ideas, no subtleties. However, there *are* subtleties—several of them. In the first place, what do we mean by *time* and *space?* It turns out that these deep philosophical questions have to be analyzed very carefully in physics, and this is not so easy to do. The theory of relativity shows that our ideas of space and time are not as simple as one might think at first sight. However, for our present purposes, for the accuracy that we need at first, we need not be very careful about defining things precisely. Perhaps you say, "That's a terrible thing—I learned that in science we have to define *everything* precisely." We cannot define *anything* precisely! If we attempt to, we get into that paralysis of thought that comes to philosophers, who sit opposite each other, one saying to the other, "You don't know what you are talking about!" The second one says, "What do you mean by *know?* What do you mean by *talking?* What do you mean by *you?*," and so on. In order to be able to talk constructively, we just have to agree that we are talking about roughly the same thing. You know as much about time as we need for the present, but remember that there are some subtleties that have to be discussed; we shall discuss them later.

Another subtlety involved, and already mentioned, is that it should be possible to imagine that the moving point we are observing is always located somewhere. (Of course when we are looking at it, there it is, but maybe when we look away it isn't there.) It turns out that in the motion of atoms, that idea also is false—we cannot find a marker on an atom and watch it move. That subtlety we shall have to get around in quantum mechanics. But we are first going to learn what the problems are before introducing the complications, and *then* we shall be in a better position to make corrections, in the light of the more recent knowledge of the subject. We shall, therefore, take a simple point of view about time and space. We know what these concepts are in a rough way, and those who have driven a car know what speed means.

8–2 Speed

Even though we know roughly what "speed" means, there are still some rather deep subtleties; consider that the learned Greeks were never able to adequately describe problems involving velocity. The subtlety comes when we try to comprehend exactly what is meant by "speed." The Greeks got very confused about this, and a new branch of mathematics had to be discovered beyond the geometry and algebra of the Greeks, Arabs, and Babylonians. As an illustration of the difficulty, try to solve this problem by sheer algebra: A balloon is being inflated so

that the volume of the balloon is increasing at the rate of 100 cm^3 per second; at what speed is the radius increasing when the volume is 1000 cm^3? The Greeks were somewhat confused by such problems, being helped, of course, by some very confusing Greeks. To show that there were difficulties in reasoning about speed at the time, Zeno produced a large number of paradoxes, of which we shall mention one to illustrate his point that there are obvious difficulties in thinking about motion. "Listen," he says, "to the following argument: Achilles runs 10 times as fast as a tortoise, nevertheless he can never catch the tortoise. For, suppose that they start in a race where the tortoise is 100 meters ahead of Achilles; then when Achilles has run the 100 meters to the place where the tortoise was, the tortoise has proceeded 10 meters, having run one-tenth as fast. Now, Achilles has to run another 10 meters to catch up with the tortoise, but on arriving at the end of that run, he finds that the tortoise is still 1 meter ahead of him; running another meter, he finds the tortoise 10 centimeters ahead, and so on, *ad infinitum*. Therefore, at any moment the tortoise is always ahead of Achilles and Achilles can never catch up with the tortoise." What is wrong with that? It is that a finite amount of time can be divided into an infinite number of pieces, just as a length of line can be divided into an infinite number of pieces by dividing repeatedly by two. And so, although there are an infinite number of steps (in the argument) to the point at which Achilles reaches the tortoise, it doesn't mean that there is an infinite amount of *time*. We can see from this example that there are indeed some subtleties in reasoning about speed.

In order to get to the subtleties in a clearer fashion, we remind you of a joke which you surely must have heard. At the point where the lady in the car is caught by a cop, the cop comes up to her and says, "Lady, you were going 60 miles an hour!" She says, "That's impossible, sir, I was travelling for only seven minutes. It is ridiculous—how can I go 60 miles an hour when I wasn't going an hour?" How would you answer her if you were the cop? Of course, if you were really the cop, then no subtleties are involved; it is very simple: you say, "Tell that to the judge!" But let us suppose that we do not have that escape and we make a more honest, intellectual attack on the problem, and try to explain to this lady what we mean by the idea that she was going 60 miles an hour. Just what *do* we mean? We say, "What we mean, lady, is this: if you kept on going the same way as you are going now, in the next hour you would go 60 miles." She could say, "Well, my foot was off the accelerator and the car was slowing down, so if I kept on going that way it would not go 60 miles." Or consider the falling ball and suppose we want to know its speed at the time three seconds if the ball kept on going the way it is going. What does that mean—kept on *accelerating*, going faster? No—kept on going with the same *velocity*. But that is what we are trying to define! For if the ball keeps on going the way it is going, it will just keep on going the way it is going. Thus we need to define the velocity better. What has to be kept the same? The lady can also argue this way: "If I kept on going the way I'm going for one more hour, I would run into that wall at the end of the street!" It is not so easy to say what we mean.

Many physicists think that measurement is the only definition of anything. Obviously, then, we should use the instrument that measures the speed—the speedometer—and say, "Look, lady, your speedometer reads 60." So she says, "My speedometer is broken and didn't read at all." Does that mean the car is standing still? We believe that there is something to measure before we build the speedometer. Only then can we say, for example, "The speedometer isn't working right," or "the speedometer is broken." That would be a meaningless sentence if the velocity had no meaning independent of the speedometer. So we have in our minds, obviously, an idea that is independent of the speedometer, and the speedometer is meant only to measure this idea. So let us see if we can get a better definition of the idea. We say, "Yes, of course, before you went an hour, you would hit that wall, but if you went one second, you would go 88 feet; lady, you were going 88 feet per second, and if you kept on going, the next second it would be 88 feet, and the wall down there is farther away than that." She says, "Yes, but there's no law against going 88 feet per second! There is only a law

against going 60 miles an hour." "But," we reply, "it's the same thing." If it *is* the same thing, it should not be necessary to go into this circumlocution about 88 feet per second. In fact, the falling ball could not keep going the same way even one second because it would be changing speed, and we shall have to define speed somehow.

Now we seem to be getting on the right track; it goes something like this: If the lady kept on going for another 1/1000 of an hour, she would go 1/1000 of 60 miles. In other words, she does not have to keep on going for the whole hour; the point is that *for a moment* she is going at that speed. Now what that means is that if she went just a little bit more in time, the extra distance she goes would be the same as that of a car that goes at a *steady* speed of 60 miles an hour. Perhaps the idea of the 88 feet per second is right; we see how far she went in the last second, divide by 88 feet, and if it comes out 1 the speed was 60 miles an hour. In other words, we can find the speed in this way: We ask, how far do we go in a very short time? We divide that distance by the time, and that gives the speed. But the time should be made as short as possible, the shorter the better, because some change could take place during that time. If we take the time of a falling body as an hour, the idea is ridiculous. If we take it as a second, the result is pretty good for a car, because there is not much change in speed, but not for a falling body; so in order to get the speed more and more accurately, we should take a smaller and smaller time interval. What we should do is take a millionth of a second, and divide that distance by a millionth of a second. The result gives the distance per second, which is what we mean by the velocity, so we can define it that way. That is a successful answer for the lady, or rather, that is the definition that we are going to use.

The foregoing definition involves a new idea, an idea that was not available to the Greeks in a general form. That idea was to take an *infinitesimal distance* and the corresponding *infinitesimal time*, form the ratio, and watch what happens to that ratio as the time that we use gets smaller and smaller and smaller. In other words, take a limit of the distance travelled divided by the time required, as the time taken gets smaller and smaller, *ad infinitum*. This idea was invented by Newton and by Leibnitz, independently, and is the beginning of a new branch of mathematics, called the *differential calculus*. Calculus was invented in order to describe motion, and its first application was to the problem of defining what is meant by going "60 miles an hour."

Let us try to define velocity a little better. Suppose that in a short time, ϵ, the car or other body goes a short distance x; then the velocity, v, is defined as

$$v = x/\epsilon,$$

an approximation that becomes better and better as the ϵ is taken smaller and smaller. If a mathematical expression is desired, we can say that the velocity equals the limit as the ϵ is made to go smaller and smaller in the expression x/ϵ, or

$$v = \lim_{\epsilon \to 0} \frac{x}{\epsilon}. \tag{8.3}$$

We cannot do the same thing with the lady in the car, because the table is incomplete. We know only where she was at intervals of one minute; we can get a rough idea that she was going 5000 ft/min during the 7th minute, but we do not know, at exactly the moment 7 minutes, whether she had been speeding up and the speed was 4900 ft/min at the beginning of the 6th minute, and is now 5100 ft/min, or something else, because we do not have the exact details in between. So only if the table were completed with an infinite number of entries could we really calculate the velocity from such a table. On the other hand, when we have a complete mathematical formula, as in the case of a falling body (Eq. 8.1), then it is possible to calculate the velocity, because we can calculate the position at any time whatsoever.

Let us take as an example the problem of determining the velocity of the falling ball at the particular time 5 seconds. One way to do this is to see from

Table 8–2 what it did in the 5th second; it went $400 - 256 = 144$ ft, so it is going 144 ft/sec; however, that is wrong, because the speed is changing; *on the average* it is 144 ft/sec during this interval, but the ball is speeding up and is really going faster than 144 ft/sec. We want to find out *exactly how fast.* The technique involved in this process is the following: We know where the ball was at 5 sec. At 5.1 sec, the distance that it has gone all together is $16(5.1)^2 = 416.16$ ft (see Eq. 8.1). At 5 sec it had already fallen 400 ft; in the last tenth of a second it fell $416.16 - 400 = 16.16$ ft. Since 16.16 ft in 0.1 sec is the same as 161.6 ft/sec, that is the speed more or less, but it is not exactly correct. Is that the speed at 5, or at 5.1, or halfway between at 5.05 sec, or when *is* that the speed? Never mind —the problem was to find the speed *at 5 seconds,* and we do not have exactly that; we have to do a better job. So, we take one-thousandth of a second more than 5 sec, or 5.001 sec, and calculate the total fall as

$$s = 16(5.001)^2 = 16(25.010001) = 400.160016 \text{ ft.}$$

In the last 0.001 sec the ball fell 0.160016 ft, and if we divide this number by 0.001 sec we obtain the speed as 160.016 ft/sec. That is closer, very close, but it is *still not exact.* It should now be evident what we must do to find the speed exactly. To perform the mathematics we state the problem a little more abstractly: to find the velocity at a special time, t_0, which in the original problem was 5 sec. Now the distance at t_0, which we call s_0, is $16t_0^2$, or 400 ft in this case. In order to find the velocity, we ask, "At the time t_0 + (a little bit), or $t_0 + \epsilon$, where is the body?" The new position is $16(t_0 + \epsilon)^2 = 16t_0^2 + 32t_0\epsilon + 16\epsilon^2$. So it is farther along than it was before, because before it was only $16t_0^2$. This distance we shall call s_0 + (a little bit more), or $s_0 + x$ (if x is the extra bit). Now if we subtract the distance at t_0 from the distance at $t_0 + \epsilon$, we get x, the extra distance gone, as $x = 32t_0 \cdot \epsilon + 16\epsilon^2$. Our first approximation to the velocity is

$$v = \frac{x}{\epsilon} = 32t_0 + 16\epsilon. \tag{8.4}$$

The true velocity is the value of this ratio, x/ϵ, when ϵ becomes vanishingly small. In other words, after forming the ratio, we take the limit as ϵ gets smaller and smaller, that is, approaches 0. The equation reduces to,

$$v \text{ (at time } t_0) = 32t_0.$$

In our problem, $t_0 = 5$ sec, so the solution is $v = 32 \times 5 = 160$ ft/sec. A few lines above, where we took ϵ as 0.1 and 0.01 sec successively, the value we got for v was a little more than this, but now we see that the actual velocity is precisely 160 ft/sec.

8–3 Speed as a derivative

The procedure we have just carried out is performed so often in mathematics that for convenience special notations have been assigned to our quantities ϵ and x. In this notation, the ϵ used above becomes Δt and x becomes Δs. This Δt means "an extra bit of t," and carries an implication that it can be made smaller. The prefix Δ is not a multiplier, any more than $\sin\theta$ means $s \cdot i \cdot n \cdot \theta$—it simply defines a time increment, and reminds us of its special character. Δs has an analogous meaning for the distance s. Since Δ is not a factor, it cannot be cancelled in the ratio $\Delta s/\Delta t$ to give s/t, any more than the ratio $\sin\theta/\sin 2\theta$ can be reduced to 1/2 by cancellation. In this notation, velocity is equal to the limit of $\Delta s/\Delta t$ when Δt gets smaller, or

$$v = \lim_{\Delta t \to 0} \frac{\Delta s}{\Delta t}. \tag{8.5}$$

This is really the same as our previous expression (8.3) with ϵ and x, but it has the advantage of showing that something is changing, and it keeps track of what is changing.

Incidentally, to a good approximation we have another law, which says that the change in distance of a moving point is the velocity times the time interval, or $\Delta s = v\,\Delta t$. This statement is true only if the velocity is not changing during that time interval, and this condition is true only in the limit as Δt goes to 0. Physicists like to write it $ds = v\,dt$, because by dt they mean Δt in circumstances in which it is very small; with this understanding, the expression is valid to a close approximation. If Δt is too long, the velocity might change during the interval, and the approximation would become less accurate. For a time dt, approaching zero, $ds = v\,dt$ precisely. In this notation we can write (8.5) as

$$v = \lim_{\Delta t \to 0} \frac{\Delta s}{\Delta t} = \frac{ds}{dt}.$$

The quantity ds/dt which we found above is called the "derivative of s with respect to t" (this language helps to keep track of what was changed), and the complicated process of finding it is called finding a derivative, or differentiating. The ds's and dt's which appear separately are called *differentials*. To familiarize you with the words, we say we found the derivative of the function $16t^2$, or the derivative (with respect to t) of $16t^2$ is $32t$. When we get used to the words, the ideas are more easily understood. For practice, let us find the derivative of a more complicated function. We shall consider the formula $s = At^3 + Bt + C$, which might describe the motion of a point. The letters A, B, and C represent constant numbers, as in the familiar general form of a quadratic equation. Starting from the formula for the motion, we wish to find the velocity at any time. To find the velocity in the more elegant manner, we change t to $t + \Delta t$ and note that s is then changed to $s +$ some Δs; then we find the Δs in terms of Δt. That is to say,

$$\begin{aligned} s + \Delta s &= A(t + \Delta t)^3 + B(t + \Delta t) + C \\ &= At^3 + Bt + C + 3At^2\,\Delta t + B\,\Delta t + 3At(\Delta t)^2 + A(\Delta t)^3, \end{aligned}$$

but since

$$s = At^3 + Bt + C,$$

we find that

$$\Delta s = 3At^2\,\Delta t + B\,\Delta t + 3At(\Delta t)^2 + A(\Delta t)^3.$$

But we do not want Δs—we want Δs divided by Δt. We divide the preceding equation by Δt, getting

$$\frac{\Delta s}{\Delta t} = 3At^2 + B + 3At(\Delta t) + A(\Delta t)^2.$$

Table 8–3. A Short Table of Derivatives

s, u, v, w are arbitrary functions of t; a, b, c, and n are arbitrary constants

Function	Derivative
$s = t^n$	$\dfrac{ds}{dt} = nt^{n-1}$
$s = cu$	$\dfrac{ds}{dt} = c\dfrac{du}{dt}$
$s = u + v + w + \cdots$	$\dfrac{ds}{dt} = \dfrac{du}{dt} + \dfrac{dv}{dt} + \dfrac{dw}{dt} + \cdots$
$s = c$	$\dfrac{ds}{dt} = 0$
$s = u^a v^b w^c \ldots$	$\dfrac{ds}{dt} = s\left(\dfrac{a}{u}\dfrac{du}{dt} + \dfrac{b}{v}\dfrac{dv}{dt} + \dfrac{c}{w}\dfrac{dw}{dt} + \cdots\right)$

As Δt goes toward 0 the limit of $\Delta s/\Delta t$ is ds/dt and is equal to

$$\frac{ds}{dt} = 3At^2 + B.$$

This is the fundamental process of calculus, differentiating functions. The process is even more simple than it appears. Observe that when these expansions contain any term with a square or a cube or any higher power of Δt, such terms may be dropped at once, since they will go to 0 when the limit is taken. After a little practice the process gets easier because one knows what to leave out. There are many rules or formulas for differentiating various types of functions. These can be memorized, or can be found in tables. A short list is found in Table 8–3.

8–4 Distance as an integral

Now we have to discuss the inverse problem. Suppose that instead of a table of distances, we have a table of speeds at different times, starting from zero. For the falling ball, such speeds and times are shown in Table 8–4. A similar table could be constructed for the velocity of the car, by recording the speedometer reading every minute or half-minute. If we know how fast the car is going at any time, can we determine how far it goes? This problem is just the inverse of the one solved above; we are given the velocity and asked to find the distance. How can we find the distance if we know the speed? If the speed of the car is not constant, and the lady goes sixty miles an hour for a moment, then slows down, speeds up, and so on, how can we determine how far she has gone? That is easy. We use the same idea, and express the distance in terms of infinitesimals. Let us say, "In the first second her speed was such and such, and from the formula $\Delta s = v\,\Delta t$ we can calculate how far the car went the first second at that speed." Now in the next second her speed is nearly the same, but slightly different; we can calculate how far she went in the next second by taking the new speed times the time. We proceed similarly for each second, to the end of the run. We now have a number of little distances, and the total distance will be the sum of all these little pieces. That is, the distance will be the sum of the velocities times the times, or $s = \sum v\,\Delta t$, where the Greek letter $\sum$ (sigma) is used to denote addition. To be more precise, it is the sum of the velocity at a certain time, let us say the i-th time, multiplied by Δt.

$$s = \sum_i v(t_i)\,\Delta t. \tag{8.6}$$

The rule for the times is that $t_{i+1} = t_i + \Delta t$. However, the distance we obtain by this method will not be correct, because the velocity changes during the time interval Δt. If we take the times short enough, the sum is precise, so we take them smaller and smaller until we obtain the desired accuracy. The true s is

$$s = \lim_{\Delta t \to 0} \sum_i v(t_i)\,\Delta t. \tag{8.7}$$

The mathematicians have invented a symbol for this limit, analogous to the symbol for the differential. The Δ turns into a d to remind us that the time is as small as it can be; the velocity is then called v at the time t, and the addition is written as a sum with a great "s," $\int$ (from the Latin *summa*), which has become distorted and is now unfortunately just called an integral sign. Thus we write

$$s = \int v(t)\,dt. \tag{8.8}$$

This process of adding all these terms together is called integration, and it is the opposite process to differentiation. The derivative of this integral is v, so one operator (d) undoes the other ($\int$). One can get formulas for integrals by taking the formulas for derivatives and running them backwards, because they are related to each other inversely. Thus one can work out his own table of integrals by differentiating all sorts of functions. For every formula with a differential, we get an integral formula if we turn it around.

Table 8–4

Velocity of a Falling Ball

t (sec)	v (ft/sec)
0	0
1	32
2	64
3	96
4	128

Every function can be differentiated analytically, i.e., the process can be carried out algebraically, and leads to a definite function. But it is not possible in a simple manner to write an analytical value for any integral at will. You can calculate it, for instance, by doing the above sum, and then doing it again with a finer interval Δt and again with a finer interval until you have it nearly right. In general, given some particular function, it is not possible to find, analytically, what the integral is. One may always try to find a function which, when differentiated, gives some desired function; but one may not find it, and it may not exist, in the sense of being expressible in terms of functions that have already been given names.

8–5 Acceleration

The next step in developing the equations of motion is to introduce another idea which goes beyond the concept of velocity to that of *change* of velocity, and we now ask, "How does the velocity *change?*" In previous chapters we have discussed cases in which forces produce changes in velocity. You may have heard with great excitement about some car that can get from rest to 60 miles an hour in ten seconds flat. From such a performance we can see how fast the speed changes, but only on the average. What we shall now discuss is the next level of complexity, which is how fast the velocity is changing. In other words, by how many feet per second does the velocity change in a second, that is, how many feet per second, per second? We previously derived the formula for the velocity of a falling body as $v = 32t$, which is charted in Table 8–4, and now we want to find out how much the velocity changes per second; this quantity is called the acceleration.

Acceleration is defined as the time rate of change of velocity. From the preceding discussion we know enough already to write the acceleration as the derivative dv/dt, in the same way that the velocity is the derivative of the distance. If we now differentiate the formula $v = 32t$ we obtain, for a falling body,

$$a = \frac{dv}{dt} = 32. \tag{8.9}$$

[To differentiate the term $32t$ we can utilize the result obtained in a previous problem, where we found that the derivative of Bt is simply B (a constant). So by letting $B = 32$, we have at once that the derivative of $32t$ is 32.] This means that the velocity of a falling body is changing by 32 feet per second, per second always. We also see from Table 8–4 that the velocity increases by 32 ft/sec in each second. This is a very simple case, for accelerations are usually not constant. The reason the acceleration is constant here is that the force on the falling body is constant, and Newton's law says that the acceleration is proportional to the force.

As a further example, let us find the acceleration in the problem we have already solved for the velocity. Starting with

$$s = At^3 + Bt + C$$

we obtained, for $v = ds/dt$,

$$v = 3At^2 + B.$$

Since acceleration is the derivative of the velocity with respect to the time, we need to differentiate the last expression above. Recall the rule that the derivative of the two terms on the right equals the sum of the derivatives of the individual terms. To differentiate the first of these terms, instead of going through the fundamental process again we note that we have already differentiated a quadratic term when we differentiated $16t^2$, and the effect was to double the numerical coefficient and change the t^2 to t; let us assume that the same thing will happen this time, and you can check the result yourself. The derivative of $3At^2$ will then be $6At$. Next we differentiate B, a constant term; but by a rule stated previously, the derivative of B is zero; hence this term contributes nothing to the acceleration. The final result, therefore, is $a = dv/dt = 6At$.

For reference, we state two very useful formulas, which can be obtained by integration. If a body starts from rest and moves with a constant acceleration, g, its velocity v at any time t is given by

$$v = gt.$$

The distance it covers in the same time is

$$s = \tfrac{1}{2}gt^2.$$

Various mathematical notations are used in writing derivatives. Since velocity is ds/dt and acceleration is the time derivative of the velocity, we can also write

$$a = \frac{d}{dt}\left(\frac{ds}{dt}\right) = \frac{d^2s}{dt^2}, \tag{8.10}$$

which are common ways of writing a second derivative.

We have another law that the velocity is equal to the integral of the acceleration. This is just the opposite of $a = dv/dt$; we have already seen that distance is the integral of the velocity, so distance can be found by twice integrating the acceleration.

In the foregoing discussion the motion was in only one dimension, and space permits only a brief discussion of motion in three dimensions. Consider a particle P which moves in three dimensions in any manner whatsoever. At the beginning of this chapter, we opened our discussion of the one-dimensional case of a moving car by observing the distance of the car from its starting point at various times. We then discussed velocity in terms of changes of these distances with time, and acceleration in terms of changes in velocity. We can treat three-dimensional motion analogously. It will be simpler to illustrate the motion on a two-dimensional diagram, and then extend the ideas to three dimensions. We establish a pair of axes at right angles to each other, and determine the position of the particle at any moment by measuring how far it is from each of the two axes. Thus each position is given in terms of an x-distance and a y-distance, and the motion can be described by constructing a table in which both these distances are given as functions of time. (Extension of this process to three dimensions requires only another axis, at right angles to the first two, and measuring a third distance, the z-distance. The distances are now measured from coordinate *planes* instead of lines.) Having constructed a table with x- and y-distances, how can we determine the velocity? We first find the components of velocity in each direction. The horizontal part of the velocity, or x-component, is the derivative of the x-distance with respect to the time, or

$$v_x = dx/dt. \tag{8.11}$$

Similarly, the vertical part of the velocity, or y-component, is

$$v_y = dy/dt. \tag{8.12}$$

In the third dimension,

$$v_z = dz/dt. \tag{8.13}$$

Now, given the components of velocity, how can we find the velocity along the actual path of motion? In the two-dimensional case, consider two successive positions of the particle, separated by a short distance Δs and a short time interval $t_2 - t_1 = \Delta t$. In the time Δt the particle moves horizontally a distance $\Delta x \sim v_x \Delta t$, and vertically a distance $\Delta y \sim v_y \Delta t$. (The symbol "$\sim$" is read "is approximately.") The actual distance moved is approximately

$$\Delta s \sim \sqrt{(\Delta x)^2 + (\Delta y)^2}, \tag{8.14}$$

as shown in Fig. 8–3. The approximate velocity during this interval can be obtained by dividing by Δt and by letting Δt go to 0, as at the beginning of the chapter.

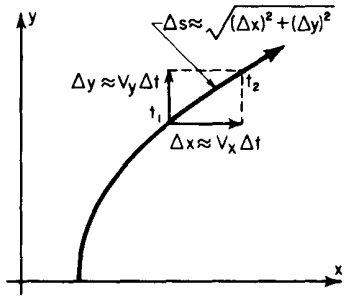


Fig. 8–3. Description of the motion of a body in two dimensions and the computation of its velocity.

We then get the velocity as

$$v = \frac{ds}{dt} = \sqrt{(dx/dt)^2 + (dy/dt^2)} = \sqrt{v_x^2 + v_y^2} \tag{8.15}$$

For three dimensions the result is

$$v = \sqrt{v_x^2 + v_y^2 + v_z^2}. \tag{8.16}$$

In the same way as we defined velocities, we can define accelerations: we have an x-component of acceleration a_x, which is the derivative of v_x, the x-component of the velocity (that is, $a_x = d^2x/dt^2$, the second derivative of x with respect to t), and so on.

Let us consider one nice example of compound motion in a plane. We shall take a motion in which a ball moves horizontally with a constant velocity u, and at the same time goes vertically downward with a constant acceleration $-g$; what is the motion? We can say $dx/dt = v_x = u$. Since the velocity v_x is constant,

$$x = ut, \tag{8.17}$$

and since the downward acceleration $-g$ is constant, the distance y the object falls can be written as

$$y = -\tfrac{1}{2}gt^2. \tag{8.18}$$

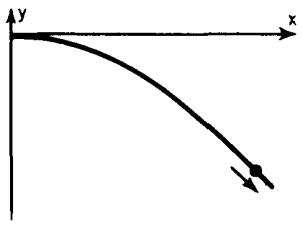


Fig. 8–4. The parabola described by a falling body with an initial horizontal velocity.

What is the curve of its path, i.e., what is the relation between y and x? We can eliminate t from Eq. (8.18), since $t = x/u$. When we make this substitution we find that

$$y = -\frac{g}{2u^2}x^2. \tag{8.19}$$

This relation between y and x may be considered as the equation of the path of the moving ball. When this equation is plotted we obtain a curve that is called a parabola; any freely falling body that is shot out in any direction will travel in a parabola, as shown in Fig. 8–4.

9

Newton's Laws of Dynamics

9–1 Momentum and force

9–2 Speed and velocity

9–3 Components of velocity, acceleration, and force

9–4 What is the force?

9–5 Meaning of the dynamical equations

9–6 Numerical solution of the equations

9–7 Planetary motions

9–1 Momentum and force

The discovery of the laws of dynamics, or the laws of motion, was a dramatic moment in the history of science. Before Newton's time, the motions of things like the planets were a mystery, but after Newton there was complete understanding. Even the slight deviations from Kepler's laws, due to the perturbations of the planets, were computable. The motions of pendulums, oscillators with springs and weights in them, and so on, could all be analyzed completely after Newton's laws were enunciated. So it is with this chapter: before this chapter we could not calculate how a mass on a spring would move; much less could we calculate the perturbations on the planet Uranus due to Jupiter and Saturn. After this chapter we *will* be able to compute not only the motion of the oscillating mass, but also the perturbations on the planet Uranus produced by Jupiter and Saturn!

Galileo made a great advance in the understanding of motion when he discovered the *principle of inertia:* if an object is left alone, is not disturbed, it continues to move with a constant velocity in a straight line if it was originally moving, or it continues to stand still if it was just standing still. Of course this never appears to be the case in nature, for if we slide a block across a table it stops, but that is because it is *not* left to itself—it is rubbing against the table. It required a certain imagination to find the right rule, and that imagination was supplied by Galileo.

Of course, the next thing which is needed is a rule for finding how an object *changes* its speed if something *is* affecting it. That is the contribution of Newton. Newton wrote down three laws: The First Law was a mere restatement of the Galilean principle of inertia just described. The Second Law gave a specific way of determining how the velocity changes under different influences called *forces.* The Third Law describes the forces to some extent, and we shall discuss that at another time. Here we shall discuss only the Second Law, which asserts that the motion of an object is changed by forces in this way: *the time-rate-of-change of a quantity called momentum is proportional to the force.* We shall state this mathematically shortly, but let us first explain the idea.

Momentum is not the same as *velocity.* A lot of words are used in physics, and they all have precise meanings in physics, although they may not have such precise meanings in everyday language. Momentum is an example, and we must define it precisely. If we exert a certain push with our arms on an object that is light, it moves easily; if we push just as hard on another object that is much heavier in the usual sense, then it moves much less rapidly. Actually, we must change the words from "light" and "heavy" to *less massive* and *more massive,* because there is a difference to be understood between the *weight* of an object and its *inertia.* (How hard it is to get it going is one thing, and how much it weighs is something else.) Weight and inertia are *proportional,* and on the earth's surface are often taken to be numerically equal, which causes a certain confusion to the student. On Mars, weights would be different but the amount of force needed to overcome inertia would be the same.

We use the term *mass* as a quantitative measure of inertia, and we may measure mass, for example, by swinging an object in a circle at a certain speed and measuring how much force we need to keep it in the circle. In this way we find a certain quantity of mass for every object. Now the *momentum* of an object is a product of two parts: its *mass* and its *velocity.* Thus Newton's Second Law may

be written mathematically this way:

$$F = \frac{d}{dt}(mv). \tag{9.1}$$

Now there are several points to be considered. In writing down any law such as this, we use many intuitive ideas, implications, and assumptions which are at first combined approximately into our "law." Later we may have to come back and study in greater detail exactly what each term means, but if we try to do this too soon we shall get confused. Thus at the beginning we take several things for granted. First, that the mass of an object is *constant*; it isn't really, but we shall start out with the Newtonian approximation that mass is constant, the same all the time, and that, further, when we put two objects together, their masses *add.* These ideas were of course implied by Newton when he wrote his equation, for otherwise it is meaningless. For example, suppose the mass varied inversely as the velocity; then the momentum would *never change* in any circumstance, so the law means nothing unless you know how the mass changes with velocity. At first we say, *it does not change.*

Then there are some implications concerning force. As a rough approximation we think of force as a kind of push or pull that we make with our muscles, but we can define it more accurately now that we have this law of motion. The most important thing to realize is that this relationship involves not only changes in the *magnitude* of the momentum or of the velocity but also in their *direction.* If the mass is constant, then Eq. (9.1) can also be written as

$$F = m\frac{dv}{dt} = ma. \tag{9.2}$$

The acceleration a is the rate of change of the velocity, and Newton's Second Law says more than that the effect of a given force varies inversely as the mass; it says also that the *direction* of the change in the velocity and the *direction* of the force are the same. Thus we must understand that a change in a velocity, or an acceleration, has a wider meaning than in common language: The velocity of a moving object can change by its speeding up, slowing down (when it slows down, we say it accelerates with a negative acceleration), or changing its direction of motion. An acceleration at right angles to the velocity was discussed in Chapter 7. There we saw that an object moving in a circle of radius R with a certain speed v along the circle falls away from a straightline path by a distance equal to $\frac{1}{2}(v^2/R)t^2$ if t is very small. Thus the formula for acceleration at right angles to the motion is

$$a = v^2/R, \tag{9.3}$$

and a force at right angles to the velocity will cause an object to move in a curved path whose radius of curvature can be found by dividing the force by the mass to get the acceleration, and then using (9.3).

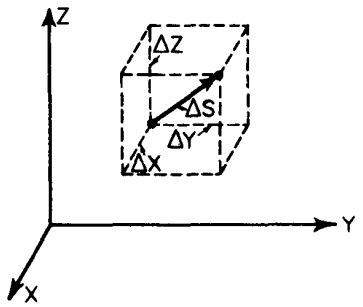


Fig. 9–1. A small displacement of an object.

9–2 Speed and velocity

In order to make our language more precise, we shall make one further definition in our use of the words *speed* and *velocity.* Ordinarily we think of speed and velocity as being the same, and in ordinary language they are the same. But in physics we have taken advantage of the fact that there *are* two words and have chosen to use them to distinguish two ideas. We carefully distinguish velocity, which has both magnitude and direction, from speed, which we choose to mean the magnitude of the velocity, but which does not include the direction. We can formulate this more precisely by describing how the x-, y-, and z-coordinates of an object change with time. Suppose, for example, that at a certain instant an object is moving as shown in Fig. 9–1. In a given small interval of time Δt it will move a certain distance Δx in the x-direction, Δy in the y-direction, and Δz in the z-direction. The total effect of these three coordinate changes is a displacement Δs along the diagonal of a parallelepiped whose sides are Δx, Δy, and Δz. In terms

of the velocity, the displacement Δx is the x-component of the velocity times Δt, and similarly for Δy and Δz:

$$\Delta x = v_x \Delta t, \qquad \Delta y = v_y \Delta t, \qquad \Delta z = v_z \Delta t. \tag{9.4}$$

9–3 Components of velocity, acceleration, and force

In Eq. (9.4) *we have resolved the velocity into components* by telling how fast the object is moving in the x-direction, the y-direction, and the z-direction. The velocity is completely specified, both as to magnitude and direction, if we give the numerical values of its three rectangular components:

$$v_x = dx/dt, \qquad v_y = dy/dt, \qquad v_z = dz/dt. \tag{9.5}$$

On the other hand, the speed of the object is

$$ds/dt = |v| = \sqrt{v_x^2 + v_y^2 + v_z^2}. \tag{9.6}$$

Next, suppose that, because of the action of a force, the velocity changes to some other direction and a different magnitude, as shown in Fig. 9–2. We can analyze this apparently complex situation rather simply if we evaluate the changes in the x-, y-, and z-components of velocity. The change in the component of the velocity in the x-direction in a time Δt is $\Delta v_x = a_x \Delta t$, where a_x is what we call the x-component of the acceleration. Similarly, we see that $\Delta v_y = a_y \Delta t$ and $\Delta v_z = a_z \Delta t$. In these terms, we see that Newton's Second Law, in saying that the force is in the same direction as the acceleration, is really three laws, in the sense that the component of the force in the x-, y-, or z-direction is equal to the mass times the rate of change of the corresponding component of velocity:

$$\begin{aligned} F_x &= m(dv_x/dt) = m(d^2x/dt^2) = ma_x, \\ F_y &= m(dv_y/dt) = m(d^2y/dt^2) = ma_y, \\ F_z &= m(dv_z/dt) = m(d^2z/dt^2) = ma_z. \end{aligned} \tag{9.7}$$

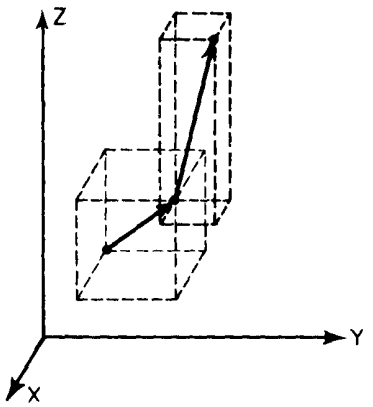


Fig. 9–2. A change in velocity in which both the magnitude and direction change.

Just as the velocity and acceleration have been resolved into components by projecting a line segment representing the quantity and its direction onto three coordinate axes, so, in the same way, a force in a given direction is represented by certain components in the x-, y-, and z-directions:

$$\begin{aligned} F_x &= F\cos(x, F), \\ F_y &= F\cos(y, F), \\ F_z &= F\cos(z, F), \end{aligned} \tag{9.8}$$

where F is the magnitude of the force and (x, F) represents the angle between the x-axis and the direction of F, etc.

Newton's Second Law is given in complete form in Eq. (9.7). If we know the forces on an object and resolve them into x-, y-, and z-components, then we can find the motion of the object from these equations. Let us consider a simple example. Suppose there are no forces in the y- and z-directions, the only force being in the x-direction, say vertically. Equation (9.7) tells us that there would be changes in the velocity in the vertical direction, but no changes in the horizontal direction. This was demonstrated with a special apparatus in Chapter 7 (see Fig. 7–3). A falling body moves horizontally without any change in horizontal motion, while it moves vertically the same way as it would move if the horizontal motion were zero. In other words, motions in the x-, y-, and z-directions are independent if the *forces* are not connected.

9–4 What is the force?

In order to use Newton's laws, we have to have some formula for the force; these laws say *pay attention to the forces.* If an object is accelerating, some agency is at work; find it. Our program for the future of dynamics must be to *find the*

laws for the force. Newton himself went on to give some examples. In the case of gravity he gave a specific formula for the force. In the case of other forces he gave some part of the information in his Third Law, which we will study in the next chapter, having to do with the equality of action and reaction.

Extending our previous example, what are the forces on objects near the earth's surface? Near the earth's surface, the force in the vertical direction due to gravity is proportional to the mass of the object and is nearly independent of height for heights small compared with the earth's radius R: $F = GmM/R^2 = mg$, where $g = GM/R^2$ is called the *acceleration of gravity*. Thus the law of gravity tells us that weight is proportional to mass; the force is in the vertical direction and is the mass times g. Again we find that the motion in the horizontal direction is at constant velocity. The interesting motion is in the vertical direction, and Newton's Second Law tells us

$$mg = m(d^2x/dt^2). \tag{9.9}$$

Cancelling the m's, we find that the acceleration in the x-direction is constant and equal to g. This is of course the well known law of free fall under gravity, which leads to the equations

$$\begin{aligned} v_x &= v_0 + gt, \\ x &= x_0 + v_0t + \tfrac{1}{2}gt^2. \end{aligned} \tag{9.10}$$

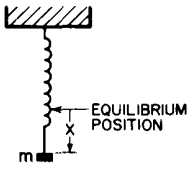


Fig. 9–3. A mass on a spring.

As another example, let us suppose that we have been able to build a gadget (Fig. 9–3) which applies a force proportional to the distance and directed oppositely —a spring. If we forget about gravity, which is of course balanced out by the initial stretch of the spring, and talk only about *excess* forces, we see that if we pull the mass down, the spring pulls up, while if we push it up the spring pulls down. This machine has been designed carefully so that the force is greater, the more we pull it up, in exact proportion to the displacement from the balanced condition, and the force upward is similarly proportional to how far we pull down. If we watch the dynamics of this machine, we see a rather beautiful motion—up, down, up, down, . . . The question is, will Newton's equations correctly describe this motion? Let us see whether we can exactly calculate how it moves with this periodic oscillation, by applying Newton's law (9.7). In the present instance, the equation is

$$-kx = m(dv_x/dt). \tag{9.11}$$

Here we have a situation where the velocity in the x-direction changes at a rate proportional to x. Nothing will be gained by retaining numerous constants, so we shall imagine either that the scale of time has changed or that there is an accident in the units, so that we happen to have $k/m = 1$. Thus we shall try to solve the equation

$$dv_x/dt = -x. \tag{9.12}$$

To proceed, we must know what v_x is, but of course we know that the velocity is the rate of change of the position.

9–5 Meaning of the dynamical equations

Now let us try to analyze just what Eq. (9.12) means. Suppose that at a given time t the object has a certain velocity v_x and position x. What is the velocity and what is the position at a slightly later time $t + \epsilon$? If we can answer this question our problem is solved, for then we can start with the given condition and compute how it changes for the first instant, the next instant, the next instant, and so on, and in this way we gradually evolve the motion. To be specific, let us suppose that at the time $t = 0$ we are given that $x = 1$ and $v_x = 0$. Why does the object move at all? Because there is a *force* on it when it is at any position except $x = 0$. If $x > 0$, that force is upward. Therefore the velocity which is zero starts to change, because of the law of motion. Once it starts to build up some velocity the object starts to move up, and so on. Now at any time t, if ϵ is very small,

we may express the position at time $t + \epsilon$ in terms of the position at time t and the velocity at time t to a very good approximation as

$$x(t + \epsilon) = x(t) + \epsilon v_x(t). \tag{9.13}$$

The smaller the ϵ, the more accurate this expression is, but it is still usefully accurate even if ϵ is not vanishingly small. Now what about the velocity? In order to get the velocity later, the velocity at the time $t + \epsilon$, we need to know how the velocity changes, the *acceleration*. And how are we going to find the acceleration? That is where the law of dynamics comes in. The law of dynamics tells us what the acceleration is. It says the acceleration is $-x$.

$$v_x(t + \epsilon) = v_x(t) + \epsilon a_x(t) \tag{9.14}$$

$$= v_x(t) - \epsilon x(t). \tag{9.15}$$

Equation (9.14) is merely kinematics; it says that a velocity changes because of the presence of acceleration. But Eq. (9.15) is *dynamics*, because it relates the acceleration to the force; it says that at this particular time for this particular problem, you can replace the acceleration by $-x(t)$. Therefore, if we know both the x and v at a given time, we know the acceleration, which tells us the new velocity, and we know the new position—this is how the machinery works. The velocity changes a little bit because of the force, and the position changes a little bit because of the velocity.

9–6 Numerical solution of the equations

Now let us really solve the problem. Suppose that we take $\epsilon = 0.100$ sec. After we do all the work if we find that this is not small enough we may have to go back and do it again with $\epsilon = 0.010$ sec. Starting with our initial value $x(0) = 1.00$, what is $x(0.1)$? It is the old position $x(0)$ plus the velocity (which is zero) times 0.10 sec. Thus $x(0.1)$ is still 1.00 because it has not yet started to move. But the new velocity at 0.10 sec will be the old velocity $v(0) = 0$ plus ϵ times the acceleration. The acceleration is $-x(0) = -1.00$. Thus

$$v(0.1) = 0.00 - 0.10 \times 1.00 = -0.10.$$

Now at 0.20 sec

$$x(0.2) = x(0.1) + \epsilon v(0.1)$$
$$= 1.00 - 0.10 \times 0.10 = 0.99$$

and

$$v(0.2) = v(0.1) + \epsilon a(0.1)$$
$$= -0.10 - 0.10 \times 1.00 = -0.20.$$

And so, on and on and on, we can calculate the rest of the motion, and that is just what we shall do. However, for practical purposes there are some little tricks by which we can increase the accuracy. If we continued this calculation as we have started it, we would find the motion only rather crudely because $\epsilon = 0.100$ sec is rather crude, and we would have to go to a very small interval, say $\epsilon = 0.01$. Then to go through a reasonable total time interval would take a lot of cycles of computation. So we shall organize the work in a way that will increase the precision of our calculations, using the same coarse interval $\epsilon = 0.10$ sec. This can be done if we make a subtle improvement in the technique of the analysis.

Notice that the new position is the old position plus the time interval ϵ times the velocity. But the velocity *when?* The velocity at the beginning of the time interval is one velocity and the velocity at the end of the time interval is another velocity. Our improvement is to use the velocity *halfway between*. If we know the speed now, but the speed is changing, then we are not going to get the right answer by going at the same speed as now. We should use some speed between the "now" speed and the "then" speed at the end of the interval. The same considerations also apply to the velocity: to compute the velocity changes, we

Table 9–1

Solution of $dv_x/dt = -x$

Interval: $\epsilon = 0.10$ sec

t	x	v_x	a_x
0.0	1.000	0.000	−1.000
		−0.050	
0.1	0.995		−0.995
		−0.150	
0.2	0.980		−0.980
		−0.248	
0.3	0.955		−0.955
		−0.343	
0.4	0.921		−0.921
		−0.435	
0.5	0.877		−0.877
		−0.523	
0.6	0.825		−0.825
		−0.605	
0.7	0.764		−0.764
		−0.682	
0.8	0.696		−0.696
		−0.751	
0.9	0.621		−0.621
		−0.814	
1.0	0.540		−0.540
		−0.868	
1.1	0.453		−0.453
		−0.913	
1.2	0.362		−0.362
		−0.949	
1.3	0.267		−0.267
		−0.976	
1.4	0.169		−0.169
		−0.993	
1.5	0.070		−0.070
		−1.000	
1.6	−0.030		+0.030

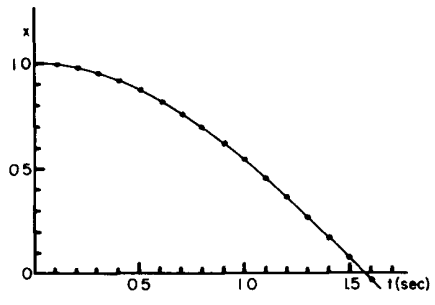


Fig. 9–4. Graph of the motion of a mass on a spring.

should use the acceleration midway between the two times at which the velocity is to be found. Thus the equations that we shall actually use will be something like this: the position later is equal to the position before plus ϵ times the velocity *at the time in the middle of the interval.* Similarly, the velocity at this halfway point is the velocity at a time ϵ before (which is in the middle of the previous interval) plus ϵ times the acceleration at the time t. That is, we use the equations

$$\begin{aligned} x(t + \epsilon) &= x(t) + \epsilon v(t + \epsilon/2), \\ v(t + \epsilon/2) &= v(t - \epsilon/2) + \epsilon a(t), \\ a(t) &= -x(t). \end{aligned} \tag{9.16}$$

There remains only one slight problem: what is $v(\epsilon/2)$? At the start, we are given $v(0)$, not $v(-\epsilon/2)$. To get our calculation started, we shall use a special equation, namely, $v(\epsilon/2) = v(0) + (\epsilon/2)a(0)$.

Now we are ready to carry through our calculation. For convenience, we may arrange the work in the form of a table, with columns for the time, the position, the velocity, and the acceleration, and the in-between lines for the velocity, as shown in Table 9–1. Such a table is, of course, just a convenient way of representing the numerical values obtained from the set of equations (9.16), and in fact the equations themselves need never be written. We just fill in the various spaces in the table one by one. This table now gives us a very good idea of the motion: it starts from rest, first picks up a little upward (negative) velocity and it loses some of its distance. The acceleration is then a little bit less but it is still gaining speed. But as it goes on it gains speed more and more slowly, until as it passes $x = 0$ at about $t = 1.50$ sec we can confidently predict that it will keep going, but now it will be on the other side; the position x will become negative, the acceleration therefore positive. Thus the speed decreases. It is interesting to compare these numbers with the function $x = \cos t$, which is done in Fig. 9–4. The agreement is within the three significant figure accuracy of our calculation! We shall see later that $x = \cos t$ is the exact mathematical solution of our equation of motion, but it is an impressive illustration of the power of numerical analysis that such an easy calculation should give such precise results.

9–7 Planetary motions

The above analysis is very nice for the motion of an oscillating spring, but can we analyze the motion of a planet around the sun? Let us see whether we can arrive at an approximation to an ellipse for the orbit. We shall suppose that the sun is infinitely heavy, in the sense that we shall not include its motion. Suppose a planet starts at a certain place and is moving with a certain velocity; it goes around the sun in some curve, and we shall try to analyze, by Newton's laws of motion and his law of gravitation, what the curve is. How? At a given moment it is at some position in space. If the radial distance from the sun to this position is called r, then we know that there is a force directed inward which, according to the law of gravity, is equal to a constant times the product of the sun's mass and the planet's mass divided by the square of the distance. To analyze this further we must find out what acceleration will be produced by this force. We shall need the *components* of the acceleration along two directions, which we call x and y. Thus if we specify the position of the planet at a given moment by giving x and y (we shall suppose that z is always zero because there is no force in the z-direction and, if there is no initial velocity v_z, there will be nothing to make z other than zero), the force is directed along the line joining the planet to the sun, as shown in Fig. 9–5.

From this figure we see that the horizontal component of the force is related to the complete force in the same manner as the horizontal distance x is to the complete hypotenuse r, because the two triangles are similar. Also, if x is positive, F_x is negative. That is, $F_x/|F| = -x/r$, or $F_x = -|F|x/r = -GMmx/r^3$. Now we use the dynamical law to find that this force component is equal to the mass of

the planet times the rate of change of its velocity in the x-direction. Thus we find the following laws:

$$\begin{aligned} m(dv_x/dt) &= -GMmx/r^3, \\ m(dv_y/dt) &= -GMmy/r^3, \\ r &= \sqrt{x^2 + y^2}. \end{aligned} \tag{9.17}$$

This, then, is the set of equations we must solve. Again, in order to simplify the numerical work, we shall suppose that the unit of time, or the mass of the sun, has been so adjusted (or luck is with us) that $GM \equiv 1$. For our specific example we shall suppose that the initial position of the planet is at $x = 0.500$ and $y = 0.000$, and that the velocity is all in the y-direction at the start, and is of magnitude 1.6300. Now how do we make the calculation? We again make a table with columns for the time, the x-position, the x-velocity v_x, and the x-acceleration a_x; then, separated by a double line, three columns for position, velocity, and acceleration in the y-direction. In order to get the accelerations we are going to need Eq. (9.17); it tells us that the acceleration in the x-direction is $-x/r^3$, and the acceleration in the y-direction is $-y/r^3$, and that r is the square root of $x^2 + y^2$. Thus, given x and y, we must do a little calculating on the side, taking the square root of the sum of the squares to find r and then, to get ready to calculate the two accelerations, it is useful also to evaluate $1/r^3$. This work can be done rather easily by using a table of squares, cubes, and reciprocals: then we need only multiply x by $1/r^3$, which we do on a slide rule.

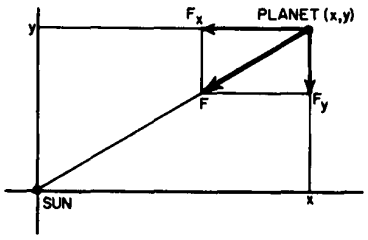


Fig. 9–5. The force of gravity on a planet.

Our calculation thus proceeds by the following steps, using time intervals $\epsilon = 0.100$: Initial values at $t = 0$:

$$\begin{aligned} x(0) &= 0.500 & y(0) &= 0.000 \\ v_x(0) &= 0.000 & v_y(0) &= +1.630 \end{aligned}$$

From these we find:

$$\begin{aligned} r(0) &= 0.500 & 1/r^3(0) &= 8.000 \\ a_x &= -4.000 & a_y &= 0.000 \end{aligned}$$

Thus we may calculate the velocities $v_x(0.05)$ and $v_y(0.05)$:

$$\begin{aligned} v_x(0.05) &= 0.000 - 4.000 \times 0.050 = -0.200; \\ v_y(0.05) &= 1.630 + 0.000 \times 0.100 = 1.630. \end{aligned}$$

Now our main calculations begin:

$$\begin{aligned} x(0.1) &= 0.500 - 0.20 \times 0.1 &&= 0.480 \\ y(0.1) &= 0.0 + 1.63 \times 0.1 &&= 0.163 \\ r &= \sqrt{0.480^2 + 0.163^2} &&= 0.507 \\ 1/r^3 &= 7.67 \\ a_x(0.1) &= 0.480 \times 7.67 &&= -3.68 \\ a_y(0.1) &= -0.163 \times 7.67 &&= -1.256 \\ v_x(0.15) &= -0.200 - 3.68 \times 0.1 &&= -0.568 \\ v_y(0.15) &= 1.630 - 1.26 \times 0.1 &&= 1.505 \\ x(0.2) &= 0.480 - 0.568 \times 0.1 &&= 0.423 \\ y(0.2) &= 0.163 + 1.50 \times 0.1 &&= 0.313 \end{aligned}$$

etc.

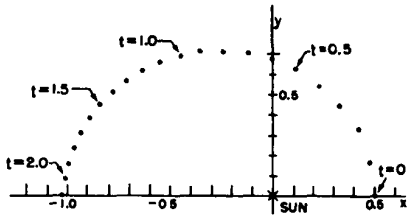


Fig. 9–6. The calculated motion of a planet around the sun.

In this way we obtain the values given in Table 9–2, and in 20 steps or so we have chased the planet halfway around the sun! In Fig. 9–6 are plotted the x- and y-coordinates given in Table 9–2. The dots represent the positions at the succession of times a tenth of a unit apart; we see that at the start the planet moves rapidly

Table 9–2

Solution of $dv_x/dt = -x/r^3$, $dv_y/dt = -y/r^3$, $r = \sqrt{x^2 + y^2}$.

Interval: $\epsilon = 0.100$

Orbit $v_y = 1.63$ $v_x = 0$ $x = 0.5$ $y = 0$ at $t = 0$

t	x	v_x	a_x	y	v_y	a_y	r	$1/r^3$
0.0	0.500		−4.00	0.000		0.00	0.500	8.000
		−0.200			1.630			
0.1	0.480		−3.68	0.163		−1.25	0.507	7.675
		−0.568			1.505			
0.2	0.423		−2.91	0.313		−2.15	0.526	6.873
		−0.859			1.290			
0.3	0.337		−1.96	0.442		−2.57	0.556	5.824
		−1.055			1.033			
0.4	0.232		−1.11	0.545		−2.62	0.592	4.81
		−1.166			0.771			
0.5	0.115		−0.453	0.622		−2.45	0.633	3.942
		−1.211			0.526			
0.6	−0.006		+0.020	0.675		−2.20	0.675	3.252
		−1.209			0.306			
0.7	−0.127		+0.344	0.706		−1.91	0.717	2.712
		−1.175			0.115			
0.8	−0.245		+0.562	0.718		−1.64	0.758	2.296
		−1.119			−0.049			
0.9	−0.357		+0.705	0.713		−1.41	0.797	1.975
		−1.048			−0.190			
1.0	−0.462		+0.796	0.694		−1.20	0.834	1.723
		−0.968			−0.310			
1.1	−0.559		+0.858	0.663		−1.02	0.867	1.535
		−0.882			−0.412			
1.2	−0.647		+0.90	0.622		−0.86	0.897	1.385
		−0.792			−0.499			
1.3	−0.726		+0.92	0.572		−0.72	0.924	1.267
		−0.700			−0.570			
1.4	−0.796		+0.93	0.515		−0.60	0.948	1.173
		−0.607			−0.630			
1.5	−0.857		+0.94	0.452		−0.50	0.969	1.099
		−0.513			−0.680			
1.6	−0.908		+0.95	0.384		−0.40	0.986	1.043
		−0.418			−0.720			
1.7	−0.950		+0.95	0.312		−0.31	1.000	1.000
		−0.323			−0.751			
1.8	−0.982		+0.95	0.237		−0.23	1.010	0.970
		−0.228			−0.773			
1.9	−1.005		+0.95	0.160		−0.15	1.018	0.948
		−0.113			−0.778			
2.0	−1.018		+0.96	0.081		−0.08	1.021	0.939
		−0.037			−0.796			
2.1	−1.022		+0.95	0.001		0.00	1.022	0.936
		+0.058			−0.796			
2.2	−1.016		+0.96	−0.079		+0.07	1.019	0.945
					−0.789			
2.3								

Crossed x-axis at 2.101 sec, $\therefore$ period = 4.20 sec.

$v_x = 0$ at 2.086 sec.

Cross x at 1.022, $\therefore$ semimajor axis $= \dfrac{1.022 + 0.500}{2} = 0.761.$

$v_y = 0.796.$

Predicted time $\pi(0.761)^{3/2} = \pi(0.663) = 2.082.$

and at the end it moves slowly, and so the shape of the curve is determined. Thus we see that we *really do* know how to calculate the motion of planets!

Now let us see how we can calculate the motion of Neptune, Jupiter, Uranus, or any other planet. If we have a great many planets, and let the sun move too, can we do the same thing? Of course we can. We calculate the force on a particular planet, let us say planet number i, which has a position x_i, y_i, z_i ($i = 1$ may represent the sun, $i = 2$ Mercury, $i = 3$ Venus, and so on). We must know the positions of all the planets. The force acting on one is due to all the other bodies which are located, let us say, at positions x_j, y_j, z_j. Therefore the equations are

$$\begin{aligned} m_i \frac{dv_{ix}}{dt} &= \sum_{j=1}^{N} -\frac{Gm_i m_j (x_i - x_j)}{r_{ij}^3}, \\ m_i \frac{dv_{iy}}{dt} &= \sum_{j=1}^{N} -\frac{Gm_i m_j (y_i - y_j)}{r_{ij}^3}, \\ m_i \frac{dv_{iz}}{dt} &= \sum_{j=1}^{N} -\frac{Gm_i m_j (z_i - z_j)}{r_{ij}^3}. \end{aligned} \tag{9.18}$$

Further, we define r_{ij} as the distance between the two planets i and j; this is equal to

$$r_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2}. \tag{9.19}$$

Also, $\sum$ means a sum over all values of j—all other bodies—except, of course, for $j = i$. Thus all we have to do is to make more columns, *lots* more columns. We need nine columns for the motions of Jupiter, nine for the motions of Saturn, and so on. Then when we have all initial positions and velocities we can calculate all the accelerations from Eq. (9.18) by first calculating all the distances, using Eq. (9.19). How long will it take to do it? If you do it at home, it will take a very long time! But in modern times we have machines which do arithmetic very rapidly; a very good computing machine may take 1 microsecond, that is, a millionth of a second, to do an addition. To do a multiplication takes longer, say 10 microseconds. It may be that in one cycle of calculation, depending on the problem, we may have 30 multiplications, or something like that, so one cycle will take 300 microseconds. That means that we can do 3000 cycles of computation per second. In order to get an accuracy, of, say, one part in a billion, we would need 4×10^5 cycles to correspond to one revolution of a planet around the sun. That corresponds to a computation time of 130 seconds or about two minutes. Thus it take only two minutes to follow Jupiter around the sun, with all the perturbations of all the planets correct to one part in a billion, by this method! (It turns out that the error varies about as the square of the interval ϵ. If we make the interval a thousand times smaller, it is a million times more accurate. So, let us make the interval 10,000 times smaller.)

So, as we said, we began this chapter not knowing how to calculate even the motion of a mass on a spring. Now, armed with the tremendous power of Newton's laws, we can not only calculate such simple motions but also, given only a machine to handle the arithmetic, even the tremendously complex motions of the planets, to as high a degree of precision as we wish!

10

Conservation of Momentum

10–1 Newton's Third Law

10–2 Conservation of momentum

10–3 Momentum *is* conserved!

10–4 Momentum and energy

10–5 Relativistic momentum

10–1 Newton's Third Law

On the basis of Newton's second law of motion, which gives the relation between the acceleration of any body and the force acting on it, any problem in mechanics can be solved in principle. For example, to determine the motion of a few particles, one can use the numerical method developed in the preceding chapter. But there are good reasons to make a further study of Newton's laws. First, there are quite simple cases of motion which can be analyzed not only by numerical methods, but also by direct mathematical analysis. For example, although we know that the acceleration of a falling body is 32 ft/sec^2, and from this fact could calculate the motion by numerical methods, it is much easier and more satisfactory to analyze the motion and find the general solution, $s = s_0 + v_0 t + 16t^2$. In the same way, although we can work out the positions of a harmonic oscillator by numerical methods, it is also possible to show analytically that the general solution is a simple cosine function of t, and so it is unnecessary to go to all that arithmetical trouble when there is a simple and more accurate way to get the result. In the same manner, although the motion of one body around the sun, determined by gravitation, can be calculated point by point by the numerical methods of Chapter 9, which show the general shape of the orbit, it is nice also to get the exact shape, which analysis reveals as a perfect ellipse.

Unfortunately, there are really very few problems which can be solved exactly by analysis. In the case of the harmonic oscillator, for example, if the spring force is not proportional to the displacement, but is something more complicated, one must fall back on the numerical method. Or if there are two bodies going around the sun, so that the total number of bodies is three, then analysis cannot produce a simple formula for the motion, and in practice the problem must be done numerically. That is the famous three-body problem, which so long challenged human powers of analysis; it is very interesting how long it took people to appreciate the fact that perhaps the powers of mathematical analysis were limited and it might be necessary to use the numerical methods. Today an enormous number of problems that cannot be done analytically are solved by numerical methods, and the old three-body problem, which was supposed to be so difficult, is solved as a matter of routine in exactly the same manner that was described in the preceding chapter, namely, by doing enough arithmetic. However, there are also situations where both methods fail: the simple problems we can do by analysis, and the moderately difficult problems by numerical, arithmetical methods, but the very complicated problems we cannot do by either method. A complicated problem is, for example, the collision of two automobiles, or even the motion of the molecules of a gas. There are countless particles in a cubic millimeter of gas, and it would be ridiculous to try to make calculations with so many variables (about 10^{17}—a hundred million billion). Anything like the motion of the molecules or atoms of a gas or a block or iron, or the motion of the stars in a globular cluster, instead of just two or three planets going around the sun—such problems we cannot do directly, so we have to seek other means.

In the situations in which we cannot follow details, we need to know some general properties, that is, general theorems or principles which are consequences of Newton's laws. One of these is the principle of conservation of energy, which was discussed in Chapter 4. Another is the principle of conservation of momentum, the subject of this chapter. Another reason for studying mechanics further is that there are certain patterns of motion that are repeated in many different circum-

stances, so it is good to study these patterns in one particular circumstance. For example, we shall study collisions; different kinds of collisions have much in common. In the flow of fluids, it does not make much difference what the fluid is, the laws of the flow are similar. Other problems that we shall study are vibrations and oscillations and, in particular, the peculiar phenomena of mechanical waves—sound, vibrations of rods, and so on.

In our discussion of Newton's laws it was explained that these laws are a kind of program that says "Pay attention to the forces," and that Newton told us only two things about the nature of forces. In the case of gravitation, he gave us the complete law of the force. In the case of the very complicated forces between atoms, he was not aware of the right laws for the forces; however, he discovered one rule, one general property of forces, which is expressed in his Third Law, and that is the total knowledge that Newton had about the nature of forces—the law of gravitation and this principle, but no other details.

This principle is that *action equals reaction.*

What is meant is something of this kind: Suppose we have two small bodies, say particles, and suppose that the first one exerts a force on the second one, pushing it with a certain force. Then, simultaneously, according to Newton's Third Law, the second particle will push on the first with an equal force, in the opposite direction; furthermore, these forces effectively act in the same line. This is the hypothesis, or law, that Newton proposed, and it seems to be quite accurate, though not exact (we shall discuss the errors later). For the moment we shall take it to be true that action equals reaction. Of course, if there is a third particle, not on the same line as the other two, the law does *not* mean that the total force on the first one is equal to the total force on the second, since the third particle, for instance, exerts its own push on each of the other two. The result is that the total effect on the first two is in some other direction, and the forces on the first two particles are, in general, neither equal nor opposite. However, the forces on each particle can be resolved into parts, there being one contribution or part due to each other interacting particle. Then each *pair* of particles has corresponding components of mutual interaction that are equal in magnitude and opposite in direction.

10–2 Conservation of momentum

Now what are the interesting consequences of the above relationship? Suppose, for simplicity, that we have just two interacting particles, possibly of different mass, and numbered 1 and 2. The forces between them are equal and opposite; what are the consequences? According to Newton's Second Law, force is the time rate of change of the momentum, so we conclude that the rate of change of momentum p_1 of particle 1 is equal to minus the rate of change of momentum p_2 of particle 2, or

$$dp_1/dt = -dp_2/dt. \tag{10.1}$$

Now if the *rate of change* is always equal and opposite, it follows that the *total change* in the momentum of particle 1 is equal and opposite to the *total change* in the momentum of particle 2; this means that if we *add* the momentum of particle 1 to the momentum of particle 2, the rate of change of the sum of these, due to the mutual forces (called internal forces) between particles, is zero; that is

$$d(p_1 + p_2)/dt = 0. \tag{10.2}$$

There is assumed to be no other force in the problem. If the rate of change of this sum is always zero, that is just another way of saying that the quantity $(p_1 + p_2)$ does not change. (This quantity is also written $m_1v_1 + m_2v_2$, and is called the *total momentum* of the two particles.) We have now obtained the result that the total momentum of the two particles does not change because of any mutual interactions between them. This statement expresses the law of conservation of

momentum in that particular example. We conclude that if there is any kind of force, no matter how complicated, between two particles, and we measure or calculate $m_1v_1 + m_2v_2$, that is, the sum of the two momenta, both before and after the forces act, the results should be equal, i.e., the total momentum is a constant.

If we extend the argument to three or more interacting particles in more complicated circumstances, it is evident that so far as internal forces are concerned, the total momentum of all the particles stays constant, since an increase in momentum of one, due to another, is exactly compensated by the decrease of the second, due to the first. That is, all the internal forces will balance out, and therefore cannot change the total momentum of the particles. Then if there are no forces from the outside (external forces), there are no forces that can change the total momentum; hence the total momentum is a constant.

It is worth describing what happens if there are forces that do *not* come from the mutual actions of the particles in question: suppose we isolate the interacting particles. If there are only mutual forces, then, as before, the total momentum of the particles does not change, no matter how complicated the forces. On the other hand, suppose there are also forces coming from the particles outside the isolated group. Any force exerted by outside bodies on inside bodies, we call an *external* force. We shall later demonstrate that the sum of all external forces equals the rate of change of the total momentum of all the particles inside, a very useful theorem.

The conservation of the total momentum of a number of interacting particles can be expressed as

$$m_1v_1 + m_2v_2 + m_3v_3 + \cdots = \text{a constant}, \tag{10.3}$$

if there are no net external forces. Here the masses and corresponding velocities of the particles are numbered 1, 2, 3, 4, . . . The general statement of Newton's Second Law for each particle,

$$f = \frac{d}{dt}(mv), \tag{10.4}$$

is true specifically for the *components* of force and momentum in any given direction; thus the x-component of the force on a particle is equal to the x-component of the rate of change of momentum of that particle, or

$$f_x = \frac{d}{dt}(mv_x), \tag{10.5}$$

and similarly for the y- and z-directions. Therefore Eq. (10.3) is really three equations, one for each direction.

In addition to the law of conservation of momentum, there is another interesting consequence of Newton's Second Law, to be proved later, but merely stated now. This principle is that the laws of physics will look the same whether we are standing still or moving with a uniform speed in a straight line. For example, a child bouncing a ball in an airplane finds that the ball bounces the same as though he were bouncing it on the ground. Even though the airplane is moving with a very high velocity, unless it changes its velocity, the laws look the same to the child as they do when the airplane is standing still. This is the so-called *relativity principle*. As we use it here we shall call it "Galilean relativity" to distinguish it from the more careful analysis made by Einstein, which we shall study later.

We have just derived the law of conservation of momentum from Newton's laws, and we could go on from here to find the special laws that describe impacts and collisions. But for the sake of variety, and also as an illustration of a kind of reasoning that can be used in physics in other circumstances where, for example, one might not know Newton's laws and might take a different approach, we shall discuss the laws of impacts and collisions from a completely different point of view. We shall base our discussion on the principle of Galilean relativity, stated above, and shall end up with the law of conservation of momentum.

We shall start by assuming that nature would look the same if we run along at a certain speed and watch it as it would if we were standing still. Before dis-

cussing collisions in which two bodies collide and stick together, or come together and bounce apart, we shall first consider two bodies that are held together by a spring or something else, and are then suddenly released and pushed by the spring or perhaps by a little explosion. Further, we shall consider motion in only one direction. First, let us suppose that the two objects are exactly the same, are nice symmetrical objects, and then we have a little explosion between them. After the explosion, one of the bodies will be moving, let us say toward the right, with a velocity v. Then it appears reasonable that the other body is moving toward the left with a velocity v, because if the objects are alike there is no reason for right or left to be preferred and so the bodies would do something that is symmetrical. This is an illustration of a kind of thinking that is very useful in many problems but would not be brought out if we just started with the formulas.

The first result from our experiment is that equal objects will have equal speed, but now suppose that we have two objects made of different materials, say copper and aluminum, and we make the two *masses* equal. We shall now suppose that if we do the experiment with two masses that are equal, even though the objects are not identical, the velocities will be equal. Someone might object: "But you know, you could do it backwards, you did not have to *suppose* that. You could *define* equal masses to mean two masses that acquire equal velocities in this experiment." We follow that suggestion and make a little explosion between the copper and a very large piece of aluminum, so heavy that the copper flies out and the aluminum hardly budges. That is too much aluminum, so we reduce the amount until there is just a very tiny piece, then when we make the explosion the aluminum goes flying away, and the *copper* hardly budges. That is not enough aluminum. Evidently there is some right amount in between; so we keep adjusting the amount until the velocities come out equal. Very well then—let us turn it around, and say that when the velocities are equal, the masses are equal. This appears to be just a definition, and it seems remarkable that we can transform physical laws into mere definitions. Nevertheless, there *are* some physical laws involved, and if we accept this definition of equal masses, we immediately find one of the laws, as follows.

Suppose we know from the foregoing experiment that two pieces of matter, A and B (of copper and aluminum), have equal masses, and we compare a third body, say a piece of gold, with the copper in the same manner as above, making sure that its mass is equal to the mass of the copper. If we now make the experiment between the aluminum and the gold, there is nothing in logic that says *these* masses must be equal; however, the *experiment* shows that they actually are. So now, by experiment, we have found a new law. A statement of this law might be: If two masses are each equal to a third mass (as determined by equal velocities in this experiment), then they are equal to each other. (This statement does *not* follow at all from a similar statement used as a postulate regarding *mathematical* quantities.) From this example we can see how quickly we start to infer things if we are careless. It is *not* just a definition to say the masses are equal when the velocities are equal, because to say the masses are equal is to imply the mathematical laws of equality, which in turn makes a prediction about an experiment.

As a second example, suppose that A and B are found to be equal by doing the experiment with one strength of explosion, which gives a certain velocity; if we then use a stronger explosion, will it be true or not true that the velocities now obtained are equal? Again, in logic there is nothing that can decide this question, but experiment shows that it *is* true. So, here is another law, which might be stated: If two bodies have equal masses, as measured by equal velocities at one velocity, they will have equal masses when measured at another velocity. From these examples we see that what appeared to be only a definition really involved some laws of physics.

In the development that follows we shall assume it is true that equal masses have equal and opposite velocities when an explosion occurs between them. We shall make another assumption in the inverse case: If two identical objects, moving in opposite directions with equal velocities, collide and stick together by some kind of glue, then which way will they be moving after the collision? This is again a

symmetrical situation, with no preference between right and left, so we assume that they stand still. We shall also suppose that any two objects of equal mass, even if the objects are made of different materials, which collide and stick together, when moving with the same velocity in opposite directions will come to rest after the collision.

10–3 Momentum *is* conserved!

We can verify the above assumptions experimentally: first, that if two stationary objects of equal mass are separated by an explosion they will move apart with the same speed, and second, if two objects of equal mass, coming together with the same speed, collide and stick together they will stop. This we can do by means of a marvelous invention called an air trough,* which gets rid of friction, the thing which continually bothered Galileo (Fig. 10–1). He could not do experiments by sliding things because they do not slide freely, but, by adding a magic touch, we can today get rid of friction. Our objects will slide without difficulty, on and on at a constant velocity, as advertised by Galileo. This is done by supporting the objects on air. Because air has very low friction, an object glides along with practically constant velocity when there is no applied force. First, we use two glide blocks which have been made carefully to have the same weight, or mass (their weight was measured really, but we know that this weight is proportional to the mass), and we place a small explosive cap in a closed cylinder between the two blocks (Fig. 10–2). We shall start the blocks from rest at the center point of the track and force them apart by exploding the cap with an electric spark. What should happen? If the speeds are equal when they fly apart, they should arrive at the ends of the trough at the same time. On reaching the ends they will both bounce back with practically opposite velocity, and will come together and stop at the center where they started. It is a good test; when it is actually done the result is just as we have described (Fig. 10–3).

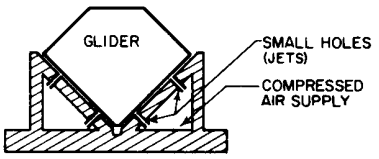


Fig. 10–1. End view of linear air trough.

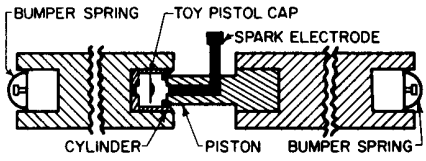


Fig. 10–2. Sectional view of gliders with explosive interaction cylinder attachment.

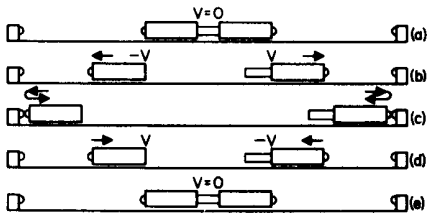


Fig. 10–3. Schematic view of action-reaction experiment with equal masses.

Now the next thing we would like to figure out is what happens in a less simple situation. Suppose we have two equal masses, one moving with velocity v and the other standing still, and they collide and stick; what is going to happen? There is a mass $2m$ altogether when we are finished, drifting with an unknown velocity. What velocity? That is the problem. To find the answer, we make the assumption that if we ride along in a car, physics will look the same as if we are standing still. We start with the knowledge that two equal masses, moving in opposite directions with equal speeds v, will stop dead when they collide. Now suppose that while this happens, we are riding by in an automobile, at a velocity $-v$. Then what does it look like? Since we are riding along with one of the two masses which are coming together, that one appears to us to have zero velocity. The other mass, however, going the other way with velocity v, will appear to be coming toward us at a velocity $2v$ (Fig. 10–4). Finally, the combined masses after collision will seem to be passing by with velocity v. We therefore conclude that an object with velocity $2v$, hitting an equal one at rest, will end up with velocity v, or what is mathematically exactly the same, an object with velocity v hitting and sticking to one at rest will produce an object moving with velocity $v/2$. Note that if we multiply the mass and the velocity beforehand and add them together, $mv + 0$, we get the same answer as when we multiply the mass and the velocity of everything afterwards, $2m$ times $v/2$. So that tells us what happens when a mass of velocity v hits one standing still.

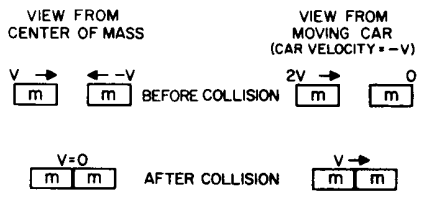


Fig. 10–4. Two views of an inelastic collision between equal masses.

In exactly the same manner we can deduce what happens when equal objects having *any* two velocities hit each other.

Suppose we have two equal bodies with velocities v_1 and v_2, respectively, which collide and stick together. What is their velocity v after the collision? Again we ride by in an automobile, say at velocity v_2, so that one body appears to be at rest. The other then appears to have a velocity $v_1 - v_2$, and we have the same case that we had before. When it is all finished they will be moving at $\frac{1}{2}(v_1 - v_2)$ with respect to the car. What then is the actual speed on the ground?

---

* H. V. Neher and R. B. Leighton, *Amer. Jour. of Phys.* **31**, 255 (1963).

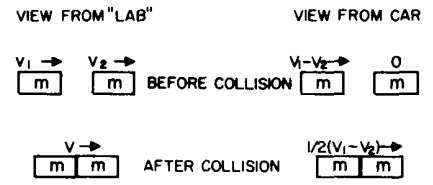


Fig. 10–5. Two views of another inelastic collision between equal masses.

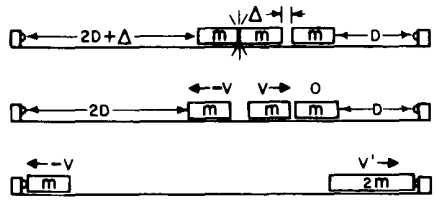


Fig. 10–6. An experiment to verify that a mass m with velocity v striking a mass m with zero velocity gives $2m$ with velocity $v/2$.

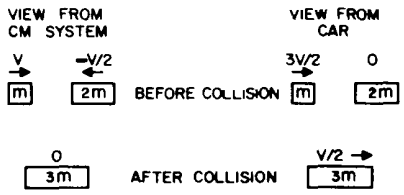


Fig. 10–7. Two views of an inelastic collision between m and $2m$.

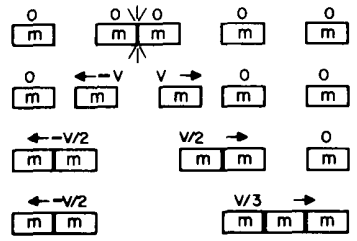


Fig. 10–8. Action and reaction between $2m$ and $3m$.

It is $v = \frac{1}{2}(v_1 - v_2) + v_2$ or $\frac{1}{2}(v_1 + v_2)$ (Fig. 10–5). Again we note that

$$mv_1 + mv_2 = 2m(v_1 + v_2)/2. \tag{10.6}$$

Thus, using this principle, we can analyze any kind of collision in which two bodies of equal mass hit each other and stick. In fact, although we have worked only in one dimension, we can find out a great deal about much more complicated collisions by imagining that we are riding by in a car in some oblique direction. The principle is the same, but the details get somewhat complicated.

In order to test experimentally whether an object moving with velocity v, colliding with an equal one at rest, forms an object moving with velocity $v/2$, we may perform the following experiment with our air-trough apparatus. We place in the trough three equally massive objects, two of which are initially joined together with our explosive cylinder device, the third being very near to but slightly separated from these and provided with a sticky bumper so that it will stick to another object which hits it. Now, a moment after the explosion, we have two objects of mass m moving with equal and opposite velocities v. A moment after that, one of these collides with the third object and makes an object of mass $2m$ moving, so we believe, with velocity $v/2$. How do we test whether it is really $v/2$? By arranging the initial positions of the masses on the trough so that the distances to the ends are not equal, but are in the ratio 2:1. Thus our first mass, which continues to move with velocity v, should cover twice as much distance in a given time as the two which are stuck together (allowing for the small distance travelled by the second object before it collided with the third). The mass m and the mass $2m$ should reach the ends at the same time, and when we try it, we find that they do (Fig. 10–6).

The next problem that we want to work out is what happens if we have two different masses. Let us take a mass m and a mass $2m$ and apply our explosive interaction. What will happen then? If, as a result of the explosion, m moves with velocity v, with what velocity does $2m$ move? The experiment we have just done may be repeated with zero separation between the second and third masses, and when we try it we get the same result, namely, the reacting masses m and $2m$ attain velocities $-v$ and $v/2$. Thus the direct reaction between m and $2m$ gives the same result as the symmetrical reaction between m and m, followed by a collision between m and a third mass m in which they stick together. Furthermore, we find that the masses m and $2m$ returning from the ends of the trough, with their velocities (nearly) exactly reversed, stop dead if they stick together.

Now the next question we may ask is this. What will happen if a mass m with velocity v, say, hits and sticks to another mass $2m$ at rest? This is very easy to answer using our principle of Galilean relativity, for we simply watch the collision which we have just described from a car moving with velocity $-v/2$ (Fig. 10–7). From the car, the velocities are

$$v_1' = v - v(\text{car}) = v + v/2 = 3v/2$$

and

$$v_2' = -v/2 - v(\text{car}) = -v/2 + v/2 = 0.$$

After the collision, the mass $3m$ appears to us to be moving with velocity $v/2$. Thus we have the answer, i.e., the ratio of velocities before and after collision is 3 to 1: if an object of mass m collides with a stationary object of mass $2m$, then the whole thing moves off, stuck together, with a velocity 1/3 as much. The general rule again is that the sum of the products of the masses and the velocities stays the same: $mv + 0$ equals $3m$ times $v/3$, so we are gradually building up the theorem of the conservation of momentum, piece by piece.

Now we have one against two. Using the same arguments, we can predict the result of one against three, two against three, etc. The case of two against three, starting from rest, is shown in Fig. 10–8.

In every case we find that the mass of the first object times its velocity, plus the mass of the second object times its velocity, is equal to the total mass of the final object times its velocity. These are all examples, then, of the conservation of

momentum. Starting from simple, symmetrical cases, we have demonstrated the law for more complex cases. We could, in fact, do it for any rational mass ratio, and since every ratio is exceedingly close to a rational ratio, we can handle every ratio as precisely as we wish.

10–4 Momentum and energy

All the foregoing examples are simple cases where the bodies collide and stick together, or were initially stuck together and later separated by an explosion. However, there are situations in which the bodies do *not* cohere, as, for example, two bodies of equal mass which collide with equal speeds and then rebound. For a brief moment they are in contact and both are compressed. At the instant of maximum compression they both have zero velocity and energy is stored in the elastic bodies, as in a compressed spring. This energy is derived from the kinetic energy the bodies had before the collision, which becomes zero at the instant their velocity is zero. The loss of kinetic energy is only momentary, however. The compressed condition is analogous to the cap that releases energy in an explosion. The bodies are immediately decompressed in a kind of explosion, and fly apart again; but we already know that case—the bodies fly apart with equal speeds. However, this speed of rebound is less, in general, than the initial speed, because not all the energy is available for the explosion, depending on the material. If the material is putty no kinetic energy is recovered, but if it is something more rigid, some kinetic energy is usually regained. In the collision the rest of the kinetic energy is transformed into heat and vibrational energy—the bodies are hot and vibrating. The vibrational energy also is soon transformed into heat. It is possible to make the colliding bodies from highly elastic materials, such as steel, with carefully designed spring bumpers, so that the collision generates very little heat and vibration. In these circumstances the velocities of rebound are practically equal to the initial velocities; such a collision is called *elastic*.

That the velocities *before* and *after* an elastic collision are equal is not a matter of conservation of momentum, but a matter of conservation of *kinetic energy*. That the speeds of the bodies rebounding after a symmetrical collision are equal to *each other*, however, is a matter of conservation of momentum.

We might similarly analyze collisions between bodies of different masses, different initial velocities, and various degrees of elasticity, and determine the final velocities and the loss of kinetic energy, but we shall not go into the details of these processes.

Elastic collisions are especially interesting for systems that have no internal "gears, wheels, or parts." Then when there is a collision there is nowhere for the energy to be impounded, because the objects that move apart are in the same condition as when they collided. Therefore, between very elementary objects, the collisions are always elastic or very nearly elastic. For instance, the collisions between atoms or molecules in a gas are said to be perfectly elastic. Although this is an excellent approximation, even such collisions are not *perfectly* elastic; otherwise one could not understand how energy in the form of light or heat radiation could come out of a gas. Once in a while, in a gas collision, a low-energy infrared ray is emitted, but this occurrence is very rare and the energy emitted is very small. So, for most purposes, collisions of molecules in gases are considered to be perfectly elastic.

As an interesting example, let us consider an *elastic* collision between two objects of *equal mass*. If they come together with the same speed, they would come apart at that same speed, by symmetry. But now look at this in another circumstance, in which one of them is moving with velocity v and the other one is at rest. What happens? We have been through this before. We watch the symmetrical collision from a car moving along with one of the objects, and we find that if a stationary body is struck elastically by another body of exactly the same mass, the moving body stops, and the one that was standing still now moves away with the same speed that the other one had; the bodies simply exchange velocities. This behavior can easily be demonstrated with a suitable impact apparatus. More

generally, if both bodies are moving, with different velocities, they simply exchange velocity at impact.

Another example of an almost elastic interaction is magnetism. If we arrange a pair of U-shaped magnets in our glide blocks, so that they repel each other, when one drifts quietly up to the other, it pushes it away and stands perfectly still, and now the other goes along, frictionlessly.

The principle of conservation of momentum is very useful, because it enables us to solve many problems without knowing the details. We did not know the details of the gas motions in the cap explosion, yet we could predict the velocities with which the bodies came apart, for example. Another interesting example is rocket propulsion. A rocket of large mass, M, ejects a small piece, of mass m, with a terrific velocity V relative to the rocket. After this the rocket, if it were originally standing still, will be moving with a small velocity, v. Using the principle of conservation of momentum, we can calculate this velocity to be

$$v = \frac{m}{M} \cdot V.$$

So long as material is being ejected, the rocket continues to pick up speed. Rocket propulsion is essentially the same as the recoil of a gun: there is no need for any air to push against.

10–5 Relativistic momentum

In modern times the law of conservation of momentum has undergone certain modifications. However, the law is still true today, the modifications being mainly in the definitions of things. In the theory of relativity it turns out that we do have conservation of momentum; the particles have mass and the momentum is still given by mv, the mass times the velocity, *but the mass changes with the velocity*, hence the momentum also changes. The mass varies with velocity according to the law

$$m = \frac{m_0}{\sqrt{1 - v^2/c^2}}, \tag{10.7}$$

where m_0 is the mass of the body at rest and c is the speed of light. It is easy to see from the formula that there is negligible difference between m and m_0 unless v is very large, and that for ordinary velocities the expression for momentum reduces to the old formula.

The components of momentum for a single particle are written as

$$p_x = \frac{m_0 v_x}{\sqrt{1 - v^2/c^2}}, \quad p_y = \frac{m_0 v_y}{\sqrt{1 - v^2/c^2}}, \quad p_z = \frac{m_0 v_z}{\sqrt{1 - v^2/c^2}}, \tag{10.8}$$

where $v^2 = v_x^2 + v_y^2 + v_z^2$. If the x-components are summed over all the interacting particles, both before and after a collision, the sums are equal; that is, momentum is conserved in the x-direction. The same holds true in any direction.

In Chapter 4 we saw that the law of conservation of energy is not valid unless we recognize that energy appears in different forms, electrical energy, mechanical energy, radiant energy, heat energy, and so on. In some of these cases, heat energy for example, the energy might be said to be "hidden." This example might suggest the question, "Are there also hidden forms of momentum—perhaps heat momentum?" The answer is that it is very hard to hide momentum for the following reasons.

The random motions of the atoms of a body furnish a measure of heat energy, if the *squares* of the velocities are summed. This sum will be a positive result, having no directional character. The heat is there, whether or not the body moves as a whole, and conservation of energy in the form of heat is not very obvious. On the other hand, if one sums the *velocities*, which have direction, and finds a result that is not zero, that means that there is a drift of the entire body in some particular direction, and such a gross momentum is readily observed. Thus there is no random internal lost momentum, because the body has net momentum only

when it moves as a whole. Therefore momentum, as a mechanical quantity, is difficult to hide. Nevertheless, momentum *can* be hidden—in the electromagnetic field, for example. This case is another effect of relativity.

One of the propositions of Newton was that interactions at a distance are instantaneous. It turns out that such is not the case; in situations involving electrical forces, for instance, if an electrical charge at one location is suddenly moved, the effects on another charge, at another place, do not appear instantaneously—there is a little delay. In those circumstances, even if the forces are equal the momentum will not check out; there will be a short time during which there will be trouble, because for a while the first charge will feel a certain reaction force, say, and will pick up some momentum, but the second charge has felt nothing and has not yet changed its momentum. It takes time for the influence to cross the intervening distance, which it does at 186,000 miles a second. In that tiny time the momentum of the particles is not conserved. Of course after the second charge has felt the effect of the first one and all is quieted down, the momentum equation will check out all right, but during that small interval momentum is not conserved. We represent this by saying that during this interval there is another kind of momentum besides that of the particle, mv, and that is momentum in the electromagnetic field. If we add the field momentum to the momentum of the particles, then momentum is conserved at any moment all the time. The fact that the electromagnetic field can possess momentum and energy makes that field very real, and so, for better understanding, the original idea that there are just the forces between particles has to be modified to the idea that a particle makes a field, and a field acts on another particle, and the field itself has such familiar properties as energy content and momentum, just as particles can have. To take another example: an electromagnetic field has waves, which we call light; it turns out that light also carries momentum with it, so when light impinges on an object it carries in a certain amount of momentum per second; this is equivalent to a force, because if the illuminated object is picking up a certain amount of momentum per second, its momentum is changing and the situation is exactly the same as if there were a force on it. Light can exert pressure by bombarding an object; this pressure is very small, but with sufficiently delicate apparatus it is measurable.

Now in quantum mechanics it turns out that momentum is a different thing—it is no longer mv. It is hard to define exactly what is meant by the velocity of a particle, but momentum still exists. In quantum mechanics the difference is that when the particles are represented as particles, the momentum is still mv, but when the particles are represented as waves, the momentum is measured by the number of waves per centimeter: the greater this number of waves, the greater the momentum. In spite of the differences, the law of conservation of momentum holds also in quantum mechanics. Even though the law $f = ma$ is false, and all the derivations of Newton were wrong for the conservation of momentum, in quantum mechanics, nevertheless, in the end, that particular law maintains itself!

11

Vectors

11–1 Symmetry in physics

11–2 Translations

11–3 Rotations

11–4 Vectors

11–5 Vector algebra

11–6 Newton's laws in vector notation

11–7 Scalar product of vectors

11–1 Symmetry in physics

In this chapter we introduce a subject that is technically known in physics as *symmetry in physical law*. The word "symmetry" is used here with a special meaning, and therefore needs to be defined. When is a thing symmetrical—how can we define it? When we have a picture that is symmetrical, one side is somehow the same as the other side. Professor Hermann Weyl has given this definition of symmetry: a thing is symmetrical if one can subject it to a certain operation and it appears exactly the same after the operation. For instance, if we look at a vase that is left-and-right symmetrical, then turn it 180° around the vertical axis, it looks the same. We shall adopt the definition of symmetry in Weyl's more general form, and in that form we shall discuss symmetry of physical laws.

Suppose we build a complex machine in a certain place, with a lot of complicated interactions, and balls bouncing around with forces between them, and so on. Now suppose we build exactly the same kind of equipment at some other place, matching part by part, with the same dimensions and the same orientation, everything the same only displaced laterally by some distance. Then, if we start the two machines in the same initial circumstances, in exact correspondence, we ask: will one machine behave exactly the same as the other? Will it follow all the motions in exact parallelism? Of course the answer may well be *no*, because if we choose the wrong place for our machine it might be inside a wall and interferences from the wall would make the machine not work.

All of our ideas in physics require a certain amount of common sense in their application; they are not purely mathematical or abstract ideas. We have to understand what we mean when we say that the phenomena are the same when we move the apparatus to a new position. We mean that we move everything that we believe is relevant; if the phenomenon is not the same, we suggest that something relevant has not been moved, and we proceed to look for it. If we never find it, then we claim that the laws of physics do not have this symmetry. On the other hand, we may find it—we expect to find it—if the laws of physics do have this symmetry; looking around, we may discover, for instance, that the wall is pushing on the apparatus. The basic question is, if we define things well enough, if all the essential forces are included inside the apparatus, if all the relevant parts are moved from one place to another, will the laws be the same? Will the machinery work the same way?

It is clear that what we want to do is to move all the equipment and *essential* influences, but not *everything* in the world—planets, stars, and all—for if we do that, we have the same phenomenon again for the trivial reason that we are right back where we started. No, we cannot move *everything*. But it turns out in practice that with a certain amount of intelligence about what to move, the machinery will work. In other words, if we do not go inside a wall, if we know the origin of the outside forces, and arrange that those are moved too, then the machinery *will* work the same in one location as in another.

11–2 Translations

We shall limit our analysis to just mechanics, for which we now have sufficient knowledge. In previous chapters we have seen that the laws of mechanics can be summarized by a set of three equations for each particle:

$$m(d^2x/dt^2) = F_x, \quad m(d^2y/dt^2) = F_y, \quad m(d^2z/dt^2) = F_z. \quad (11.1)$$

Now this means that there exists a way to *measure* x, y, and z on three perpendicular axes, and the forces along those directions, such that these laws are true. These must be measured from some origin, but *where do we put the origin?* All that Newton would tell us at first is that there *is* some place that we can measure from, perhaps the center of the universe, such that these laws are correct. But we can show immediately that we can never find the center, because if we use some other origin it would make no difference. In other words, suppose that there are two people—Joe, who has an origin in one place, and Moe, who has a parallel system whose origin is somewhere else (Fig. 11–1). Now when Joe measures the location of the point in space, he finds it at x, y, and z (we shall usually leave z out because it is too confusing to draw in a picture). Moe, on the other hand, when measuring the same point, will obtain a different x (in order to distinguish it, we will call it x'), and in principle a different y, although in our example they are numerically equal. So we have

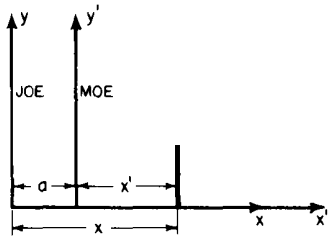


Fig. 11–1. Two parallel coordinate systems.

$$x' = x - a, \qquad y' = y, \qquad z' = z. \tag{11.2}$$

Now in order to complete our analysis we must know what Moe would obtain for the forces. The force is supposed to act along some line, and by the force in the x-direction we mean the part of the total which is in the x-direction, which is the magnitude of the force times this cosine of its angle with the x-axis. Now we see that Moe would use exactly the same projection as Joe would use, so we have a set of equations

$$F_{x'} = F_x, \qquad F_{y'} = F_y, \qquad F_{z'} = F_z. \tag{11.3}$$

These would be the relationships between quantities as seen by Joe and Moe.

The question is, if Joe knows Newton's laws, and if Moe tries to write down Newton's laws, will they also be correct for him? Does it make any difference from which origin we measure the points? In other words, assuming that equations (11.1) are true, and the Eqs. (11.2) and (11.3) give the relationship of the measurements, is it or is it not true that

$$\begin{aligned} &\text{(a)} \quad m(d^2x'/dt^2) = F_{x'}, \\ &\text{(b)} \quad m(d^2y'/dt^2) = F_{y'}, \\ &\text{(c)} \quad m(d^2z'/dt^2) = F_{z'}? \end{aligned} \tag{11.4}$$

In order to test these equations we shall differentiate the formula for x' twice. First of all

$$\frac{dx'}{dt} = \frac{d}{dt}(x - a) = \frac{dx}{dt} - \frac{da}{dt}.$$

Now we shall assume that Moe's origin is fixed (not moving) relative to Joe's; therefore a is a constant and $da/dt = 0$, so we find that

$$dx'/dt = dx/dt'$$

and therefore

$$d^2x'/dt^2 = d^2x/dt^2;$$

therefore we know that Eq. (11.4a) becomes

$$m(d^2x/dt^2) = F_{x'}.$$

(We also suppose that the masses measured by Joe and Moe are equal.) Thus the acceleration times the mass is the same as the other fellow's. We have also found the formula for $F_{x'}$, for, substituting from Eq. (11.1), we find that

$$F_{x'} = F_x.$$

Therefore the laws as seen by Moe appear the same; he can write Newton's laws too, with different coordinates, and they will still be right. That means that

there is no unique way to define the origin of the world, because the laws will appear the same, from whatever position they are observed.

This is also true: if there is a piece of equipment in one place with a certain kind of machinery in it, the same equipment in another place will behave in the same way. Why? Because one machine, when analyzed by Moe, has exactly the same equations as the other one, analyzed by Joe. Since the *equations* are the same, the *phenomena* appear the same. So the proof that an apparatus in a new position behaves the same as it did in the old position is the same as the proof that the equations when displaced in space reproduce themselves. Therefore we say that *the laws of physics are symmetrical for translational displacements*, symmetrical in the sense that the laws do not change when we make a translation of our coordinates. Of course it is quite obvious intuitively that this is true, but it is interesting and entertaining to discuss the mathematics of it.

11–3 Rotations

The above is the first of a series of ever more complicated propositions concerning the symmetry of a physical law. The next proposition is that it should make no difference in which *direction* we choose the axes. In other words, if we build a piece of equipment in some place and watch it operate, and nearby we build the same kind of apparatus but put it up on an angle, will it operate in the same way? Obviously it will not if it is a Grandfather clock, for example! If a pendulum clock stands upright, it works fine, but if it is tilted the pendulum falls against the side of the case and nothing happens. The theorem is then false in the case of the pendulum clock, unless we include the earth, which is pulling on the pendulum. Therefore we can make a prediction about pendulum clocks if we believe in the symmetry of physical law for rotation: something else is involved in the operation of a pendulum clock besides the machinery of the clock, something outside it that we should look for. We may also predict that pendulum clocks will not work the same way when located in different places relative to this mysterious source of asymmetry, perhaps the earth. Indeed, we know that a pendulum clock up in an artificial satellite, for example, would not tick either, because there is no effective force, and on Mars it would go at a different rate. Pendulum clocks *do* involve something more than just the machinery inside, they involve something on the outside. Once we recognize this factor, we see that we must turn the earth along with the apparatus. Of course we do not have to worry about that, it is easy to do; one simply waits a moment or two and the earth turns; then the pendulum clock ticks again in the new position the same as it did before. While we are rotating in space our angles are always changing, absolutely; this change does not seem to bother us very much, for in the new position we seem to be in the same condition as in the old. This has a certain tendency to confuse one, because it is true that in the new turned position the laws are the same as in the unturned position, but it is *not* true that *as we turn* a thing it follows the same laws as it does when we are not turning it. If we perform sufficiently delicate experiments, we can tell that the earth *is rotating*, but not that it *had rotated*. In other words, we cannot locate its angular position, but we can tell that it is changing.

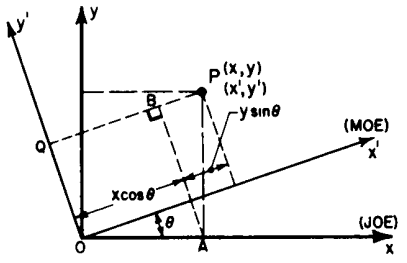


Fig. 11–2. Two coordinate systems having different angular orientations.

Now we may discuss the effects of angular orientation upon physical laws. Let us find out whether the same game with Joe and Moe works again. This time, to avoid needless complication, we shall suppose that Joe and Moe use the same origin (we have already shown that the axes can be moved by translation to another place). Assume that Moe's axes have rotated relative to Joe's by an angle θ. The two coordinate systems are shown in Fig. 11–2, which is restricted to two dimensions. Consider any point P having coordinates (x, y) in Joe's system and (x', y') in Moe's system. We shall begin, as in the previous case, by expressing the coordinates x' and y' in terms of x, y, and θ. To do so, we first drop perpendiculars from P to all four axes and draw AB perpendicular to PQ. Inspection of the figure shows that x' can be written as the sum of two lengths along the x'-axis, and y' as the difference of two lengths along AB. All these lengths are expressed

in terms of x, y, and θ in equations (11.5), to which we have added an equation for the third dimension.

$$\begin{aligned} x' &= x\cos\theta + y\sin\theta, \\ y' &= y\cos\theta - x\sin\theta, \\ z' &= z. \end{aligned} \tag{11.5}$$

The next step is to analyze the relationship of forces as seen by the two observers, following the same general method as before. Let us assume that a force F, which has already been analyzed as having components F_x and F_y (as seen by Joe), is acting on a particle of mass m, located at point P in Fig. 11–2. For simplicity, let us move both sets of axes so that the origin is at P, as shown in Fig. 11–3. Moe sees the components of F along his axes as $F_{x'}$ and $F_{y'}$. F_x has components along both the x'- and y'-axes, and F_y likewise has components along both these axes. To express $F_{x'}$ in terms of F_x and F_y, we sum these components along the x'-axis, and in a like manner we can express $F_{y'}$ in terms of F_x and F_y. The results are

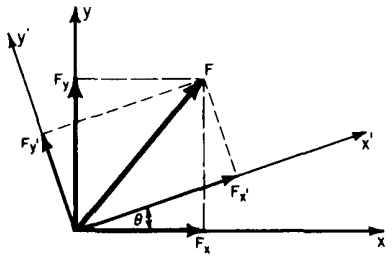


Fig. 11–3. Components of a force in the two systems.

$$\begin{aligned} F_{x'} &= F_x\cos\theta + F_y\sin\theta, \\ F_{y'} &= F_y\cos\theta - F_x\sin\theta, \\ F_{z'} &= F_z. \end{aligned} \tag{11.6}$$

It is interesting to note an accident of sorts, which is of extreme importance: the formulas (11.5) and (11.6), for coordinates of P and components of F, respectively, *are of identical form.*

As before, Newton's laws are assumed to be true in Joe's system, and are expressed by equations (11.1). The question, again, is whether Moe can apply Newton's laws—will the results be correct for his system of rotated axes? In other words, if we assume that Eqs. (11.5) and (11.6) give the relationship of the measurements, is it true or not true that

$$\begin{aligned} m(d^2x'/dt^2) &= F_{x'}, \\ m(d^2y'/dt^2) &= F_{y'}, \\ m(d^2z'/dt^2) &= F_{z'}? \end{aligned} \tag{11.7}$$

To test these equations, we calculate the left and right sides independently, and compare the results. To calculate the left sides, we multiply equations (11.5) by m, and differentiate twice with respect to time, assuming the angle θ to be constant. This gives

$$\begin{aligned} m(d^2x'/dt^2) &= m(d^2x/dt^2)\cos\theta + m(d^2y/dt^2)\sin\theta, \\ m(d^2y'/dt^2) &= m(d^2y/dt^2)\cos\theta - m(d^2x/dt^2)\sin\theta, \\ m(d^2z'/dt^2) &= m(d^2z/dt^2). \end{aligned} \tag{11.8}$$

We calculate the right sides of equations (11.7) by substituting equations (11.1) into equations (11.6). This gives

$$\begin{aligned} F_{x'} &= m(d^2x/dt^2)\cos\theta + m(d^2y/dt^2)\sin\theta, \\ F_{y'} &= m(d^2y/dt^2)\cos\theta - m(d^2x/dt^2)\sin\theta, \\ F_{z'} &= m(d^2z/dt^2). \end{aligned} \tag{11.9}$$

Behold! The right sides of Eqs. (11.8) and (11.9) are identical, so we conclude that if Newton's laws are correct on one set of axes, they are also valid on any other set of axes. This result, which has now been established for both translation and rotation of axes, has certain consequences: first, no one can claim his particular axes are unique, but of course they can be more *convenient* for certain particular problems. For example, it is handy to have gravity along one axis, but this is not physically necessary. Second, it means that any piece of equipment which is completely self-contained, with all the force-generating equipment completely inside the apparatus, would work the same when turned at an angle.

11–4 Vectors

Not only Newton's laws, but also the other laws of physics, so far as we know today, have the two properties which we call invariance (or symmetry) under translation of axes and rotation of axes. These properties are so important that a mathematical technique has been developed to take advantage of them in writing and using physical laws.

The foregoing analysis involved considerable tedious mathematical work. To reduce the details to a minimum in the analysis of such questions, a very powerful mathematical machinery has been devised. This system, called *vector analysis*, supplies the title of this chapter; strictly speaking, however, this is a chapter on the symmetry of physical laws. By the methods of the preceding analysis we were able to do everything required for obtaining the results that we sought, but in practice we should like to do things more easily and rapidly, so we employ the vector technique.

We began by noting some characteristics of two kinds of quantities that are important in physics. (Actually there are more than two, but let us start out with two.) One of them, like the number of potatoes in a sack, we call an ordinary quantity, or an undirected quantity, or a *scalar*. Temperature is an example of such a quantity. Other quantities that are important in physics do have direction, for instance velocity: we have to keep track of which way a body is going, not just its speed. Momentum and force also have direction, as does displacement: when someone steps from one place to another in space, we can keep track of how far he went, but if we wish also to know *where* he went, we have to specify a direction.

All quantities that have a direction, like a step in space, are called *vectors*.

A vector is three numbers. In order to represent a step in space, say from the origin to some particular point P whose location is (x, y, z), we really need three numbers, but we are going to invent a single mathematical symbol, $\mathbf{r}$, which is unlike any other mathematical symbols we have so far used.* It is *not* a single number, it represents *three* numbers: x, y, and z. It means three numbers, but not really only *those* three numbers, because if we were to use a different coordinate system, the three numbers would be changed to x', y', and z'. However, we want to keep our mathematics simple and so we are going to use the *same mark* to represent the three numbers (x, y, z) and the three numbers (x', y', z'). That is, we use the same mark to represent the first set of three numbers for one coordinate system, but the second set of three numbers if we are using the other coordinate system. This has the advantage that when we change the coordinate system, we do not have to change the letters of our equations. If we write an equation in terms of x, y, z, and then use another system, we have to change to x', y', z', but we shall just write $\mathbf{r}$, with the convention that it represents (x, y, z) if we use one set of axes, or (x', y', z') if we use another set of axes, and so on. The three numbers which describe the quantity in a given coordinate system are called the *components* of the vector in the direction of the coordinate axes of that system. That is, we use the same symbol for the three letters that correspond to the *same object, as seen from different axes*. The very fact that we can say "the same object" implies a physical intuition about the reality of a step in space, that is independent of the components in terms of which we measure it. So the symbol $\mathbf{r}$ will represent the same thing no matter how we turn the axes.

Now suppose there is another directed physical quantity, any other quantity, which also has three numbers associated with it, like force, and these three numbers change to three other numbers by a certain mathematical rule, if we change the axes. It must be the same rule that changes (x, y, z) into (x', y', z'). In other words, any physical quantity associated with three numbers which transform as do the components of a step in space is a vector. An equation like

$$\mathbf{F} = \mathbf{r}$$

would thus be true in *any* coordinate system if it were true in one. This equation,

---

* In type, vectors are represented by boldface; in handwritten form an arrow is used: $\vec{r}$.

of course, stands for the three equations

$$F_x = x, \quad F_y = y, \quad F_z = z,$$

or, alternatively, for

$$F_{x'} = x', \quad F_{y'} = y', \quad F_{z'} = z'.$$

The fact that a physical relationship can be expressed as a vector equation assures us the relationship is unchanged by a mere rotation of the coordinate system. That is the reason why vectors are so useful in physics.

Now let us examine some of the properties of vectors. As examples of vectors we may mention velocity, momentum, force, and acceleration. For many purposes it is convenient to represent a vector quantity by an arrow that indicates the direction in which it is acting. Why can we represent force, say, by an arrow? Because it has the same mathematical transformation properties as a "step in space." We thus represent it in a diagram as if it were a step, using a scale such that one unit of force, or one newton, corresponds to a certain convenient length. Once we have done this, all forces can be represented as lengths, because an equation like

$$\mathbf{F} = k\mathbf{r},$$

where k is some constant, is a perfectly legitimate equation. Thus we can always represent forces by lines, which is very convenient, because once we have drawn the line we no longer need the axes. Of course, we can quickly calculate the three components as they change upon turning the axes, because that is just a geometric problem.

11–5 Vector algebra

Now we must describe the laws, or rules, for combining vectors in various ways. The first such combination is the *addition* of two vectors: suppose that **a** is a vector which in some particular coordinate system has the three components (a_x, a_y, a_z), and that **b** is another vector which has the three components (b_x, b_y, b_z). Now let us invent three new numbers $(a_x + b_x, a_y + b_y, a_z + b_z)$. Do these form a vector? "Well," we might say, "they are three numbers, and every three numbers form a vector." No, *not* every three numbers form a vector! In order for it to be a vector, not only must there be three numbers, but these must be associated with a coordinate system in such a way that if we turn the coordinate system, the three numbers "revolve" on each other, get "mixed up" in each other, by the precise laws we have already described. So the question is, if we now rotate the coordinate system so that (a_x, a_y, a_z) become $(a_{x'}, a_{y'}, a_{z'})$ and (b_x, b_y, b_z) become $(b_{x'}, b_{y'}, b_{z'})$, what do $(a_x + b_x, a_y + b_y, a_z + b_z)$ become? Do they become $(a_{x'} + b_{x'}, a_{y'} + b_{y'}, a_{z'} + b_{z'})$ or not? The answer is, of course, yes, because the prototype transformations of Eq. (11.5) constitute what we call a *linear* transformation. If we apply those transformations to a_x and b_x to get $a_{x'} + b_{x'}$, we find that the transformed $a_x + b_x$ is indeed the same as $a_{x'} + b_{x'}$. When **a** and **b** are "added together" in this sense, they will form a vector which we may call **c**. We would write this as

$$\mathbf{c} = \mathbf{a} + \mathbf{b}.$$

Now **c** has the interesting property

$$\mathbf{c} = \mathbf{b} + \mathbf{a},$$

as we can immediately see from its components. Thus also,

$$\mathbf{a} + (\mathbf{b} + \mathbf{c}) = (\mathbf{a} + \mathbf{b}) + \mathbf{c}.$$

We can add vectors in any order.

What is the geometric significance of $\mathbf{a} + \mathbf{b}$? Suppose that **a** and **b** were represented by lines on a piece of paper, what would **c** look like? This is shown in

Fig. 11–4. We see that we can add the components of **b** to those of **a** most conveniently if we place the rectangle representing the components of **b** next to that representing the components of **a** in the manner indicated. Since **b** just "fits" into its rectangle, as does **a** into its rectangle, this is the same as putting the "tail" of **b** on the "head" of **a**, the arrow from the "tail" of **a** to the "head" of **b** being the vector **c**. Of course, if we added **a** to **b** the other way around, we would put the "tail" of **a** on the "head" of **b**, and by the geometrical properties of parallelograms we would get the same result for **c**. Note that vectors can be added in this way without reference to any coordinate axes.

Suppose we multiply a vector by a number α, what does this mean? We *define* it to mean a new vector whose components are αa_x, αa_y, and αa_z. We leave it as a problem for the student to prove that it *is* a vector.

Now let us consider vector subtraction. We may define subtraction in the same way as addition, but instead of adding, we subtract the components. Or we might define subtraction by defining a negative vector, $-\mathbf{b} = -1\mathbf{b}$, and then we would add the components. It comes to the same thing. The result is shown in Fig. 11–5. This figure shows $\mathbf{d} = \mathbf{a} - \mathbf{b} = \mathbf{a} + (-\mathbf{b})$; we also note that the difference $\mathbf{a} - \mathbf{b}$ can be found very easily from **a** and **b** by using the equivalent relation $\mathbf{a} = \mathbf{b} + \mathbf{d}$. Thus the difference is even easier to find than the sum: we just draw the vector from **b** to **a**, to get $\mathbf{a} - \mathbf{b}$!

Next we discuss velocity. Why is velocity a vector? If position is given by the three coordinates (x, y, z), what is the velocity? The velocity is given by dx/dt, dy/dt, and dz/dt Is that a vector, or not? We can find out by differentiating the expressions in Eq. (11.5) to find out whether dx'/dt *transforms* in the right way. We see that the components dx/dt and dy/dt *do* transform according to the same law as x and y, and therefore the time derivative is a vector. So the velocity *is* a vector. We can write the velocity in an interesting way as

$$\mathbf{v} = d\mathbf{r}/dt.$$

What the velocity is, and why it is a vector, can also be understood more pictorially: How far does a particle move in a short time Δt? Answer: $\Delta\mathbf{r}$, so if a particle is "here" at one instant and "there" at another instant, then the vector difference of the positions $\Delta\mathbf{r} = \mathbf{r}_2 - \mathbf{r}_1$, which is in the direction of motion shown in Fig. 11–6, divided by the time interval $\Delta t = t_2 - t_1$, is the "average velocity" vector.

In other words, by vector velocity we mean the limit, as Δt goes to 0, of the difference between the radius vectors at the time $t + \Delta t$ and the time t, divided by Δt:

$$\mathbf{v} = \lim_{\Delta t \to 0} (\Delta\mathbf{r}/\Delta t) = d\mathbf{r}/dt. \tag{11.10}$$

Thus velocity is a vector because it is the difference of two vectors. It is also the right definition of velocity because its components are dx/dt, dy/dt, and dz/dt. In fact, we see from this argument that if we differentiate *any* vector with respect to time we produce a new vector. So we have several ways of producing new vectors: (1) multiply by a constant, (2) differentiate with respect to time, (3) add or subtract two vectors.

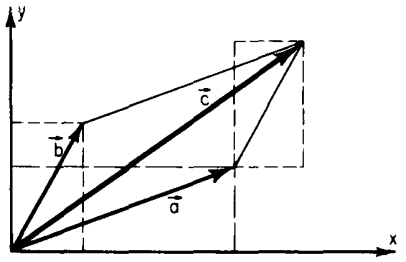


Fig. 11–4. The addition of vectors.

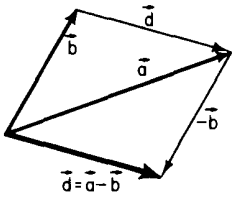


Fig. 11–5. The subtraction of vectors.

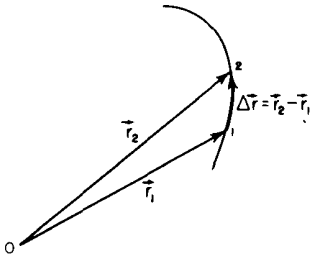


Fig. 11–6. The displacement of a particle in a short time interval $\Delta t = t_2 - t_1$.

11–6 Newton's laws in vector notation

In order to write Newton's laws in vector form, we have to go just one step further, and define the acceleration vector. This is the time derivative of the velocity vector, and it is easy to demonstrate that its components are the second derivatives of x, y, and z with respect to t:

$$\mathbf{a} = \frac{d\mathbf{v}}{dt} = \left(\frac{d}{dt}\right)\left(\frac{d\mathbf{r}}{dt}\right) = \frac{d^2\mathbf{r}}{dt^2}, \tag{11.11}$$

$$a_x = \frac{dv_x}{dt} = \frac{d^2x}{dt^2}, \quad a_y = \frac{dv_y}{dt} = \frac{d^2y}{dt^2}, \quad a_z = \frac{dv_z}{dt} = \frac{d^2z}{dt^2}. \tag{11.12}$$

With this definition, then, Newton's laws can be written in this way:

$$m\mathbf{a} = \mathbf{F} \tag{11.13}$$

or

$$m(d^2\mathbf{r}/dt^2) = \mathbf{F}. \tag{11.14}$$

Now the problem of proving the invariance of Newton's laws under rotation of coordinates is this: prove that **a** is a vector; this we have just done. Prove that **F** is a vector; we *suppose* it is. So if force is a vector, then, since we know acceleration is a vector, Eq. (11.13) will look the same in any coordinate system. Writing it in a form which does not explicitly contain x's, y's, and z's has the advantage that from now on we need not write *three* laws every time we write Newton's equations or other laws of physics. We write what looks like *one* law, but really, of course, it is the three laws for any particular set of axes, because any vector equation involves the statement that *each of the components is equal.*

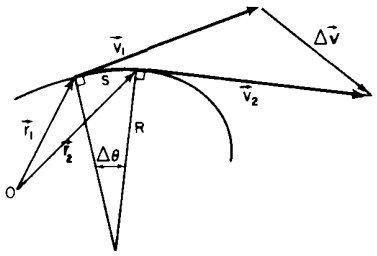


Fig. 11–7. A curved trajectory.

The fact that the acceleration is the rate of change of the vector velocity helps us to calculate the acceleration in some rather complicated circumstances. Suppose, for instance, that a particle is moving on some complicated curve (Fig. 11–7) and that, at a given instant t, it had a certain velocity $\mathbf{v}_1$, but that when we go to another instant t_2 a little later, it has a different velocity $\mathbf{v}_2$. What is the acceleration? Answer: Acceleration is the difference in the velocity divided by the small time interval, so we need the difference of the two velocities. How do we get the difference of the velocities? To subtract two vectors, we put the vector across the ends of $\mathbf{v}_2$ and $\mathbf{v}_1$; that is, we draw $\boldsymbol{\Delta}$ as the difference of the two vectors, right? *No!* That only works when the *tails* of the vectors are in the same place! It has no meaning if we move the vector somewhere else and then draw a line across, so watch out! We have to draw a new diagram to subtract the vectors. In Fig. 11–8, $\mathbf{v}_1$ and $\mathbf{v}_2$ are both drawn parallel and equal to their counterparts in Fig. 11–7, and now we can discuss the acceleration. Of course the acceleration is simply $\Delta\mathbf{v}/\Delta t$. It is interesting to note that we can compose the velocity difference out of two parts; we can think of acceleration as having *two components*, $\Delta\mathbf{v}_{||}$ in the direction tangent to the path and $\Delta\mathbf{v}_{\perp}$ at right angles to the path, as indicated in Fig. 11–8. The acceleration tangent to the path is, of course, just the change in the *length* of the vector, i.e., the change in the *speed* v:

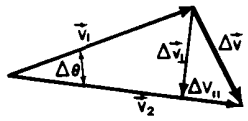


Fig. 11–8. Diagram for calculating the acceleration.

$$a_{||} = dv/dt. \tag{11.15}$$

The other component of acceleration, at right angles to the curve, is easy to calculate, using Figs. 11–7 and 11–8. In the short time Δt let the change in angle between $\mathbf{v}_1$ and $\mathbf{v}_2$ be the small angle $\Delta\theta$. If the magnitude of the velocity is called v, then of course

$$\Delta v_{\perp} = v\,\Delta\theta$$

and the acceleration a will be

$$a_{\perp} = v\,(\Delta\theta/\Delta t).$$

Now we need to know $\Delta\theta/\Delta t$, which can be found this way: If, at the given moment, the curve is approximated as a circle of a certain radius R, then in a time Δt the distance s is, of course, $v\,\Delta t$, where v is the speed.

$$\Delta\theta = v\,(\Delta t/R), \qquad \text{or} \qquad \Delta\theta/\Delta t = v/R.$$

Therefore, we find

$$a = v^2/R, \tag{11.16}$$

as we have seen before.

11–7 Scalar product of vectors

Now let us examine a little further the properties of vectors. It is easy to see that the *length* of a step in space would be the same in any coordinate system. That is, if a particular step **r** is represented by x, y, z, in one coordinate system,

and by x', y', z' in another coordinate system, surely the distance $r = |\mathbf{r}|$ would be the same in both. Now

$$r = \sqrt{x^2 + y^2 + z^2}$$

and also

$$r' = \sqrt{x'^2 + y'^2 + z'^2}.$$

So what we wish to verify is that these two quantities are equal. It is much more convenient not to bother to take the square root, so let us talk about the square of the distance; that is, let us find out whether

$$x^2 + y^2 + z^2 = x'^2 + y'^2 + z'^2. \tag{11.17}$$

It had better be—and if we substitute Eq. (11.5) we do indeed find that it is. So we see that there are other kinds of equations which are true for any two coordinate systems.

Something new is involved. We can produce a new quantity, a function of x, y, and z, called a *scalar function*, a quantity which has no direction but which is the same in both systems. Out of a vector we can make a scalar. We have to find a general rule for that. It is clear what the rule is for the case just considered: add the squares of the components. Let us now define a new thing, which we call $\mathbf{a} \cdot \mathbf{a}$. This is not a vector, but a scalar; it is a number that is the same in all coordinate systems, and it is defined to be the sum of the squares of the three components of the vector:

$$\mathbf{a} \cdot \mathbf{a} = a_x^2 + a_y^2 + a_z^2. \tag{11.18}$$

Now you say, "But with what axes?" It does not depend on the axes, the answer is the same in *every* set of axes. So we have a new *kind* of quantity, a new *invariant* or *scalar* produced by one vector "squared." If we now define the following quantity for any two vectors $\mathbf{a}$ and $\mathbf{b}$:

$$\mathbf{a} \cdot \mathbf{b} = a_x b_x + a_y b_y + a_z b_z, \tag{11.19}$$

we find that this quantity, calculated in the primed and unprimed systems, also stays the same. To prove it we note that it is true of $\mathbf{a} \cdot \mathbf{a}$, $\mathbf{b} \cdot \mathbf{b}$, and $\mathbf{c} \cdot \mathbf{c}$, where $\mathbf{c} = \mathbf{a} + \mathbf{b}$. Therefore the sum of the squares $(a_x + b_x)^2 + (a_y + b_y)^2 + (a_z + b_z)^2$ will be invariant:

$$(a_x + b_x)^2 + (a_y + b_y)^2 + (a_z + b_z)^2 = (a_{x'} + b_{x'})^2 + (a_{y'} + b_{y'})^2 + (a_{z'} + b_{z'})^2. \tag{11.20}$$

If both sides of this equation are expanded, there will be cross products of just the type appearing in Eq. (11.19), as well as the sums of squares of the components of $\mathbf{a}$ and $\mathbf{b}$. The invariance of terms of the form of Eq. (11.18) then leaves the cross product terms (11.19) invariant also.

The quantity $\mathbf{a} \cdot \mathbf{b}$ is called the *scalar product* of two vectors, $\mathbf{a}$ and $\mathbf{b}$, and it has many interesting and useful properties. For instance, it is easily proved that

$$\mathbf{a} \cdot (\mathbf{b} + \mathbf{c}) = \mathbf{a} \cdot \mathbf{b} + \mathbf{a} \cdot \mathbf{c}. \tag{11.21}$$

Also, there is a simple geometrical way to calculate $\mathbf{a} \cdot \mathbf{b}$, without having to calculate the components of $\mathbf{a}$ and $\mathbf{b}$: $\mathbf{a} \cdot \mathbf{b}$ is the product of the length of $\mathbf{a}$ and the length of $\mathbf{b}$ times the cosine of the angle between them. Why? Suppose that we choose a special coordinate system in which the x-axis lies along $\mathbf{a}$; in those circumstances, the only component of $\mathbf{a}$ that will be there is a_x, which is of course the whole length of $\mathbf{a}$. Thus Eq. (11.19) reduces to $\mathbf{a} \cdot \mathbf{b} = a_x b_x$ for this case, and this is the length of $\mathbf{a}$ times the component of $\mathbf{b}$ in the direction of $\mathbf{a}$, that is, $b \cos \theta$:

$$\mathbf{a} \cdot \mathbf{b} = ab \cos \theta.$$

Therefore, in that special coordinate system, we have proved that $\mathbf{a} \cdot \mathbf{b}$ is the

length of **a** times the length of **b** times cos θ. But *if it is true in one coordinate system, it is true in all*, because $\mathbf{a} \cdot \mathbf{b}$ is independent of the coordinate system; that is our argument.

What good is the dot product? Are there any cases in physics where we need it? Yes, we need it all the time. For instance, in Chapter 4 the kinetic energy was called $\frac{1}{2}mv^2$, but if the object is moving in space it should be the velocity squared in the x-direction, the y-direction, and the z-direction, and so the formula for kinetic energy according to vector analysis is

$$\text{K.E.} = \tfrac{1}{2}m(\mathbf{v} \cdot \mathbf{v}) = \tfrac{1}{2}m(v_x^2 + v_y^2 + v_z^2). \tag{11.22}$$

Energy does not have direction. Momentum has direction; it is a vector, and it is the mass times the velocity vector.

Another example of a dot product is the work done by a force when something is pushed from one place to the other. We have not yet defined work, but it is equivalent to the energy change, the weights lifted, when a force **F** acts through a distance **s**:

$$\text{Work} = \mathbf{F} \cdot \mathbf{s} \tag{11.23}$$

It is sometimes very convenient to talk about the component of a vector in a certain direction (say the vertical direction because that is the direction of gravity). For such purposes, it is useful to invent what we call a *unit vector* in the direction that we want to study. By a unit vector we mean one whose dot product with itself is equal to unity. Let us call this unit vector **i**; then $\mathbf{i} \cdot \mathbf{i} = 1$. Then, if we want the component of some vector in the direction of **i**, we see that the dot product $\mathbf{a} \cdot \mathbf{i}$ will be $a \cos \theta$, i.e., the component of **a** in the direction of **i**. This is a nice way to get the component; in fact, it permits us to get *all* the components and to write a rather amusing formula. Suppose that in a given system of coordinates, x, y, and z, we invent three vectors: **i**, a unit vector in the direction x; **j**, a unit vector in the direction y; and **k**, a unit vector in the direction z. Note first that $\mathbf{i} \cdot \mathbf{i} = 1$. What is $\mathbf{i} \cdot \mathbf{j}$? When two vectors are at right angles, their dot product is zero. Thus

$$\begin{aligned} &\mathbf{i} \cdot \mathbf{i} = 1 \\ &\mathbf{i} \cdot \mathbf{j} = 0 \qquad \mathbf{j} \cdot \mathbf{j} = 1 \\ &\mathbf{i} \cdot \mathbf{k} = 0 \qquad \mathbf{j} \cdot \mathbf{k} = 0 \qquad \mathbf{k} \cdot \mathbf{k} = 1 \end{aligned} \tag{11.24}$$

Now with these definitions, any vector whatsoever can be written this way:

$$\mathbf{a} = a_x\mathbf{i} + a_y\mathbf{j} + a_z\mathbf{k}. \tag{11.25}$$

By this means we can go from the components of a vector to the vector itself.

This discussion of vectors is by no means complete. However, rather than try to go more deeply into the subject now, we shall first learn to use in physical situations some of the ideas so far discussed. Then, when we have properly mastered this basic material, we shall find it easier to penetrate more deeply into the subject without getting too confused. We shall later find that it is useful to define another kind of product of two vectors, called the vector product, and written as $\mathbf{a} \times \mathbf{b}$. However, we shall undertake a discussion of such matters in a later chapter.

12

Characteristics of Force

12–1 What is a force?

12–2 Friction

12–3 Molecular forces

12–4 Fundamental forces. Fields

12–5 Pseudo forces

12–6 Nuclear forces

12–1 What is a force?

Although it is interesting and worth while to study the physical laws simply because they help us to understand and to use nature, one ought to stop every once in a while and think, "What do they really mean?" The meaning of any statement is a subject that has interested and troubled philosophers from time immemorial, and the meaning of physical laws is even more interesting, because it is generally believed that these laws represent some kind of real knowledge. The meaning of knowledge is a deep problem in philosophy, and it is always important to ask, "What does it mean?"

Let us ask, "What is the meaning of the physical laws of Newton, which we write as $F = ma$? What is the meaning of force, mass, and acceleration?" Well, we can intuitively sense the meaning of mass, and we can *define* acceleration if we know the meaning of position and time. We shall not discuss those meanings, but shall concentrate on the new concept of *force*. The answer is equally simple: "If a body is accelerating, then there is a force on it." That is what Newton's laws say, so the most precise and beautiful definition of force imaginable might simply be to say that force is the mass of an object times the acceleration. Suppose we have a law which says that the conservation of momentum is valid if the sum of all the external forces is zero; then the question arises, "What does it *mean*, that the sum of all the external forces is zero?" A pleasant way to define that statement would be: "When the total momentum is a constant, then the sum of the external forces is zero." There must be something wrong with that, because it is just not saying anything new. If we have discovered a fundamental law, which asserts that the force is equal to the mass times the acceleration, and then *define* the force to be the mass times the acceleration, we have found out nothing. We could also define force to mean that a moving object with no force acting on it continues to move with constant velocity in a straight line. If we then observe an object *not* moving in a straight line with a constant velocity, we might say that there is a force on it. Now such things certainly cannot be the content of physics, because they are definitions going in a circle. The Newtonian statement above, however, seems to be a most precise definition of force, and one that appeals to the mathematician; nevertheless, it is completely useless, because no prediction whatsoever can be made from a definition. One might sit in an armchair all day long and define words at will, but to find out what happens when two balls push against each other, or when a weight is hung on a spring, is another matter altogether, because the way the bodies *behave* is something completely outside any choice of definitions.

For example, if we were to choose to say that an object left to itself keeps its position and does not move, then when we see something drifting, we could say that must be due to a "gorce"—a gorce is the rate of change of position. Now we have a wonderful new law, everything stands still except when a gorce is acting. You see, that would be analogous to the above definition of force, and it would contain no information. The real content of Newton's laws is this: that the force is supposed to have some *independent properties*, in addition to the law $F = ma$; but the *specific* independent properties that the force has were not completely described by Newton or by anybody else, and therefore the physical law $F = ma$ is an incomplete law. It implies that if we study the mass times the acceleration and call the product the force, i.e., if we study the characteristics of force as a program

of interest, then we shall find that forces have some simplicity; the law is a good program for analyzing nature, it is a suggestion that the forces will be simple.

Now the first example of such forces was the complete law of gravitation, which was given by Newton, and in stating the law he answered the question, "What is the force?" If there were nothing but gravitation, then the combination of this law and the force law (second law of motion) would be a complete theory, but there is much more than gravitation, and we want to use Newton's laws in many different situations. Therefore in order to proceed we have to tell something about the properties of force.

For example, in dealing with force the tacit assumption is always made that the force is equal to zero unless some physical body is present, that if we find a force that is not equal to zero we also find something in the neighborhood that is a source of the force. This assumption is entirely different from the case of the "gorce" that we introduced above. One of the most important characteristics of force is that it has a material origin, and this is *not* just a definition.

Newton also gave one rule about the force: that the forces between interacting bodies are equal and opposite—action equals reaction; that rule, it turns out, is not exactly true. In fact, the law $F = ma$ is not exactly true; if it were a definition we should have to say that it is *always* exactly true; but it is not.

The student may object, "I do not like this imprecision, I should like to have everything defined exactly; in fact, it says in some books that any science is an exact subject, in which *everything* is defined." If you insist upon a precise definition of force, you will never get it! First, because Newton's Second Law is not exact, and second, because in order to understand physical laws you must understand that they are all some kind of approximation.

Any simple idea is approximate; as an illustration, consider an object, . . . what *is* an object? Philosophers are always saying, "Well, just take a chair for example." The moment they say that, you know that they do not know what they are talking about any more. What *is* a chair? Well, a chair is a certain thing over there . . . certain?, how certain? The atoms are evaporating from it from time to time—not many atoms, but a few—dirt falls on it and gets dissolved in the paint; so to define a chair precisely, to say exactly which atoms are chair, and which atoms are air, or which atoms are dirt, or which atoms are paint that belongs to the chair is impossible. So the mass of a chair can be defined only approximately. In the same way, to define the mass of a single object is impossible, because there are not any single, left-alone objects in the world—every object is a mixture of a lot of things, so we can deal with it only as a series of approximations and idealizations.

The trick is the idealizations. To an excellent approximation of perhaps one part in 10^{10}, the number of atoms in the chair does not change in a minute, and if we are not too precise we may idealize the chair as a definite thing; in the same way we shall learn about the characteristics of force, in an ideal fashion, if we are not too precise. One may be dissatisfied with the approximate view of nature that physics tries to obtain (the attempt is always to increase the accuracy of the approximation), and may prefer a mathematical definition; but mathematical definitions can never work in the real world. A mathematical definition will be good for mathematics, in which all the logic can be followed out completely, but the physical world is complex, as we have indicated in a number of examples, such as those of the ocean waves and a glass of wine. When we try to isolate pieces of it, to talk about one mass, the wine and the glass, how can we know which is which, when one dissolves in the other? The forces on a single thing already involve approximation, and if we have a system of discourse about the real world, then that system, at least for the present day, must involve approximations of some kind.

This system is quite unlike the case of mathematics, in which everything can be defined, and then we do not *know* what we are talking about. In fact, the glory of mathematics is that *we do not have to say what we are talking about.* The glory is that the laws, the arguments, and the logic are independent of what "it" is. If we have any other set of objects that obey the same system of axioms as Euclid's

geometry, then if we make new definitions and follow them out with correct logic, all the consequences will be correct, and it makes no difference what the subject was. In nature, however, when we draw a line or establish a line by using a light beam and a theodolite, as we do in surveying, are we measuring a line in the sense of Euclid? No, we are making an approximation; the cross hair has some width, but a geometrical line has no width, and so, whether Euclidean geometry can be used for surveying or not is a physical question, not a mathematical question. However, from an experimental standpoint, not a mathematical standpoint, we need to know whether the laws of Euclid apply to the kind of geometry that we use in measuring land; so we make a hypothesis that it does, and it works pretty well; but it is not precise, because our surveying lines are not really geometrical lines. Whether or not those lines of Euclid, which are really abstract, apply to the lines of experience is a question for experience; it is not a question that can be answered by sheer reason.

In the same way, we cannot just call $F = ma$ a definition, deduce everything purely mathematically, and make mechanics a mathematical theory, when mechanics is a description of nature. By establishing suitable postulates it is always possible to make a system of mathematics, just as Euclid did, but we cannot make a mathematics of the world, because sooner or later we have to find out whether the axioms are valid for the objects of nature. Thus we immediately get involved with these complicated and "dirty" objects of nature, but with approximations ever increasing in accuracy.

12–2 Friction

The foregoing considerations show that a true understanding of Newton's laws requires a discussion of forces, and it is the purpose of this chapter to introduce such a discussion, as a kind of completion of Newton's laws. We have already studied the definitions of acceleration and related ideas, but now we have to study the properties of force, and this chapter, unlike the previous chapters, will not be very precise, because forces are quite complicated.

To begin with a particular force, let us consider the drag on an airplane flying through the air. What is the law for that force? (Surely there is a law for every force, we *must* have a law!) One can hardly think that the law for that force will be simple. Try to imagine what makes a drag on an airplane flying through the air—the air rushing over the wings, the swirling in the back, the changes going on around the fuselage, and many other complications, and you see that there is not going to be a simple law. On the other hand, it is a remarkable fact that the drag force on an airplane is approximately a constant times the square of the velocity, or $F \sim cv^2$.

Now what is the status of such a law, is it analogous to $F = ma$? Not at all, because in the first place this law is an empirical thing that is obtained roughly by tests in a wind tunnel. You say, "Well $F = ma$ might be empirical too." That is not the reason that there is a difference. The difference is not that it is empirical, but that, as we understand nature, this law is the result of an enormous complexity of events and is not, fundamentally, a simple thing. If we continue to study it more and more, measuring more and more accurately, the law will continue to become *more* complicated, not *less*. In other words, as we study this law of the drag on an airplane more and more closely, we find out that it is "falser" and "falser," and the more deeply we study it, and the more accurately we measure, the more complicated the truth becomes; so in that sense we consider it not to result from a simple, fundamental process, which agrees with our original surmise. For example, if the velocity is extremely low, so low that an ordinary airplane is not flying, as when the airplane is dragged slowly through the air, then the law changes, and the drag friction depends more nearly linearly on the velocity. To take another example, the frictional drag on a ball or a bubble or anything that is moving slowly through a viscous liquid like honey, is proportional to the velocity, but for motion so fast that the fluid swirls around (honey does not but water and air do) then the drag becomes more nearly proportional to the square of the velocity ($F = cv^2$), and

if the velocity continues to increase, then even this law begins to fail. People who say, "Well the coefficient changes slightly," are dodging the issue. Second, there are other great complications: can this force on the airplane be divided or analyzed as a force on the wings, a force on the front, and so on? Indeed, this can be done, if we are concerned about the torques here and there, but then we have to get special laws for the force on the wings, and so on. It is an amazing fact that the force on a wing depends upon the other wing: in other words, if we take the airplane apart and put just one wing in the air, then the force is not the same as if the rest of the plane were there. The reason, of course, is that some of the wind that hits the front goes around to the wings and changes the force on the wings. It seems a miracle that there is such a simple, rough, empirical law that can be used in the design of airplanes, but this law is not in the same class as the *basic* laws of physics, and further study of it will only make it more and more complicated. A study of how the coefficient c depends on the shape of the front of the airplane is, to put it mildly, frustrating. There just is no simple law for determining the coefficient in terms of the shape of the airplane. In contrast, the law of gravitation is simple, and further study only indicates its greater simplicity.

We have just discussed two cases of friction, resulting from fast movement in air and slow movement in honey. There is another kind of friction, called dry friction or sliding friction, which occurs when one solid body slides on another. In this case a force is needed to maintain motion. This is called a frictional force, and its origin, also, is a very complicated matter. Both surfaces of contact are irregular, on an atomic level. There are many points of contact where the atoms seem to cling together, and then, as the sliding body is pulled along, the atoms snap apart and vibration ensues; something like that has to happen. Formerly the mechanism of this friction was thought to be very simple, that the surfaces were merely full of irregularities and the friction originated in lifting the slider over the bumps; but this cannot be, for there is no loss of energy in that process, whereas power is in fact consumed. The mechanism of power loss is that as the slider snaps over the bumps, the bumps deform and then generate waves and atomic motions and, after a while, heat, in the two bodies. Now it is very remarkable that again, empirically, this friction can be described approximately by a simple law. This law is that the force needed to overcome friction and to drag one object over another depends upon the normal force (i.e., perpendicular to the surface) between the two surfaces that are in contact. Actually, to a fairly good approximation, the frictional force is proportional to this normal force, and has a more or less constant coefficient; that is,

$$F = \mu N, \tag{12.1}$$

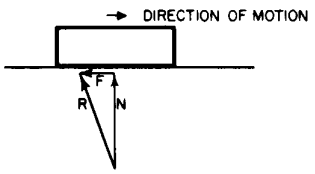


Fig. 12–1. The relation between frictional force and the normal force for sliding contact.

where μ is called the *coefficient of friction* (Fig. 12–1). Although this coefficient is not exactly constant, the formula is a good empirical rule for judging approximately the amount of force that will be needed in certain practical or engineering circumstances. If the normal force or the speed of motion gets too big, the law fails because of the excessive heat generated. It is important to realize that each of these empirical laws has its limitations, beyond which it does not really work.

That the formula $F = \mu N$ is approximately correct can be demonstrated by a simple experiment. We set up a plane, inclined at a small angle θ, and place a block of weight W on the plane. We then tilt the plane at a steeper angle, until the block just begins to slide from its own weight. The component of the weight downward along the plane is $W \sin \theta$, and this must equal the frictional force F when the block is sliding uniformly. The component of the weight normal to the plane is $W \cos \theta$, and this is the normal force N. With these values, the formula becomes $W \sin \theta = \mu W \cos \theta$, from which we get $\mu = \sin \theta / \cos \theta = \tan \theta$. If this law were exactly true, an object would start to slide at some definite inclination. If the same block is loaded by putting extra weight on it, then, although W is increased, all the forces in the formula are increased in the same proportion, and W cancels out. If μ stays constant, the loaded block will slide again at the same slope. When the angle θ is determined by trial with the original weight, it is found

that with the greater weight the block will slide at about the same angle. This will be true even when one weight is many times as great as the other, and so we conclude that the coefficient of friction is independent of the weight.

In performing this experiment it is noticeable that when the plane is tilted at about the correct angle θ, the block does not slide steadily but in a halting fashion. At one place it may stop, at another it may move with acceleration. This behavior indicates that the coefficient of friction is only roughly a constant, and varies from place to place along the plane. The same erratic behavior is observed whether the block is loaded or not. Such variations are caused by different degrees of smoothness or hardness of the plane, and perhaps dirt, oxides, or other foreign matter. The tables that list purported values of μ for "steel on steel," "copper on copper," and the like, are all false, because they ignore the factors mentioned above, which really determine μ. The friction is never due to "copper on copper," etc., but to the impurities clinging to the copper.

In experiments of the type described above, the friction is nearly independent of the velocity. Many people believe that the friction to be overcome to get something started (static friction) exceeds the force required to keep it sliding (sliding friction), but with dry metals it is very hard to show any difference. The opinion probably arises from experiences where small bits of oil or lubricant are present, or where blocks, for example, are supported by springs or other flexible supports so that they appear to bind.

It is quite difficult to do accurate quantitative experiments in friction, and the laws of friction are still not analyzed very well, in spite of the enormous engineering value of an accurate analysis. Although the law $F = \mu N$ is fairly accurate once the surfaces are standardized, the reason for this form of the law is not really understood. To show that the coefficient μ is nearly independent of velocity requires some delicate experimentation, because the apparent friction is much reduced if the lower surface vibrates very fast. When the experiment is done at very high speed, care must be taken that the objects do not vibrate relative to one another, since apparent decreases of the friction at high speed are often due to vibrations. At any rate, this friction law is another of those semiempirical laws that are not thoroughly understood, and in view of all the work that has been done it is surprising that more understanding of this phenomenon has not come about. At the present time, in fact, it is impossible even to estimate the coefficient of friction between two substances.

It was pointed out above that attempts to measure μ by sliding pure substances such as copper on copper will lead to spurious results, because the surfaces in contact are not pure copper, but are mixtures of oxides and other impurities. If we try to get absolutely pure copper, if we clean and polish the surfaces, outgas the materials in a vacuum, and take every conceivable precaution, we still do not get μ. For if we tilt the apparatus even to a vertical position, the slider will not fall off—the two pieces of copper stick together! The coefficient μ, which is ordinarily less than unity for reasonably hard surfaces, becomes several times unity! The reason for this unexpected behavior is that when the atoms in contact are all of the same kind, there is no way for the atoms to "know" that they are in different pieces of copper. When there are other atoms, in the oxides and greases and more complicated thin surface layers of contaminants in between, the atoms "know" when they are not on the same part. When we consider that it is forces between atoms that hold the copper together as a solid, it should become clear that it is impossible to get the right coefficient of friction for pure metals.

The same phenomenon can be observed in a simple home-made experiment with a flat glass plate and a glass tumbler. If the tumbler is placed on the plate and pulled along with a loop of string, it slides fairly well and one can feel the coefficient of friction; it is a little irregular, but it is a coefficient. If we now wet the glass plate and the bottom of the tumbler and pull again, we find that it binds, and if we look closely we shall find scratches, because the water is able to lift the grease and the other contaminants off the surface, and then we really have a glass-to-glass contact; this contact is so good that it holds tight and resists separation so much that the glass is torn apart; that is, it makes scratches.

12–3 Molecular forces

We shall next discuss the characteristics of molecular forces. These are forces between the atoms, and are the ultimate origin of friction. Molecular forces have never been satisfactorily explained on a basis of classical physics; it takes quantum mechanics to understand them fully. Empirically, however, the force between atoms is illustrated schematically in Fig. 12–2, where the force F between two atoms is plotted as a function of the distance r between them. There are different cases: in the water molecule, for example, the negative charges sit more on the oxygen, and the mean positions of the negative charges and of the positive charges are not at the same point; consequently, another molecule nearby feels a relatively large force, which is called a dipole-dipole force. However, for many systems the charges are very much better balanced, in particular for oxygen gas, which is perfectly symmetrical. In this case, although the minus charges and the plus charges are dispersed over the molecule, the distribution is such that the center of the minus charges and the center of the plus charges coincide. A molecule where the centers do not coincide is called a polar molecule, and charge times the separation between centers is called the dipole moment. A nonpolar molecule is one in which the centers of the charges coincide. For all nonpolar molecules, in which all the electrical forces are neutralized, it nevertheless turns out that the force at very large distances is an attraction and varies inversely as the seventh power of the distance, or $F = k/r^7$, where k is a constant that depends on the molecules. Why this is we shall learn only when we learn quantum mechanics. When there are dipoles the forces are greater. When atoms or molecules get too close they repel with a very large repulsion; that is what keeps us from falling through the floor!

These molecular forces can be demonstrated in a fairly direct way: one of these is the friction experiment with a sliding glass tumbler; another is to take two very carefully ground and lapped surfaces which are very accurately flat, so that the surfaces can be brought very close together. An example of such surfaces is the Johansson blocks that are used in machine shops as standards for making accurate length measurements. If one such block is slid over another very carefully and the upper one is lifted, the other one will adhere and also be lifted by the molecular forces, exemplifying the direct attraction between the atoms on one block for the atoms on the other block.

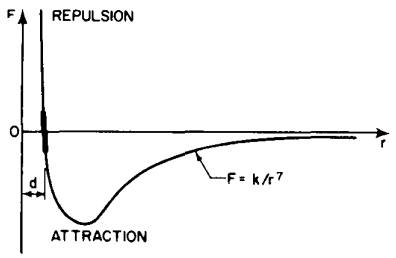


Fig. 12–2. The force between two atoms as a function of their distance of separation.

Nevertheless these molecular forces of attraction are still not fundamental in the sense that gravitation is fundamental; they are due to the vastly complex interactions of all the electrons and nuclei in one molecule with all the electrons and nuclei in another. Any simple-looking formula we get represents a summation of complications, so we still have not got the fundamental phenomena.

Since the molecular forces attract at large distances and repel at short distances, as shown in Fig. 12–2, we can make up solids in which all the atoms are held together by their attractions and held apart by the repulsion that sets in when they are too close together. At a certain distance d (where the graph in Fig. 12–2 crosses the axis) the forces are zero, which means that they are all balanced, so that the molecules stay that distance apart from one another. If the molecules are pushed closer together than the distance d they all show a repulsion, represented by the portion of the graph above the r-axis. To push the molecules only slightly closer together requires a great force, because the molecular repulsion rapidly becomes very great at distances less than d. If the molecules are pulled slightly apart there is a slight attraction, which increases as the separation increases. If they are pulled sufficiently hard, they will separate permanently—the bond is broken.

If the molecules are pushed only a *very small* distance closer, or pulled only a *very small* distance farther than d, the corresponding distance along the curve of Fig. 12–2 is also very small, and can then be approximated by a straight line. Therefore, in many circumstances, if the displacement is not too great *the force is proportional to the displacement.* This principle is known as Hooke's law, or the law of elasticity, which says that the force in a body which tries to restore the body

to its original condition when it is distorted is proportional to the distortion. This law, of course, holds true only if the distortion is relatively small; when it gets too large the body will be torn apart or crushed, depending on the kind of distortion. The amount of force for which Hooke's law is valid depends upon the material; for instance, for dough or putty the force is very small, but for steel it is relatively large. Hooke's law can be nicely demonstrated with a long coil spring, made of steel and suspended vertically. A suitable weight hung on the lower end of the spring produces a tiny twist throughout the length of the wire, which results in a small vertical deflection in each turn and adds up to a large displacement if there are many turns. If the total elongation produced, say, by a 100-gram weight, is measured, it is found that additional weights of 100 grams will each produce an additional elongation that is very nearly equal to the stretch that was measured for the first 100 grams. This constant ratio of force to displacement begins to change when the spring is overloaded, i.e., Hooke's law no longer holds.

12–4 Fundamental forces. Fields

We shall now discuss the only remaining forces that are fundamental. We call them fundamental in the sense that their laws are fundamentally simple. We shall first discuss electrical force. Objects carry electrical charges which consist simply of electrons or protons. If any two bodies are electrically charged, there is an electrical force between them, and if the magnitudes of the charges are q_1 and q_2, respectively, the force varies inversely as the square of the distance between the charges, or $F = (\text{const})\, q_1 q_2 / r^2$. For unlike charges, this law is like the law of gravitation, but for *like* charges the force is repulsive and the sign (direction) is reversed. The charges q_1 and q_2 can be intrinsically either positive or negative, and in any specific application of the formula the direction of the force will come out right if the q's are given the proper plus or minus sign; the force is directed along the line between the two charges. The constant in the formula depends, of course, upon what units are used for the force, the charge, and the distance. In current practice the charge is measured in coulombs, the distance in meters, and the force in newtons. Then, in order to get the force to come out properly in newtons, the constant (which for historical reasons is written $1/4\pi\epsilon_0$) takes the numerical value

$$\epsilon_0 = 8.854 \times 10^{-12}\ \text{coul}^2/\text{newton} \cdot \text{m}^2$$

or

$$1/4\pi\epsilon_0 = 8.99 \times 10^9\ \text{n} \cdot \text{m}^2/\text{coul}^2.$$

Thus the force law for static charges is

$$\mathbf{F} = q_1 q_2 \mathbf{r} / 4\pi\epsilon_0 r^3. \tag{12.2}$$

In nature, the most important charge of all is the charge on a single electron, which is 1.60×10^{-19} coulomb. In working with electrical forces between fundamental particles rather than with large charges, many people prefer the combination $(q_{el})^2/4\pi\epsilon_0$, in which q_{el} is defined as the charge on an electron. This combination occurs frequently, and to simplify calculations it has been defined by the symbol e^2; its numerical value in the mks system of units turns out to be $(1.52 \times 10^{-14})^2$. The advantage of using the constant in this form is that the force between two electrons in newtons can then be written simply as e^2/r^2, with r in meters, without all the individual constants. Electrical forces are much more complicated than this simple formula indicates, since the formula gives the force between two objects only when the objects are standing still. We shall consider the more general case shortly.

In the analysis of forces of the more fundamental kinds (not such forces as friction, but the electrical force or the gravitational force), an interesting and very important concept has been developed. Since at first sight the forces are very much more complicated than is indicated by the inverse-square laws and these laws hold true only when the interacting bodies are standing still, an improved

method is needed to deal with the very complex forces that ensue when the bodies start to move in a complicated way. Experience has shown that an approach known as the concept of a "field" is of great utility for the analysis of forces of this type. To illustrate the idea for, say, electrical force, suppose we have two electrical charges, q_1 and q_2, located at points P and R respectively. Then the force between the charges is given by

$$\mathbf{F} = q_1 q_2 \mathbf{r}/r^3. \tag{12.3}$$

To analyze this force by means of the field concept, we say that the charge q_1 at P produces a "condition" at R, such that when the charge q_2 is placed at R it "feels" the force. This is one way, strange perhaps, of describing it; we say that the force $\mathbf{F}$ on q_2 at R can be written in two parts. It is q_2 multiplied by a quantity $\mathbf{E}$ that would be there whether q_2 were there or not (provided we keep all the other charges in their right places). $\mathbf{E}$ is the "condition" produced by q_1, we say, and $\mathbf{F}$ is the response of q_2 to $\mathbf{E}$. $\mathbf{E}$ is called an *electric field*, and it is a vector. The formula for the electric field $\mathbf{E}$ that is produced at R by a charge q_1 at P is the charge q_1 times the constant $1/4\pi\epsilon_0$ divided by r^2 (r is the distance from P to R), and it is acting in the direction of the radius vector (the radius vector $\mathbf{r}$ divided by its own length). The expression for $\mathbf{E}$ is thus

$$\mathbf{E} = q_1 \mathbf{r}/4\pi\epsilon_0 r^3. \tag{12.4}$$

We then write

$$\mathbf{F} = q_2 \mathbf{E}, \tag{12.5}$$

which expresses the force, the field, and the charge in the field. What is the point of all this? The point is to divide the analysis into two parts. One part says that something *produces* a field. The other part says that something is *acted on* by the field. By allowing us to look at the two parts independently, this separation of the analysis simplifies the calculation of a problem in many situations. If many charges are present, we first work out the total electric field produced at R by all the charges, and then, knowing the charge that is placed at R, we find the force on it.

In the case of gravitation, we can do exactly the same thing. In this case, where the force $\mathbf{F} = -Gm_1 m_2 \mathbf{r}/r^3$, we can make an analogous analysis, as follows: the force on a body in a gravitational field is the mass of that body times the field $\mathbf{C}$. The force on m_2 is the mass m_2 times the field $\mathbf{C}$ produced by m_1; that is, $\mathbf{F} = m_2\mathbf{C}$. Then the field $\mathbf{C}$ produced by a body of mass m_1 is $\mathbf{C} = -Gm_1\mathbf{r}/r^3$ and it is directed radially, as in the electrical case.

In spite of how it might at first seem, this separation of one part from another is not a triviality. It would be trivial, just another way of writing the same thing, if the laws of force were simple, but the laws of force are so complicated that it turns out that the fields have a reality that is almost independent of the objects which create them. One can do something like shake a charge and produce an effect, a field, at a distance; if one then stops moving the charge, the field keeps track of all the past, because the interaction between two particles is not instantaneous. It is desirable to have some way to remember what happened previously. If the force upon some charge depends upon where another charge was yesterday, which it does, then we need machinery to keep track of what went on yesterday, and that is the character of a field. So when the forces get more complicated, the field becomes more and more real, and this technique becomes less and less of an artificial separation.

In analyzing forces by the use of fields, we need two kinds of laws pertaining to fields. The first is the response to a field, and that gives the equations of motion. For example, the law of response of a mass to a gravitational field is that the force is equal to the mass times the gravitational field; or, if there is also a charge on the body, the response of the charge to the electric field equals the charge times the electric field. The second part of the analysis of nature in these situations is to formulate the laws which determine the strength of the field and how it is produced. These laws are sometimes called the *field equations*. We shall learn more about them in due time, but shall write down a few things about them now.

First, the most remarkable fact of all, which is true exactly and which can be easily understood, is that the total electric field produced by a number of sources is the vector sum of the electric fields produced by the first source, the second source, and so on. In other words, if we have numerous charges making a field, and if all by itself one of them would make the field $\mathbf{E}_1$, another would make the field $\mathbf{E}_2$, and so on, then we merely add the vectors to get the total field. This principle can be expressed as

$$\mathbf{E} = \mathbf{E}_1 + \mathbf{E}_2 + \mathbf{E}_3 + \cdots \tag{12.6}$$

or, in view of the definition given above,

$$\mathbf{E} = \sum_i \frac{q_i \mathbf{r}_i}{4\pi\epsilon_0 r_i^3}. \tag{12.7}$$

Can the same methods be applied to gravitation? The force between two masses m_1 and m_2 was expressed by Newton as $\mathbf{F} = Gm_1m_2\mathbf{r}/r^3$. But according to the field concept, we may say that m_1 creates a field $\mathbf{C}$ in all the surrounding space, such that the force on m_2 is given by

$$\mathbf{F} = m_2\mathbf{C}. \tag{12.8}$$

By complete analogy with the electrical case,

$$\mathbf{C} = -Gm_i\mathbf{r}_i/r_i^3 \tag{12.9}$$

and the gravitational field produced by several masses is

$$\mathbf{C} = \mathbf{C}_1 + \mathbf{C}_2 + \mathbf{C}_3 + \cdots \tag{12.10}$$

In Chapter 7, in working out a case of planetary motion, we used this principle in essence. We simply added all the force vectors to get the resultant force on a planet. If we divide out the mass of the planet in question, we get Eq. (12.10).

Equations (12.6) and (12.10) express what is known as *the principle of superposition* of fields. This principle states that the total field due to all the sources is the sum of the fields due to each source. So far as we know today, for electricity this is an absolutely guaranteed law, which is true even when the force law is complicated because of the motions of the charges. There are apparent violations, but more careful analysis has always shown these to be due to the overlooking of certain moving charges. However, although the principle of superposition applies exactly for electrical forces, it is not exact for gravity if the field is too strong, and Newton's equation (12.10) is only approximate, according to Einstein's gravitational theory.

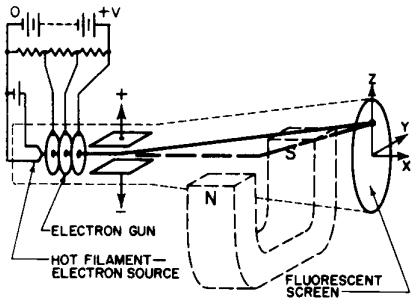


Fig. 12–3. An electron-beam tube.

Closely related to electrical force is another kind, called magnetic force, and this too is analyzed in terms of a field. Some of the qualitative relations between electrical and magnetic forces can be illustrated by an experiment with an electron-ray tube (Fig. 12–3). At one end of such a tube is a source that emits a stream of electrons. Within the tube are arrangements for accelerating the electrons to a high speed and sending some of them in a narrow beam to a fluorescent screen at the other end of the tube. A spot of light glows in the center of the screen where the electrons strike, and this enables us to trace the electron path. On the way to the screen the electron beam passes through a narrow space between a pair of parallel metal plates, which are arranged, say, horizontally. A voltage can be applied across the plates, so that either plate can be made negative at will. When such a voltage is present, there is an electric field between the plates.

The first part of the experiment is to apply a negative voltage to the lower plate, which means that extra electrons have been placed on the lower plate. Since like charges repel, the light spot on the screen instantly shifts upward. (We could also say this in another way—that the electrons "felt" the field, and responded by deflecting upward.) We next reverse the voltage, making the *upper* plate negative. The light spot on the screen now jumps below the center, showing that the electrons in the beam were repelled by those in the plate above them. (Or we could say again

that the electrons had "responded" to the field, which is now in the reverse direction.)

The second part of the experiment is to disconnect the voltage from the plates and test the effect of a magnetic field on the electron beam. This is done by means of a horseshoe magnet, whose poles are far enough apart to more or less straddle the tube. Suppose we hold the magnet below the tube in the same orientation as the letter U, with its poles up and part of the tube in between. We note that the light spot is deflected, say, upward, as the magnet approaches the tube from below. So it appears that the magnet repels the electron beam. However, it is not that simple, for if we invert the magnet without reversing the poles side-for-side, and now approach the tube from above, the spot still moves *upward*, so the electron beam is *not* repelled; instead, it appears to be attracted this time. Now we start again, restoring the magnet to its original U orientation and holding it below the tube, as before. Yes, the spot is still deflected upward; but now turn the magnet 180 degrees around a vertical axis, so that it is still in the U position but the poles are reversed side-for-side. Behold, the spot now jumps downward, and stays down, even if we invert the magnet and approach from above, as before.

To understand this peculiar behavior, we have to have a new combination of forces. We explain it thus: Across the magnet from one pole to the other there is a *magnetic field.* This field has a direction which is always away from one particular pole (which we could mark) and toward the other. Inverting the magnet did not change the direction of the field, but reversing the poles side-for-side did reverse its direction. For example, if the electron velocity were horizontal in the x-direction and the magnetic field were also horizontal but in the y-direction, the magnetic force *on the moving electrons* would be in the z-direction, i.e., up or down, depending on whether the field was in the positive or negative y-direction.

Although we shall not at the present time give the correct law of force between charges moving in an arbitrary manner, one relative to the other, because it is too complicated, we shall give one aspect of it: the complete law of the forces *if the fields are known.* The force on a charged object depends upon its motion; if, when the object is standing still at a given place, there is some force, this is taken to be proportional to the charge, the coefficient being what we call the *electric field.* When the object moves the force may be different, and the correction, the new "piece" of force, turns out to be dependent exactly *linearly on the velocity*, but at *right angles* to $\mathbf{v}$ and to another vector quantity which we call the *magnetic induction* $\mathbf{B}$. If the components of the electric field $\mathbf{E}$ and the magnetic induction $\mathbf{B}$ are, respectively, (E_x, E_y, E_z) and (B_x, B_y, B_z), and if the velocity $\mathbf{v}$ has the components (v_x, v_y, v_z), then the total electric and magnetic force on a moving charge q has the components

$$\begin{aligned} F_x &= q(E_x + v_yB_z - v_zB_y), \\ F_y &= q(E_y + v_zB_x - v_xB_z), \\ F_z &= q(E_z + v_xB_y - v_yB_x). \end{aligned} \tag{12.11}$$

If, for instance, the only component of the magnetic field were B_y and the only component of the velocity were v_x, then the only term left in the magnetic force would be a force in the z-direction, at right angles to both $\mathbf{B}$ and $\mathbf{v}$.

12–5 Pseudo forces

The next kind of force we shall discuss might be called a pseudo force. In Chapter 11 we discussed the relationship between two people, Joe and Moe, who use different coordinate systems. Let us suppose that the positions of a particle as measured by Joe are x and by Moe are x'; then the laws are as follows:

$$x = x' + s, \qquad y = y', \qquad z = z',$$

where s is the displacement of Moe's system relative to Joe's. If we suppose that

the laws of motion are correct for Joe, how do they look for Moe? We find first, that

$$dx/dt = dx'/dt + ds/dt.$$

Previously, we considered the case where s was constant, and we found that s made no difference in the laws of motion, since $ds/dt = 0$; ultimately, therefore, the laws of physics were the same in both systems. But another case we can take is that $s = ut$, where u is a uniform velocity in a straight line. Then s is not constant, and ds/dt is not zero, but is u, a constant. However, the acceleration d^2x/dt^2 is still the same as d^2x'/dt^2, because $du/dt = 0$. This proves the law that we used in Chapter 10, namely, that if we move in a straight line with uniform velocity the laws of physics will look the same to us as when we are standing still. That is the Galilean transformation. But we wish to discuss the interesting case where s is still more complicated, say $s = at^2/2$. Then $ds/dt = at$ and $d^2s/dt^2 = a$, a uniform acceleration; or in a still more complicated case, the acceleration might be a function of time. This means that although the laws of force from the point of view of Joe would look like

$$m\frac{d^2x}{dt^2} = F_x,$$

the laws of force as looked upon by Moe would appear as

$$m\frac{d^2x'}{dt^2} = F_{x'} - ma.$$

That is, since Moe's coordinate system is accelerating with respect to Joe's, the extra term ma comes in, and Moe will have to correct his forces by that amount in order to get Newton's laws to work. In other words, here is an apparent, mysterious new force of unknown origin which arises, of course, because Moe has the wrong coordinate system. This is an example of a pseudo force; other examples occur in coordinate systems that are *rotating*.

Another example of pseudo force is what is often called "centrifugal force." An observer in a rotating coordinate system, e.g., in a rotating box, will find mysterious forces, not accounted for by any known origin of force, throwing things outward toward the walls. These forces are due merely to the fact that the observer does not have Newton's coordinate system, which is the simplest coordinate system.

Pseudo force can be illustrated by an interesting experiment in which we push a jar of water along a table, with acceleration. Gravity, of course, acts downward on the water, but because of the horizontal acceleration there is also a pseudo force acting horizontally and in a direction opposite to the acceleration. The resultant of gravity and pseudo force makes an angle with the vertical, and during the acceleration the surface of the water will be perpendicular to the resultant force, i.e., inclined at an angle with the table, with the water standing higher in the rearward side of the jar. When the push on the jar stops and the jar decelerates because of friction, the pseudo force is reversed, and the water stands higher in the forward side of the jar (Fig. 12–4).

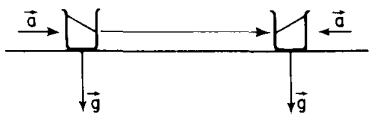


Fig. 12–4. Illustration of a pseudo force

One very important feature of pseudo forces is that they are always proportional to the masses; the same is true of gravity. The possibility exists, therefore, that *gravity itself is a pseudo force*. Is it not possible that perhaps gravitation is due simply to the fact that we do not have the right coordinate system? After all, we can always get a force proportional to the mass if we imagine that a body is accelerating. For instance, a man shut up in a box that is standing still on the earth finds himself held to the floor of the box with a certain force that is proportional to his mass. But if there were no earth at all and the box were standing still, the man inside would float in space. On the other hand, if there were no earth at all and something were *pulling* the box along with an acceleration g, then the man in the box, analyzing physics, would find a pseudo force which would pull him to the floor, just as gravity does.

Einstein put forward the famous hypothesis that accelerations give an imitation of gravitation, that the forces of acceleration (the pseudo forces) *cannot be distinguished* from those of gravity; it is not possible to tell how much of a given force is gravity and how much is pseudo force.

It might seem all right to consider gravity to be a pseudo force, to say that we are all held down because we are accelerating upward, but how about the people in Madagascar, on the other side of the earth—are they accelerating too? Einstein found that gravity could be considered a pseudo force only at one point at a time, and was led by his considerations to suggest that the *geometry of the world* is more complicated than ordinary Euclidean geometry. The present discussion is only qualitative, and does not pretend to convey anything more than the general idea. To give a rough idea of how gravitation could be the result of pseudo forces, we present an illustration which is purely geometrical and does not represent the real situation. Suppose that we all lived in two dimensions, and knew nothing of a third. We think we are on a plane, but suppose we are really on the surface of a sphere. And suppose that we shoot an object along the ground, with no forces on it. Where will it go? It will appear to go in a straight line, but it has to remain on the surface of a sphere, where the shortest distance between two points is along a great circle; so it goes along a great circle. If we shoot another object similarly, but in another direction, it goes along another great circle. Because we think we are on a plane, we expect that these two bodies will continue to diverge linearly with time, but careful observation will show that if they go far enough they move closer together again, as though they were attracting each other. But they are *not* attracting each other—there is just something "weird" about this geometry. This particular illustration does not describe correctly the way in which Euclid's geometry is "weird," but it illustrates that if we distort the geometry sufficiently it is possible that all gravitation is related in some way to pseudo forces; that is the general idea of the Einsteinian theory of gravitation.

12–6 Nuclear forces

We conclude this chapter with a brief discussion of the only other known forces, which are called *nuclear forces*. These forces are within the nuclei of atoms, and although they are much discussed, no one has ever calculated the force between two nuclei, and indeed at present there is no known law for nuclear forces. These forces have a very tiny range which is just about the same as the size of the nucleus, perhaps 10^{-13} centimeter. With particles so small and at such a tiny distance, only the quantum-mechanical laws are valid, not the Newtonian laws. In nuclear analysis we no longer think in terms of forces, and in fact we can replace the force concept with a concept of the energy of interaction of two particles, a subject that will be discussed later. Any formula that can be written for nuclear forces is a rather crude approximation which omits many complications; one might be somewhat as follows: forces within a nucleus do not vary inversely as the square of the distance, but die off exponentially over a certain distance r, as expressed by $F = (1/r^2)\exp(-r/r_0)$, where the distance r_0 is of the order of 10^{-13} centimeter. In other words, the forces disappear as soon as the particles are any great distance apart, although they are very strong within the 10^{-13} centimeter range. So far as they are understood today, the laws of nuclear force are very complex; we do not understand them in any simple way, and the whole problem of analyzing the fundamental machinery behind nuclear forces is unsolved. Attempts at a solution have led to the discovery of numerous strange particles, the π-mesons, for example, but the origin of these forces remains obscure.

13

Work and Potential Energy (A)

13–1 Energy of a falling body

13–2 Work done by gravity

13–3 Summation of energy

13–4 Gravitational field of large objects

13–1 Energy of a falling body

In Chapter 4 we discussed the conservation of energy. In that discussion, we did not use Newton's laws, but it is, of course, of great interest to see how it comes about that energy is in fact conserved in accordance with these laws. For clarity we shall start with the simplest possible example, and then develop harder and harder examples.

The simplest example of the conservation of energy is a vertically falling object, one that moves only in a vertical direction. An object which changes its height under the influence of gravity alone has a kinetic energy T (or K.E.) due to its motion during the fall, and a potential energy mgh, abbreviated U or P.E., whose sum is constant:

$$\underset{\text{K E}}{\tfrac{1}{2}mv^2} + \underset{\text{P E}}{mgh} = \text{const},$$

or

$$T + U = \text{const}. \tag{13.1}$$

Now we would like to show that this statement is true. What do we mean, show it is true? From Newton's Second Law we can easily tell how the object moves, and it is easy to find out how the velocity varies with time, namely, that it increases proportionally with the time, and that the height varies as the square of the time. So if we measure the height from a zero point where the object is stationary, it is no miracle that the height turns out to be equal to the square of the velocity times a number of constants. However, let us look at it a little more closely.

Let us find out *directly* from Newton's Second Law how the kinetic energy should change, by taking the derivative of the kinetic energy with respect to time and then using Newton's laws. When we differentiate $\frac{1}{2}mv^2$ with respect to time, we obtain

$$\frac{dT}{dt} = \frac{d}{dt}(\tfrac{1}{2}mv^2) = \tfrac{1}{2}m2v\frac{dv}{dt} = mv\frac{dv}{dt}, \tag{13.2}$$

since m is assumed constant. But from Newton's Second Law, $m(dv/dt) = F$, so that

$$dT/dt = Fv. \tag{13.3}$$

In general, it will come out to be $\mathbf{F} \cdot \mathbf{v}$, but in our one-dimensional case let us leave it as the force times the velocity.

Now in our simple example the force is constant, equal to $-mg$, a vertical force (the minus sign means that it acts downward), and the velocity, of course, is the rate of change of the vertical position, or height h, with time. Thus the rate of change of the kinetic energy is $-mg(dh/dt)$, which quantity, miracle of miracles, is the rate of change of something else! It is the time rate of change of mgh! Therefore, as time goes on, the changes in kinetic energy and in the quantity mgh are equal and opposite, so that the sum of the two quantities remains constant. Q.E.D.

We have shown, from Newton's second law of motion, that energy is conserved for constant forces when we add the potential energy mgh to the kinetic energy $\frac{1}{2}mv^2$. Now let us look into this further and see whether it can be generalized, and thus advance our understanding. Does it work only for a freely falling body, or is it more general? We expect from our discussion of the conservation of energy

that it would work for an object moving from one point to another in some kind of frictionless curve, under the influence of gravity (Fig. 13–1). If the object reaches a certain height h from the original height H, then the same formula should again be right, even though the velocity is now in some direction other than the vertical. We would like to understand *why* the law is still correct. Let us follow the same analysis, finding the time rate of change of the kinetic energy. This will again be $mv(dv/dt)$, but $m(dv/dt)$ is the rate of change of the magnitude of the momentum, i.e., the *force in the direction of motion*—the tangential force F_t. Thus

$$\frac{dT}{dt} = mv\frac{dv}{dt} = F_t v.$$

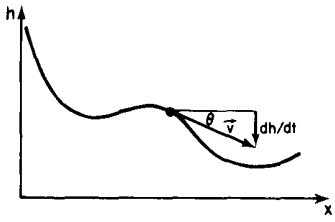


Fig. 13–1. An object moving on a frictionless curve under the influence of gravity.

Now the speed is the rate of change of distance along the curve, ds/dt, and the tangential force F_t is not mg but is weaker by the ratio of the distance ds along the path to the vertical distance dh. In other words,

$$F_t = -mg\sin\theta = -mg\frac{dh}{ds},$$

so that

$$F_t\frac{ds}{dt} = -mg\left(\frac{dh}{ds}\right)\left(\frac{ds}{dt}\right) = -mg\frac{dh}{dt},$$

since the ds's cancel. Thus we get $-mg(dh/dt)$, which is equal to the rate of change of mgh, as before.

In order to understand exactly how the conservation of energy works in general in mechanics, we shall now discuss a number of concepts which will help us to analyze it.

First, we discuss the rate of change of kinetic energy in general in three dimensions. The kinetic energy in three dimensions is

$$T = \tfrac{1}{2}m(v_x^2 + v_y^2 + v_z^2).$$

When we differentiate this with respect to time, we get three terrifying terms:

$$\frac{dT}{dt} = m\left(v_x\frac{dv_x}{dt} + v_y\frac{dv_y}{dt} + v_z\frac{dv_z}{dt}\right). \tag{13.4}$$

But $m(dv_x/dt)$ is the force F_x acting on the object in the x-direction. Thus the right side of Eq. (13.4) is $F_xv_x + F_yv_y + F_zv_z$. We recall our vector analysis and recognize this as $\mathbf{F}\cdot\mathbf{v}$; therefore

$$dT/dt = \mathbf{F}\cdot\mathbf{v}. \tag{13.5}$$

This result can be derived more quickly as follows: if $\mathbf{a}$ and $\mathbf{b}$ are two vectors, both of which may depend upon the time, the derivative of $\mathbf{a}\cdot\mathbf{b}$ is, in general,

$$d(\mathbf{a}\cdot\mathbf{b})/dt = \mathbf{a}\cdot d\mathbf{b}/dt + (d\mathbf{a}/dt)\cdot\mathbf{b}. \tag{13.6}$$

We then use this in the form $\mathbf{a} = \mathbf{b} = \mathbf{v}$:

$$\frac{d(\frac{1}{2}mv^2)}{dt} = \frac{d(\frac{1}{2}m\mathbf{v}\cdot\mathbf{v})}{dt} = m\frac{d\mathbf{v}}{dt}\cdot\mathbf{v} = \mathbf{F}\cdot\mathbf{v} = \mathbf{F}\cdot\frac{d\mathbf{s}}{dt}. \tag{13.7}$$

Because the concepts of kinetic energy, and energy in general, are so important, various names have been given to the important terms in equations such as these. $\frac{1}{2}mv^2$ is, as we know, called *kinetic energy*. $\mathbf{F}\cdot\mathbf{v}$ is called *power*: the force acting on an object times the velocity of the object (vector "dot" product) is the power being delivered to the object by that force. We thus have a marvelous theorem: *the rate of change of kinetic energy of an object is equal to the power expended by the forces acting on it.*

However, to study the conservation of energy, we want to analyze this still more closely. Let us evaluate the change in kinetic energy in a very short time dt. If we multiply both sides of Eq. (13.7) by dt, we find that the differential change in

the kinetic energy is the force "dot" the differential distance moved:

$$dT = \mathbf{F} \cdot d\mathbf{s}. \tag{13.8}$$

If we now integrate, we get

$$\Delta T = \int_1^2 \mathbf{F} \cdot d\mathbf{s}. \tag{13.9}$$

What does this mean? It means that if an object is moving *in any way* under the influence of a force, moving in some kind of curved path, then the change in K.E. when it goes from one point to another along the curve is equal to the integral of the component of the force along the curve times the differential displacement $d\mathbf{s}$, the integral being carried out from one point to the other. This integral also has a name; it is called the *work done by the force on the object.* We see immediately that *power equals work done per second.* We also see that it is only a component of force *in the direction of motion* that contributes to the work done. In our simple example the forces were only vertical, and had only a single component, say F_z, equal to $-mg$. No matter how the object moves in those circumstances, falling in a parabola for example, $\mathbf{F} \cdot d\mathbf{s}$, which can be written as $F_x\,dx + F_y\,dy + F_z\,dz$, has nothing left of it but $F_z\,dz = -mg\,dz$, because the other components of force are zero. Therefore, in our simple case,

$$\int_1^2 \mathbf{F} \cdot d\mathbf{s} = \int_{z_1}^{z_2} -mg\,dz = -mg(z_2 - z_1), \tag{13.10}$$

so again we find that it is only the *vertical height* from which the object falls that counts toward the potential energy.

A word about units. Since forces are measured in newtons, and we multiply by a distance in order to obtain work, work is measured in *newton · meters* (n · m), but people do not like to say newton-meters, they prefer to say *joules* (j). A newton-meter is called a joule; work is measured in joules. Power, then, is joules per second, and that is also called a *watt* (w). If we multiply watts by time, the result is the work done. The work done by the electrical company in our houses, technically, is equal to the watts times the time. That is where we get things like kilowatt hours, 1000 watts times 3600 seconds, or 3.6×10^6 joules.

Now we take another example of the law of conservation of energy. Consider an object which initially has kinetic energy and is moving very fast, and which slides against the floor with friction. It stops. At the start the kinetic energy is *not* zero, but at the end it *is* zero; there is work done by the forces, because whenever there is friction there is always a component of force in a direction opposite to that of the motion, and so energy is steadily lost. But now let us take a mass on the end of a pivot swinging in a vertical plane in a gravitational field with no friction. What happens here is different, because when the mass is going up the force is downward, and when it is coming down, the force is also downward. Thus $\mathbf{F} \cdot d\mathbf{s}$ has one sign going up and another sign coming down. At each corresponding point of the downward and upward paths the values of $\mathbf{F} \cdot d\mathbf{s}$ are exactly equal in size but of opposite sign, so the net result of the integral will be zero for this case. Thus the kinetic energy with which the mass comes back to the bottom is the same as it had when it left, that is the principle of the conservation of energy. (Note that when there are friction forces the conservation of energy seems at first sight to be invalid. We have to find another *form* of energy. It turns out, in fact, that *heat* is generated in an object when it rubs another with friction, but at the moment we supposedly do not know that.)

13–2 Work done by gravity

The next problem to be discussed is much more difficult than the above; it has to do with the case when the forces are not constant, or simply vertical, as they were in the cases we have worked out. We want to consider a planet, for example, moving around the sun, or a satellite in the space around the earth.

We shall first consider the motion of an object which starts at some point 1 and falls, say, *directly* toward the sun or toward the earth (Fig. 13–2). Will there be a law of conservation of energy in these circumstances? The only difference is that in this case, the force is *changing* as we go along, it is not just a constant. As we know, the force is GM/r^2 times the mass m, where m is the mass that moves. Now certainly when a body falls toward the earth, the kinetic energy increases as the distance fallen increases, just as it does when we do not worry about the variation of force with height. The question is whether it is possible to find another formula for potential energy different from mgh, a different function of distance away from the earth, so that conservation of energy will still be true.

This one-dimensional case is easy to treat because we know that the change in the kinetic energy is equal to the integral, from one end of the motion to the other, of $-GMm/r^2$ times the displacement dr:

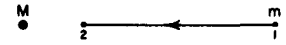


Fig. 13–2. A small mass m falls under the influence of gravity toward a large mass M.

$$T_2 - T_1 = -\int_1^2 GMm\,\frac{dr}{r^2}. \tag{13.11}$$

There are no cosines needed for this case because the force and the displacement are in the same direction. It is easy to integrate dr/r^2; the result is $-1/r$, so Eq. (13.11) becomes

$$T_2 - T_1 = +GMm\left(\frac{1}{r_2} - \frac{1}{r_1}\right). \tag{13.12}$$

Thus we have a different formula for potential energy. Equation (13.12) tells us that the quantity $(\frac{1}{2}mv^2 - GMm/r)$ calculated at point 1, at point 2, or at any other place, has a constant value.

We now have the formula for the potential energy in a gravitational field for vertical motion. Now we have an interesting problem. Can we make *perpetual motion* in a gravitational field? The gravitational field varies; in different places it is in different directions and has different strengths. Could we do something like this, using a fixed, frictionless track: start at some point and lift an object out to some other point, then move it around an arc to a third point, then lower it a certain distance, then move it in at a certain slope and pull it out some other way, so that when we bring it back to the starting point, a certain amount of work has been done by the gravitational force, and the kinetic energy of the object is increased? Can we design the curve so that it comes back moving a little bit faster than it did before, so that it goes around and around and around, and gives us perpetual motion? Since perpetual motion is impossible, we ought to find out that this is also impossible. We ought to discover the following proposition: since there is no friction the object should come back with neither higher nor lower velocity—it should be able to keep going around and around any closed path. Stated in another way, *the total work done in going around a complete cycle should be zero* for gravity forces, because if it is not zero we can get energy out by going around. (If the work turns out to be less than zero, so that we get less speed when we go around one way, then we merely go around the other way, because the forces, of course, depend only upon the position, not upon the direction; if one way is plus, the other way would be minus, so unless it is zero we will get perpetual motion by going around either way.)

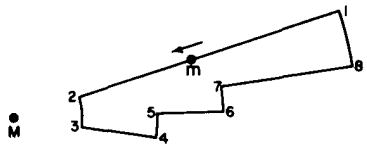


Fig. 13–3. A closed path in a gravitational field.

Is the work really zero? Let us try to demonstrate that it is. First we shall explain more or less why it is zero, and then we shall examine it a little better mathematically. Suppose that we use a simple path such as that shown in Fig. 13–3, in which a small mass is carried from point 1 to point 2, and then is made to go around a circle to 3, back to 4, then to 5, 6, 7, and 8, and finally back to 1. All of the lines are either purely radial or circular, with M as the center. How much work is done in carrying m around this path? Between points 1 and 2, it is GMm times the difference of $1/r$ between these two points:

$$W_{12} = \int_1^2 \mathbf{F}\cdot d\mathbf{s} = \int_1^2 -GMm\,\frac{dr}{r^2} = -GMm\left(\frac{1}{r_2} - \frac{1}{r_1}\right).$$

From 2 to 3 the force is exactly at right angles to the curve, so that $W_{23} \equiv 0$. The work from 3 to 4 is

$$W_{34} = \int_3^4 \mathbf{F} \cdot d\mathbf{s} = -GMm\left(\frac{1}{r_4} - \frac{1}{r_3}\right).$$

In the same fashion, we find that $W_{45} = 0$, $W_{56} = -GMm(1/r_6 - 1/r_5)$, $W_{67} = 0$, $W_{78} = -GMm(1/r_8 - 1/r_7)$, and $W_{81} = 0$. Thus

$$W = GMm\left(\frac{1}{r_1} - \frac{1}{r_2} + \frac{1}{r_3} - \frac{1}{r_4} + \frac{1}{r_5} - \frac{1}{r_6} + \frac{1}{r_7} - \frac{1}{r_8}\right).$$

But we note that $r_2 = r_3$, $r_4 = r_5$, $r_6 = r_7$, and $r_8 = r_1$. Therefore $W = 0$.

Of course we may wonder whether this is too trivial a curve. What if we use a *real* curve? Let us try it on a real curve. First of all, we might like to assert that a real curve could always be imitated sufficiently well by a series of sawtooth jiggles like those of Fig. 13–4, and that therefore, etc., Q.E.D., but without a little analysis, it is not obvious at first that the work done going around even a small triangle is zero. Let us magnify one of the triangles, as shown in Fig. 13–4. Is the work done in going from a to b and b to c on a triangle the same as the work done in going directly from a to c? Suppose that the force is acting in a certain direction; let us take the triangle such that the side bc is in this direction, just as an example. We also suppose that the triangle is so small that the force is essentially constant over the entire triangle. What is the work done in going from a to c? It is

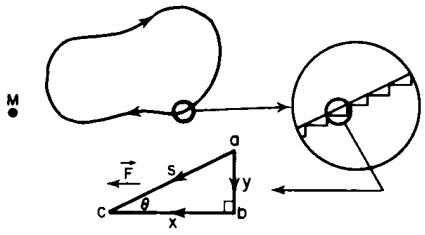


Fig. 13–4. A "smooth" closed path, showing a magnified segment of it approximated by a series of radial and circumferential steps, and an enlarged view of one step.

$$W_{ac} = \int_a^c \mathbf{F} \cdot d\mathbf{s} = Fs\cos\theta,$$

since the force is constant. Now let us calculate the work done in going around the other two sides of the triangle. On the vertical side ab the force is perpendicular to ds, so that here the work is zero. On the horizontal side bc,

$$W_{bc} = \int_b^c \mathbf{F} \cdot d\mathbf{s} = Fx.$$

Thus we see that the work done in going along the sides of a small triangle is the same as that done going on a slant, because $s\cos\theta$ is equal to x. We have proved previously that the answer is zero for any path composed of a series of notches like those of Fig. 13–3, and also that we do the same work if we cut across the corners instead of going along the notches (so long as the notches are fine enough, and we can always make them very fine); therefore, *the work done in going around any path in a gravitational field is zero.*

This is a very remarkable result. It tells us something we did not previously know about planetary motion. It tells us that when a planet moves around the sun (without any other objects around, no other forces) it moves in such a manner that the square of the speed at any point minus some constants divided by the radius at that point is always the same at every point on the orbit. For example, the closer the planet is to the sun, the faster it is going, but by how much? By the following amount: if instead of letting the planet go around the sun, we were to change the direction (but not the magnitude) of its velocity and make it move radially, and then we let it fall from some special radius to the radius of interest, the new speed would be the same as the speed it had in the actual orbit, because this is just another example of a complicated path. So long as we come back to the same distance, the kinetic energy will be the same. So, whether the motion is the real, undisturbed one, or is changed in direction by channels, by frictionless constraints, the kinetic energy with which the planet arrives at a point will be the same.

Thus, when we make a numerical analysis of the motion of the planet in its orbit, as we did earlier, we can check whether or not we are making appreciable errors by calculating this constant quantity, the energy, at every step, and it should not change. For the orbit of Table 9–2 the energy does change,* it changes by

* The energy is $\frac{1}{2}(v_x^2 + v_y^2) - 1/r$ in the units of Table 9–2.

some 1.5 percent from the beginning to the end. Why? Either because for the numerical method we use finite intervals, or else because we made a slight mistake somewhere in arithmetic.

Let us consider the energy in another case: the problem of a mass on a spring. When we displace the mass from its balanced position, the restoring force is proportional to the displacement. In those circumstances, can we work out a law for conservation of energy? Yes, because the work done by such a force is

$$W = \int_0^x F\,dx = \int_0^x -kx\,dx = -\tfrac{1}{2}kx^2. \tag{13.13}$$

Therefore, for a mass on a spring we have that the kinetic energy of the oscillating mass plus $\frac{1}{2}kx^2$ is a constant. Let us see how this works. We pull the mass down; it is standing still and so its speed is zero. But x is not zero, x is at its maximum, so there is some energy, the potential energy, of course. Now we release the mass and things begin to happen (the details not to be discussed), but at any instant the kinetic plus potential energy must be a constant. For example, after the mass is on its way past the original equilibrium point, the position x equals zero, but that is when it has its biggest v^2, and as it gets more x^2 it gets less v^2, and so on. So the balance of x^2 and v^2 is maintained as the mass goes up and down. Thus we have another rule now, that the potential energy for a spring is $\frac{1}{2}kx^2$, if the force is $-kx$.

13–3 Summation of energy

Now we go on to the more general consideration of what happens when there are large numbers of objects. Suppose we have the complicated problem of many objects, which we label $i = 1, 2, 3, \ldots$, all exerting graviational pulls on each other. What happens then? We shall prove that if we add the kinetic energies of all the particles, and add to this the sum, over all *pairs* of particles, of their mutual gravitational potential energy, $-GMm/r_{ij}$, the total is a constant:

$$\sum_i \tfrac{1}{2}m_i v_i^2 + \sum_{(\text{pairs } ij)} -\frac{Gm_i m_j}{r_{ij}} = \text{const.} \tag{13.14}$$

How do we prove it? We differentiate each side with respect to time and get zero. When we differentiate $\frac{1}{2}m_i v_i^2$, we find derivatives of the velocity that are the forces, just as in Eq. (13.5). We replace these forces by the law of force that we know from Newton's law of gravity and then we notice that what is left is the same as the time derivative of

$$\sum_{\text{pairs}} -\frac{Gm_i m_j}{r_{ij}}.$$

The time derivative of the kinetic energy is

$$\begin{aligned}\frac{d}{dt}\sum_i \tfrac{1}{2}m_i v_i^2 &= \sum_i m_i \mathbf{v}_i \cdot \frac{d\mathbf{v}_i}{dt} \\ &= \sum_i \mathbf{F}_i \cdot \mathbf{v}_i \\ &= \sum_i \left(\sum_j -\frac{Gm_i m_j \mathbf{r}_{ij}}{r_{ij}^3}\right)\cdot \mathbf{v}_i.\end{aligned} \tag{13.15}$$

The time derivative of the potential energy is

$$\frac{d}{dt}\sum_{\text{pairs}} -\frac{Gm_i m_j}{r_{ij}} = \sum_{\text{pairs}} \left(+\frac{Gm_i m_j}{r_{ij}^2}\right)\left(\frac{dr_{ij}}{dt}\right).$$

But

$$r_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2},$$

so that

$$\frac{dr_{ij}}{dt} = \frac{1}{2r_{ij}}\left[2(x_i - x_j)\left(\frac{dx_i}{dt} - \frac{dx_j}{dt}\right) + 2(y_i - y_j)\left(\frac{dy_i}{dt} - \frac{dy_j}{dt}\right) + 2(z_i - z_j)\left(\frac{dz_i}{dt} - \frac{dz_j}{dt}\right)\right]$$

$$= \mathbf{r}_{ij} \cdot \frac{\mathbf{v}_i - \mathbf{v}_j}{r_{ij}}$$

$$= \mathbf{r}_{ij} \cdot \frac{\mathbf{v}_i}{r_{ij}} + \mathbf{r}_{ji} \cdot \frac{\mathbf{v}_j}{r_{ji}},$$

since $\mathbf{r}_{ij} = -\mathbf{r}_{ji}$, while $r_{ij} = r_{ji}$. Thus

$$\frac{d}{dt}\sum_{\text{pairs}} -\frac{Gm_im_j}{r_{ij}} = \sum_{\text{pairs}}\left[\frac{Gm_im_j\,\mathbf{r}_{ij}}{r_{ij}^3} \cdot \mathbf{v}_i + \frac{Gm_jm_i\,\mathbf{r}_{ji}}{r_{ji}^3} \cdot \mathbf{v}_j\right]. \quad (13.16)$$

Now we must note carefully what $\sum_i \{\sum_j\}$ and $\sum_{\text{pairs}}$ mean. In Eq. (13.15), $\sum_i \{\sum_j\}$ means that i takes on all values $i = 1, 2, 3, \ldots$ in turn, and for each value of i, the index j takes on all values except i. Thus if $i = 3$, j takes on the values $1, 2, 4, \ldots$

In Eq. (13.16), on the other hand, $\sum_{\text{pairs}}$ means that given values of i and j occur only once. Thus the particle pair 1 and 3 contributes only one term to the sum. To keep track of this, we might agree to let i range over all values $1, 2, 3, \ldots$, and for each i let j range only over values *greater* than i. Thus if $i = 3$, j could only have values $4, 5, 6, \ldots$ But we notice that for each i, j value there are two contributions to the sum, one involving $\mathbf{v}_i$, and the other $\mathbf{v}_j$, and that these terms have the same appearance as those of Eq. (13.14), where *all* values of i and j (except $i = j$) are included in the sum. Therefore, by matching the terms one by one, we see that Eqs. (13.16) and (13.15) are precisely the same, but of opposite sign, so that the time derivative of the kinetic plus potential energy is indeed zero. Thus we see that, for many objects, *the kinetic energy is the sum of the contributions from each individual object*, and that the potential energy is also simple, it being also just a sum of contributions, the energies between all the pairs. We can understand *why* it should be the energy of every pair this way: Suppose that we want to find the total amount of work that must be done to bring the objects to certain distances from each other. We may do this in several steps, bringing them in from infinity where there is no force, one by one. First we bring in number one, which requires no work, since no other objects are yet present to exert force on it. Next we bring in number two, which does take some work, namely $W_{12} = -Gm_1m_2/r_{12}$. Now, and this is an important point, suppose we bring in the next object to position three. At any moment the force on number 3 can be written as the sum of two forces—the force exerted by number 1 and that exerted by number 2. Therefore *the work done is the sum of the works done by each*, because if $\mathbf{F}_3$ can be resolved into the sum of two forces,

$$\mathbf{F}_3 = \mathbf{F}_{13} + \mathbf{F}_{23},$$

then the work is

$$\int \mathbf{F}_3 \cdot d\mathbf{s} = \int \mathbf{F}_{13} \cdot d\mathbf{s} + \int \mathbf{F}_{23} \cdot d\mathbf{s} = W_{13} + W_{23}.$$

That is, the work done is the sum of the work done against the first force and the second force, as if each acted independently. Proceeding in this way, we see that the total work required to assemble the given configuration of objects is precisely the value given in Eq. (13.14) as the potential energy. It is because gravity obeys the principle of superposition of forces that we can write the potential energy as a sum over each pair of particles.

13–4 Gravitational field of large objects

Now we shall calculate the fields which are met in a few physical circumstances involving *distributions of mass*. We have not so far considered distributions of mass, only particles, so it is interesting to calculate the forces when they are produced by more than just one particle. First we shall find the gravitational force on a mass that is produced by a plane sheet of material, infinite in extent. The force on a unit mass at a given point P, produced by this sheet of material (Fig. 13–5), will of course be directed toward the sheet. Let the distance of the point from the sheet be a, and let the amount of mass per unit area of this huge sheet be μ. We shall suppose μ to be constant; it is a uniform sheet of material. Now, what small field $d\mathbf{C}$ is produced by the mass dm lying between ρ and $\rho + d\rho$ from the point O of the sheet nearest point P? Answer: $d\mathbf{C} = G(dm\mathbf{r}/r^3)$. But this field is directed along $\mathbf{r}$, and we know that only the x-component of it will remain when we add all the little vector $d\mathbf{C}$'s to produce $\mathbf{C}$. The x-component of $d\mathbf{C}$ is

$$dC_x = G\frac{dmr_x}{r^3} = G\frac{dma}{r^3}.$$

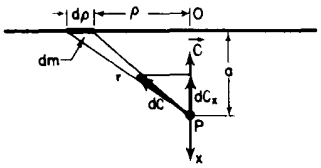


Fig. 13–5. The gravitational force F on a mass point produced by an infinite plane sheet of matter.

Now all masses dm which are at the same distance r from P will yield the same dC_x, so we may at once write for dm the total mass in the *ring* between ρ and $\rho + d\rho$, namely $dm = \mu 2\pi\rho\, d\rho$ ($2\pi\rho\, d\rho$ is the area of a ring of radius ρ and width $d\rho$, if $d\rho \ll \rho$). Thus

$$dC_x = G\mu 2\pi\rho\frac{d\rho a}{r^3}.$$

Then, since $r^2 = \rho^2 + a^2$, $\rho\, d\rho = r\, dr$. Therefore,

$$C_x = 2\pi G\mu a\int_a^\infty \frac{dr}{r^2} = 2\pi G\mu a\left(\frac{1}{a} - \frac{1}{\infty}\right) = 2\pi G\mu. \qquad (13.17)$$

Thus the force is independent of distance a! Why? Have we made a mistake? One might think that the farther away we go, the weaker the force would be. But no! If we are close, most of the matter is pulling at an unfavorable angle; if we are far away, more of the matter is situated more favorably to exert a pull toward the plane. At any distance, the matter which is most effective lies in a certain cone. When we are farther away the force is smaller by the inverse square, but in the same cone, in the same angle, there is much *more matter*, larger by just the square of the distance! This analysis can be made rigorous by just noticing that the differential contribution in any given cone is in fact independent of the distance, because of the reciprocal variation of the strength of the force from a given mass, and the amount of mass included in the cone, with changing distance. The force is not really constant of course, because when we go on the other side of the sheet it is reversed in sign.

We have also, in effect, solved an electrical problem: if we have an electrically charged plate, with an amount σ of charge per unit area, then the electric field at a point outside the sheet is equal to $\sigma/2\epsilon_0$, and is in the outward direction if the sheet is positively charged, and inward if the sheet is negatively charged. To prove this, we merely note that G, gravity, plays the same role as $1/4\pi\epsilon_0$ for electricity.

Now suppose that we have two plates, with a positive charge $+\sigma$ on one and a negative charge $-\sigma$ on another at a distance D from the first. What is the field? Outside the two plates it is zero. Why? Because one attracts and the other repels, the force being *independent of distance*, so that the two balance out! Also, the force *between* the two plates is clearly twice as great as that from one plate, namely $E = \sigma/\epsilon_0$, and is directed from the positive plate to the negative one.

Now we come to a most interesting and important problem, whose solution we have been assuming all the time, namely, that the force produced by the earth at a point on the surface or outside it is the same as if all the mass of the earth were located at its center. The validity of this assumption is not obvious, because when we are close, some of the mass is very close to us, and some is farther away,

and so on. When we add the effects all together, it seems a miracle that the net force is exactly the same as we would get if we put all the mass in the middle!

We now demonstrate the correctness of this miracle. In order to do so, however, we shall consider a thin uniform hollow shell instead of the whole earth. Let the total mass of the shell be m, and let us calculate the *potential energy* of a particle of mass m' a distance R away from the sphere (Fig. 13–6) and show that the potential energy is the same as it would be if the mass m were a point at the center. (The potential energy is easier to work with than is the field because we do not have to worry about angles, we merely add the potential energies of all the pieces of mass.) If we call x the distance of a certain plane section from the center, then all the mass that is in a slice dx is at the same distance r from P, and the potential energy due to this ring is $-Gm'\,dm/r$. How much mass is in the small slice dx? An amount

$$dm = 2\pi y\mu\,ds = \frac{2\pi y\mu\,dx}{\sin\theta} = \frac{2\pi y\mu\,dxa}{y} = 2\pi a\mu\,dx,$$

where $\mu = m/4\pi a^2$ is the surface density of mass on the spherical shell. (It is a general rule that the area of a zone of a sphere is proportional to its axial width.) Therefore the potential energy due to dm is

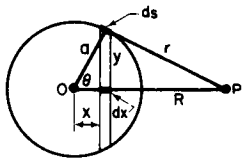


Fig. 13–6. A thin spherical shell of mass or charge.

$$dW = -\frac{Gm'\,dm}{r} = -\frac{Gm'2\pi a\mu\,dx}{r}.$$

But we see that

$$\begin{aligned} r^2 = y^2 + (R - x)^2 &= y^2 + x^2 + R^2 - 2Rx \\ &= a^2 + R^2 - 2Rx. \end{aligned}$$

Thus

$$2r\,dr = -2R\,dx$$

or

$$\frac{dx}{r} = \frac{dr}{R}.$$

Therefore,

$$dW = -\frac{Gm'2\pi a\mu\,dr}{R},$$

and so

$$\begin{aligned} W &= -\frac{Gm'2\pi a\mu}{R}\int_{R-a}^{R+a} dr \\ &= -\frac{Gm'2\pi a\mu}{R}\,2a = -\frac{Gm'(4\pi a^2\mu)}{R} \\ &= -\frac{Gm'm}{R}. \end{aligned} \tag{13.18}$$

Thus, for a thin spherical shell, the potential energy of a mass m', external to the shell, is the same as though the mass of the shell were concentrated at its center. The earth can be imagined as a series of spherical shells, each one of which contributes an energy which depends only on its mass and the distance from the center; adding them all together we get the *total mass*, and therefore the earth acts as though all the material were at the center!

But notice what happens if our point is on the *inside* of the shell. Making the same calculation, but with P on the inside, we still get the difference of the two r's, but now in the form $a + R - (a - R) = 2R$, or twice the distance from the center. In other words, W comes out to be $W = -Gm'm/a$, which is *independent of* R and independent of position, i.e., the same energy no matter *where* we are inside. Therefore no force; no work is done when we move about inside. If the potential energy is the same no matter where an object is placed inside the sphere, there can be no force on it. So there is no force inside, there is only a force outside, and the force outside is the same as though the mass were all at the center.

14

Work and Potential Energy (conclusion)

14–1 Work

14–2 Constrained motion

14–3 Conservative forces

14–4 Nonconservative forces

14–5 Potentials and fields

14–1 Work

In the preceding chapter we have presented a great many new ideas and results that play a central role in physics. These ideas are so important that it seems worth while to devote a whole chapter to a closer examination of them. In the present chapter we shall not repeat the "proofs" or the specific tricks by which the results were obtained, but shall concentrate instead upon a discussion of the ideas themselves.

In learning any subject of a technical nature where mathematics plays a role, one is confronted with the task of understanding and storing away in the memory a huge body of facts and ideas, held together by certain relationships which can be "proved" or "shown" to exist between them. It is easy to confuse the proof itself with the relationship which it establishes. Clearly, the important thing to learn and to remember is the relationship, not the proof. In any particular circumstance we can either say "it can be shown that" such and such is true, or we can show it. In almost all cases, the particular proof that is used is concocted, first of all, in such form that it can be written quickly and easily on the chalkboard or on paper, and so that it will be as smooth-looking as possible. Consequently, the proof may look deceptively simple, when in fact, the author might have worked for hours trying different ways of calculating the same thing until he has found the neatest way, so as to be able to show that it can be shown in the shortest amount of time! The thing to be remembered, when seeing a proof, is not the proof itself, but rather that it *can be shown* that such and such is true. Of course, if the proof involves some mathematical procedures or "tricks" that one has not seen before, attention should be given not to the trick exactly, but to the mathematical idea involved.

It is certain that in all the demonstrations that are made in a course such as this, not one has been remembered from the time when the author studied freshman physics. Quite the contrary: he merely remembers that such and such is true, and to explain how it can be shown he invents a demonstration at the moment it is needed. Anyone who has really learned a subject should be able to follow a similar procedure, but it is no use remembering the proofs. That is why, in this chapter, we shall avoid the proofs of the various statements made previously, and merely summarize the results.

The first idea that has to be digested is *work done by a force*. The physical word "work" is not the word in the ordinary sense of "Workers of the world unite!," but is a different idea. Physical work is expressed as $\int \mathbf{F} \cdot d\mathbf{s}$, called "the line integral of F dot ds," which means that if the force, for instance, is in one direction and the object on which the force is working is displaced in a certain direction, then *only the component of force in the direction of the displacement* does any work. If, for instance, the force were constant and the displacement were a finite distance $\Delta\mathbf{s}$, then the work done in moving the constant force through that distance is only the component of force along $\Delta\mathbf{s}$ times Δs. The rule is "force times distance," but we really mean only the component of force in the direction of the displacement times Δs or, equivalently, the component of displacement in the direction of force times $\mathbf{F}$. It is evident that no work whatsoever is done by a force which is at right angles to the displacement.

Now if the vector displacement $\Delta\mathbf{s}$ is resolved into components, in other words, if the actual displacement is $\Delta\mathbf{s}$ and we want to consider it effectively as a component of displacement Δx in the x-direction, Δy in the y-direction, and Δz

in the z-direction, then the work done in carrying an object from one place to another can be calculated in three parts, by calculating the work done along x, along y, and along z. The work done in going along x involves only that component of force, namely F_x, and so on, so the work is $F_x \Delta x + F_y \Delta y + F_z \Delta z$. When the force is not constant, and we have a complicated curved motion, then we must resolve the path into a lot of little $\Delta\mathbf{s}$'s, add the work done in carrying the object along each $\Delta\mathbf{s}$, and take the limit as $\Delta\mathbf{s}$ goes to zero. This is the meaning of the "line integral."

Everything we have just said is contained in the formula $W = \int \mathbf{F} \cdot d\mathbf{s}$. It is all very well to say that it is a marvelous formula, but it is another thing to understand what it means, or what some of the consequences are.

The word "work" in physics has a meaning so different from that of the word as it is used in ordinary circumstances that it must be observed carefully that there are some peculiar circumstances in which it appears not to be the same. For example, according to the physical definition of work, if one holds a hundred-pound weight off the ground for a while, he is doing no work. Nevertheless, everyone knows that he begins to sweat, shake, and breathe harder, as if he were running up a flight of stairs. Yet running upstairs *is* considered as doing work (in running *downstairs*, one gets work out of the world, according to physics), but in simply holding an object in a fixed position, no work is done. Clearly, the physical definition of work differs from the physiological definition, for reasons we shall briefly explore.

It is a fact that when one holds a weight he has to do "physiological" work. Why should he sweat? Why should he need to consume food to hold the weight up? Why is the machinery inside him operating at full throttle, just to hold the weight up? Actually, the weight could be held up with no effort by just placing it on a table; then the table, quietly and calmly, without any supply of energy, is able to maintain the same weight at the same height! The physiological situation is something like the following. There are two kinds of muscles in the human body and in other animals: one kind, called *striated* or *skeletal* muscle, is the type of muscle we have in our arms, for example, which is under voluntary control; the other kind, called *smooth* muscle, is like the muscle in the intestines or, in the clam, the greater adductor muscle that closes the shell. The smooth muscles work very slowly, but they can hold a "set"; that is to say, if the clam tries to close its shell in a certain position, it will hold that position, even if there is a very great force trying to change it. It will hold a position under load for hours and hours without getting tired because it is very much like a table holding up a weight, it "sets" into a certain position, and the molecules just lock there temporarily with no work being done, no effort being generated by the clam. The fact that we have to generate effort to hold up a weight is simply due to the design of striated muscle. What happens is that when a nerve impulse reaches a muscle fiber, the fiber gives a little twitch and then relaxes, so that when we hold something up, enormous volleys of nerve impulses are coming in to the muscle, large numbers of twitches are maintaining the weight, while the other fibers relax. We can see this, of course: when we hold a heavy weight and get tired, we begin to shake. The reason is that the volleys are coming irregularly, and the muscle is tired and not reacting fast enough. Why such an inefficient scheme? We do not know exactly why, but evolution has not been able to develop *fast* smooth muscle. Smooth muscle would be much more effective for holding up weights because you could just stand there and it would lock in; there would be no work involved and no energy would be required. However, it has the disadvantage that it is very slow-operating.

Returning now to physics, we may ask *why* we want to calculate the work done. The answer is that it is interesting and useful to do so, since the work done on a particle by the resultant of all the forces acting on it is exactly equal to the change in kinetic energy of that particle. That is, if an object is being pushed, it picks up speed, and

$$\Delta(v^2) = \frac{2}{m} \mathbf{F} \cdot \Delta\mathbf{s} \cdot$$

14–2 Constrained motion

Another interesting feature of forces and work is this: suppose that we have a sloping or a curved track, and a particle that must move along the track, but without friction. Or we may have a pendulum with a string and a weight; the string constrains the weight to move in a circle about the pivot point. The pivot point may be changed by having the string hit a peg, so that the path of the weight is along two circles of different radii. These are examples of what we call *fixed, frictionless constraints.*

In motion with a fixed frictionless constraint, no work is done by the constraint because the forces of constraint are always at right angles to the motion. By the "forces of constraint" we mean those forces which are applied to the object directly by the constraint itself—the contact force with the track, or the tension in the string.

The forces involved in the motion of a particle on a slope moving under the influence of gravity are quite complicated, since there is a constraint force, a gravitational force, and so on. However, if we base our calculation of the motion on conservation of energy and *the gravitational force alone*, we get the right result. This seems rather strange, because it is not strictly the right way to do it—we should use the *resultant* force. Nevertheless, the work done by the gravitational force alone will turn out to be the change in the kinetic energy, because the work done by the constraint part of the force is zero (Fig. 14–1).

The important feature here is that if a force can be analyzed as the sum of two or more "pieces" then the work done by the resultant force in going along a certain curve is the sum of the works done by the various "component" forces into which the force is analyzed. Thus if we analyze the force as being the vector sum of several effects, gravitational plus constraint forces, etc., or the x-component of all forces and the y-component of all forces, or any other way that we wish to split it up, then the work done by the net force is equal to the sum of the works done by all the parts into which we have divided the force in making the analysis.

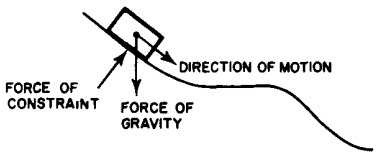


Fig. 14–1. Forces acting on a sliding body (no friction).

14–3 Conservative forces

In nature there are certain forces, that of gravity, for example, which have a very remarkable property which we call "conservative" (no political ideas involved, it is again one of those "crazy words"). If we calculate how much work is done by a force in moving an object from one point to another along some curved path, in general the work depends upon the curve, but in special cases it does not. If it does not depend upon the curve, we say that the force is a conservative force. In other words, if the integral of the force times the distance in going from position 1 to position 2 in Fig. 14–2 is calculated along curve A and then along B, we get the same number of joules, and if this is true for this pair of points on *every curve*, and if the same proposition works *no matter which pair of points* we use, then we say the force is conservative. In such circumstances, the work integral going from 1 to 2 can be evaluated in a simple manner, and we can give a formula for the result. Ordinarily it is not this easy, because we also have to specify the curve, but when we have a case where the work does not depend on the curve, then, of course, the work depends only upon the *positions* of 1 and 2.

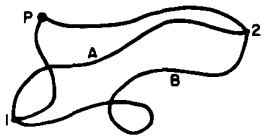


Fig. 14–2. Possible paths between two points in a field of force.

To demonstrate this idea, consider the following. We take a "standard" point P, at an arbitrary location (Fig. 14–2). Then, the work line-integral from 1 to 2, which we want to calculate, can be evaluated as the work done in going from 1 to P plus the work done in going from P to 2, because the forces are conservative and the work does not depend upon the curve. Now, the work done in going from position P to a particular position in space is a function of that position in space. Of course it really depends on P also, but we hold the arbitrary point P fixed permanently for the analysis. If that is done, then the work done in going from point P to point 2 is some function of the final position of 2. It depends upon where 2 is; if we go to some other point we get a different answer.

We shall call this function of position $-U(x, y, z)$, and when we wish to refer to some particular point 2 whose coordinates are (x_2, y_2, z_2), we shall write $U(2)$,

as an abbreviation for $U(x_2, y_2, z_2)$. The work done in going from point 1 to point P can be written also by going the *other way* along the integral, reversing all the $d\mathbf{s}$'s. That is, the work done in going from 1 to P is *minus* the work done in going from the point P to 1:

$$\int_1^P \mathbf{F} \cdot d\mathbf{s} = \int_P^1 \mathbf{F} \cdot (-d\mathbf{s}) = -\int_P^1 \mathbf{F} \cdot d\mathbf{s}.$$

Thus the work done in going from P to 1 is $-U(1)$, and from P to 2 the work is $-U(2)$. Therefore the integral from 1 to 2 is equal to $-U(2)$ plus $[-U(1)$ backwards], or $+U(1) - U(2)$:

$$U(1) = -\int_P^1 \mathbf{F} \cdot d\mathbf{s}, \qquad U(2) = -\int_P^2 \mathbf{F} \cdot d\mathbf{s},$$

$$\int_1^2 \mathbf{F} \cdot d\mathbf{s} = U(1) - U(2). \tag{14.1}$$

The quantity $U(1) - U(2)$ is called the change in the potential energy, and we call U the potential energy. We shall say that when the object is located at position 2, it has potential energy $U(2)$ and at position 1 it has potential energy $U(1)$. If it is located at position P, it has zero potential energy. If we had used any other point, say Q, instead of P, it would turn out (and we shall leave it to you to demonstrate) that the *potential energy is changed only by the addition of a constant.* Since the conservation of energy depends only upon *changes*, it does not matter if we add a constant to the potential energy. Thus the point P is arbitrary.

Now, we have the following two propositions: (1) that the work done by a force is equal to the change in kinetic energy of the particle, but (2) mathematically, for a conservative force, the work done is minus the change in a function U which we call the potential energy. As a consequence of these two, we arrive at the proposition that *if only conservative forces act, the kinetic energy T plus the potential energy U remains constant:*

$$T + U = \text{constant}. \tag{14.2}$$

Let us now discuss the formulas for the potential energy for a number of cases. If we have a gravitational field that is uniform, if we are not going to heights comparable with the radius of the earth, then the force is a constant vertical force and the work done is simply the force times the vertical distance. Thus

$$U(z) = mgz, \tag{14.3}$$

and the point P which corresponds to zero potential energy happens to be any point in the plane $z = 0$. We could also have said that the potential energy is $mg(z - 6)$ if we had wanted to—all the results would, of course, be the same in our analysis except that the value of the potential energy at $z = 0$ would be $-mg6$. It makes no difference, because only *differences* in potential energy count.

The energy needed to compress a linear spring a distance x from an equilibrium point is

$$U(x) = \tfrac{1}{2}kx^2, \tag{14.4}$$

and the zero of potential energy is at the point $x = 0$, the equilibrium position of the spring. Again we could add any constant we wish.

The potential energy of gravitation for point masses M and m, a distance r apart, is

$$U(r) = -GMm/r. \tag{14.5}$$

The constant has been chosen here so that the potential is zero at infinity. Of course the same formula applies to electrical charges, because it is the same law:

$$U(r) = q_1q_2/4\pi\epsilon_0 r. \tag{14.6}$$

Now let us actually use one of these formulas, to see whether we understand what it means. *Question*: How fast do we have to shoot a rocket away from the

earth in order for it to leave? *Solution*: The kinetic plus potential energy must be a constant; when it "leaves," it will be millions of miles away, and if it is just barely able to leave, we may suppose that it is moving with zero speed out there, just barely going. Let a be the radius of the earth, and M its mass. The kinetic plus potential energy is then initially given by $\frac{1}{2}mv^2 - GmM/a$. At the end of the motion the two energies must be equal. The kinetic energy is taken to be zero at the end of the motion, because it is supposed to be just barely drifting away at essentially zero speed, and the potential energy is GmM divided by infinity, which is zero. So everything is zero on one side and that tells us that the square of the velocity must be $2GM/a$. But GM/a^2 is what we call the acceleration of gravity, g. Thus

$$v^2 = 2ga.$$

At what speed must a satellite travel in order to keep going around the earth? We worked this out long ago and found that $v^2 = GM/a$. Therefore to go *away* from the earth, we need $\sqrt{2}$ times the velocity we need to just go *around* the earth near its surface. We need, in other words, *twice as much energy* (because energy goes as the square of the velocity) to leave the earth as we do to go around it. Therefore the first thing that was done historically with satellites was to get one to go around the earth, which requires a speed of five miles per second. The next thing was to send a satellite away from the earth permanently; this required twice the energy, or about seven miles per second.

Now, continuing our discussion of the characteristics of potential energy, let us consider the interaction of two molecules, or two atoms, two oxygen atoms for instance. When they are very far apart, the force is one of attraction, which varies as the inverse seventh power of the distance, and when they are very close the force is a very large repulsion. If we integrate the inverse seventh power to find the work done, we find that the potential energy U, which is a function of the radial distance between the two oxygen atoms, varies as the inverse sixth power of the distance for large distances.

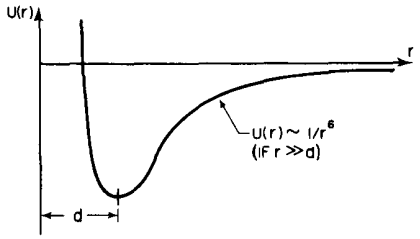


Fig. 14–3. The potential energy between two atoms as a function of the distance between them.

If we sketch the curve of the potential energy $U(r)$ as in Fig. 14–3, we thus start out at large r with an inverse sixth power, but if we come in sufficiently near we reach a point d where there is a minimum of potential energy. The minimum of potential energy at $r = d$ means this: if we start at d and move a small distance, a very small distance, the work done, which is the change in potential energy when we move this distance, is nearly zero, because there is very little change in potential energy at the bottom of the curve. Thus there is no force at this point, and so it is the equilibrium point. Another way to see that it is the equilibrium point is that it takes work to move away from d in either direction. When the two oxygen atoms have settled down, so that no more energy can be liberated from the force between them, they are in the lowest energy state, and they will be at this separation d. This is the way an oxygen molecule looks when it is cold. When we heat it up, the atoms shake and move farther apart, and we can in fact break them apart, but to do so takes a certain amount of work or energy, which is the potential energy difference between $r = d$ and $r = \infty$. When we try to push the atoms very close together the energy goes up very rapidly, because they repel each other.

The reason we bring this out is that the idea of *force* is not particularly suitable for quantum mechanics; there the idea of *energy* is most natural. We find that although forces and velocities "dissolve" and disappear when we consider the more advanced forces between nuclear matter and between molecules and so on, the energy concept remains. Therefore we find curves of potential energy in quantum mechanics books, but very rarely do we ever see a curve for the force between two molecules, because by that time people who are doing analyses are thinking in terms of energy rather than of force.

Next we note that if several conservative forces are acting on an object at the same time, then the potential energy of the object is the sum of the potential energies from each of the separate forces. This is the same proposition that we mentioned before, because if the force can be represented as a vector sum of forces, then the work done by the total force is the sum of the works done by the partial

forces, and it can therefore be analyzed as changes in the potential energies of each of them separately. Thus the total potential energy is the sum of all the little pieces.

We could generalize this to the case of a system of many objects interacting with one another, like Jupiter, Saturn, Uranus, etc., or oxygen, nitrogen, carbon, etc., which are acting with respect to one another in pairs due to forces all of which are conservative. In these circumstances the kinetic energy in the entire system is simply the sum of the kinetic energies of all of the particular atoms or planets or whatever, and the potential energy of the system is the sum, over the pairs of particles, of the potential energy of mutual interaction of a single pair, as though the others were not there. (This is really not true for molecular forces, and the formula is somewhat more complicated; it certainly is true for Newtonian gravitation, and it is true as an approximation for molecular forces. For molecular forces there is a potential energy, but it is sometimes a more complicated function of the positions of the atoms than simply a sum of terms from pairs.) In the special case of gravity, therefore, the potential energy is the sum, over all the pairs i and j, of $-Gm_im_j/r_{ij}$, as was indicated in Eq. (13.14). Equation (13.14) expressed mathematically the following proposition: that the total kinetic energy plus the total potential energy does not change with time. As the various planets wheel about, and turn and twist and so on, if we calculate the total kinetic energy and the total potential energy we find that the total remains constant.

14–4 Nonconservative forces

We have spent a considerable time discussing conservative forces; what about nonconservative forces? We shall take a deeper view of this than is usual, and state that there are no nonconservative forces! As a matter of fact, all the fundamental forces in nature appear to be conservative. This is not a consequence of Newton's laws. In fact, so far as Newton himself knew, the forces could be nonconservative, as friction apparently is. When we say friction *apparently* is, we are taking a modern view, in which it has been discovered that all the deep forces, the forces between the particles at the most fundamental level, are conservative.

If, for example, we analyze a system like that great globular star cluster that we saw a picture of, with the thousands of stars all interacting, then the formula for the total potential energy is simply one term plus another term, etc., summed over all pairs of stars, and the kinetic energy is the sum of the kinetic energies of all the individual stars. But the globular cluster as a whole is drifting in space too, and, if we were far enough away from it and did not see the details, could be thought of as a single object. Then if forces were applied to it, some of those forces might end up driving it forward as a whole, and we would see the center of the whole thing moving. On the other hand, some of the forces can be, so to speak, "wasted" in increasing the kinetic or potential energy of the "particles" inside. Let us suppose, for instance, that the action of these forces expands the whole cluster and makes the particles move faster. The total energy of the whole thing is really conserved, but seen from the outside with our crude eyes which cannot see the confusion of motions inside, and just thinking of the kinetic energy of the motion of the whole object as though it were a single particle, it would appear that energy is not conserved, but this is due to a lack of appreciation of what it is that we see. And that, it turns out, is the case: the total energy of the world, kinetic plus potential, is a constant when we look closely enough.

When we study matter in the finest detail at the atomic level, it is not always *easy* to separate the total energy of a thing into two parts, kinetic energy and potential energy, and such separation is not always necessary. It is *almost* always possible to do it, so let us say that it *is* always possible, and that the potential-plus-kinetic energy of the world is constant. Thus the total potential-plus-kinetic energy inside the whole world is constant, and if the "world" is a piece of isolated material, the energy is constant if there are no external forces. But as we have seen, some of the kinetic and potential energy of a thing may be internal, for instance the internal molecular motions, in the sense that we do not notice it. We know that in a glass of water everything is jiggling around, all the parts are moving

all the time, so there is a certain kinetic energy inside, which we ordinarily may not pay any attention to. We do not notice the motion of the atoms, which produces heat, and so we do not call it kinetic energy, but heat is primarily kinetic energy. Internal potential energy may also be in the form, for instance, of chemical energy: when we burn gasoline energy is liberated because the potential energies of the atoms in the new atomic arrangement are lower than in the old arrangement. It is not strictly possible to treat heat as being pure kinetic energy, for a little of the potential gets in, and vice versa for chemical energy, so we put the two together and say that the total kinetic and potential energy inside an object is partly heat, partly chemical energy, and so on. Anyway, all these different forms of internal energy are sometimes considered as "lost" energy in the sense described above; this will be made clearer when we study thermodynamics.

As another example, when friction is present it is not true that kinetic energy is lost, even though a sliding object stops and the kinetic energy seems to be lost. The kinetic energy is not lost because, of course, the atoms inside are jiggling with a greater amount of kinetic energy than before, and although we cannot see that, we can measure it by determining the temperature. Of course if we disregard the heat energy, then the conservation of energy theorem will appear to be false.

Another situation in which energy conservation appears to be false is when we study only part of a system. Naturally, the conservation of energy theorem will appear not to be true if something is interacting with something else on the outside and we neglect to take that interaction into account.

In classical physics potential energy involved only gravitation and electricity, but now we have nuclear energy and other energies also. Light, for example, would involve a new form of energy in the classical theory, but we can also, if we want to, imagine that the energy of light is the kinetic energy of a photon, and then our formula (14.2) would still be right.

14–5 Potentials and fields

We shall now discuss a few of the ideas associated with potential energy and with the idea of a *field*. Suppose we have two large objects A and B and a third very small one which is attracted gravitationally by the two, with some resultant force $\mathbf{F}$. We have already noted in Chapter 12 that the gravitational force on a particle can be written as its mass, m, times another vector, $\mathbf{C}$, which is dependent only upon the *position* of the particle:

$$\mathbf{F} = m\mathbf{C}.$$

We can analyze gravitation, then, by imagining that there is a certain vector $\mathbf{C}$ at every position in space which "acts" upon a mass which we may place there, but which is there itself whether we actually supply a mass for it to "act" on or not. $\mathbf{C}$ has three components, and each of those components is a function of (x, y, z), a function of position in space. Such a thing we call a *field*, and we say that the objects A and B *generate* the field, i.e., they "make" the vector $\mathbf{C}$. When an object is put in a field, the force on it is equal to its mass times the value of the field vector at the point where the object is put.

We can also do the same with the potential energy. Since the potential energy, the integral of **(force)** $\cdot$ $(d\mathbf{s})$ can be written as m times the integral of the **(field)** $\cdot$ $(d\mathbf{s})$, a mere change of scale, we see that the potential energy $U(x, y, z)$ of an object located at a point $(x, y, z,)$ in space can be written as m times another function which we may call the *potential* Ψ. The integral $\int \mathbf{C} \cdot d\mathbf{s} = -\Psi$, just as $\int \mathbf{F} \cdot d\mathbf{s} = -U$; there is only a scale factor between the two:

$$U = -\int \mathbf{F} \cdot d\mathbf{s} = -m\int \mathbf{C} \cdot d\mathbf{s} = m\Psi. \tag{14.7}$$

By having this function $\Psi(x, y, z)$ at every point in space, we can immediately calculate the potential energy of an object at any point in space, namely, $U(x, y, z) = m\Psi(x, y, z)$—rather a trivial business, it seems. But it is not really trivial, because it is sometimes much nicer to describe the field by giving the value

of Ψ everywhere in space instead of having to give **C**. Instead of having to write three complicated components of a vector function, we can give instead the scalar function Ψ. Furthermore, it is much easier to calculate Ψ than any given component of **C** when the field is produced by a number of masses, for since the potential is a scalar we merely add, without worrying about direction. Also, the field **C** can be recovered easily from Ψ, as we shall shortly see. Suppose we have point masses $m_1, m_2, \ldots$ at the points $1, 2, \ldots$ and we wish to know the potential Ψ at some arbitrary point p. This is simply the sum of the potentials at P due to the individual masses taken one by one:

$$\Psi(p) = \sum_i -\frac{Gm_i}{r_{ip}}, \qquad i = 1, 2, \ldots \tag{14.8}$$

In the last chapter we used this formula, that the potential is the sum of the potentials from all the different objects, to calculate the potential due to a spherical shell of matter by adding the contributions to the potential at a point from all parts of the shell. The result of this calculation is shown graphically in Fig. 14–4. It is negative, having the value zero at $r = \infty$ and varying as $1/r$ down to the radius a, and then is constant inside the shell. Outside the shell the potential is $-Gm/r$, where m is the mass of the shell, which is exactly the same as it would have been if all the mass were located at the center. But it is not *everywhere* exactly the same, for inside the shell the potential turns out to be $-Gm/a$, and is a constant! *When the potential is constant, there is no field,* or when the potential energy is constant there is no force, because if we move an object from one place to another anywhere inside the sphere the work done by the force is exactly zero. Why? Because the work done in moving the object from one place to the other is equal to minus the change in the potential energy (or, the corresponding field integral is the change of the potential). But the potential energy is the *same* at any two points inside, so there is zero change in potential energy, and therefore no work is done in going between any two points inside the shell. The only way the work can be zero for all directions of displacement is that there is no force at all.

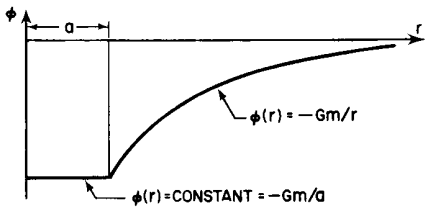


Fig. 14–4. Potential due to a spherical shell of radius a.

This gives us a clue as to how we can obtain the force or the field, given the potential energy. Let us suppose that the potential energy of an object is known at the position (x, y, z) and we want to know what the force on the object is. It will not do to know the potential at only this one point, as we shall see; it requires knowledge of the potential at neighboring points as well. Why? How can we calculate the x-component of the force? (If we can do this, of course, we can also find the y- and z-components, and we will then know the whole force.) Now, if we were to move the object a small distance Δx, the work done by the force on the object would be the x-component of the force times Δx, if Δx is sufficiently small, and this should equal the change in potential energy in going from one point to the other:

$$\Delta W = -\Delta U = F_x \, \Delta x. \tag{14.9}$$

We have merely used the formula $\int \mathbf{F} \cdot d\mathbf{s} = -\Delta U$, but for a *very short* path. Now we divide by Δx and so find that the force is

$$F_x = -\Delta U/\Delta x. \tag{14.10}$$

Of course this is not exact. What we really want is the limit of (14.10) as Δx gets smaller and smaller, because it is only *exactly* right in the limit of infinitesimal Δx. This we recognize as the derivative of U with respect to x, and we would be inclined, therefore, to write $-dU/dx$. But U depends on x, y, and z, and the mathematicians have invented a different symbol to remind us to be very careful when we are differentiating such a function, so as to remember that we are considering that *only* x *varies*, and y and z do not vary. Instead of a d they simply make a "backwards 6," or ∂. (A ∂ should have been used in the beginning of calculus because we always want to cancel that d, but we never want to cancel a ∂!) So they write $\partial U/\partial x$, and furthermore, in moments of duress, if they want to be *very* careful, they put a line beside it with a little yz at the bottom $(\partial U/\partial x|_{yz})$,

which means "Take the derivative of U with respect to x, keeping y and z constant." Most often we leave out the remark about what is kept constant because it is usually evident from the context, so we usually do not use the line with the y and z. However, *always* use a ∂ instead of a d as a warning that it is a derivative with some other variables kept constant. This is called a *partial derivative*; it is a derivative in which we vary only x.

Therefore, we find that the force in the x-direction is minus the partial derivative of U with respect to x:

$$F_x = -\partial U/\partial x. \tag{14.11}$$

In a similar way, the force in the y-direction can be found by differentiating U with respect to y, keeping x and z constant, and the third component, of course, is the derivative with respect to z, keeping y and x constant:

$$F_y = -\partial U/\partial y, \qquad F_z = -\partial U/\partial z. \tag{14.12}$$

This is the way to get from the potential energy to the force. We get the *field* from the *potential* in exactly the same way:

$$C_x = -\partial\Psi/\partial x, \qquad C_y = -\partial\Psi/\partial y, \qquad C_z = -\partial\Psi/\partial z. \tag{14.13}$$

Incidentally, we shall mention here another notation, which we shall not actually use for quite a while: Since $\mathbf{C}$ is a vector and has x-, y-, and z-components, the symbolized $\partial/\partial x$, $\partial/\partial y$, and $\partial/\partial z$ which produce the x-, y-, and z-components are something like vectors. The mathematicians have invented a glorious new symbol, ∇, called "grad" or "gradient" which is not a quantity but an *operator* which makes a vector from a scalar. It has the following "components": The x-component of this "grad" is $\partial/\partial x$, the y-component is $\partial/\partial y$, and the z-component is $\partial/\partial z$, and then we have the fun of writing our formulas this way:

$$\mathbf{F} = -\nabla U, \qquad \mathbf{C} = -\nabla\Psi. \tag{14.14}$$

Using ∇ gives us a quick way of testing whether we have a real vector equation or not, but actually Eq. (14.14) means precisely the same as Eqs. (14.11) and (14.12); it is just another way of writing them, and since we do not want to write three equations every time, we just write ∇U instead.

One more example of fields and potentials has to do with the electrical case. In the case of electricity the force on a stationary object is the charge times the electric field: $\mathbf{F} = q\mathbf{E}$. (In general, of course, the x-component of force in an electrical problem has also a part which depends on the magnetic field. It is easy to show from Eq. (12.10) that the force on a particle due to magnetic fields is always at right angles to its velocity, and also at right angles to the field. Since the force due to magnetism on a moving charge is at right angles to the velocity, *no work is done* by the magnetism on the moving charge because the motion is at right angles to the force. Therefore, in calculating theorems of kinetic energy in electric and magnetic fields we can disregard the contribution from the magnetic field, since it does not change the kinetic energy.) We suppose that there is only an electric field. Then we can calculate the energy, or work done, in the same way as for gravity, and calculate a quantity ϕ which is minus the integral of $\mathbf{E}\cdot d\mathbf{s}$, from the arbitrary fixed point to the point where we make the calculation, and then the potential energy in an electric field is just charge times this quantity ϕ:

$$\phi(\mathbf{r}) = \int \mathbf{E}\cdot d\mathbf{s},$$
$$U = q\phi.$$

Let us take, as an example, the case of two parallel metal plates, each with a surface charge of $\pm\sigma$ per unit area. This is called a parallel-plate capacitor. We found previously that there is zero force outside the plates and that there is a constant electric field between them, directed from $+$ to $-$ and of magnitude σ/ϵ_0 (Fig. 14–5). We would like to know how much work would be done in

carrying a charge from one plate to the other. The work would be the **(force)** · (*d***s**) integral, which can be written as charge times the potential value at plate 1 minus that at plate 2:

$$W = \int_1^2 \mathbf{F} \cdot d\mathbf{s} = q(\phi_1 - \phi_2).$$

We can actually work out the integral because the force is constant, and if we call the separation of the plates d, then the integral is easy:

$$\int_1^2 \mathbf{F} \cdot d\mathbf{s} = \frac{q\sigma}{\epsilon_0} \int_1^2 dx = \frac{q\sigma d}{\epsilon_0}.$$

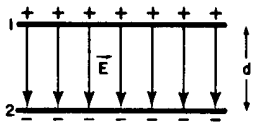


Fig. 14–5. Field between parallel plates.

The difference in potential, $\Delta\phi = \sigma d/\epsilon_0$, is called the *voltage difference*, and ϕ is measured in volts. When we say a pair of plates is charged to a certain voltage, what we mean is that the difference in electrical potential of the two plates is so-and-so many volts. For a capacitor made of two parallel plates carrying a surface charge $\pm\sigma$, the voltage, or difference in potential, of the pair of plates is $\sigma d/\epsilon_0$.

15

The Special Theory of Relativity

15–1 The principle of relativity

15–2 The Lorentz transformation

15–3 The Michelson-Morley experiment

15–4 Transformation of time

15–5 The Lorentz contraction

15–6 Simultaneity

15–7 Four-vectors

15–8 Relativistic dynamics

15–9 Equivalence of mass and energy

15–1 The principle of relativity

For over 200 years the equations of motion enunciated by Newton were believed to describe nature correctly, and the first time that an error in these laws was discovered, the way to correct it was also discovered. Both the error and its correction were discovered by Einstein in 1905.

Newton's Second Law, which we have expressed by the equation

$$F = d(mv)/dt,$$

was stated with the tacit assumption that m is a constant, but we now know that this is not true, and that the mass of a body increases with velocity. In Einstein's corrected formula m has the value

$$m = \frac{m_0}{\sqrt{1 - v^2/c^2}}, \tag{15.1}$$

where the "rest mass" m_0 represents the mass of a body that is not moving and c is the speed of light, which is about 3×10^5 km · sec^{-1} or about 186,000 mi · sec^{-1}.

For those who want to learn just enough about it so they can solve problems, that is all there is to the theory of relativity—it just changes Newton's laws by introducing a correction factor to the mass. From the formula itself it is easy to see that this mass increase is very small in ordinary circumstances. If the velocity is even as great as that of a satellite, which goes around the earth at 5 mi/sec, then $v/c = 5/186{,}000$: putting this value into the formula shows that the correction to the mass is only one part in two to three billion, which is nearly impossible to observe. Actually, the correctness of the formula has been amply confirmed by the observation of many kinds of particles, moving at speeds ranging up to practically the speed of light. However, because the effect is ordinarily so small, it seems remarkable that it was discovered theoretically before it was discovered experimentally. Empirically, at a sufficiently high velocity, the effect is very large, but it was not discovered that way. Therefore it is interesting to see how a law that involved so delicate a modification (at the time when it was first discovered) was brought to light by a combination of experiments and physical reasoning. Contributions to the discovery were made by a number of people, the final result of whose work was Einstein's discovery.

There are really two Einstein theories of relativity. This chapter is concerned with the Special Theory of Relativity, which dates from 1905. In 1915 Einstein published an additional theory, called the General Theory of Relativity. This latter theory deals with the extension of the Special Theory to the case of the law of gravitation; we shall not discuss the General Theory here.

The principle of relativity was first stated by Newton, in one of his corollaries to the laws of motion: "The motions of bodies included in a given space are the same among themselves, whether that space is at rest or moves uniformly forward in a straight line." This means, for example, that if a space ship is drifting along at a uniform speed, all experiments performed in the space ship and all the phenomena in the space ship will appear the same as if the ship were not moving, provided, of course, that one does not look outside. That is the meaning of the principle of relativity. This is a simple enough idea, and the only question is whether it is *true* that in all experiments performed inside a moving system the laws of physics

will appear the same as they would if the system were standing still. Let us first investigate whether Newton's laws appear the same in the moving system.

Suppose that Moe is moving in the x-direction with a uniform velocity u, and he measures the position of a certain point, shown in Fig. 15–1. He designates the "x-distance" of the point in his coordinate system as x'. Joe is at rest, and measures the position of the same point, designating its x-coordinate in his system as x. The relationship of the coordinates in the two systems is clear from the diagram. After time t Moe's origin has moved a distance ut, and if the two systems originally coincided,

$$\begin{aligned} x' &= x - ut, \\ y' &= y, \\ z' &= z, \\ t' &= t. \end{aligned} \tag{15.2}$$

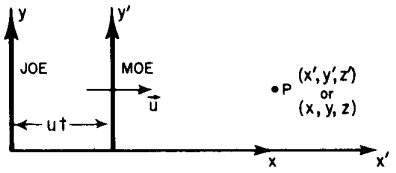


Fig. 15–1. Two coordinate systems in uniform relative motion along their x-axes.

If we substitute this transformation of coordinates into Newton's laws we find that these laws transform to the same laws in the primed system; that is, the laws of Newton are of the same form in a moving system as in a stationary system, and therefore it is impossible to tell, by making mechanical experiments, whether the system is moving or not.

The principle of relativity has been used in mechanics for a long time. It was employed by various people, in particular Huygens, to obtain the rules for the collision of billiard balls, in much the same way as we used it in Chapter 10 to discuss the conservation of momentum. In the past century interest in it was heightened as the result of investigations into the phenomena of electricity, magnetism, and light. A long series of careful studies of these phenomena by many people culminated in Maxwell's equations of the electromagnetic field, which describe electricity, magnetism, and light in one uniform system. However, the Maxwell equations did *not* seem to obey the principle of relativity. That is, if we transform Maxwell's equations by the substitution of equations 15.2, *their form does not remain the same;* therefore, in a moving space ship the electrical and optical phenomena should be different from those in a stationary ship. Thus one could use these optical phenomena to determine the speed of the ship; in particular, one could determine the absolute speed of the ship by making suitable optical or electrical measurements. One of the consequences of Maxwell's equations is that if there is a disturbance in the field such that light is generated, these electromagnetic waves go out in all directions equally and at the same speed c, or 186,000 mi/sec. Another consequence of the equations is that if the source of the disturbance is moving, the light emitted goes through space at the same speed c. This is analogous to the case of sound, the speed of sound waves being likewise independent of the motion of the source.

This independence of the motion of the source, in the case of light, brings up an interesting problem:

Suppose we are riding in a car that is going at a speed u, and light from the rear is going past the car with speed c. Differentiating the first equation in (15.2) gives

$$dx'/dt = dx/dt - u,$$

which means that according to the Galilean transformation the apparent speed of the passing light, as we measure it in the car, should not be c but should be $c - u$. For instance, if the car is going 100,000 mi/sec, and the light is going 186,000 mi/sec, then apparently the light going past the car should go 86,000 mi/sec. In any case, by measuring the speed of the light going past the car (if the Galilean transformation is correct for light), one could determine the speed of the car. A number of experiments based on this general idea were performed to determine the velocity of the earth, but they all failed—they gave *no velocity at all.* We shall discuss one of these experiments in detail, to show exactly what was done and what was the matter; something *was* the matter, of course, something was wrong with the equations of physics. What could it be?

15–2 The Lorentz transformation

When the failure of the equations of physics in the above case came to light, the first thought that occurred was that the trouble must lie in the new Maxwell equations of electrodynamics, which were only 20 years old at the time. It seemed almost obvious that these equations must be wrong, so the thing to do was to change them in such a way that under the Galilean transformation the principle of relativity would be satisfied. When this was tried, the new terms that had to be put into the equations led to predictions of new electrical phenomena that did not exist at all when tested experimentally, so this attempt had to be abandoned. Then it gradually became apparent that Maxwell's laws of electrodynamics were correct, and the trouble must be sought elsewhere.

In the meantime, H. A. Lorentz noticed a remarkable and curious thing when he made the following substitutions in the Maxwell equations:

$$\begin{aligned} x' &= \frac{x - ut}{\sqrt{1 - u^2/c^2}}, \\ y' &= y, \\ z' &= z, \\ t' &= \frac{t - ux/c^2}{\sqrt{1 - u^2/c^2}}, \end{aligned} \tag{15.3}$$

namely, Maxwell's equations remain in the same form when this transformation is applied to them! Equations (15.3) are known as a *Lorentz transformation.* Einstein, following a suggestion originally made by Poincaré, then proposed that *all the physical laws* should be of such a kind that they *remain unchanged under a Lorentz transformation.* In other words, we should change, not the laws of electrodynamics, but the laws of mechanics. How shall we change Newton's laws so that *they* will remain unchanged by the Lorentz transformation? If this goal is set, we then have to rewrite Newton's equations in such a way that the conditions we have imposed are satisfied. As it turned out, the only requirement is that the mass m in Newton's equations must be replaced by the form shown in Eq. (15.1). When this change is made, Newton's laws and the laws of electrodynamics will harmonize. Then if we use the Lorentz transformation in comparing Moe's measurements with Joe's, we shall never be able to detect whether either is moving, because the form of all the equations will be the same in both coordinate systems!

It is interesting to discuss what it means that we replace the old transformation between the coordinates and time with a new one, because the old one (Galilean) seems to be self-evident, and the new one (Lorentz) looks peculiar. We wish to know whether it is logically and experimentally possible that the new, and not the old, transformation can be correct. To find that out, it is not enough to study the laws of mechanics but, as Einstein did, we too must analyze our ideas of *space* and *time* in order to understand this transformation. We shall have to discuss these ideas and their implications for mechanics at some length, so we say in advance that the effort will be justified, since the results agree with experiment.

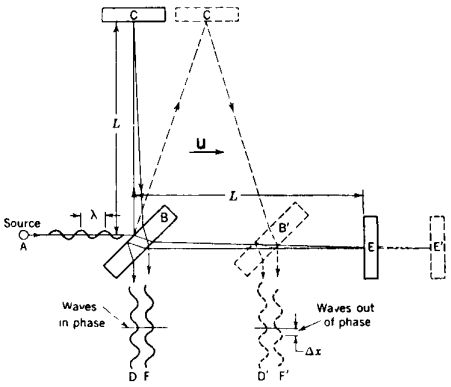


Fig. 15–2. Schematic diagram of the Michelson-Morley experiment.

15–3 The Michelson-Morley experiment

As mentioned above, attempts were made to determine the absolute velocity of the earth through the hypothetical "ether" that was supposed to pervade all space. The most famous of these experiments is one performed by Michelson and Morley in 1887. It was 18 years later before the negative results of the experiment were finally explained, by Einstein.

The Michelson-Morley experiment was performed with an apparatus like that shown schematically in Fig. 15–2. This apparatus is essentially comprised of a light source A, a partially silvered glass plate B, and two mirrors C and E, all mounted on a rigid base. The mirrors are placed at equal distances L from B. The plate B splits an oncoming beam of light, and the two resulting beams con-

tinue in mutually perpendicular directions to the mirrors, where they are reflected back to B. On arriving back at B, the two beams are recombined as two superposed beams, D and F. If the time taken for the light to go from B to E and back is the same as the time from B to C and back, the emerging beams D and F will be in phase and will reinforce each other, but if the two times differ slightly, the beams will be slightly out of phase and interference will result. If the apparatus is "at rest" in the ether, the times should be precisely equal, but if it is moving toward the right with a velocity u, there should be a difference in the times. Let us see why.

First, let us calculate the time required for the light to go from B to E and back. Let us say that the time for light to go from plate B to mirror E is t_1, and the time for the return is t_2. Now, while the light is on its way from B to the mirror, the apparatus moves a distance ut_1, so the light must traverse a distance $L + ut_1$, at the speed c. We can also express this distance as ct_1, so we have

$$ct_1 = L + ut_1, \qquad \text{or} \qquad t_1 = L/(c - u).$$

(This result is also obvious from the point of view that the velocity of light relative to the apparatus is $c - u$, so the time is the length L divided by $c - u$.) In a like manner, the time t_2 can be calculated. During this time the plate B advances a distance ut_2, so the return distance of the light is $L - ut_2$. Then we have

$$ct_2 = L - ut_2, \qquad \text{or} \qquad t_2 = L/(c + u).$$

Then the total time is

$$t_1 + t_2 = 2Lc/(c^2 - u^2).$$

For convenience in later comparison of times we write this as

$$t_1 + t_2 = \frac{2L/c}{1 - u^2/c^2}. \tag{15.4}$$

Our second calculation will be of the time t_3 for the light to go from B to the mirror C. As before, during time t_3 the mirror C moves to the right a distance ut_3 to the position C'; in the same time, the light travels a distance ct_3 along the hypotenuse of a triangle, which is BC'. For this right triangle we have

$$(ct_3)^2 = L^2 + (ut_3)^2$$

or

$$L^2 = c^2t_3^2 - u^2t_3^2 = (c^2 - u^2)t_3^2,$$

from which we get

$$t_3 = L/\sqrt{c^2 - u^2}.$$

For the return trip from C' the distance is the same, as can be seen from the symmetry of the figure; therefore the return time is also the same, and the total time is $2t_3$. With a little rearrangement of the form we can write

$$2t_3 = \frac{2L}{\sqrt{c^2 - u^2}} = \frac{2L/c}{\sqrt{1 - u^2/c^2}}. \tag{15.5}$$

We are now able to compare the times taken by the two beams of light. In expressions (15.4) and (15.5) the numerators are identical, and represent the time that would be taken if the apparatus were at rest. In the denominators, the term u^2/c^2 will be small, unless u is comparable in size to c. The denominators represent the modifications in the times caused by the motion of the apparatus. And behold, these modifications are *not the same*—the time to go to C and back is a little less than the time to E and back, even though the mirrors are equidistant from B, and all we have to do is to measure that difference with precision.

Here a minor technical point arises—suppose the two lengths L are not exactly equal? In fact, we surely cannot make them exactly equal. In that case we simply turn the apparatus 90 degrees, so that BC is in the line of motion and BE is perpendicular to the motion. Any small difference in length then becomes

unimportant, and what we look for is a *shift* in the interference fringes when we rotate the apparatus.

In carrying out the experiment, Michelson and Morley oriented the apparatus so that the line BE was nearly parallel to the earth's motion in its orbit (at certain times of the day and night). This orbital speed is about 18 miles per second, and any "ether drift" should be at least that much at some time of the day or night and at some time during the year. The apparatus was amply sensitive to observe such an effect, but no time difference was found—the velocity of the earth through the ether could not be detected. The result of the experiment was null.

The result of the Michelson-Morley experiment was very puzzling and most disturbing. The first fruitful idea for finding a way out of the impasse came from Lorentz. He suggested that material bodies contract when they are moving, and that this foreshortening is only in the direction of the motion, and also, that if the length is L_0 when a body is at rest, then when it moves with speed u parallel to its length, the new length, which we call $L_{||}$ (L-parallel), is given by

$$L_{||} = L_0 \sqrt{1 - u^2/c^2}. \tag{15.6}$$

When this modification is applied to the Michelson-Morley interferometer apparatus the distance from B to C does not change, but the distance from B to E is shortened to $L\sqrt{1 - u^2/c^2}$. Therefore Eq. (15.5) is not changed, but the L of Eq. (15.4) must be changed in accordance with Eq. (15.6). When this is done we obtain

$$t_1 + t_2 = \frac{(2L/c)\sqrt{1 - u^2/c^2}}{1 - u^2/c^2} = \frac{2L/c}{\sqrt{1 - u^2/c^2}}. \tag{15.7}$$

Comparing this result with Eq. (15.5), we see that $t_1 + t_2 = 2t_3$. So if the apparatus shrinks in the manner just described, we have a way of understanding why the Michelson-Morley experiment gives no effect at all. Although the contraction hypothesis successfully accounted for the negative result of the experiment, it was open to the objection that it was invented for the express purpose of explaining away the difficulty, and was too artificial. However, in many other experiments to discover an ether wind, similar difficulties arose, until it appeared that nature was in a "conspiracy" to thwart man by introducing some new phenomenon to undo every phenomenon that he thought would permit a measurement of u.

It was ultimately recognized, as Poincaré pointed out, that *a complete conspiracy is itself a law of nature!* Poincaré then proposed that there *is* such a law of nature, that it is not possible to discover an ether wind by *any* experiment; that is, there is no way to determine an absolute velocity.

15–4 Transformation of time

In checking out whether the contraction idea is in harmony with the facts in other experiments, it turns out that everything is correct provided that the *times* are also modified, in the manner expressed in the fourth equation of the set (15.3). That is because the time t_3, calculated for the trip from B to C and back, is not the same when calculated by a man performing the experiment in a moving space ship as when calculated by a stationary observer who is watching the space ship. To the man in the ship the time is simply $2L/c$, but to the other observer it is $(2L/c)/\sqrt{1 - u^2/c^2}$ (Eq. 15.5). In other words, when the outsider sees the man in the space ship lighting a cigar, all the actions appear to be slower than normal, while to the man inside, everything moves at a normal rate. So not only must the lengths shorten, but also the time-measuring instruments ("clocks") must apparently slow down. That is, when the clock in the space ship records 1 second elapsed, as seen by the man in the ship, it shows $1/\sqrt{1 - u^2/c^2}$ second to the man outside.

This slowing of the clocks in a moving system is a very peculiar phenomenon, and is worth an explanation. In order to understand this, we have to watch the machinery of the clock and see what happens when it is moving. Since that is

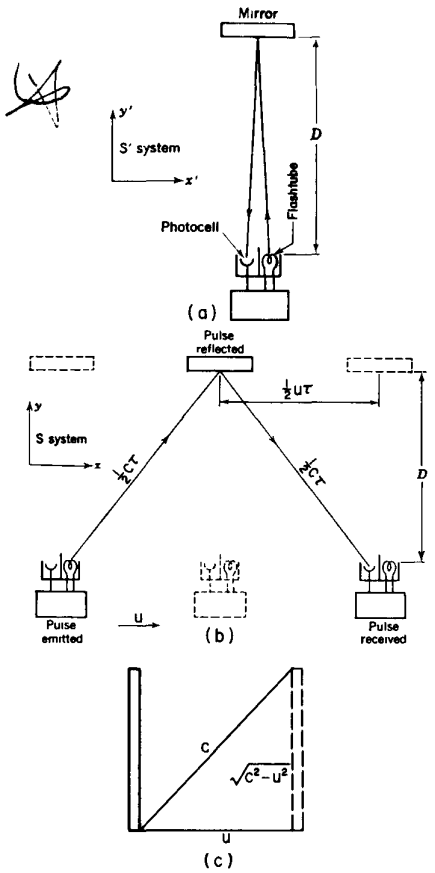


Fig. 15–3. (a) A "light clock" at rest in the S′ system. (b) The same clock, moving through the S system. (c) Illustration of the diagonal path taken by the light beam in a moving "light clock."

rather difficult, we shall take a very simple kind of clock. The one we choose is rather a silly kind of clock, but it will work in principle: it is a rod (meter stick) with a mirror at each end, and when we start a light signal between the mirrors, the light keeps going up and down, making a click every time it comes down, like a standard ticking clock. We build two such clocks, with exactly the same lengths, and synchronize them by starting them together; then they agree always thereafter, because they are the same in length, and light always travels with speed c. We give one of these clocks to the man to take along in his space ship, and he mounts the rod perpendicular to the direction of motion of the ship; then the length of the rod will not change. How do we know that perpendicular lengths do not change? The men can agree to make marks on each other's y-meter stick as they pass each other. By symmetry, the two marks must come at the same y- and y'-coordinates, since otherwise, when they get together to compare results, one mark will be above or below the other, and so we could tell who was really moving.

Now let us see what happens to the moving clock. Before the man took it aboard, he agreed that it was a nice, standard clock, and when he goes along in the space ship he will not see anything peculiar. If he did, he would know he was moving—if anything at all changed because of the motion, he could tell he was moving. But the principle of relativity says this is impossible in a uniformly moving system, so nothing has changed. On the other hand, when the external observer looks at the clock going by, he sees that the light, in going from mirror to mirror, is "really" taking a zigzag path, since the rod is moving sidewise all the while. We have already analyzed such a zigzag motion in connection with the Michelson-Morley experiment. If in a given time the rod moves forward a distance proportional to u in Fig. 15–3, the distance the light travels in the same time is proportional to c, and the vertical distance is therefore proportional to $\sqrt{c^2 - u^2}$.

That is, it takes a *longer time* for light to go from end to end in the moving clock than in the stationary clock. Therefore the apparent time between clicks is longer for the moving clock, in the same proportion as shown in the hypotenuse of the triangle (that is the source of the square root expressions in our equations). From the figure it is also apparent that the greater u is, the more slowly the moving clock appears to run. Not only does this particular kind of clock run more slowly, but if the theory of relativity is correct, any other clock, operating on any principle whatsoever, would also appear to run slower, and in the same proportion—we can say this without further analysis. Why is this so?

To answer the above question, suppose we had two other clocks made exactly alike with wheels and gears, or perhaps based on radioactive decay, or something else. Then we adjust these clocks so they both run in precise synchronism with our first clocks. When light goes up and back in the first clocks and announces its arrival with a click, the new models also complete some sort of cycle, which they simultaneously announce by some doubly coincident flash, or bong, or other signal. One of these clocks is taken into the space ship, along with the first kind. Perhaps *this* clock will not run slower, but will continue to keep the same time as its stationary counterpart, and thus disagree with the other moving clock. Ah no, if that should happen, the man in the ship could use this mismatch between his two clocks to determine the speed of his ship, which we have been supposing is impossible. *We need not know anything about the machinery* of the new clock that might cause the effect—we simply know that whatever the reason, it will appear to run slow, just like the first one.

Now if *all* moving clocks run slower, if no way of measuring time gives anything but a slower rate, we shall just have to say, in a certain sense, that *time itself* appears to be slower in a space ship. All the phenomena there—the man's pulse rate, his thought processes, the time he takes to light a cigar, how long it takes to grow up and get old—all these things must be slowed down in the same proportion, because he cannot tell he is moving. The biologists and medical men sometimes say it is not quite certain that the time it takes for a cancer to develop will be longer in a space ship, but from the viewpoint of a modern physicist it is nearly certain; otherwise one could use the rate of cancer development to determine the speed of the ship!

A very interesting example of the slowing of time with motion is furnished by mu-mesons (muons), which are particles that disintegrate spontaneously after an average lifetime of 2.2×10^{-6} sec. They come to the earth in cosmic rays, and can also be produced artificially in the laboratory. Some of them disintegrate in midair, but the remainder disintegrate only after they encounter a piece of material and stop. It is clear that in its short lifetime a muon cannot travel, even at the speed of light, much more than 600 meters. But although the muons are created at the top of the atmosphere, some 10 kilometers up, yet they are actually found in a laboratory down here, in cosmic rays. How can that be? The answer is that different muons move at various speeds, some of which are very close to the speed of light. While from their own point of view they live only about 2 μsec, from our point of view they live considerably longer—enough longer that they may reach the earth. The factor by which the time is increased has already been given as $1/\sqrt{1 - u^2/c^2}$. The average life has been measured quite accurately for muons of different velocities, and the values agree closely with the formula.

We do not know why the meson disintegrates or what its machinery is, but we do know its behavior satisfies the principle of relativity. That is the utility of the principle of relativity—it permits us to make predictions, even about things that otherwise we do not know much about. For example, before we have any idea at all about what makes the meson disintegrate, we can still predict that when it is moving at nine-tenths of the speed of light, the apparent length of time that it lasts is $(2.2 \times 10^{-6})/\sqrt{1 - 9^2/10^2}$ sec; and our prediction works—that is the good thing about it.

15–5 The Lorentz contraction

Now let us return to the Lorentz transformation (15.3) and try to get a better understanding of the relationship between the (x, y, z, t) and the (x', y', z', t') coordinate systems, which we shall call the S and S' systems, or Joe and Moe systems, respectively. We have already noted that the first equation is based on the Lorentz suggestion of contraction along the x-direction; how can we prove that a contraction takes place? In the Michelson-Morley experiment, we now appreciate that the *transverse* arm BC cannot change length, by the principle of relativity; yet the null result of the experiment demands that the *times* must be equal. So, in order for the experiment to give a null result, the longitudinal arm BE must appear shorter, by the square root $\sqrt{1 - u^2/c^2}$. What does this contraction mean, in terms of measurements made by Joe and Moe? Suppose that Moe, moving with the S' system in the x-direction, is measuring the x'-coordinate of some point with a meter stick. He lays the stick down x' times, so he thinks the distance is x' meters. From the viewpoint of Joe in the S system, however, Moe is using a foreshortened ruler, so the "real" distance measured is $x'\sqrt{1 - u^2/c^2}$ meters. Then if the S' system has travelled a distance ut away from the S system, the S observer would say that the same point, measured in his coordinates, is at a distance $x = x'\sqrt{1 - u^2/c^2} + ut$, or

$$x' = \frac{x - ut}{\sqrt{1 - u^2/c^2}},$$

which is the first equation of the Lorentz transformation.

15–6 Simultaneity

In an analogous way, because of the difference in time scales, the denominator expression is introduced into the fourth equation of the Lorentz transformation. The most interesting term in that equation is the ux/c^2 in the numerator, because that is quite new and unexpected. Now what does that mean? If we look at the situation carefully we see that events that occur at two separated places at the same time, as seen by Moe in S', do *not* happen at the same time as viewed by Joe in S. If one event occurs at point x_1 at time t_0 and the other event at x_2 and t_0 (the same

time), we find that the two corresponding times t_1' and t_2' differ by an amount

$$t_2' - t_1' = \frac{u(x_1 - x_2)/c^2}{\sqrt{1 - u^2/c^2}}.$$

This circumstance is called "failure of simultaneity at a distance," and to make the idea a little clearer let us consider the following experiment.

Suppose that a man moving in a space ship (system S') has placed a clock at each end of the ship and is interested in making sure that the two clocks are in synchronism. How can the clocks be synchronized? There are many ways. One way, involving very little calculation, would be first to locate exactly the midpoint between the clocks. Then from this station we send out a light signal which will go both ways at the same speed and will arrive at both clocks, clearly, at the same time. This simultaneous arrival of the signals can be used to synchronize the clocks. Let us then suppose that the man in S' synchronizes his clocks by this particular method. Let us see whether an observer in system S would agree that the two clocks are synchronous. The man in S' has a right to believe they are, because he does not know that he is moving. But the man in S reasons that since the ship is moving forward, the clock in the front end was running away from the light signal, hence the light had to go more than halfway in order to catch up; the rear clock, however, was advancing to meet the light signal, so this distance was shorter. Therefore the signal reached the rear clock first, although the man in S' thought that the signals arrived simultaneously. We thus see that when a man in a space ship thinks the times at two locations are simultaneous, equal values of t' in his coordinate system must correspond to *different* values of t in the other coordinate system!

15–7 Four-vectors

Let us see what else we can discover in the Lorentz transformation. It is interesting to note that the transformation between the x's and t's is analogous in form to the transformation of the x's and y's that we studied in Chapter 11 for a rotation of coordinates. We then had

$$\begin{aligned} x' &= x\cos\theta + y\sin\theta, \\ y' &= y\cos\theta - x\sin\theta, \end{aligned} \tag{15.8}$$

in which the new x' mixes the old x and y, and the new y' also mixes the old x and y; similarly, in the Lorentz transformation we find a new x' which is a mixture of x and t, and a new t' which is a mixture of t and x. So the Lorentz transformation is analogous to a rotation, only it is a "rotation" in *space and time*, which appears to be a strange concept. A check of the analogy to rotation can be made by calculating the quantity

$$x'^2 + y'^2 + z'^2 - c^2t'^2 = x^2 + y^2 + z^2 - c^2t^2. \tag{15.9}$$

In this equation the first three terms on each side represent, in three-dimensional geometry, the square of the distance between a point and the origin (surface of a sphere) which remains unchanged (invariant) regardless of rotation of the coordinate axes. Similarly, Eq. (15.9) shows that there is a certain combination which includes time, that is invariant to a Lorentz transformation. Thus, the analogy to a rotation is complete, and is of such a kind that vectors, i.e., quantities involving "components" which transform the same way as the coordinates and time, are also useful in connection with relativity.

Thus we contemplate an extension of the idea of vectors, which we have so far considered to have only space components, to include a time component. That is, we expect that there will be vectors with four components, three of which are like the components of an ordinary vector, and with these will be associated a fourth component, which is the analog of the time part.

This concept will be analyzed further in the next chapters, where we shall find that if the ideas of the preceding paragraph are applied to momentum, the transformation gives three space parts that are like ordinary momentum components, and a fourth component, the time part, which is the *energy*.

15–8 Relativistic dynamics

We are now ready to investigate, more generally, what form the laws of mechanics take under the Lorentz transformation. [We have thus far explained how length and time change, but not how we get the modified formula for m (Eq. 15.1). We shall do this in the next chapter.] To see the consequences of Einstein's modification of m for Newtonian mechanics, we start with the Newtonian law that force is the rate of change of momentum, or

$$\mathbf{F} = d(m\mathbf{v})/dt.$$

Momentum is still given by mv, but when we use the new m this becomes

$$\mathbf{p} = m\mathbf{v} = \frac{m_0\mathbf{v}}{\sqrt{1 - v^2/c^2}}. \tag{15.10}$$

This is Einstein's modification of Newton's laws. Under this modification, if action and reaction are still equal (which they may not be in detail, but are in the long run), there will be conservation of momentum in the same way as before, but the quantity that is being conserved is not the old $m\mathbf{v}$ with its constant mass, but instead is the quantity shown in (15.10), which has the modified mass. When this change is made in the formula for momentum, conservation of momentum still works.

Now let us see how momentum varies with speed. In Newtonian mechanics it is proportional to the speed and, according to (15.10), over a considerable range of speed, but small compared with c, it is nearly the same in relativistic mechanics, because the square-root expression differs only slightly from 1. But when v is almost equal to c, the square-root expression approaches zero, and the momentum therefore goes toward infinity.

What happens if a constant force acts on a body for a long time? In Newtonian mechanics the body keeps picking up speed until it goes faster than light. But this is impossible in relativistic mechanics. In relativity, the body keeps picking up, not speed, but momentum, which can continually increase because the mass is increasing. After a while there is practically no acceleration in the sense of a change of velocity, but the momentum continues to increase. Of course, whenever a force produces very little change in the velocity of a body, we say that the body has a great deal of inertia, and that is exactly what our formula for relativistic mass says (see Eq. 15.10)—it says that the inertia is very great when v is nearly as great as c. As an example of this effect, to deflect the high-speed electrons in the synchrotron that is used here at Caltech, we need a magnetic field that is 2000 times stronger than would be expected on the basis of Newton's laws. In other words, the mass of the electrons in the synchrotron is 2000 times as great as their normal mass, and is as great as that of a proton! That m should be 2000 times m_0 means that $1 - v^2/c^2$ must be 1/4,000,000, and that means that v^2/c^2 differs from 1 by one part in 4,000,000, or that v differs from c by one part in 8,000,000, so the electrons are getting pretty close to the speed of light. If the electrons and light were both to start from the synchrotron (estimated as 700 feet away) and rush out to Bridge Lab, which would arrive first? The light, of course, because light always travels faster.* How much earlier? That is too hard to tell—instead, we tell by what distance the light is ahead: it is about 1/1000 of an inch, or $\frac{1}{4}$ the thickness of a piece of paper! When the electrons are going that fast their masses are enormous, but their speed cannot exceed the speed of light.

---

* The electrons would actually win the race versus *visible* light because of the index of refraction of air. A gamma ray would make out better.

Now let us look at some further consequences of relativistic change of mass. Consider the motion of the molecules in a small tank of gas. When the gas is heated, the speed of the molecules is increased, and therefore the mass is also increased and the gas is heavier. An approximate formula to express the increase of mass, for the case when the velocity is small, can be found by expanding $m_0/\sqrt{1 - v^2/c^2} = m_0(1 - v^2/c^2)^{-1/2}$ in a power series, using the binomial theorem. We get

$$m_0(1 - v^2/c^2)^{-1/2} = m_0(1 + \tfrac{1}{2}v^2/c^2 + \tfrac{3}{8}v^4/c^4 + \cdots).$$

We see clearly from the formula that the series converges rapidly when v is small, and the terms after the first two or three are negligible. So we can write

$$m \cong m_0 + \tfrac{1}{2}m_0v^2\left(\frac{1}{c^2}\right) \tag{15.11}$$

in which the second term on the right expresses the increase of mass due to molecular velocity. When the temperature increases the v^2 increases proportionately, so we can say that the increase in mass is proportional to the increase in temperature. But since $\frac{1}{2}m_0v^2$ is the kinetic energy in the old-fashioned Newtonian sense, we can also say that the increase in mass of all this body of gas is equal to the increase in kinetic energy divided by c^2, or $\Delta m = \Delta(\text{K.E.})/c^2$.

15–9 Equivalence of mass and energy

The above observation led Einstein to the suggestion that the mass of a body can be expressed more simply than by the formula (15.1), if we say that the mass is equal to the total energy content divided by c^2. If Eq. (15.11) is multiplied by c^2 the result is

$$mc^2 = m_0c^2 + \tfrac{1}{2}m_0v^2 + \cdots \tag{15.12}$$

Here, the term on the left expresses the total energy of a body, and we recognize the last term as the ordinary kinetic energy. Einstein interpreted the large constant term, m_0c^2, to be part of the total energy of the body, an intrinsic energy known as the "rest energy."

Let us follow out the consequences of assuming, with Einstein, that *the energy of a body always equals* mc^2. As an interesting result, we shall find the formula (15.1) for the variation of mass with speed, which we have merely assumed up to now. We start with the body at rest, when its energy is m_0c^2. Then we apply a force to the body, which starts it moving and gives it kinetic energy; therefore, since the energy has increased, the mass has increased—this is implicit in the original assumption. So long as the force continues, the energy and the mass both continue to increase. We have already seen (Chapter 13) that the rate of change of energy with time equals the force times the velocity, or

$$\frac{dE}{dt} = \mathbf{F} \cdot \mathbf{v}. \tag{15.13}$$

We also have (Chapter 9, Eq. 9.1) that $F = d(mv)/dt$. When these relations are put together with the definition of E, Eq. (15.13) becomes

$$\frac{d(mc^2)}{dt} = \mathbf{v} \cdot \frac{d(m\mathbf{v})}{dt}. \tag{15.14}$$

We wish to solve this equation for m. To do this we first use the mathematical trick of multiplying both sides by $2m$, which changes the equation to

$$c^2(2m)\frac{dm}{dt} = 2mv\frac{d(mv)}{dt}. \tag{15.15}$$

We need to get rid of the derivatives, which can be accomplished by integrating

both sides. The quantity $(2m)\ dm/dt$ can be recognized as the time derivative of m^2, and $(2m\mathbf{v}) \cdot d(m\mathbf{v})/dt$ is the time derivative of $(mv)^2$. So, Eq. (15.15) is the same as

$$c^2 \frac{d(m^2)}{dt} = \frac{d(m^2v^2)}{dt}. \tag{15.16}$$

If the derivatives of two quantities are equal, the quantities themselves differ at most by a constant, say C. This permits us to write

$$m^2c^2 = m^2v^2 + C. \tag{15.17}$$

We need to define the constant C more explicitly. Since Eq. (15.17) must be true for all velocities, we can choose a special case where $v = 0$, and say that in this case the mass is m_0. Substituting these values into Eq. (15.17) gives

$$m_0^2c^2 = 0 + C.$$

We can now use this value of C in Eq. (15.17), which becomes

$$m^2c^2 = m^2v^2 + m_0^2c^2. \tag{15.18}$$

Dividing by c^2 and rearranging terms gives

$$m^2(1 - v^2/c^2) = m_0^2,$$

from which we get

$$m = m_0/\sqrt{1 - v^2/c^2}. \tag{15.19}$$

This is the formula (15.1), and is exactly what is necessary for the agreement between mass and energy in Eq. (15.12).

Ordinarily these energy changes represent extremely slight changes in mass, because most of the time we cannot generate much energy from a given amount of material; but in an atomic bomb of explosive energy equivalent to 20 kilotons of TNT, for example, it can be shown that the dirt after the explosion is lighter by 1 gram than the initial mass of the reacting material, because of the energy that was released, i.e., the released energy had a mass of 1 gram, according to the relationship $\Delta E = \Delta(mc^2)$. This theory of equivalence of mass and energy has been beautifully verified by experiments in which matter is annihilated—converted totally to energy: An electron and a positron come together at rest, each with a rest mass m_0. When they come together they disintegrate and two gamma rays emerge, each with the measured energy of m_0c^2. This experiment furnishes a direct determination of the energy associated with the existence of the rest mass of a particle.

16

Relativistic Energy and Momentum

16–1 Relativity and the philosophers

16–2 The twin paradox

16–3 Transformation of velocities

16–4 Relativistic mass

16–5 Relativistic energy

16–1 Relativity and the philosophers

In this chapter we shall continue to discuss the principle of relativity of Einstein and Poincaré, as it affects our ideas of physics and other branches of human thought.

Poincaré made the following statement of the principle of relativity: "According to the principle of relativity, the laws of physical phenomena must be the same for a fixed observer as for an observer who has a uniform motion of translation relative to him, so that we have not, nor can we possibly have, any means of discerning whether or not we are carried along in such a motion."

When this idea descended upon the world, it caused a great stir among philosophers, particularly the "cocktail-party philosophers," who say, "Oh, it is very simple: Einstein's theory says all is relative!" In fact, a surprisingly large number of philosophers, not only those found at cocktail parties (but rather than embarrass them, we shall just call them "cocktail-party philosophers"), will say, "That all is relative is a consequence of Einstein, and it has profound influences on our ideas." In addition, they say "It has been demonstrated in physics that phenomena depend upon your frame of reference." We hear that a great deal, but it is difficult to find out what it means. Probably the frames of reference that were originally referred to were the coordinate systems which we use in the analysis of the theory of relativity. So the fact that "things depend upon your frame of reference" is supposed to have had a profound effect on modern thought. One might well wonder why, because, after all, that things depend upon one's point of view is so simple an idea that it certainly cannot have been necessary to go to all the trouble of the physical relativity theory in order to discover it. That what one sees depends upon his frame of reference is certainly known to anybody who walks around, because he sees an approaching pedestrian first from the front and then from the back; there is nothing deeper in most of the philosophy which is said to have come from the theory of relativity than the remark that "A person looks different from the front than from the back." The old story about the elephant that several blind men describe in different ways is another example, perhaps, of the theory of relativity from the philosopher's point of view.

But certainly there must be deeper things in the theory of relativity than just this simple remark that "A person looks different from the front than from the back." Of course relativity is deeper than this, because *we can make definite predictions with it.* It certainly would be rather remarkable if we could predict the behavior of nature from such a simple observation alone.

There is another school of philosophers who feel very uncomfortable about the theory of relativity, which asserts that we cannot determine our absolute velocity without looking at something outside, and who would say, "It is obvious that one cannot measure his velocity without looking outside. It is self-evident that it is *meaningless* to talk about the velocity of a thing without looking outside; the physicists are rather stupid for having thought otherwise, but it has just dawned on them that this is the case. If only we philosophers had realized what the problems were that the physicists had, we could have decided immediately by brainwork that it is impossible to tell how fast one is moving without looking outside, and we could have made an enormous contribution to physics." These philosophers are always with us, struggling in the periphery to try to tell us something, but they never really understand the subtleties and depths of the problem.

Our inability to detect absolute motion is a result of *experiment* and not a result of plain thought, as we can easily illustrate. In the first place, Newton believed that it was true that one could not tell how fast he is going if he is moving with uniform velocity in a straight line. In fact, Newton first stated the principle of relativity, and one quotation made in the last chapter was a statement of Newton's. Why then did the philosophers not make all this fuss about "all is relative," or whatever, in Newton's time? Because it was not until Maxwell's theory of electrodynamics was developed that there were physical laws that suggested that one *could* measure his velocity without looking outside; soon it was found *experimentally* that one could *not*.

Now, *is* it absolutely, definitely, philosophically *necessary* that one should not be able to tell how fast he is moving without looking outside? One of the consequences of relativity was the development of a philosophy which said, "You can only define what you can measure! Since it is self-evident that one cannot measure a velocity without seeing what he is measuring it relative to, therefore it is clear that there is no *meaning* to absolute velocity. The physicists should have realized that they can talk only about what they can measure." But *that is the whole problem:* whether or not one *can define* absolute velocity is the same as the problem of whether or not one *can detect in an experiment*, without looking outside, whether he is moving. In other words, whether or not a thing is measurable is not something to be decided *a priori* by thought alone, but something that can be decided only by experiment. Given the fact that the velocity of light is 186,000 mi/sec, one will find few philosophers who will calmly state that it is self-evident that if light goes 186,000 mi/sec inside a car, and the car is going 100,000 mi/sec, that the light also goes 186,000 mi/sec past an observer on the ground. That is a shocking fact to them; the very ones who claim it is obvious find, when you give them a specific fact, that it is not obvious.

Finally, there is even a philosophy which says that one cannot detect *any* motion except by looking outside. It is simply not true in physics. True, one cannot perceive a *uniform* motion in a *straight line*, but if the whole room were *rotating* we would certainly know it, for everybody would be thrown to the wall—there would be all kinds of "centrifugal" effects. That the earth is turning on its axis can be determined without looking at the stars, by means of the so-called Foucault pendulum, for example. Therefore it is not true that "all is relative"; it is only *uniform velocity* that cannot be detected without looking outside. Uniform *rotation* about a fixed axis *can* be. When this is told to a philosopher, he is very upset that he did not really understand it, because to him it seems impossible that one should be able to determine rotation about an axis without looking outside. If the philosopher is good enough, after some time he may come back and say, "I understand. We really do not have such a thing as absolute rotation; we are really rotating *relative to the stars*, you see. And so some influence exerted by the stars on the object must cause the centrifugal force."

Now, for all we know, that is true; we have no way, at the present time, of telling whether there would have been centrifugal force if there were no stars and nebulae around. We have not been able to do the experiment of removing all the nebulae and then measuring our rotation, so we simply do not know. We must admit that the philosopher may be right. He comes back, therefore, in delight and says, "It is absolutely necessary that the world ultimately turn out to be this way: *absolute* rotation means nothing; it is only *relative* to the nebulae." Then we say to him, "*Now*, my friend, is it or is it not obvious that uniform velocity in a straight line, *relative to the nebulae* should produce no effects inside a car?" Now that the motion is no longer absolute, but is a motion *relative to the nebulae*, it becomes a mysterious question, and a question that can be answered only by experiment.

What, then, *are* the philosophic influences of the theory of relativity? If we limit ourselves to influences in the sense of *what kind of new ideas and suggestions* are made to the physicist by the principle of relativity, we could describe some of them as follows. The first discovery is, essentially, that even those ideas which have been held for a very long time and which have been very accurately verified might be wrong. It was a shocking discovery, of course, that Newton's laws are

wrong, after all the years in which they seemed to be accurate. Of course it is clear, not that the experiments were wrong, but that they were done over only a limited range of velocities, so small that the relativistic effects would not have been evident. But nevertheless, we now have a much more humble point of view of our physical laws—everything *can* be wrong!

Secondly, if we have a set of "strange" ideas, such as that time goes slower when one moves, and so forth, whether we *like* them or do *not* like them is an irrelevant question. The only relevant question is whether the ideas are consistent with what is found experimentally. In other words, the "strange ideas" need only agree with *experiment*, and the only reason that we have to discuss the behavior of clocks and so forth is to demonstrate that although the notion of the time dilation is strange, it is *consistent* with the way we measure time.

Finally, there is a third suggestion which is a little more technical but which has turned out to be of enormous utility in our study of other physical laws, and that is to *look at the symmetry of the laws* or, more specifically, to look for the ways in which the laws can be transformed and leave their form the same. When we discussed the theory of vectors, we noted that the fundamental laws of motion are not changed when we rotate the coordinate system, and now we learn that they are not changed when we change the space and time variables in a particular way, given by the Lorentz transformation. So this idea of studying the patterns or operations under which the fundamental laws are not changed has proved to be a very useful one.

16–2 The twin paradox

To continue our discussion of the Lorentz transformation and relativistic effects, we consider a famous so-called "paradox" of Peter and Paul, who are supposed to be twins, born at the same time. When they are old enough to drive a space ship, Paul flies away at very high speed. Because Peter, who is left on the ground, sees Paul going so fast, all of Paul's clocks appear to go slower, his heart beats go slower, his thoughts go slower, everything goes slower, from Peter's point of view. Of course, Paul notices nothing unusual, but if he travels around and about for a while and then comes back, he will be younger than Peter, the man on the ground! That is actually right; it is one of the consequences of the theory of relativity which has been clearly demonstrated. Just as the mu-mesons last longer when they are moving, so also will Paul last longer when he is moving. This is called a "paradox" only by the people who believe that the principle of relativity means that *all motion* is relative; they say, "Heh, heh, heh, from the point of view of Paul, can't we say that *Peter* was moving and should therefore appear to age more slowly? By symmetry, the only possible result is that both should be the same age when they meet." But in order for them to come back together and make the comparison, Paul must either stop at the end of the trip and make a comparison of clocks or, more simply, he has to come back, and the one who comes back must be the man who was moving, and he knows this, because he had to turn around. When he turned around, all kinds of unusual things happened in his space ship—the rockets went off, things jammed up against one wall, and so on—while Peter felt nothing.

So the way to state the rule is to say that *the man who has felt the accelerations*, who has seen things fall against the walls, and so on, is the one who would be the younger; that is the difference between them in an "absolute" sense, and it is certainly correct. When we discussed the fact that moving mu-mesons live longer, we used as an example their straight-line motion in the atmosphere. But we can also make mu-mesons in a laboratory and cause them to go in a curve with a magnet, and even under this accelerated motion, they last exactly as much longer as they do when they are moving in a straight line. Although no one has arranged an experiment explicitly so that we can get rid of the paradox, one could compare a mu-meson which is left standing with one that had gone around a complete circle, and it would surely be found that the one that went around the circle lasted longer. Although we have not actually carried out an experiment using a complete

circle, it is really not necessary, of course, because everything fits together all right. This may not satisfy those who insist that every single fact be demonstrated directly, but we confidently predict the result of the experiment in which Paul goes in a complete circle.

16–3 Transformation of velocities

The main difference between the relativity of Einstein and the relativity of Newton is that the laws of transformation connecting the coordinates and times between relatively moving systems are different. The correct transformation law, that of Lorentz, is

$$\begin{aligned} x' &= \frac{x - ut}{\sqrt{1 - u^2/c^2}}, \\ y' &= y, \\ z' &= z, \\ t' &= \frac{t - ux/c^2}{\sqrt{1 - u^2/c^2}}. \end{aligned} \tag{16.1}$$

These equations correspond to the relatively simple case in which the relative motion of the two observers is along their common x-axes. Of course other directions of motion are possible, but the most general Lorentz transformation is rather complicated, with all four quantities mixed up together. We shall continue to use this simpler form, since it contains all the essential features of relativity.

Let us now discuss more of the consequences of this transformation. First, it is interesting to solve these equations in reverse. That is, here is a set of linear equations, four equations with four unknowns, and they can be solved in reverse, for x, y, z, t in terms of x', y', z', t'. The result is very interesting, since it tells us how a system of coordinates "at rest" looks from the point of view of one that is "moving." Of course, since the motions are relative and of uniform velocity, the man who is "moving" can say, if he wishes, that it is really the other fellow who is moving and he himself who is at rest. And since he is moving in the opposite direction, he should get the same transformation, but with the opposite sign of velocity. That is precisely what we find by manipulation, so that is consistent. If it did not come out that way, we would have real cause to worry!

$$\begin{aligned} x &= \frac{x' + ut'}{\sqrt{1 - u^2/c^2}}, \\ y &= y', \\ z &= z', \\ t &= \frac{t' + ux'/c^2}{\sqrt{1 - u^2/c^2}}. \end{aligned} \tag{16.2}$$

Next we discuss the interesting problem of the addition of velocities in relativity. We recall that one of the original puzzles was that light travels at 186,000 mi/sec in all systems, even when they are in relative motion. This is a special case of the more general problem exemplified by the following. Suppose that an object inside a space ship is going at 100,000 mi/sec and the space ship itself is going at 100,000 mi/sec; how fast is the object inside the space ship moving from the point of view of an observer outside? We might want to say 200,000 mi/sec, which is faster than the speed of light. This is very unnerving, because it is not supposed to be going faster than the speed of light! The general problem is as follows.

Let us suppose that the object inside the ship, from the point of view of the man inside, is moving with velocity v, and that the space ship itself has a velocity u with respect to the ground. We want to know with what velocity v_x this object is moving from the point of view of the man on the ground. This is, of course, still but a special case in which the motion is in the x-direction. There will also be a

transformation for velocities in the y-direction, or for any angle; these can be worked out as needed. Inside the space ship the velocity is $v_{x'}$, which means that the displacement x is equal to the velocity times the time:

$$x' = v_{x'}t'. \tag{16.3}$$

Now we have only to calculate what the position and time are from the point of view of the outside observer for an object which has the relation (16.2) between x' and t'. So we simply substitute (16.3) into (16.2), and obtain

$$x = \frac{v_{x'}t' + ut'}{\sqrt{1 - u^2/c^2}}. \tag{16.4}$$

But here we find x expressed in terms of t'. In order to get the velocity as seen by the man on the outside, we must divide *his distance* by *his time*, not by the *other man's time!* So we must also calculate the *time* as seen from the outside, which is

$$t = \frac{t' + u(v_{x'}t')/c^2}{\sqrt{1 - u^2/c^2}}. \tag{16.5}$$

Now we must find the ratio of x to t, which is

$$v_x = \frac{x}{t} = \frac{u + v_{x'}}{1 + uv_{x'}/c^2}, \tag{16.6}$$

the square roots having cancelled. This is the law that we seek: the resultant velocity, the "summing" of two velocities, is not just the algebraic sum of two velocities (we know that it cannot be or we get in trouble), but is "corrected" by $1 + uv/c^2$.

Now let us see what happens. Suppose that you are moving inside the space ship at half the speed of light, and that the space ship itself is going at half the speed of light. Thus u is $\frac{1}{2}c$ and v is $\frac{1}{2}c$, but in the denominator uv is one-fourth, so that

$$v = \frac{\frac{1}{2}c + \frac{1}{2}c}{1 + \frac{1}{4}} = \frac{4c}{5}.$$

So, in relativity, "half" and "half" does not make "one," it makes only "4/5." Of course low velocities can be added quite easily in the familiar way, because so long as the velocities are small compared with the speed of light we can forget about the $(1 + uv/c^2)$ factor; but things are quite different and quite interesting at high velocity.

Let us take a limiting case. Just for fun, suppose that inside the space ship the man was observing *light itself.* In other words, $v = c$, and yet the space ship is moving. How will it look to the man on the ground? The answer will be

$$v = \frac{u + c}{1 + uc/c^2} = c\frac{u + c}{u + c} = c.$$

Therefore, if something is moving at the speed of light inside the ship, it will appear to be moving at the speed of light from the point of view of the man on the ground too! This is good, for it is, in fact, what the Einstein theory of relativity was designed to do in the first place—so it had *better* work!

Of course, there are cases in which the motion is not in the direction of the uniform translation. For example, there may be an object inside the ship which is just moving "upward" with the velocity $v_{y'}$ with respect to the ship, and the ship is moving "horizontally." Now, we simply go through the same thing, only using y's instead of x's, with the result

$$y = y' = v_{y'}t',$$

so that if $v_{x'} = 0$,

$$v_y = \frac{y}{t} = v_{y'}\sqrt{1 - u^2/c^2}. \tag{16.7}$$

Thus a sidewise velocity is no longer $v_{y'}$, but $v_{y'}\sqrt{1 - u^2/c^2}$. We found this result by substituting and combining the transformation equations, but we can also see the result directly from the principle of relativity for the following reason (it is always good to look again to see whether we can see the reason). We have already (Fig. 15–3) discussed how a possible clock might work when it is moving; the light appears to travel at an angle at the speed c in the fixed system, while it simply goes vertically with the same speed in the moving system. We found that the *vertical component* of the velocity in the fixed system is less than that of light by the factor $\sqrt{1 - u^2/c^2}$ (see Eq. 15–3). But now suppose that we let a material particle go back and forth in this same "clock," but at some integral fraction $1/n$ of the speed of light (Fig. 16–1). Then when the particle has gone back and forth once, the light will have gone exactly n times. That is, each "click" of the "particle" clock will coincide with each nth "click" of the light clock. *This fact must still be true when the whole system is moving*, because the physical phenomenon of coincidence will be a coincidence in any frame. Therefore, since the speed c_y is less than the speed of light, the speed v_y of the particle must be slower than the corresponding speed by the same square-root ratio! That is why the square root appears in any vertical velocity.

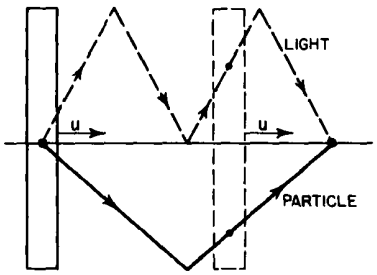


Fig 16–1. Trajectories described by a light ray and particle inside a moving clock.

16–4 Relativistic mass

We learned in the last chapter that the mass of an object increases with velocity, but no demonstration of this was given, in the sense that we made no arguments analogous to those about the way clocks have to behave. However, we *can* show that, as a consequence of relativity plus a few other reasonable assumptions, the mass must vary in this way. (We have to say "a few other assumptions" because we cannot prove anything unless we have some laws which we assume to be true, if we expect to make meaningful deductions.) To avoid the need to study the transformation laws of force, we shall analyze a *collision*, where we need know nothing about the laws of force, except that we shall assume the conservation of momentum and energy. Also, we shall assume that the momentum of a particle which is moving is a vector and is always directed in the direction of the velocity. However, we shall not assume that the momentum is a *constant* times the velocity, as Newton did, but only that it is some *function* of velocity. We thus write the momentum vector as a certain coefficient times the vector velocity:

$$\mathbf{p} = m_v\mathbf{v}. \tag{16.8}$$

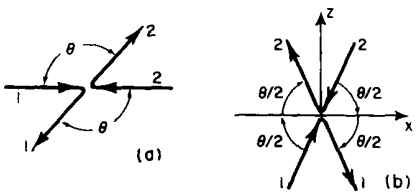


Fig. 16–2. Two views of an elastic collision between equal objects moving at the same speed in opposite directions.

We put a subscript v on the coefficient to remind us that it is a function of velocity, and we shall agree to call this coefficient m_v the "mass." Of course, when the velocity is small, it is the same mass that we would measure in the slow-moving experiments that we are used to. Now we shall try to demonstrate that the formula for m_v must be $m_0/\sqrt{1 - v^2/c^2}$, by arguing from the principle of relativity that the laws of physics must be the same in every coordinate system.

Suppose that we have two particles, like two protons, that are absolutely equal, and they are moving toward each other with exactly equal velocities. Their total momentum is zero. Now what can happen? After the collision, their directions of motion must be exactly opposite to each other, because if they are not exactly opposite, there will be a nonzero total vector momentum, and momentum would not have been conserved. Also they must have the same speeds, since they are exactly similar objects; in fact, they must have the same speed they started with, since we suppose that the energy is conserved in these collisions. So the diagram of an elastic collision, a reversible collision, will look like Fig. 16–2(a): all the arrows are the same length, all the speeds are equal. We shall suppose that such collisions can always be arranged, that any angle θ can occur, and that any speed could be used in such a collision. Next, we notice that this same collision can be viewed differently by turning the axes, and just for convenience we *shall* turn the axes, so that the horizontal splits it evenly, as in Fig. 16–2(b). It is the same collision redrawn, only with the axes turned.

Now here is the real trick: let us look at this collision from the point of view of someone riding along in a car that is moving with a speed equal to the horizontal component of the velocity of one particle. Then how does the collision look? It looks as though particle 1 is just going straight up, because it has lost its horizontal component, and it comes straight down again, also because it does not have that component. That is, the collision appears as shown in Fig. 16–3(a). Particle 2, however, was going the other way, and as we ride past it appears to fly by at some terrific speed and at a smaller angle, but we can appreciate that the angles before and after the collision are the *same*. Let us denote by u the horizontal component of the velocity of particle 2, and by w the vertical velocity of particle 1.

Now the question is, what is the vertical velocity $u \tan \alpha$? If we knew that, we could get the correct expression for the momentum, using the law of conservation of momentum in the vertical direction. Clearly, the horizontal component of the momentum is conserved: it is the same before and after the collision for both particles, and is zero for particle 1. So we need use the conservation law only for the upward velocity $u \tan \alpha$. But we *can* get the upward velocity, simply by looking at the same collision going the other way! If we look at the collision of Fig. 16–3(a) from a car to the left moving with speed u, we see the same collision, except "turned over," as shown in Fig. 16–3(b). Now particle 2 is the one that goes up and down with speed w, and particle 1 has picked up the horizontal speed u. Of course, now we *know* what the velocity $u \tan \alpha$ is: it is $w\sqrt{1 - u^2/c^2}$ (see Eq. 16.7). We know that the change in the vertical momentum of the vertically moving particle is

$$\Delta p = 2m_w w$$

(2, because it moves up and back down). The obliquely moving particle has a certain velocity v whose components we have found to be u and $w\sqrt{1 - u^2/c^2}$, and whose mass is m_v. The change in *vertical* momentum of this particle is therefore $\Delta p' = 2m_v w\sqrt{1 - u^2/c^2}$ because, in accordance with our assumed law (16.8), the momentum component is always the mass corresponding to the magnitude of the velocity times the component of the velocity in the direction of interest. Thus in order for the total momentum to be zero the vertical momenta must cancel and the ratio of the mass moving with speed v and the mass moving with speed w must therefore be

$$\frac{m_w}{m_v} = \sqrt{1 - u^2/c^2}. \tag{16.9}$$

Let us take the limiting case that w is infinitesimal. If w is very tiny indeed, it is clear that v and u are practically equal. In this case, $m_w \to m_0$ and $m_v \to m_u$. The grand result is

$$m_u = \frac{m_0}{\sqrt{1 - u^2/c^2}}. \tag{16.10}$$

It is an interesting exercise now to check whether or not Eq. (16.9) is indeed true for arbitrary values of w, assuming that Eq. (16.10) is the right formula for the mass. Note that the velocity v needed in Eq. (16.9) can be calculated from the right-angle triangle:

$$v^2 = u^2 + w^2(1 - u^2/c^2).$$

It will be found to check out automatically, although we used it only in the limit of small w.

Now, let us accept that momentum is conserved and that the mass depends upon the velocity according to (16.10) and go on to find what else we can conclude. Let us consider what is commonly called an *inelastic collision*. For simplicity, we shall suppose that two objects of the same kind, moving oppositely with equal speeds w, hit each other and stick together, to become some new, stationary object, as shown in Fig. 16–4(a). The mass m of each corresponds to w, which, as we know, is $m_0/\sqrt{1 - w^2/c^2}$. If we assume the conservation of momentum and the principle of relativity, we can demonstrate an interesting fact about the mass of the new object which has been formed. We imagine an infinitesimal velocity

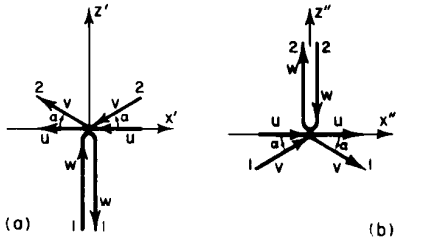


Fig. 16–3. Two more views of the collision, from moving cars.

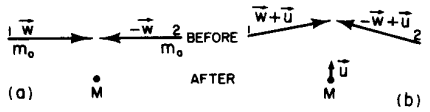


Fig. 16–4. Two views of an inelastic collision between equally massive objects.

u at right angles to w (we can do the same with finite values of u, but it is easier to understand with an infinitesimal velocity), then look at this same collision as we ride by in an elevator at the velocity $-u$. What we see is shown in Fig. 16–4(b). The composite object has an unknown mass M. Now object 1 moves with an upward component of velocity u and a horizontal component which is practically equal to w, and so also does object 2. After the collision we have the mass M moving upward with velocity u, considered very small compared with the speed of light, and also small compared with w. Momentum must be conserved, so let us estimate the momentum in the upward direction before and after the collision. Before the collision we have $p \sim 2m_w u$, and after the collision the momentum is evidently $p' = M_u u$, but M_u is essentially the same as M_0 because u is so small. These momenta must be equal because of the conservation of momentum, and therefore

$$M_0 = 2m_w. \tag{16.11}$$

The mass of the object which is formed when two equal objects collide must be twice the mass of the objects which come together. You might say, "Yes, of course, that is the conservation of mass." But not "Yes, of course," so easily, because *these masses have been enhanced* over the masses that they would be if they were standing still, yet they still contribute, to the total M, not the mass they have when standing still, but *more*. Astonishing as that may seem, in order for the conservation of momentum to work when two objects come together, the mass that they form must be greater than the rest masses of the objects, even though the objects are at rest after the collision!

16–5 Relativistic energy

In the last chapter we demonstrated that as a result of the dependence of the mass on velocity and Newton's laws, the changes in the kinetic energy of an object resulting from the total work done by the forces on it always comes out to be

$$\Delta T = (m_u - m_0)c^2 = \frac{m_0 c^2}{\sqrt{1 - u^2/c^2}} - m_0 c^2. \tag{16.12}$$

We even went further, and guessed that the total energy is the total mass times c^2. Now we continue this discussion.

Suppose that our two equally massive objects that collide can still be "seen" inside M. For instance, a proton and a neutron are "stuck together," but are still moving about inside of M. Then, although we might at first expect the mass M to be $2m_0$, we have found that it is not $2m_0$, but $2m_w$. Since $2m_w$ is what is put in, but $2m_0$ are the rest masses of the things inside, the *excess* mass of the composite object is equal to the kinetic energy brought in. This means, of course, that *energy has inertia.* In the last chapter we discussed the heating of a gas, and showed that because the gas molecules are moving and moving things are heavier, when we put energy into the gas its molecules move faster and so the gas gets heavier. But in fact the argument is completely general, and our discussion of the inelastic collision shows that the mass is there whether or not it is *kinetic* energy. In other words, if two particles come together and produce potential or any other form of energy; if the pieces are slowed down by climbing hills, doing work against internal forces, or whatever; then it is still true that the mass is the total energy that has been put in. So we see that the conservation of mass which we have deduced above is equivalent to the conservation of energy, and therefore there is no place in the theory of relativity for strictly inelastic collisions, as there was in Newtonian mechanics. According to Newtonian mechanics it is all right for two things to collide and so form an object of mass $2m_0$ which is in no way distinct from the one that would result from putting them together slowly. Of course we know from the law of conservation of energy that there is more kinetic energy inside, but that does not affect the mass, according to Newton's laws. But now we see that this is impossible; because of the kinetic energy involved in the collision, the resulting

object will be *heavier;* therefore, it will be a *different* object. When we put the objects together gently they make something whose mass is $2m_0$; when we put them together forcefully, they make something whose mass is greater. When the mass is different, we can *tell* that it is different. So, necessarily, the conservation of energy must go along with the conservation of momentum in the theory of relativity.

This has interesting consequences. For example, suppose that we have an object whose mass M is measured, and suppose something happens so that it flies into two equal pieces moving with speed w, so that they each have a mass m_w. Now suppose that these pieces encounter enough material to slow them up until they stop; then they will have mass m_0. How much energy will they have given to the material when they have stopped? Each will give an amount $(m_w - m_0)c^2$, by the theorem that we proved before. This much energy is left in the material in some form, as heat, potential energy, or whatever. Now $2m_w = M$, so the liberated energy is $E = (M - 2m_0)c^2$. This equation was used to estimate how much energy would be liberated under fission in the atomic bomb, for example. (Although the fragments are not exactly equal, they are nearly equal.) The mass of the uranium atom was known—it had been measured ahead of time—and the atoms into which it split, iodine, xenon, and so on, all were of known mass. By masses, we do not mean the masses while the atoms are moving, we mean the masses when the atoms are *at rest.* In other words, both M and m_0 are known. So by subtracting the two numbers one can calculate how much energy will be released if M can be made to split in "half." For this reason poor old Einstein was called the "father" of the atomic bomb in all the newspapers. Of course, all that meant was that he could tell us ahead of time how much energy would be released if we told him what process would occur. The energy that should be liberated when an atom of uranium undergoes fission was estimated about six months before the first direct test, and as soon as the energy was in fact liberated, someone measured it directly (and if Einstein's formula had not worked, they would have measured it anyway), and the moment they measured it they no longer needed the formula. Of course, we should not belittle Einstein, but rather should criticize the newspapers and many popular descriptions of what causes what in the history of physics and technology. The problem of how to get the thing to occur in an effective and rapid manner is a completely different matter.

The result is just as significant in chemistry. For instance, if we were to weigh the carbon dioxide molecule and compare its mass with that of the carbon and the oxygen, we could find out how much energy would be liberated when carbon and oxygen form carbon dioxide. The only trouble here is that the differences in masses are so small that it is technically very difficult to do.

Now let us turn to the question of whether we should add m_0c^2 to the kinetic energy and say from now on that the total energy of an object is mc^2. First, if we can still *see* the component pieces of rest mass m_0 inside M, then we could say that some of the mass M of the compound object is the mechanical rest mass of the parts, part of it is kinetic energy of the parts, and part of it is potential energy of the parts. But we have discovered, in nature, particles of various kinds which undergo reactions just like the one we have treated above, in which with all the study in the world, *we cannot see the parts inside.* For instance, when a K-meson disintegrates into two pions it does so according to the law (16.11), but the idea that a K is made out of 2 π's is a useless idea, because it also disintegrates into 3 π's!

Therefore we have a *new idea*: we do not have to know what things are made of inside; we cannot and need not identify, inside a particle, which of the energy is rest energy of the parts into which it is going to disintegrate. It is not convenient and often not possible to separate the total mc^2 energy of an object into rest energy of the inside pieces, kinetic energy of the pieces, and potential energy of the pieces; instead, we simply speak of the *total energy* of the particle. We "shift the origin" of energy by adding a constant m_0c^2 to everything, and say that the total energy of a particle is the mass in motion times c^2, and when the object is standing still, the energy is the mass at rest times c^2.

Finally, we find that the velocity v, momentum P, and total energy E are related in a rather simple way. That the mass in motion at speed v is the mass m_0 at rest divided by $\sqrt{1 - v^2/c^2}$, surprisingly enough, is rarely used. Instead, the following relations are easily proved, and turn out to be very useful:

$$E^2 - P^2c^2 = m_0^2c^4 \tag{16.13}$$

and

$$Pc = Ev/c. \tag{16.14}$$

17

Space-Time

17–1 The geometry of space-time

17–2 Space-time intervals

17–3 Past, present, and future

17–4 More about four-vectors

17–5 Four-vector algebra

17–1 The geometry of space-time

The theory of relativity shows us that the relationships of positions and times as measured in one coordinate system and another are not what we would have expected on the basis of our intuitive ideas. It is very important that we thoroughly understand the relations of space and time implied by the Lorentz transformation, and therefore we shall consider this matter more deeply in this chapter.

The Lorentz transformation between the positions and times (x, y, z, t) as measured by an observer "standing still," and the corresponding coordinates and time (x', y', z', t') measured inside a "moving" space ship, moving with velocity u are

$$\begin{aligned} x' &= \frac{x - ut}{\sqrt{1 - u^2/c^2}}, \\ y' &= y, \\ z' &= z, \\ t' &= \frac{t - ux/c^2}{\sqrt{1 - u^2/c^2}}. \end{aligned} \tag{17.1}$$

Let us compare these equations with Eq. (11.5), which also relates measurements in two systems, one of which in this instance is *rotated* relative to the other:

$$\begin{aligned} x' &= x\cos\theta + y\sin\theta, \\ y' &= y\cos\theta - x\sin\theta, \\ z' &= z. \end{aligned} \tag{17.2}$$

In this particular case, Moe and Joe are measuring with axes having an angle θ between the x'- and x-axes. In each case, we note that the "primed" quantities are "mixtures" of the "unprimed" ones: the new x' is a mixture of x and y, and the new y' is also a mixture of x and y.

An analogy is useful: When we look at an object, there is an obvious thing we might call the "apparent width," and another we might call the "depth." But the two ideas, width and depth, are not *fundamental* properties of the object, because if we step aside and look at the same thing from a different angle, we get a different width and a different depth, and we may develop some formulas for computing the new ones from the old ones and the angles involved. Equations (17.2) are these formulas. One might say that a given depth is a kind of "mixture" of all depth and all width. If it were impossible ever to move, and we always saw a given object from the same position, then this whole business would be irrelevant—we would always see the "true" width and the "true" depth, and they would appear to have quite different qualities, because one appears as a subtended optical angle and the other involves some focusing of the eyes or even intuition; they would seem to be very different things and would never get mixed up. It is because we *can* walk around that we realize that depth and width are, somehow or other, just two different aspects of the same thing.

Can we not look at the Lorentz transformations in the same way? Here also we have a mixture—of positions and the time. A difference between a space measurement and a time measurement produces a new space measurement. In other words, in the space measurements of one man there is mixed in a little bit of the time, as seen by the other. Our analogy permits us to generate this idea: The "reality" of

an object that we are looking at is somehow greater (speaking crudely and intuitively) than its "width" and its "depth" because *they* depend upon *how* we look at it; when we move to a new position, our brain immediately recalculates the width and the depth. But our brain does not immediately recalculate coordinates and time when we move at high speed, because we have had no effective experience of going nearly as fast as light to appreciate the fact that time and space are also of the same nature. It is as though we were always stuck in the position of having to look at just the width of something, not being able to move our heads appreciably one way or the other; if we could, we understand now, we would see some of the other man's time—we would see "behind," so to speak, a little bit.

Thus we shall try to think of objects in a new kind of world, of space and time mixed together, in the same sense that the objects in our ordinary space-world are real, and can be looked at from different directions. We shall then consider that objects occupying space and lasting for a certain length of time occupy a kind of a "blob" in a new kind of world, and that we look at this "blob" from different points of view when we are moving at different velocities. This new world, this geometrical entity in which the "blobs" exist by occupying position and taking up a certain amount of time, is called *space-time*. A given point (x, y, z, t) in space-time is called an *event*. Imagine, for example, that we plot the x-positions horizontally, y and z in two other directions, both mutually at "right angles" and at "right angles" to the paper (!), and time, vertically. Now, how does a moving particle, say, look on such a diagram? If the particle is standing still, then it has a certain x, and as time goes on, it has the same x, the same x, the same x; so its "path" is a line that runs parallel to the t-axis (Fig. 17–1 a). On the other hand, if it drifts outward, then as the time goes on x increases (Fig. 17–1 b). So a particle, for example, which starts to drift out and then slows up should have a motion something like that shown in Fig. 17–1(c). A patricle, in other words, which is permanent and does not disintegrate is represented by a line in space-time. A particle which disintegrates would be represented by a forked line, because it would turn into two other things which would start from that point.

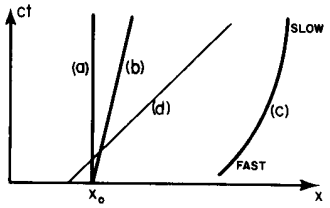


Fig. 17–1. Three particle paths in space-time: (a) a particle at rest at $x = x_0$; (b) a particle which starts at $x = x_0$ and moves with constant speed; (c) a particle which starts at high speed but slows down.

What about light? Light travels at the speed c, and that would be represented by a line having a certain fixed slope (Fig. 17–1 d).

Now according to our new idea, if a given event occurs to a particle, say if it suddenly disintegrates at a certain space-time point into two new ones which follow some new tracks, and this interesting event occurred at a certain value of x and a certain value of t, then we would expect that, if this makes any sense, we just have to take a new pair of axes and turn them, and that will give us the new t and the new x in our new system, as shown in Fig. 17–2(a). But this is wrong, because Eq. (17.1) is not *exactly* the same mathematical transformation as Eq. (17.2). Note, for example, the difference in sign between the two, and the fact that one is written in terms of cos θ and sin θ, while the other is written with algebraic quantities. (Of course, it is not impossible that the algebraic quantities could be written as cosine and sine, but actually they cannot.) But still, the two expressions *are* very similar. As we shall see, it is not really possible to think of space-time as a real, ordinary geometry because of that difference in sign. In fact, although we shall not emphasize this point, it turns out that a man who is moving has to use a set of axes which are inclined equally to the light ray, using a special kind of projection parallel to the x'- and t'-axes, for his x' and t', as shown in Fig. 17–2(b). We shall not deal with the geometry, since it does not help much; it is easier to work with the equations.

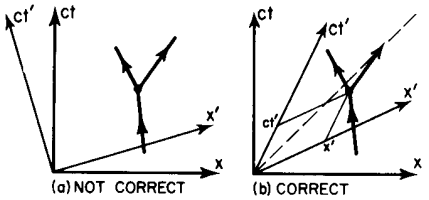


Fig. 17–2. Two views of a disintegrating particle.

17–2 Space-time intervals

Although the geometry of space-time is not Euclidean in the ordinary sense, there *is* a geometry which is very similar, but peculiar in certain respects. If this idea of geometry is right, there ought to be some functions of coordinates and time which are independent of the coordinate system. For example, under ordinary rotations, if we take two points, one at the origin, for simplicity, and the other one somewhere else, both systems would have the same origin, and the distance from

here to the other point is the same in both. That is one property that is independent of the particular way of measuring it. The square of the distance is $x^2 + y^2 + z^2$. Now what about space-time? It is not hard to demonstrate that we have here, also, something which stays the same, namely, the combination $c^2t^2 - x^2 - y^2 - z^2$ is the same before and after the transformation:

$$c^2t'^2 - x'^2 - y'^2 - z'^2 = c^2t^2 - x^2 - y^2 - z^2. \tag{17.3}$$

This quantity is therefore something which, like the distance, is "real" in some sense; it is called the *interval* between the two space-time points, one of which is, in this case, at the origin. (Actually, of course, it is the interval squared, just as $x^2 + y^2 + z^2$ is the distance squared.) We give it a different name because it is in a different geometry, but the interesting thing is only that some signs are reversed and there is a c in it.

Let us get rid of the c; that is an absurdity if we are going to have a wonderful space with x's and y's that can be interchanged. One of the confusions that could be caused by someone with no experience would be to measure widths, say, by the angle subtended at the eye, and measure depth in a different way, like the strain on the muscles needed to focus them, so that the depths would be measured in feet and the widths in meters. Then one would get an enormously complicated mess of equations in making transformations such as (17.2), and would not be able to see the clarity and simplicity of the thing for a very simple technical reason, that the same thing is being measured in two different units. Now in Eqs. (17.1) and (17.3) nature is telling us that time and space are equivalent; time becomes space; *they should be measured in the same units.* What distance is a "second"? It is easy to figure out from (17.3) what it is. It is 3×10^8 meters, *the distance that light would go in one second.* In other words, if we were to measure all distances and times in the same units, seconds, then our unit of distance would be 3×10^8 meters, and the equations would be simpler. Or another way that we could make the units equal is to measure time in meters. What is a meter of time? A meter of time is the time it takes for light to go one meter, and is therefore $1/3 \times 10^{-8}$ sec, or 3.3 billionths of a second! We would like, in other words, to put all our equations in a system of units in which $c = 1$. If time and space are measured in the same units, as suggested, then the equations are obviously much simplified. They are

$$\begin{aligned} x' &= \frac{x - ut}{\sqrt{1 - u^2}}, \\ y' &= y, \\ z' &= z, \\ t' &= \frac{t - ux}{\sqrt{1 - u^2}}. \end{aligned} \tag{17.4}$$

$$t'^2 - x'^2 - y'^2 - z'^2 = t^2 - x^2 - y^2 - z^2. \tag{17.5}$$

If we are ever unsure or "frightened" that after we have this system with $c = 1$ we shall never be able to get our equations right again, the answer is quite the opposite. It is much easier to remember them without the c's in them, and it is always easy to put the c's back, by looking after the dimensions. For instance, in $\sqrt{1 - u^2}$, we know that we cannot subtract a velocity squared, which has units, from the pure number 1, so we know that we must divide u^2 by c^2 in order to make that unitless, and that is the way it goes.

The difference between space-time and ordinary space, and the character of an interval as related to the distance, is very interesting. According to formula (17.5), if we consider a point which in a given coordinate system had zero time, and only space, then the interval squared would be negative and we would have an imaginary interval, the square root of a negative number. Intervals can be either real or imaginary in the theory. The square of an interval may be either positive or negative, unlike distance, which has a positive square. When an interval is imaginary, we say that the two points have a *space-like interval* between them

(instead of imaginary), because the interval is more like space than like time. On the other hand, if two objects are at the same place in a given coordinate system, but differ only in time, then the square of the time is positive and the distances are zero and the interval squared is positive; this is called a *time-like interval.* In our diagram of space-time, therefore, we would have a representation something like this: at 45° there are two lines (actually, in four dimensions these will be "cones," called light cones) and points on these lines are all at zero interval from the origin. Where light goes from a given point is always separated from it by a zero interval, as we see from Eq. (17.5). Incidentally, we have just proved that if light travels with speed c in one system, it travels with speed c in another, for if the interval is the same in both systems, i.e., zero in one and zero in the other, then to state that the propagation speed of light is invariant is the same as saying that the interval is zero.

17–3 Past, present, and future

The space-time region surrounding a given space-time point can be separated into three regions, as shown in Fig. 17–3. In one region we have space-like intervals, and in two regions, time-like intervals. Physically, these three regions into which space-time around a given point is divided have an interesting physical relationship to that point: a physical object or a signal can get from a point in region 2 to the event O by moving along at a speed less than the speed of light. Therefore events in this region can affect the point O, can have an influence on it from the past. In fact, of course, an object at P on the negative t-axis is precisely in the "past" with respect to O; it is the same space-point as O, only earlier. What happened there then, affects O now. (Unfortunately, that is the way life is.) Another object at Q can get to O by moving with a certain speed less than c, so if this object were in a space ship and moving, it would be, again, the past of the same space-point. That is, in another coordinate system, the axis of time might go through both O and Q. So all points of region 2 are in the "past" of O, and anything that happens in this region *can* affect O. Therefore region 2 is sometimes called the *affective past*, or affecting past; it is the locus of all events which can affect point O in any way.

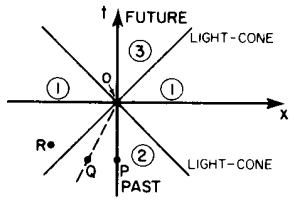


Fig. 17–3. The space-time region surrounding a point at the origin.

Region 3, on the other hand, is a region which we can affect *from O*, we can "hit" things by shooting "bullets" out at speeds less than c. So this is the world whose future can be affected by us, and we may call that the *affective future.* Now the interesting thing about all the rest of space-time, i.e., region 1, is that we can neither affect it now *from O*, nor can it affect us now *at O*, because nothing can go faster than the speed of light. Of course, what happens at R *can* affect us *later*; that is, if the sun is exploding "right now," it takes eight minutes before we know about it, and it cannot possibly affect us before then.

What we mean by "right now" is a mysterious thing which we cannot define and we cannot affect, but it can affect us later, or we could have affected it if we had done something far enough in the past. When we look at the star Alpha Centauri, we see it as it was four years ago; we might wonder what it is like "now." "Now" means at the same time from our special coordinate system. We can only see Alpha Centauri by the light that has come from our past, up to four years ago, but we do not know what it is doing "now"; it will take four years before what it is doing "now" can affect us. Alpha Centauri "now" is an idea or concept of our mind; it is not something that is really definable physically at the moment, because we have to wait to observe it; we cannot even define it right "now." Furthermore, the "now" depends on the coordinate system. If, for example, Alpha Centauri were moving, an observer there would not agree with us because he would put his axes at an angle, and his "now" would be a *different* time. We have already talked about the fact that simultaneity is not a unique thing.

There are fortune tellers, or people who tell us they can know the future, and there are many wonderful stories about the man who suddenly discovers that he has knowledge about the affective future. Well, there are lots of paradoxes produced by that because if we know something is going to happen, then we can make

sure we will avoid it by doing the right thing at the right time, and so on. But actually there is no fortune teller who can even tell us the *present*! There is no one who can tell us what is really happening right now, at any reasonable distance, because that is unobservable. We might ask ourselves this question, which we leave to the student to try to answer: Would any paradox be produced if it were suddenly to become possible to know things that are in the space-like intervals of region 1?

17–4 More about four-vectors

Let us now return to our consideration of the analogy of the Lorentz transformation and rotations of the space axes. We have learned the utility of collecting together other quantities which have the same transformation properties as the coordinates, to form what we call *vectors*, directed lines. In the case of ordinary rotations, there are many quantities that transform the same way as x, y, and z under rotation: for example, the velocity has three components, an x, y, and z-component; when seen in a different coordinate system, none of the components is the same, instead they are all transformed to new values. But, somehow or other, the velocity "itself" has a greater reality than do any of its particular components, and we represent it by a directed line.

We therefore ask: Is it or is it not true that there are quantities which transform, or which are related, in a moving system and in a nonmoving system, in the same way as x, y, z, and t? From our experience with vectors, we know that three of the quantities, like x, y, z, would constitute the three components of an ordinary space-vector, but the fourth quantity would look like an ordinary scalar under space rotation, because it does not change so long as we do not go into a moving coordinate system. Is it possible, then, to associate with some of our known "three-vectors" a fourth object, that we could call the "time component," in such a manner that the four objects together would "rotate" the same way as position and time in space-time? We shall now show that there is, indeed, at least one such thing (there are many of them, in fact): *the three components of momentum, and the energy as the time component, transform together* to make what we call a "four-vector." In demonstrating this, since it is quite inconvenient to have to write c's everywhere, we shall use the same trick concerning units of the energy, the mass, and the momentum, that we used in Eq. (17.4). Energy and mass, for example, differ only by a factor c^2 which is merely a question of units, so we can say energy *is* the mass. Instead of having to write the c^2, we put $E = m$, and then, of course, if there were any trouble we would put in the right amounts of c so that the units would straighten out in the last equation, but not in the intermediate ones.

Thus our equations for energy and momentum are

$$\begin{aligned} E &= m = m_0/\sqrt{1 - v^2}, \\ \mathbf{p} &= m\mathbf{v} = m_0\mathbf{v}/\sqrt{1 - v^2}. \end{aligned} \tag{17.6}$$

Also in these units, we have

$$E^2 - p^2 = m_0^2. \tag{17.7}$$

For example, if we measure energy in electron volts, what does a mass of 1 electron volt mean? It means the mass whose rest energy is 1 electron volt, that is, m_0c^2 is one electron volt. For example, the rest mass of an electron is 0.511×10^6 ev.

Now what would the momentum and energy look like in a new coordinate system? To find out, we shall have to transform Eq. (17.6), which we can do because we know how the velocity transforms. Suppose that, as we measure it, an object has a velocity v, but we look upon the same object from the point of view of a space ship which itself is moving with a velocity u, and in that system we use a prime to designate the corresponding thing. In order to simplify things at first, we shall take the case that the velocity v is in the direction of u. (Later, we can do the more general case.) What is v', the velocity as seen from the space ship? It is the

composite velocity, the "difference" between v and u. By the law which we worked out before,

$$v' = \frac{v - u}{1 - uv}. \tag{17.8}$$

Now let us calculate the new energy E', the energy as the fellow in the space ship would see it. He would use the same rest mass, of course, but he would use v' for the velocity. What we have to do is square v', subtract it from one, take the square root, and take the reciprocal:

$$v'^2 = \frac{v^2 - 2uv + u^2}{1 - 2uv + u^2v^2},$$

$$1 - v'^2 = \frac{1 - 2uv + u^2v^2 - v^2 + 2uv - u^2}{1 - 2uv + u^2v^2}$$

$$= \frac{1 - v^2 - u^2 + u^2v^2}{1 - 2uv + u^2v^2}$$

$$= \frac{(1 - v^2)(1 - u^2)}{(1 - uv)^2}.$$

Therefore

$$\frac{1}{\sqrt{1 - v'^2}} = \frac{1 - uv}{\sqrt{1 - v^2}\sqrt{1 - u^2}}. \tag{17.9}$$

The energy E' is then simply m_0 times the above expression. But we want to express the energy in terms of the unprimed energy and momentum, and we note that

$$E' = \frac{m_0 - m_0uv}{\sqrt{1 - v^2}\sqrt{1 - u^2}} = \frac{(m_0/\sqrt{1 - v^2}) - (m_0v/\sqrt{1 - v^2})\,u}{\sqrt{1 - u^2}},$$

or

$$E' = \frac{E - up_x}{\sqrt{1 - u^2}}, \tag{17.10}$$

which we recognize as being exactly of the same form as

$$t' = \frac{t - ux}{\sqrt{1 - u^2}}.$$

Next we must find the new momentum p'_x. This is just the energy E times v', and is also simply expressed in terms of E and p:

$$p'_x = E'v' = \frac{m_0(1 - uv)}{\sqrt{1 - v^2}\sqrt{1 - u^2}} \cdot \frac{v - u}{(1 - uv)} = \frac{m_0v - m_0u}{\sqrt{1 - v^2}\sqrt{1 - u^2}}.$$

Thus

$$p'_x = \frac{p_x - uE}{\sqrt{1 - u^2}}, \tag{17.11}$$

which we recognize as being of precisely the same form as

$$x' = \frac{x - ut}{\sqrt{1 - u^2}}.$$

Thus the transformations for the new energy and momentum in terms of the old energy and momentum are exactly the same as the transformations for t' in terms of t and x, and x' in terms of x and t: all we have to do is, every time we see t in (17.4) substitute E, and every time we see x substitute p_x, and then the equations (17.4) will become the same as Eqs. (17.10) and (17.11). This would imply, if everything works right, an additional rule that $p'_y = p_y$ and that $p'_z = p_z$. To prove this would require our going back and studying the case of motion up and down. Actually, we did study the case of motion up and down in the last

chapter. We analyzed a complicated collision and we noticed that, in fact, the transverse momentum is *not* changed when viewed from a moving system; so we have already verified that $p'_y = p_y$ and $p'_z = p_z$. The complete transformation, then, is

$$\begin{aligned} p'_x &= \frac{p_x - uE}{\sqrt{1-u^2}}, \\ p'_y &= p_y, \\ p'_z &= p_z \\ E' &= \frac{E - up_x}{\sqrt{1-u^2}}. \end{aligned} \tag{17.12}$$

In these transformations, therefore, we have discovered four quantities which transform like x, y, z, and t, and which we call the *four-vector momentum*. Since the momentum is a four-vector, it can be represented on a space-time diagram of a moving particle as an "arrow" tangent to the path, as shown in Fig. 17–4. This arrow has a time component equal to the energy, and its space components represent its three-vector momentum; this arrow is more "real" than either the energy or the momentum, because those just depend on how we look at the diagram.

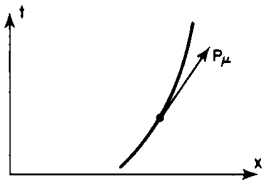


Fig. 17–4. The four-vector momentum of a particle.

17–5 Four-vector algebra

The notation for four-vectors is different than it is for three-vectors. In the case of three-vectors, if we were to talk about the ordinary three-vector momentum we would write it **p**. If we wanted to be more specific, we could say it has three components which are, for the axes in question, p_x, p_y, and p_z, or we could simply refer to a general component as p_i, and say that i could either be x, y, or z, and that these are the three components; that is, imagine that i is any one of three directions, x, y, or z. The notation that we use for four-vectors is analogous to this: we write p_μ for the four-vector, and μ stands for the *four* possible directions t, x, y, or z.

We could, of course, use any notation we want; do not laugh at notations; invent them, they are powerful. In fact, mathematics is, to a large extent, invention of better notations. The whole idea of a four-vector, in fact, is an improvement in notation so that the transformations can be remembered easily. A_μ, then, is a general four-vector, but for the special case of momentum, the p_t is identified as the energy, p_x is the momentum in the x-direction, p_y is that in the y-direction, and p_z is that in the z-direction. To add four-vectors, we add the corresponding components.

If there is an equation among four-vectors, then the equation is true for *each component*. For instance, if the law of conservation of three-vector momentum is to be true in particle collisions, i.e., if the sum of the momenta for a large number of interacting or colliding particles is to be a constant, that must mean that the sums of all momenta in the x-direction, in the y-direction, and in the z-direction, for all the particles, must each be constant. This law alone would be impossible in relativity because it is *incomplete;* it is like talking about only two of the components of a three-vector. It is incomplete because if we rotate the axes, we mix the various components, so we must include all three components in our law. Thus, in relativity, we must complete the law of conservation of momentum by extending it to include the *time* component. This is *absolutely necessary* to go with the other three, or there cannot be relativistic invariance. The *conservation of energy* is the fourth equation which goes with the conservation of momentum to make a valid four-vector relationship in the geometry of space and time. Thus the law of conservation of energy and momentum in four-dimensional notation is

$$\sum_{\substack{\text{particles}\\ \text{in}}} p_\mu = \sum_{\substack{\text{particles}\\ \text{out}}} p_\mu \tag{17.13}$$

or, in a slightly different notation

$$\sum_i p_{i\mu} = \sum_j p_{j\mu}, \tag{17.14}$$

where $i = 1, 2, \ldots$ refers to the particles going into the collision, $j = 1, 2, \ldots$ refers to the particles coming out of the collision, and $\mu = x, y, z,$ or t. You say, "In which axes?" It makes no difference. The law is true for each component, using *any* axes.

In vector analysis we discussed one other thing, the dot product of two vectors. Let us now consider the corresponding thing in space-time. In ordinary rotation we discovered there was an unchanged quantity $x^2 + y^2 + z^2$. In four dimensions, we find that the corresponding quantity is $t^2 - x^2 - y^2 - z^2$ (Eq. 17.3). How can we write that? One way would be to write some kind of four-dimensional thing with a square dot between, like $A_\mu \diamondsuit B_\mu$; one of the notations which is actually used is

$$\sum_\mu{}' A_\mu A_\mu = A_t^2 - A_x^2 - A_y^2 - A_z^2. \tag{17.15}$$

The prime on $\sum$ means that the first term, the "time" term, is positive, but the other three terms have minus signs. This quantity, then, will be the same in any coordinate system, and we may call it the square of the length of the four-vector. For instance, what is the square of the length of the four-vector momentum of a single particle? This will be equal to $p_t^2 - p_x^2 - p_y^2 - p_z^2$ or, in other words, $E^2 - p^2$, because we know that p_t is E. What is $E^2 - p^2$? It must be something which is the same in every coordinate system. In particular, it must be the same for a coordinate system which is moving right along with the particle, in which the particle is standing still. If the particle is standing still, it would have no momentum. So in that coordinate system, it is purely its energy, which is the same as its rest mass. Thus $E^2 - p^2 = m_0^2$. So we see that the square of the length of this vector, the four-vector momentum, is equal to m_0^2.

From the square of a vector, we can go on to invent the "dot product," or the product which is a scalar: if a_μ is one four-vector and b_μ is another four-vector, then the scalar product is

$$\sum{}' a_\mu b_\mu = a_t b_t - a_x b_x - a_y b_y - a_z b_z. \tag{17.16}$$

It is the same in all coordinate systems.

Finally, we shall mention certain things whose rest mass m_0 is zero. A photon of light, for example. A photon is like a particle, in that it carries an energy and a momentum. The energy of a photon is a certain constant, called Planck's constant, times the frequency of the photon: $E = h\nu$. Such a photon also carries a momentum, and the momentum of a photon (or of any other particle, in fact) is h divided by the wavelength: $p = h/\lambda$. But, for a photon, there is a definite relationship between the frequency and the wavelength: $\nu = c/\lambda$. (The number of waves per second, times the wavelength of each, is the distance that the light goes in one second, which, of course, is c.) Thus we see immediately that the energy of a photon must be the momentum times c, or if $c = 1$, *the energy and momentum are equal.* That is to say, the rest mass is zero. Let us look at that again; that is quite curious. If it is a particle of zero rest mass, what happens when it stops? *It never stops!* It always goes at the speed c. The usual formula for energy is $m_0/\sqrt{1 - v^2}$. Now can we say that $m_0 = 0$ and $v = 1$, so the energy is 0? We *cannot* say that it is zero; the photon really can (and does) have energy even though it has no rest mass, but this it possesses by perpetually going at the speed of light!

We also know that the momentum of any particle is equal to its total energy times its velocity: if $c = 1$, $p = vE$ or, in ordinary units, $p = vE/c^2$. For any particle moving at the speed of light, $p = E$ if $c = 1$. The formulas for the energy of a photon as seen from a moving system are, of course, given by Eq. (17.12), but for the momentum we must substitute the energy times c (or times 1 in this case). The different energies after transformation means that there are different frequencies. This is called the Doppler effect, and one can calculate it easily from Eq. (17.12), using also $E = p$ and $E = h\nu$.

As Minkowski said, "Space of itself, and time of itself will sink into mere shadows, and only a kind of union between them shall survive."

18

Rotation in Two Dimensions

18–1 The center of mass

18–2 Rotation of a rigid body

18–3 Angular momentum

18–4 Conservation of angular momentum

18–1 The center of mass

In the previous chapters we have been studying the mechanics of points, or small particles whose internal structure does not concern us. For the next few chapters we shall study the application of Newton's laws to more complicated things. When the world becomes more complicated, it also becomes more interesting, and we shall find that the phenomena associated with the mechanics of a more complex object than just a point are really quite striking. Of course these phenomena involve nothing but combinations of Newton's laws, but it is sometimes hard to believe that only $F = ma$ is at work.

The more complicated objects we deal with can be of several kinds: water flowing, galaxies whirling, and so on. The simplest "complicated" object to analyze, at the start, is what we call a *rigid body*, a solid object that is turning as it moves about. However, even such a simple object may have a most complex motion, and we shall therefore first consider the simplest aspects of such motion, in which an extended body rotates about a *fixed axis*. A given point on such a body then moves in a plane perpendicular to this axis. Such rotation of a body about a fixed axis is called *plane rotation* or rotation in two dimensions. We shall later generalize the results to three dimensions, but in doing so we shall find that, unlike the case of ordinary particle mechanics, rotations are subtle and hard to understand unless we first get a solid grounding in two dimensions.

The first interesting theorem concerning the motion of complicated objects can be observed at work if we throw an object made of a lot of blocks and spokes, held together by strings, into the air. Of course we know it goes in a parabola, because we studied that for a particle. But now our object is *not* a particle; it wobbles and it jiggles, and so on. It does go in a parabola though; one can see that. *What* goes in a parabola? Certainly not the point on the corner of the block, because that is jiggling about; neither is it the end of the wooden stick, or the middle of the wooden stick, or the middle of the block. But *something* goes in a parabola, there is an effective "center" which moves in a parabola. So our first theorem about complicated objects is to demonstrate that there *is* a mean position which is mathematically definable, but not necessarily a point of the material itself, which goes in a parabola. That is called the theorem of the center of the mass, and the proof of it is as follows.

We may consider any object as being made of lots of little particles, the atoms, with various forces among them. Let i represent an index which defines one of the particles. (There are millions of them, so i goes to 10^{23}, or something.) Then the force on the ith particle is, of course, the mass times the acceleration of that particle:

$$\mathbf{F}_i = \mathrm{m}_i(d^2\mathbf{r}_i/dt^2). \tag{18.1}$$

In the next few chapters our moving objects will be ones in which all the parts are moving at speeds very much slower than the speed of light, and we shall use the nonrelativistic approximation for all quantities. In these circumstances the mass is constant, so that

$$\mathbf{F}_i = d^2(m_i\mathbf{r}_i)/dt^2 \tag{18.2}$$

If we now add the force on all the particles, that is, if we take the sum of all the $\mathbf{F}_i$'s for all the different indexes, we get the total force, $\mathbf{F}$. On the other side of the

equation, we get the same thing as though we added before the differentiation:

$$\sum_i \mathbf{F}_i = \mathbf{F} = \frac{d^2(\sum_i m_i \mathbf{r}_i)}{dt^2}. \tag{18.3}$$

Therefore the total force is the second derivative of the masses times their positions, added together.

Now the total force on all the particles is the same as the *external* force. Why? Although there are all kinds of forces on the particles because of the strings, the wigglings, the pullings and pushings, and the atomic forces, and who knows what, and we have to add all these together, we are rescued by Newton's Third Law. Between any two particles the action and reaction are equal, so that when we add all the equations together, if any two particles have forces between them it cancels out in the sum; therefore the net result is only those forces which arise from other particles which are not included in whatever object we decide to sum over. So if Eq. (18.3) is the sum over a certain number of the particles, which together are called "the object," then the *external* force on the total object is equal to the sum of *all* the forces on all its constituent particles.

Now it would be nice if we could write Eq. (18.3) as the total mass times some acceleration. We can. Let us say M is the sum of all the masses, i.e., the total mass. Then if we *define* a certain vector $\mathbf{R}$ to be

$$\mathbf{R} = \sum_i m_i \mathbf{r}_i / M, \tag{18.4}$$

Eq. (18.3) will be simply

$$\mathbf{F} = d^2(M\mathbf{R})/dt^2 = M(d^2\mathbf{R}/dt^2), \tag{18.5}$$

since M is a constant. Thus we find that the external force is the total mass times the acceleration of an imaginary point whose location is $\mathbf{R}$. This point is called the *center of mass* of the body. It is a point somewhere in the "middle" of the object, a kind of average $\mathbf{r}$ in which the different $\mathbf{r}_i$'s have weights or importances proportional to the masses.

We shall discuss this important theorem in more detail in a later chapter, and we shall therefore limit our remarks to two points: First, if the external forces are zero, if the object were floating in empty space, it might whirl, and jiggle, and twist, and do all kinds of things. But the *center of mass*, this artificially invented, calculated position, somewhere in the middle, *will move with a constant velocity*. In particular, if it is initially at rest, it will stay at rest. So if we have some kind of a box, perhaps a space ship, with people in it, and we calculate the location of the center of mass and find it is standing still, then the center of mass will continue to stand still if no external forces are acting on the box. Of course, the space ship may move a little in space, but that is because the people are walking back and forth inside; when one walks toward the front, the ship goes toward the back so as to keep the average position of all the masses in exactly the same place.

Is rocket propulsion therefore absolutely impossible because one cannot move the center of mass? No; but of course we find that to propel an interesting part of the rocket, an uninteresting part must be thrown away. In other words, if we start with a rocket at zero velocity and we spit some gas out the back end, then this little blob of gas goes one way as the rocket ship goes the other, but the center of mass is still exactly where it was before. So we simply move the part that we are interested in against the part we are not interested in.

The second point concerning the center of mass, which is the reason we introduced it into our discussion at this time, is that it may be treated separately from the "internal" motions of an object, and may therefore be ignored in our discussion of rotation.

18–2 Rotation of a rigid body

Now let us discuss rotations. Of course an ordinary object does not simply rotate, it wobbles, shakes, and bends, so to simplify matters we shall discuss the motion of a nonexistent ideal object which we call a *rigid body*. This means an

object in which the forces between the atoms are so strong, and of such character, that the little forces that are needed to move it do not bend it. Its shape stays essentially the same as it moves about. If we wish to study the motion of such a body, and agree to ignore the motion of its center of mass, there is only one thing left for it to do, and that is to *turn*. We have to describe that. How? Suppose there is some line in the body which stays put (perhaps it includes the center of mass and perhaps not), and the body is rotating about this particular line as an axis. How do we define the rotation? That is easy enough, for if we mark a point somewhere on the object, anywhere except on the axis, we can always tell exactly where the object is, if we only know where this point has gone to. The only thing needed to describe the position of that point is an *angle*. So rotation consists of a study of the variations of the angle with time.

In order to study rotation, we observe the angle through which a body has turned. Of course, we are not referring to any particular angle *inside* the object itself; it is not that we draw some angle *on* the object. We are talking about the *angular change of the position* of the whole thing, from one time to another.

First, let us study the kinematics of rotations. The angle will change with time, and just as we talked about position and velocity in one dimension, we may talk about angular position and angular velocity in plane rotation. In fact, there is a very interesting relationship between rotation in two dimensions and one-dimensional displacement, in which almost every quantity has its analog. First, we have the angle θ which defines how far the body has gone *around*; this replaces the distance y, which defines how far it has gone *along*. In the same manner, we have a velocity of turning, $\omega = d\theta/dt$, which tells us how much the angle changes in a second, just as $v = ds/dt$ describes how fast a thing moves, or how far it moves in a second. If the angle is measured in radians, then the angular velocity ω will be so and so many radians per second. The greater the angular velocity, the faster the object is turning, the faster the angle changes. We can go on: we can differentiate the angular velocity with respect to time, and we can call $\alpha = d\omega/dt = d^2\theta/dt^2$ the angular acceleration. That would be the analog of the ordinary acceleration.

Now of course we shall have to relate the dynamics of rotation to the laws of dynamics of the particles of which the object is made, so we must find out how a particular particle moves when the angular velocity is such and such. To do this, let us take a certain particle which is located at a distance r from the axis and say it is in a certain location $P(x, y)$ at a given instant, in the usual manner (Fig. 18–1). If at a moment Δt later the angle of the whole object has turned through $\Delta\theta$, then this particle is carried with it. It is at the same radius away from O as it was before, but is carried to Q. The first thing we would like to know is how much the distance x changes and how much the distance y changes. If OP is called r, then the length PQ is $r\,\Delta\theta$, because of the way angles are defined. The change in x, then, is simply the projection of $r\,\Delta\theta$ in the x-direction:

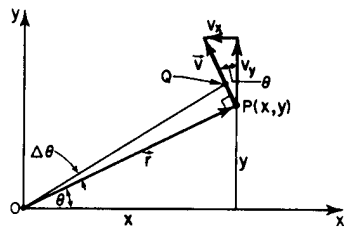


Fig. 18–1. Kinematics of two-dimensional rotation.

$$\Delta x = -PQ \sin\theta = -r\,\Delta\theta \cdot (y/r) = -y\,\Delta\theta. \tag{18.6}$$

Similarly,

$$\Delta y = +x\,\Delta\theta. \tag{18.7}$$

If the object is turning with a given angular velocity ω, we find, by dividing both sides of (18.6) and (18.7) by Δt, that the velocity of the particle is

$$v_x = -\omega y \qquad \text{and} \qquad v_y = +\omega x. \tag{18.8}$$

Of course if we want to find the magnitude of the velocity, we just write

$$v = \sqrt{v_x^2 + v_y^2} = \sqrt{\omega^2 y^2 + \omega^2 x^2} = \omega\sqrt{x^2 + y^2} = \omega r. \tag{18.9}$$

It should not be mysterious that the value of the magnitude of this velocity is ωr; in fact, it should be self-evident, because the distance that it moves is $r\,\Delta\theta$ and the distance it moves per second is $r\,\Delta\theta/\Delta t$, or $r\omega$.

Let us now move on to consider the *dynamics* of rotation. Here a new concept, *force*, must be introduced. Let us inquire whether we can invent something which we shall call the *torque* (L. *torquere*, to twist) which bears the same relationship to rotation as force does to linear movement. A force is the thing that is needed to make linear motion, and the thing that makes something rotate is a "rotary force" or a "twisting force," i.e., a torque. Qualitatively, a torque is a "twist"; what is a torque quantitatively? We shall get to the theory of torques quantitatively by studying the *work* done in turning an object, for one very nice way of defining a force is to say how much work it does when it acts through a given displacement. We are going to try to maintain the analogy between linear and angular quantities by equating the work that we do when we turn something a little bit when there are forces acting on it, to the *torque* times the *angle* it turns through. In other words, the definition of the torque is going to be so arranged that the theorem of work has an absolute analog: force times distance is work, and torque times angle is going to be work. That tells us what torque is. Consider, for instance, a rigid body of some kind with various forces acting on it, and an axis about which the body rotates. Let us at first concentrate on one force and suppose that this force is applied at a certain point (x, y). How much work would be done if we were to turn the object through a very small angle? That is easy. The work done is

$$\Delta W = F_x \Delta x + F_y \Delta y. \tag{18.10}$$

We need only to substitute Eqs. (18.6) and (18.7) for Δx and Δy to obtain

$$\Delta W = (xF_y - yF_x)\,\Delta\theta. \tag{18.11}$$

That is, the amount of work that we have done is, in fact, equal to the angle through which we have turned the object, multiplied by a strange-looking combination of the force and the distance. This "strange combination" is what we call the torque. So, defining the change in work as the torque times the angle, we now have the formula for torque in terms of the forces. (Obviously, torque is not a completely new idea independent of Newtonian mechanics—torque must have a definite definition in terms of the force.)

When there are several forces acting, the work that is done is, of course, the sum of the works done by all the forces, so that ΔW will be a whole lot of terms, all added together, for all the forces, *each of which is proportional, however, to* $\Delta\theta$. We can take the $\Delta\theta$ outside and therefore can say that the change in the work is equal to the sum of all the torques due to all the different forces that are acting, times $\Delta\theta$. This sum we might call the total torque, τ. Thus torques add by the ordinary laws of algebra, but we shall later see that this is only because we are working in a plane. It is like one-dimensional kinematics, where the forces simply add algebraically, but only because they are all in the same direction. It is more complicated in three dimensions. Thus, for two-dimensional rotation,

$$\tau_i = x_i F_{yi} - y_i F_{xi} \tag{18.12}$$

and

$$\tau = \sum \tau_i. \tag{18.13}$$

It must be emphasized that the torque is about a given axis. If a different axis is chosen, so that all the x_i and y_i are changed, the value of the torque is (usually) changed too.

Now we pause briefly to note that our foregoing introduction of torque, through the idea of work, gives us a most important result for an object in equilibrium: if all the forces on an object are in balance both for translation and rotation, not only is the net *force* zero, but the total of all the *torques* is also zero, because if an object is in equilibrium, *no work is done by the forces for a small displacement.* Therefore, since $\Delta W = \tau\,\Delta\theta = 0$, the sum of all the torques must be zero. So there are two conditions for equilibrium: that the sum of the forces is zero, and that the sum of the torques is zero. Prove that it suffices to be sure that the sum of torques about any *one* axis (in two dimensions) is zero.

Now let us consider a single force, and try to figure out, geometrically, what this strange thing $xF_y - yF_x$ amounts to. In Fig. 18–2 we see a force **F** acting at a point **r**. When the object has rotated through a small angle $\Delta\theta$, the work done, of course, is the component of force in the direction of the displacement times the displacement. In other words, it is only the tangential component of the force that counts, and this must be multiplied by the distance $r\,\Delta\theta$. Therefore we see that the torque is also equal to the tangential component of force (perpendicular to the radius) times the radius. That makes sense in terms of our ordinary idea of the torque, because if the force were completely radial, it would not put any "twist" on the body; it is evident that the twisting effect should involve only the part of the force which is not pulling out from the center, and that means the tangential component. Furthermore, it is clear that a given force is more effective on a long arm than near the axis. In fact, if we take the case where we push right *on* the axis, we are not twisting at all! So it makes sense that the amount of twist, or torque, is proportional both to the radial distance and to the tangential component of the force.

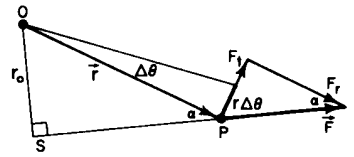


Fig. 18–2. The torque produced by a force.

There is still a third formula for the torque which is very interesting. We have just seen that the torque is the force times the radius times the sine of the angle α, in Fig. 18–2. But if we extend the line of action of the force and draw the line OS, the perpendicular distance to the line of action of the force (the *lever arm* of the force) we notice that this lever arm is shorter than r in just the same proportion as the tangential part of the force is less than the total force. Therefore the formula for the torque can also be written as the magnitude of the force times the length of the lever arm.

The torque is also often called the *moment* of the force. The origin of this term is obscure, but it may be related to the fact that "moment" is derived from the Latin *movimentum*, and that the capability of a force to move an object (using the force on a lever or crowbar) increases with the length of the lever arm. In mathematics 'moment" means weighted by how far away it is from an axis.

18–3 Angular momentum

Although we have so far considered only the special case of a rigid body, the properties of torques and their mathematical relationships are interesting also even when an object is not rigid. In fact, we can prove a very remarkable theorem: just as external force is the rate of change of a quantity p, which we call the total momentum of a collection of particles, so the external torque is the rate of change of a quantity L which we call the *angular momentum* of the group of particles.

To prove this, we shall suppose that there is a system of particles on which there are some forces acting and find out what happens to the system as a result of the torques due to these forces. First, of course, we should consider just *one* particle. In Fig. 18–3 is one particle of mass m, and an axis O; the particle is not necessarily rotating in a circle about O, it may be moving in an ellipse, like a planet going around the sun, or in some other curve. It is moving somehow, and there are forces on it, and it accelerates according to the usual formula that the x-component of force is the mass times the x-component of acceleration, etc. But let us see what the *torque* does. The torque equals $xF_y - yF_x$, and the force in the x- or y-direction is the mass times the acceleration in the x- or y-direction:

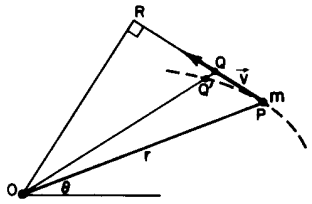


Fig. 18–3. A particle moves about an axis O.

$$\begin{aligned}\tau &= xF_y - yF_x \\ &= xm(d^2y/dt^2) - ym(d^2x/dt^2).\end{aligned} \tag{18.14}$$

Now, although this does not appear to be the derivative of any simple quantity, it is in fact the derivative of the quantity $xm(dy/dt) - ym(dx/dt)$:

$$\begin{aligned}\frac{d}{dt}\left[xm\left(\frac{dy}{dt}\right) - ym\left(\frac{dx}{dt}\right)\right] &= xm\left(\frac{d^2y}{dt^2}\right) + \left(\frac{dx}{dt}\right)m\left(\frac{dy}{dt}\right) \\ - ym\left(\frac{d^2x}{dt^2}\right) - \left(\frac{dy}{dt}\right)m\left(\frac{dx}{dt}\right) &= xm\left(\frac{d^2y}{dt^2}\right) - ym\left(\frac{d^2x}{dt^2}\right).\end{aligned} \tag{18.15}$$

So it is true that the torque is the rate of change of something with time! So we pay attention to the "something," we give it a name: we call it L, the angular momentum:

$$\begin{aligned} L &= xm(dy/dt) - ym(dx/dt) \\ &= xp_y - yp_x. \end{aligned} \tag{18.16}$$

Although our present discussion is nonrelativistic, the second form for L given above is relativistically correct. So we have found that there is also a rotational analog for the momentum, and that this analog, the angular momentum, is given by an expression in terms of the components of linear momentum that is just like the formula for torque in terms of the force components! Thus, if we want to know the angular momentum of a particle about an axis, we take only the component of the momentum that is tangential, and multiply it by the radius. In other words, what counts for angular momentum is not how fast it is going *away* from the origin, but how much it is going *around* the origin. Only the tangential part of the momentum counts for angular momentum. Furthermore, the farther out the line of the momentum extends, the greater the angular momentum. And also, because the geometrical facts are the same whether the quantity is labeled p or F, it is true that there is a lever arm (*not* the same as the lever arm of the force on the particle!) which is obtained by extending the line of the *momentum* and finding the perpendicular distance to the axis. Thus the angular momentum is the magnitude of the momentum times the momentum lever arm. So we have three formulas for angular momentum, just as we have three formulas for the torque:

$$\begin{aligned} L &= xp_y - yp_x \\ &= rp_{\text{tang}} \\ &= p \cdot \text{lever arm}. \end{aligned} \tag{18.17}$$

Like torque, angular momentum depends upon the position of the axis about which it is to be calculated.

Before proceeding to a treatment of more than one particle, let us apply the above results to a planet going around the sun. In which direction is the force? The force is toward the sun. What, then, is the torque on the object? Of course, this depends upon where we take the axis, but we get a very simple result if we take it at the sun itself, for the torque is the force times the lever arm, or the component of force perpendicular to r, times r. But there is no tangential force, so there is no torque about an axis at the sun! Therefore, the angular momentum of the planet going around the sun must remain constant. Let us see what that means. The tangential component of velocity, times the mass, times the radius, will be constant, because that is the angular momentum, and the rate of change of the angular momentum is the torque, and, in this problem, the torque is zero. Of course since the mass is also a constant, this means that the tangential velocity times the radius is a constant. But this is something we already knew for the motion of a planet. Suppose we consider a small amount of time Δt. How far will the planet move when it moves from P to Q (Fig. 18–3)? How much *area* will it sweep through? Disregarding the very tiny area $QQ'P$ compared with the much larger area OPQ, it is simply half the base PQ times the height, OR. In other words, the area that is swept through in unit time will be equal to the velocity times the lever arm of the velocity (times one-half). Thus the rate of change of area is proportional to the angular momentum, which is constant. So Kepler's law about equal areas in equal times is a word description of the statement of the law of conservation of angular momentum, when there is no torque produced by the force.

18–4 Conservation of angular momentum

Now we shall go on to consider what happens when there is a large number of particles, when an object is made of many pieces with many forces acting between them and on them from the outside. Of course, we already know that, about any given fixed axis, the torque on the ith particle (which is the force on the ith particle

times the lever arm of that force) is equal to the rate of change of the angular momentum of that particle, and that the angular momentum of the ith particle is its momentum times its momentum lever arm. Now suppose we add the torques τ_i for all the particles and call it the total torque τ. Then this will be the rate of change of the sum of the angular momenta of all the particles L_i, and that defines a new quantity which we call the total angular momentum L. Just as the total *momentum* of an object is the sum of the momenta of all the parts, so the angular momentum is the sum of the angular momenta of all the parts. Then the rate of change of the total L is the total torque:

$$\tau = \sum \tau_i = \sum \frac{dL_i}{dt} = \frac{dL}{dt}. \tag{18.18}$$

Now it might seem that the total torque is a complicated thing. There are all those internal forces and all the outside forces to be considered. But, if we take Newton's law of action and reaction to say, not simply that the action and reaction are equal, but also that they are *directed exactly oppositely along the same line* (Newton may or may not actually have said this, but he tacitly assumed it), then the two *torques* on the reacting objects, due to their mutual interaction, will be equal and opposite because the lever arms for any axis are equal. Therefore the internal torques balance out pair by pair, and so we have the remarkable theorem that the *rate of change of the total angular momentum about any axis is equal to the external torque about that axis!*

$$\tau = \sum \tau_i = \tau_{\text{ext}} = dL/dt. \tag{18.19}$$

Thus we have a very powerful theorem concerning the motion of large collections of particles, which permits us to study the over-all motion without having to look at the detailed machinery inside. This theorem is true for any collection of objects, whether they form a rigid body or not.

One extremely important case of the above theorem is the *law of conservation of angular momentum:* if no external torques act upon a system of particles, the angular momentum remains constant.

A special case of great importance is that of a rigid body, that is, an object of a definite shape that is just turning around. Consider an object that is fixed in its geometrical dimensions, and which is rotating about a fixed axis. Various parts of the object bear the same relationship to one another at all times. Now let us try to find the total angular momentum of this object. If the mass of one of its particles is m_i, and its position or location is at (x_i, y_i), then the problem is to find the angular momentum of that particle, because the total angular momentum is the sum of the angular momenta of all such particles in the body. For an object going around in a circle, the angular momentum, of course, is the mass times the velocity times the distance from the axis, and the velocity is equal to the angular velocity times the distance from the axis:

$$L_i = m_i v_i r_i = m_i r_i^2 \omega, \tag{18.20}$$

or, summing over all the particles i, we get

$$L = I\omega, \tag{18.21}$$

where

$$I = \sum_i m_i r_i^2. \tag{18.22}$$

This is the analog of the law that the momentum is mass times velocity. Velocity is replaced by angular velocity, and we see that the mass is replaced by a new thing which we call the *moment of inertia I*, which is analogous to the mass. Equations (18.21) and (18.22) say that a body has inertia for turning which depends, not just on the masses, but on *how far away they are* from the axis. So, if we have two objects of the same mass, when we put the masses farther away from the axis, the inertia for turning will be higher. This is easily demonstrated by the apparatus

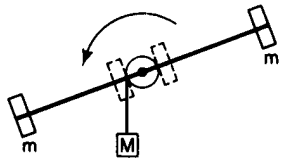


Fig. 18–4. The "inertia for turning" depends upon the lever arm of the masses.

shown in Fig. 18–4, where a weight M is kept from falling very fast because it has to turn the large weighted rod. At first, the masses m are close to the axis, and M speeds up at a certain rate. But when we change the moment of inertia by putting the two masses m much farther away from the axis, then we see that M accelerates much less rapidly than it did before, because the body has much more inertia against turning. The moment of inertia is the inertia against turning, and is the sum of the contributions of all the masses, times their distances *squared*, from the axis.

There is one important difference between mass and moment of inertia which is very dramatic. The mass of an object never changes, but its moment of inertia *can* be changed. If we stand on a frictionless rotatable stand with our arms outstretched, and hold some weights in our hands as we rotate slowly, we may change our moment of inertia by drawing our arms in, but our mass does not change. When we do this, all kinds of wonderful things happen, because of the law of the conservation of angular momentum: If the external torque is zero, then the angular momentum, the moment of inertia times omega, remains constant. Initially, we were rotating with a large moment of inertia I_1 at a low angular velocity ω_1, and the angular momentum was $I_1\omega_1$. Then we changed our moment of inertia by pulling our arms in, say to a smaller value I_2. Then the product $I\omega$, which has to stay the same because the total angular momentum has to stay the same, was $I_2\omega_2$. So $I_1\omega_1 = I_2\omega_2$. That is, if we *reduce* the moment of inertia, we have to *increase* the angular velocity.

19

Center of Mass; Moment of Inertia

19–1 Properties of the center of mass

19–2 Locating the center of mass

19–3 Finding the moment of inertia

19–4 Rotational kinetic energy

19–1 Properties of the center of mass

In the previous chapter we found that if a great many forces are acting on a complicated mass of particles, whether the particles comprise a rigid or a nonrigid body, or a cloud of stars, or anything else, and we find the sum of all the forces (that is, of course, the external forces, because the internal forces balance out), then if we consider the body as a whole, and say it has a total mass M, there is a certain point "inside" the body, called the *center of mass*, such that the net resulting external force produces an acceleration of this point, just as though the whole mass were concentrated there. Let us now discuss the center of mass in a little more detail.

The location of the center of mass (abbreviated CM) is given by the equation

$$\mathbf{R}_{\text{CM}} = \frac{\sum m_i \mathbf{r}_i}{\sum m_i}. \tag{19.1}$$

This is, of course, a vector equation which is really three equations, one for each of the three directions. We shall consider only the x-direction, because if we can understand that one, we can understand the other two. What does $X_{\text{CM}} = \sum m_i x_i / \sum m_i$ mean? Suppose for a moment that the object is divided into little pieces, all of which have the same mass m; then the total mass is simply the number N of pieces times the mass of one piece, say one gram, or any unit. Then this equation simply says that we add all the x's, and then divide by the number of things that we have added: $X_{\text{CM}} = m\sum x_i / mN = \sum x_i / N$. In other words, X_{CM} is the average of all the x's, if the masses are equal. But suppose one of them were twice as heavy as the others. Then in the sum, that x would come in twice. This is easy to understand, for we can think of this double mass as being split into two equal ones, just like the others; then in taking the average, of course, we have to count that x twice because there are two masses there. Thus X is the average position, in the x-direction, of all the masses, every mass being counted a number of times proportional to the mass, as though it were divided into "little grams." From this it is easy to prove that X must be somewhere between the largest and the smallest x, and, therefore lies inside the envelope including the entire body. It does not have to be in the *material* of the body, for the body could be a circle, like a hoop, and the center of mass is in the center of the hoop, not in the hoop itself.

Of course, if an object is symmetrical in some way, for instance, a rectangle, so that it has a plane of symmetry, the center of mass lies somewhere on the plane of symmetry. In the case of a rectangle there are two planes, and that locates it uniquely. But if it is just any symmetrical object, then the center of gravity lies somewhere on the axis of symmetry, because in those circumstances there are as many positive as negative x's.

Another interesting proposition is the following very curious one. Suppose that we imagine an object to be made of two pieces, A and B (Fig. 19–1). Then the center of mass of the whole object can be calculated as follows. First, find the center of mass of piece A, and then of piece B. Also, find the total mass of each piece, M_A and M_B. Then consider a new problem, in which a *point* mass M_A is at the center of mass of object A, and another *point* mass M_B is at the center of mass of object B. The center of mass of these two point masses is then the center of mass of the whole object. In other words, if the centers of mass of various parts

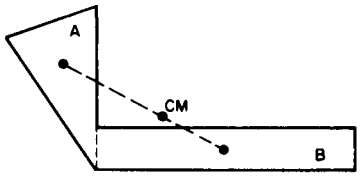


Fig. 19–1. The CM of a compound body lies on the line joining the CM's of the two composite parts.

of an object have been worked out, we do not have to start all over again to find the center of mass of the whole object; we just have to put the pieces together, treating each one as a point mass situated at the center of mass of that piece. Let us see why that is. Suppose that we wanted to calculate the center of mass of a complete object, some of whose particles are considered to be members of object A and some members of object B. The total sum $\sum m_i x_i$ can then be split into two pieces—the sum $\sum_A m_i x_i$ for the A object only, and the sum $\sum_B m_i x_i$ for object B only. Now if we were computing the center of mass of object A alone, we would have exactly the first of these sums, and we know that this by itself is $M_A X_A$, the total mass of all the particles in A times the position of the center of mass of A, because that is the theorem of the center of mass, applied to object A. In the same manner, just by looking at object B, we get $M_B X_B$, and of course, adding the two yields MX:

$$\begin{aligned} MX_{\text{CM}} &= \sum_A m_i x_i + \sum_B m_i x_i \\ &= M_A X_A + M_B X_B. \end{aligned} \tag{19.2}$$

Now since M is evidently the sum of M_A and M_B, we see that Eq. (19.2) can be interpreted as a special example of the center of mass formula for two point objects, one of mass M_A located at X_A and the other of mass M_B located at X_B.

The theorem concerning the motion of the center of mass is very interesting, and has played an important part in the development of our understanding of physics. Suppose we assume that Newton's law is right for the small component parts of a much larger object. Then this theorem shows that Newton's law is also correct for the larger object, even if we do not study the details of the object, but only the total force acting on it and its mass. In other words, Newton's law has the peculiar property that if it is right on a certain small scale, then it will be right on a larger scale. If we do not consider a baseball as a tremendously complex thing, made of myriads of interacting particles, but study only the motion of the center of mass and the external forces on the ball, we find $\mathbf{F} = m\mathbf{a}$, where $\mathbf{F}$ is the external force on the baseball, m is its mass, and $\mathbf{a}$ is the acceleration of its center of mass. So $\mathbf{F} = m\mathbf{a}$ is a law which reproduces itself on a larger scale. (There ought to be a good word, out of the Greek, perhaps, to describe a law which reproduces the same law on a larger scale.)

Of course, one might suspect that the first laws that would be discovered by human beings would be those that would reproduce themselves on a larger scale. Why? Because the actual scale of the fundamental gears and wheels of the universe are of atomic dimensions, which are so much finer than our observations that we are nowhere near that scale in our ordinary observations. So the first things that we would discover must be true for objects of no special size relative to an atomic scale. If the laws for small particles did not reproduce themselves on a larger scale, we would not discover those laws very easily. What about the reverse problem? Must the laws on a small scale be the same as those on a larger scale? Of course it is not necessarily so in nature, that at an atomic level the laws have to be the same as on a large scale. Suppose that the true laws of motion of atoms were given by some strange equation which does *not* have the property that when we go to a larger scale we reproduce the same law, but instead has the property that if we go to a larger scale, we can *approximate it by a certain expression* such that, if we extend that expression up and up, *it* keeps reproducing itself on a larger and larger scale. That is possible, and in fact that is the way it works. Newton's laws are the "tail end" of the atomic laws, extrapolated to a very large size. The actual laws of motion of particles on a fine scale are very peculiar, but if we take large numbers of them and compound them, they approximate, but *only* approximate, Newton's laws. Newton's laws then permit us to go on to a higher and higher scale, and it still seems to be the same law. In fact, it becomes more and more accurate as the scale gets larger and larger. This self-reproducing factor of Newton's laws is thus really not a fundamental feature of nature, but is an important historical feature. We would never discover the fundamental laws of the atomic particles at first observation because the first observations are much too crude. In fact, it turns

out that the fundamental atomic laws, which we call quantum mechanics, are quite different from Newton's laws, and are difficult to understand because all our direct experiences are with large-scale objects and the small-scale atoms behave like nothing we see on a large scale. So we cannot say, "An atom is just like a planet going around the sun," or anything like that. It is like *nothing* we are familiar with because there is *nothing like it.* As we apply quantum mechanics to larger and larger things, the laws about the behavior of many atoms together do *not* reproduce themselves, but produce *new laws*, which are Newton's laws, which then continue to reproduce themselves from, say, micro-microgram size, which still is billions and billions of atoms, on up to the size of the earth, and above.

Let us now return to the center of mass. The center of mass is sometimes called the center of gravity, for the reason that, in many cases, gravity may be considered uniform. Let us suppose that we have small enough dimensions that the gravitational force is not only proportional to the mass, but is everywhere parallel to some fixed line. Then consider an object in which there are gravitational forces on each of its constituent masses. Let m_i be the mass of one part. Then the gravitational force on that part is m_i times g. Now the question is, where can we apply a single force to balance the gravitational force on the whole thing, so that the entire object, if it is a rigid body, will not turn? The answer is that this force must go through the center of mass, and we show this in the following way. In order that the body will not turn, the torque produced by all the forces must add up to zero, because if there is a torque, there is a change of angular momentum, and thus a rotation. So we must calculate the total of all the torques on all the particles, and see how much torque there is about any given axis; it should be zero if this axis is at the center of mass. Now, measuring x horizontally and y vertically, we know that the torques are the forces in the y-direction, times the lever arm x (that is to say, the force times the lever arm around which we want to measure the torque). Now the total torque is the sum

$$\tau = \sum m_i g x_i = g \sum m_i x_i, \tag{19.3}$$

so if the total torque is to be zero, the sum $\sum m_i x_i$ must be zero. But $\sum m_i x_i = MX$, the total mass times the distance of the center of mass from the axis. Thus the x-distance of the center of mass from the axis is zero.

Of course, we have checked the result only for the x-distance, but if we use the true center of mass the object will balance in any position, because if we turned it 90 degrees, we would have y's instead of x's. In other words, when an object is supported at its center of mass, there is no torque on it because of a parallel gravitational field. In case the object is so large that the nonparallelism of the gravitational forces is significant, then the center where one must apply the balancing force is not simple to describe, and it departs slightly from the center of mass. That is why one must distinguish between the center of mass and the center of gravity. The fact that an object supported exactly at the center of mass will balance in all positions has another interesting consequence. If, instead of gravitation, we have a pseudoforce due to acceleration, we may use exactly the same mathematical procedure to find the position to support it so that there are no torques produced by the inertial force of acceleration. Suppose that the object is held in some manner inside a box, and that the box, and everything contained in it, is accelerating. We know that, from the point of view of someone at rest relative to this accelerating box, there will be an effective force due to inertia. That is, to make the object go along with the box, we have to push on it to accelerate it, and this force is "balanced" by the "force of inertia," which is a pseudoforce equal to the mass times the acceleration of the box. To the man in the box, this is the same situation as if the object were in a uniform gravitational field whose "g" value is equal to the acceleration a. Thus the inertial force due to accelerating an object has no torque about the center of mass.

This fact has a very interesting consequence. In an inertial frame that is not accelerating, the torque is always equal to the rate of change of the angular momentum. However, about an axis through the center of mass of an object which

is accelerating, it is *still true* that the torque is equal to the rate of change of the angular momentum. Even if the center of mass is accelerating, we may still choose one special axis, namely, one passing through the center of mass, such that it will still be true that the torque is equal to the rate of change of angular momentum around that axis. Thus the theorem that torque equals the rate of change of angular momentum is true in two general cases: (1) a fixed axis in inertial space, (2) an axis through the center of mass, even though the object may be accelerating.

19–2 Locating the center of mass

The mathematical techniques for the calculation of centers of mass are in the province of a mathematics course, and such problems provide good exercise in integral calculus. After one has learned calculus, however, and wants to know how to locate centers of mass, it is nice to know certain tricks which can be used to do so. One such trick makes use of what is called the theorem of Pappus. It works like this: if we take any closed area in a plane and generate a solid by moving it through space such that each point is always moved perpendicular to the plane of the area, the resulting solid has a total volume equal to the area of the cross section times the distance that the center of mass moved! Certainly this is true if we move the area in a straight line perpendicular to itself, but if we move it in a circle or in some other curve, then it generates a rather peculiar volume. For a curved path, the outside goes around farther, and the inside goes around less, and these effects balance out. So if we want to locate the center of mass of a plane sheet of uniform density, we can remember that the volume generated by spinning it about an axis is the distance that the center of mass goes around, times the area of the sheet.

For example, if we wish to find the center of mass of a right triangle of base D and height H (Fig. 19–2), we might solve the problem in the following way. Imagine an axis along H, and rotate the triangle about that axis through a full 360 degrees. This generates a cone. The distance that the x-coordinate of the center of mass has moved is $2\pi x$. The area which is being moved is the area of the triangle, $\frac{1}{2}HD$. So the x-distance of the center of mass times the area of the triangle is the volume swept out, which is of course $\pi D^2H/3$. Thus $(2\pi x)(\frac{1}{2}HD) = 1/3\pi D^2H$, or $x = D/3$. In a similar manner, by rotating about the other axis, or by symmetry, we find $y = H/3$. In fact, the center of mass of any uniform triangular area is where the three medians, the lines from the vertices through the centers of the opposite sides, all meet. That point is 1/3 of the way along each median. *Clue:* Slice the triangle up into a lot of little pieces, each parallel to a base. Note that the median line bisects every piece, and therefore the center of mass must lie on this line.

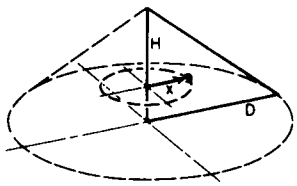


Fig. 19–2. A right triangle and a right circular cone generated by rotating the triangle.

Now let us try a more complicated figure. Suppose that it is desired to find the position of the center of mass of a uniform semicircular disc—a disc sliced in half. Where is the center of mass? For a full disc, it is at the center, of course, but a half-disc is more difficult. Let r be the radius and x be the distance of the center of mass from the straight edge of the disc. Spin it around this edge as axis to generate a sphere. Then the center of mass has gone around $2\pi x$, the area is $\pi r^2/2$ (because it is only half a circle). The volume generated is, of course, $4\pi r^3/3$, from which we find that

$$(2\pi x)(\tfrac{1}{2}\pi r^2) = 4\pi r^3/3,$$

or

$$x = 4r/3\pi.$$

There is another theorem of Pappus which is a special case of the above one, and therefore equally true. Suppose that, instead of the solid semicircular disc, we have a semicircular piece of wire with uniform mass density along the wire, and we want to find its center of mass. In this case there is no mass in the interior, only on the wire. Then it turns out that the *area* which is swept by a plane curved line, when it moves as before, is the distance that the center of mass moves times the *length* of the line. (The line can be thought of as a very narrow area, and the previous theorem can be applied to it.)

19–3 Finding the moment of inertia

Now let us discuss the problem of finding the *moments of inertia* of various objects. The formula for the moment of inertia about the z-axis of an object is

$$I = \sum m_i(x_i^2 + y_i^2)$$

or

$$I = \int (x^2 + y^2)\, dm = \int (x^2 + y^2)\rho\, dv. \tag{19.4}$$

That is, we must sum the masses, each one multiplied by the square of its distance $(x_i^2 + y_i^2)$ from the axis. Note that it is not the three-dimensional distance, only the two-dimensional distance squared, even for a three-dimensional object. For the most part, we shall restrict ourselves to two-dimensional objects, but the formula for rotation about the z-axis is just the same in three dimensions.

As a simple example, consider a rod rotating about a perpendicular axis through one end (Fig. 19–3). Now we must sum all the masses times the x-distances squared (the y's being all zero in this case). What we mean by "the sum," of course, is the integral of x^2 times the little elements of mass. If we divide the rod into small elements of length dx, the corresponding elements of mass are proportional to dx, and if dx were the length of the whole rod the mass would be M. Therefore

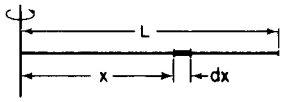


Fig. 19–3. A straight rod of length L rotating about an axis through one end.

$$dm = M\, dx/L$$

and so

$$I = \int_0^L x^2 \frac{M\, dx}{L} = \frac{M}{L}\int_0^L x^2\, dx = \frac{ML^2}{3}. \tag{19.5}$$

The dimensions of moment of inertia are always mass times length squared, so all we really had to work out was the factor 1/3.

Now what is I if the rotation axis is at the center of the rod? We could just do the integral over again, letting x range from $-\frac{1}{2}L$ to $+\frac{1}{2}L$. But let us notice a few things about the moment of inertia. We can imagine the rod as two rods, each of mass $M/2$ and length $L/2$; the moments of inertia of the two small rods are equal, and are both given by the formula (19.5). Therefore the moment of inertia is

$$I = \frac{2(M/2)(L/2)^2}{3} = \frac{ML^2}{12}. \tag{19.6}$$

Thus it is much easier to turn a rod about its center, than to swing it around an end.

Of course, we could go on to compute the moments of inertia of various other bodies of interest. However, while such computations provide a certain amount of important exercise in the calculus, they are not basically of interest to us as such. There is, however, an interesting theorem which is very useful. Suppose we have an object, and we want to find its moment of inertia around some axis. That means we want the inertia needed to carry it by rotation about that axis. Now if we support the object on pivots at the center of mass, so that the object does not turn as it rotates about the axis (because there is no torque on it from inertial effects, and therefore it will not turn when we start moving it), then the forces needed to swing it around are the same as though all the mass were concentrated at the center of mass, and the moment of inertia would be simply $I_1 = MR_{\text{CM}}^2$, where R_{CM} is the distance from the axis to the center of mass. But of course that is not the right formula for the moment of inertia of an object which is really being rotated as it revolves, because not only is the center of it moving in a circle, which would contribute an amount I_1 to the moment of inertia, but also we must turn it about its center of mass. So it is not unreasonable that we must add to I_1 the moment of inertia I_c about the center of mass. So it is a good guess that the total moment of inertia about any axis will be

$$I = I_c + MR_{\text{CM}}^2. \tag{19.7}$$

This theorem is called the *parallel-axis theorem*, and may be easily proved. The moment of inertia about any axis is the mass times the sum of the x_i's and the y_i's, each squared: $I = \sum (x_i^2 + y_i^2)m_i$. We shall concentrate on the x's, but of course the y's work the same way. Now x is the distance of a particular point mass from the origin, but let us consider how it would look if we measured x' from the CM, instead of x from the origin. To get ready for this analysis, we write

$$x_i = x'_i + X_{\text{CM}}.$$

Then we just square this to find

$$x_i^2 = x_i'^2 + 2X_{\text{CM}}x'_i + X_{\text{CM}}^2.$$

So, when this is multiplied by m_i and summed over all i, what happens? Taking the constants outside the summation sign, we get

$$I_x = \sum m_i x_i'^2 + 2X_{\text{CM}} \sum m_i x'_i + X_{\text{CM}}^2 \sum m_i.$$

The third sum is easy; it is just MX_{CM}^2. In the second sum there are two pieces, one of them is $\sum m_i x'_i$, which is the total mass times the x'-coordinate of the center of mass. But this contributes nothing, because x' is *measured* from the center of mass, and in these axes the average position of all the particles, weighted by the masses, is zero. The first sum, of course, is the x part of I_c. Thus we arrive at Eq. (19.7), just as we guessed.

Let us check (19.7) for one example. Let us just see whether it works for the rod. For an axis through one end, the moment of inertia should be $mL^2/3$, for we calculated that. The center of mass of a rod, of course, is in the center of the rod, at a distance $L/2$. Therefore we should find that $ML^2/3 = ML^2/12 + M(L/2)^2$. Since one-quarter plus one-twelfth is one-third, we have made no fundamental error.

Incidentally, we did not really need to use an integral to find the moment of inertia (19.5). If we simply assume that it is ML^2 times γ, an unknown coefficient, and then use the argument about the two halves to get $\frac{1}{4}\gamma$ for (19.6), then from our argument about transferring the axes we could prove that $\gamma = \frac{1}{4}\gamma + \frac{1}{4}$, so γ must be 1/3. There is always another way to do it!

In applying the parallel-axis theorem, it is of course important to remember that the axis for I_c *must be parallel* to the axis about which the moment of inertia is wanted.

One further property of the moment of inertia is worth mentioning because it is often helpful in finding the moment of inertia of certain kinds of objects. This property is that if one has a *plane figure* and a set of coordinate axes with origin in the plane and z-axis perpendicular to the plane, then the moment of inertia of this figure about the z-axis is equal to the sum of the moments of inertia about the x- and y-axes. This is easily proved by noting that

$$I_x = \sum m_i(y_i^2 + z_i^2) = \sum m_i y_i^2$$

(since $z_i = 0$). Similarly,

$$I_y = \sum m_i(x_i^2 + z_i^2) = \sum m_i x_i^2,$$

but

$$I_z = \sum m_i(x_i^2 + y_i^2) = \sum m_i x_i^2 + \sum m_i y_i^2$$
$$= I_x + I_y.$$

As an example, the moment of inertia of a uniform rectangular plate of mass M, width w, and length L, about an axis perpendicular to the plate and through its center is simply

$$I = M(w^2 + L^2)/12,$$

because its moment of inertia about an axis in its plane and parallel to its length is $Mw^2/12$, i.e., just as for a rod of length w, and the moment of inertia about the other axis in its plane is $ML^2/12$, just as for a rod of length L.

To summarize, the moment of inertia of an object about a given axis, which we shall call the z-axis, has the following properties:

(1) The moment of inertia is

$$I_z = \sum_i m_i(x_i^2 + y_i^2) = \int (x^2 + y^2)\, dm.$$

(2) If the object is made of a number of parts, each of whose moment of inertia is known, the total moment of inertia is the sum of the moments of inertia of the pieces.

(3) The moment of inertia about any given axis is equal to the moment of inertia about a parallel axis through the CM plus the total mass times the square of the distance from the axis to the CM.

(4) If the object is a plane figure, the moment of inertia about an axis perpendicular to the plane is equal to the sum of the moments of inertia about any two mutually perpendicular axes lying in the plane and intersecting at the perpendicular axis.

The moments of inertia of a number of elementary shapes having uniform mass densities are given in Table 19–1, and the moments of inertia of some other objects, which may be deduced from Table 19–1, using the above properties, are given in Table 19–2.

Table 19–1

Object	z-axis	I_z
Thin rod, length L	$\perp$ rod at center	$ML^2/12$
Thin concentric circular ring, radii r_1 and r_2	$\perp$ ring at center	$M(r_1^2 + r_2^2)/2$
Sphere, radius r	through center	$2Mr^2/5$

Table 19–2

Object	z-axis	I_z
Rect. sheet, sides a, b	$\parallel b$ at center	$Ma^2/12$
Rect. sheet, sides a, b	$\perp$ sheet at center	$M(a^2 + b^2)/12$
Thin annular ring, radii r_1, r_2	any diameter	$M(r_1^2 + r_2^2)/4$
Rect. parallelepiped, sides a, b, c	$\parallel c$, through center	$M(a^2 + b^2)/12$
Rt. circ. cyl., radius r, length L	$\parallel L$, through center	$Mr^2/2$
Rt. circ. cyl., radius r, length L	$\perp L$, through center	$M(r^2/4 + L^2/12)$

19–4 Rotational kinetic energy

Now let us go on to discuss dynamics further. In the analogy between linear motion and angular motion that we discussed in Chapter 18, we used the work theorem, but we did not talk about kinetic energy. What is the kinetic energy of a rigid body, rotating about a certain axis with an angular velocity ω? We can immediately guess the correct answer by using our analogies. The moment of inertia corresponds to the mass, angular velocity corresponds to velocity, and so the kinetic energy ought to be $\frac{1}{2}I\omega^2$, and indeed it is, as will now be demonstrated. Suppose the object is rotating about some axis so that each point has a velocity whose magnitude is ωr_i, where r_i is the radius from the particular point to the axis.

Then if m_i is the mass of that point, the total kinetic energy of the whole thing is just the sum of the kinetic energies of all of the little pieces:

$$T = \tfrac{1}{2}\sum m_i v_i^2 = \tfrac{1}{2}\sum m_i (r_i\omega)^2.$$

Now ω^2 is a constant, the same for all points. Thus

$$T = \tfrac{1}{2}\omega^2 \sum m_i r_i^2 = \tfrac{1}{2} I\omega^2. \tag{19.8}$$

At the end of Chapter 18 we pointed out that there are some interesting phenomena associated with an object which is not rigid, but which changes from one rigid condition with a definite moment of inertia, to another rigid condition. Namely, in our example of the turntable, we had a certain moment of inertia I_1 with our arms stretched out, and a certain angular velocity ω_1. When we pulled our arms in, we had a different moment of inertia, I_2, and a different angular velocity, ω_2, but again we were "rigid." The angular momentum remained constant, since there was no torque about the vertical axis of the turntable. This means that $I_1\omega_1 = I_2\omega_2$. Now what about the energy? That is an interesting question. With our arms pulled in, we turn faster, but our moment of inertia is less, and it looks as though the energies might be equal. But they are not, because what *does* balance is $I\omega$, not $I\omega^2$. So if we compare the kinetic energy before and after, the kinetic energy before is $\frac{1}{2}I_1\omega_1^2 = \frac{1}{2}L\omega_1$, where $L = I_1\omega_1 = I_2\omega_2$ is the angular momentum. Afterward, by the same argument, we have $T = \frac{1}{2}L\omega_2$, and since $\omega_2 > \omega_1$ the kinetic energy of rotation is greater than it was before. So we had a certain energy when our arms were out, and when we pulled them in, we were turning faster and had more kinetic energy. What happened to the theorem of the conservation of energy? Somebody must have done some work. We did work! When did we do any work? When we move a weight horizontally, we do not do any work. If we hold a thing out and pull it in, we do not do any work. But that is when we are not rotating! When we *are* rotating, there is centrifugal force on the weights. They are trying to fly out, so when we are going around we have to pull the weights in against the centrifugal force. So, the work we do against the centrifugal force ought to agree with the difference in rotational energy, and of course it does. That is where the extra kinetic energy comes from.

There is still another interesting feature which we can treat only descriptively, as a matter of general interest. This feature is a little more advanced, but is worth pointing out because it is quite curious and produces many interesting effects.

Consider that turntable experiment again. Consider the body and the arms separately, from the point of view of the man who is rotating. After the weights are pulled in, the whole object is spinning faster, but observe, *the central part of the body is not changed*, yet it is turning faster after the event than before. So, if we were to draw a circle around the inner body, and consider only objects inside the circle, *their* angular momentum would *change*; they are going faster. Therefore there must be a torque exerted on the body while we pull in our arms. No torque can be exerted by the centrifugal force, because that is radial. So that means that among the forces that are developed in a rotating system, centrifugal force is not the entire story, *there is another force*. This other force is called *Coriolis force*, and it has the very strange property that when we move something in a rotating system, it seems to be pushed sidewise. Like the centrifugal force, it is an apparent force. But if we live in a system that is rotating, and move something radially, we find that we must also push it sidewise to move it radially. This sidewise push which we have to exert is what turned our body around.

Now let us develop a formula to show how this Coriolis force really works. Suppose Moe is sitting on a carousel that appears to him to be stationary. But from the point of view of Joe, who is standing on the ground and who knows the right laws of mechanics, the carousel is going around. Suppose that we have drawn a radial line on the carousel, and that Moe is moving some mass radially along this line. We would like to demonstrate that a sidewise force is required to do that. We can do this by paying attention to the angular momentum of the mass. It is

always going around with the same angular velocity ω, so that the angular momentum is

$$L = mv_{\text{tang}}r = m\omega r \cdot r = m\omega r^2.$$

So when the mass is close to the center, it has relatively little angular momentum, but if we move it to a new position farther out, if we increase r, m has more angular momentum, so *a torque must be exerted* in order to move it along the radius. (To walk along the radius in a carousel, one has to lean over and push sidewise. Try it sometime.) The torque that is required is the rate of change of L with time as m moves along the radius. If m moves only along the radius, omega stays constant, so that the torque is

$$\tau = F_c r = \frac{dL}{dt} = \frac{d(m\omega r^2)}{dt} = 2m\omega r \frac{dr}{dt},$$

where F_c is the Coriolis force. What we really want to know is what sidewise *force* has to be exerted by Moe in order to move m out at speed $v_r = dr/dt$. This is $F_c = \tau/r = 2m\omega v_r$.

Now that we have a formula for the Coriolis force, let us look at the situation a little more carefully, to see whether we can understand the origin of this force from a more elementary point of view. We note that the Coriolis force is the same at every radius, and is evidently present even at the origin! But it is especially easy to understand it at the origin, just by looking at what happens from the inertial system of Joe, who is standing on the ground. Figure 19–4 shows three successive views of m just as it passes the origin at $t = 0$. Because of the rotation of the carousel, we see that m does not move in a straight line, but in a *curved path* tangent to a diameter of the carousel where $r = 0$. In order for m to go in a curve, there must be a force to accelerate it in absolute space. This is the Coriolis force.

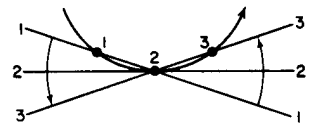


Fig. 19–4. Three successive views of a point moving radially on a rotating turntable.

This is not the only case in which the Coriolis force occurs. We can also show that if an object is moving with constant speed around the circumference of a circle, there is also a Coriolis force. Why? Moe sees a velocity v_M around the circle. On the other hand, Joe sees m going around the circle with the velocity $v_J = v_M + \omega r$, because m is also carried by the carousel. Therefore we know what the force really is, namely, the total centripetal force due to the velocity v_J, or mv_J^2/r; that is the actual force. Now from Moe's point of view, this centripetal force has three pieces. We may write it all out as follows:

$$F_r = -\frac{mv_J^2}{r} = -\frac{mv_M^2}{r} - 2mv_M\omega - m\omega^2 r.$$

Now, F_r is the force that Moe would see. Let us try to understand it. Would Moe appreciate the first term? "Yes," he would say, "even if I were not turning, there would be a centripetal force if I were to run around a circle with velocity v_M." This is simply the centripetal force that Moe would expect, having nothing to do with rotation. In addition, Moe is quite aware that there is another centripetal force that would act even on objects which are standing still on his carousel. This is the third term. But there is another term in addition to these, namely the second term, which is again $2m\omega v$. The Coriolis force F_c was tangential when the velocity was radial, and now it is radial when the velocity is tangential. In fact, one expression has a minus sign relative to the other. The force is always in the same direction, relative to the velocity, no matter in which direction the velocity is. The force is at right angles to the velocity, and of magnitude $2m\omega v$.

20

Rotation in space

20–1 Torques in three dimensions

20–2 The rotation equations using cross products

20–3 The gyroscope

20–4 Angular momentum of a solid body

20–1 Torques in three dimensions

In this chapter we shall discuss one of the most remarkable and amusing consequences of mechanics, the behavior of a rotating wheel. In order to do this we must first extend the mathematical formulation of rotational motion, the principles of angular momentum, torque, and so on, to three-dimensional space. We shall not *use* these equations in all their generality and study all their consequences, because this would take many years, and we must soon turn to other subjects. In an introductory course we can present only the fundamental laws and apply them to a very few situations of special interest.

First, we notice that if we have a rotation in three dimensions, whether of a rigid body or any other system, what we deduced for two dimensions is still right. That is, it is still true that $xF_y - yF_x$ is the torque "in the xy-plane," or the torque "around the z-axis." It also turns out that this torque is still equal to the rate of change of $xp_y - yp_x$, for if we go back over the derivation of Eq. (18.15) from Newton's laws we see that we did not have to assume that the motion was in a plane; when we differentiate $xp_y - yp_x$, we get $xF_y - yF_x$, so this theorem is still right. The quantity $xp_y - yp_x$, then, we call the angular momentum belonging to the xy-plane, or the angular momentum about the z-axis. This being true, we can use any other pair of axes and get another equation. For instance, we can use the yz-plane, and it is clear from symmetry that if we just substitute y for x and z for y, we would find $yF_z - zF_y$ for the torque and $yp_z - zp_y$ would be the angular momentum associated with the yz-plane. Of course we could have another plane, the zx-plane, and for this we would find $zF_x - xF_z = d/dt\,(zp_x - xp_z)$.

That these three equations can be deduced for the motion of a single particle is quite clear. Furthermore, if we added such things as $xp_y - yp_x$ together for many particles and called it the total angular momentum, we would have three kinds for the three planes xy, yz, and zx, and if we did the same with the forces, we would talk about the torque in the planes xy, yz, and zx also. Thus we would have laws that the external torque associated with any plane is equal to the rate of change of the angular momentum associated with that plane. This is just a generalization of what we wrote in two dimensions.

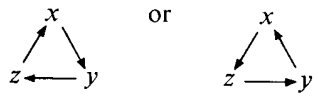
But now one may say, "Ah, but there are more planes; after all, can we not take some other plane at some angle, and calculate the torque on that plane from the forces? Since we would have to write another set of equations for every such plane, we would have a lot of equations!" Interestingly enough, it turns out that if we were to work out the combination $x'F_{y'} - y'F_{x'}$ for another plane, measuring the x', $F_{y'}$, etc., in that plane, the result can be written as some *combination* of the three expressions for the xy-, yz- and zx-planes. There is nothing new. In other words, if we know what the three torques in the xy-, yz-, and zx-planes are, then the torque in any other plane, and correspondingly the angular momentum also, can be written as some combination of these: six percent of one and ninety-two percent of another, and so on. This property we shall now analyze.

Suppose that in the xyz-axes, Joe has worked out all his torques and his angular momenta in his planes. But Moe has axes x', y', z' in some other direction. To make it a little easier, we shall suppose that only the x- and y-axes have been turned. Moe's x' and y' are new, but his z' happens to be the same. That is, he has new planes, let us say, for yz and zx. He therefore has new torques and angular momenta which he would work out. For example, his torque in the $x'y'$-plane would be equal to $x'F_{y'} - y'F_{x'}$ and so forth. What we must now do is to find the relationship between the new torques and the old torques, so we will be able to make a

connection from one set of axes to the other. Someone may say, "That looks just like what we did with vectors." And indeed, that is exactly what we are intending to do. Then he may say, "Well, isn't torque just a vector?" It *does* turn out to be a vector, but we do not know that right away without making an analysis. So in the following steps we shall make the analysis. We shall not discuss every step in detail, since we only want to illustrate how it works. The torques calculated by Joe are

$$\begin{aligned} \tau_{xy} &= xF_y - yF_x, \\ \tau_{yz} &= yF_z - zF_y, \\ \tau_{zx} &= zF_x - xF_z. \end{aligned} \tag{20.1}$$

We digress at this point to note that in such cases as this one may get the wrong sign for some quantity if the coordinates are not handled in the right way. Why not write $\tau_{yz} = zF_y - yF_z$? The problem arises from the fact that a coordinate system may be either "right-handed" or "left-handed." Having chosen (arbitrarily) a sign for, say τ_{xy}, then the correct expressions for the other two quantities may always be found by interchanging the letters xyz in either order



Moe now calculates the torques in his system:

$$\begin{aligned} \tau_{x'y'} &= x'F_{y'} - y'F_{x'}, \\ \tau_{y'z'} &= y'F_{z'} - z'F_{y'}, \\ \tau_{z'x'} &= z'F_{x'} - x'F_{z'}. \end{aligned} \tag{20.2}$$

Now we suppose that one coordinate system is rotated by a fixed angle θ, such that the z- and z'-axes are the same. (This angle θ has nothing to do with rotating objects or what is going on inside the coordinate system. It is merely the relationship between the axes used by one man and the axes used by the other, and is supposedly constant.) Thus the coordinates of the two systems are related by

$$\begin{aligned} x' &= x\cos\theta + y\sin\theta, \\ y' &= y\cos\theta - x\sin\theta, \\ z' &= z. \end{aligned} \tag{20.3}$$

Likewise, because force is a vector it transforms into the new system in the same way as do x, y, and z, since a thing is a vector if and only if the various components transform in the same way as x, y, and z:

$$\begin{aligned} F_{x'} &= F_x\cos\theta + F_y\sin\theta, \\ F_{y'} &= F_y\cos\theta - F_x\sin\theta, \\ F_{z'} &= F_z. \end{aligned} \tag{20.4}$$

Now we can find out how the torque transforms by merely substituting for x', y', and z' the expressions (20.3), and for $F_{x'}$, $F_{y'}$, $F_{z'}$ those given by (20.4), all into (20.2). So, we have a rather long string of terms for $\tau_{x'y'}$ and (rather surprisingly at first) it turns out that it comes right down to $xF_y - yF_x$, which we recognize to be the torque in the xy-plane:

$$\begin{aligned} \tau_{x'y'} &= (x\cos\theta + y\sin\theta)(F_y\cos\theta - F_x\sin\theta) \\ &\quad -(y\cos\theta - x\sin\theta)(F_x\cos\theta + F_y\sin\theta) \\ &= xF_y(\cos^2\theta + \sin^2\theta) - yF_x(\sin^2\theta + \cos^2\theta) \\ &\quad +xF_x(-\sin\theta\cos\theta + \sin\theta\cos\theta) \\ &\quad +yF_y(\sin\theta\cos\theta - \sin\theta\cos\theta) \\ &= xF_y - yF_x = \tau_{xy}. \end{aligned} \tag{20.5}$$

That result is clear, for if we only turn our axes *in the plane*, the twist around z in that plane is no different than it was before, because it is the same plane! What will be more interesting is the expression for $\tau_{y'z'}$, because that is a new plane. We now do exactly the same thing with the $y'z'$-plane, and it comes out as follows:

$$\begin{aligned}\tau_{y'z'} &= (y\cos\theta - x\sin\theta)F_z \\ &\quad -z(F_y\cos\theta - F_x\sin\theta) \\ &= (yF_z - zF_y)\cos\theta + (zF_x - xF_z)\sin\theta \\ &= \tau_{yz}\cos\theta + \tau_{zx}\sin\theta. \end{aligned} \tag{20.6}$$

Finally, we do it for $z'x'$:

$$\begin{aligned}\tau_{z'x'} &= z(F_x\cos\theta + F_y\sin\theta) \\ &\quad -(x\cos\theta + y\sin\theta)F_z \\ &= (zF_x - xF_z)\cos\theta - (yF_z - zF_y)\sin\theta \\ &= \tau_{zx}\cos\theta - \tau_{yz}\sin\theta. \end{aligned} \tag{20.7}$$

We wanted to get a rule for finding torques in new axes in terms of torques in old axes, and now we have the rule. How can we ever remember that rule? If we look carefully at (20.5), (20.6), and (20.7), we see that there is a close relationship between these equations and the equations for x, y, and z. If, somehow, we could call τ_{xy} the *z-component* of something, let us call it the z-component of $\boldsymbol{\tau}$, then it would be all right; we would understand (20.5) as a vector transformation, since the z-component would be unchanged, as it should be. Likewise, if we associate with the yz-plane the x-component of our newly invected vector, and with the zx-plane, the y-component, then these transformation expressions would read

$$\begin{aligned}\tau_{z'} &= \tau_z, \\ \tau_{x'} &= \tau_x\cos\theta + \tau_y\sin\theta, \\ \tau_{y'} &= \tau_y\cos\theta - \tau_x\sin\theta, \end{aligned} \tag{20.8}$$

which is just the rule for vectors!

Therefore we have proved that we may identify the combination of $xF_y - yF_x$ with what we ordinarily call the z-component of a certain artificially invented vector. Although a torque is a twist on a plane, and it has no *a priori* vector character, mathematically it does behave like a vector. This vector is at right angles to the plane of the twist, and its length is proportional to the strength of the twist. The three components of such a quantity will transform like a real vector.

So we represent torques by vectors; with each plane on which the torque is supposed to be acting, we associate a line at right angles, by a rule. But "at right angles" leaves the sign unspecified. To get the sign right, we must adopt a rule which will tell us that if the torque were in a certain sense on the xy-plane, then the axis that we want to associate with it is in the "up" z-direction. That is, somebody has to define "right" and "left" for us. Supposing that the coordinate system is x, y, z in a right-hand system, then the rule will be the following: if we think of the twist as if we were turning a screw having a right-hand thread, then the direction of the vector that we will associate with that twist is in the direction that the screw would advance.

Why is torque a vector? It is a miracle of good luck that we can associate a single axis with a plane, and therefore that we can associate a vector with the torque; it is a special property of three-dimensional space. In two dimensions, the torque is an ordinary scalar, and there need be no direction associated with it. In three dimensions, it is a vector. If we had four dimensions, we would be in great difficulty, because (if we had time, for example, as the fourth dimension) we would not only have planes like xy, yz, and zx, we would also have tx-, ty-, and tz-planes. There would be *six* of them, and one cannot represent six quantities as one vector in four dimensions.

We will be living in three dimensions for a long time, so it is well to notice that the foregoing mathematical treatment did not depend upon the fact that x

was position and F was force; it only depended on the transformation laws for vectors. Therefore if, instead of x, we used the x-component of some other vector, it is not going to make any difference. In other words, if we were to calculate $a_xb_y - a_yb_x$, where $\mathbf{a}$ and $\mathbf{b}$ are vectors, and call it the z-component of some new quantity c, then these new quantities form a vector $\mathbf{c}$. We need a mathematical notation for the relationship of the new vector, with its three components, to the vectors $\mathbf{a}$ and $\mathbf{b}$. The notation that has been devised for this is $\mathbf{c} = \mathbf{a} \times \mathbf{b}$. We have then, in addition to the ordinary scalar product in the theory of vector analysis, a new kind of product, called the *vector product*. Thus, if $\mathbf{c} = \mathbf{a} \times \mathbf{b}$, this is the same as writing

$$\begin{aligned} c_x &= a_yb_z - a_zb_y, \\ c_y &= a_zb_x - a_xb_z, \\ c_z &= a_xb_y - a_yb_x. \end{aligned} \tag{20.9}$$

If we reverse the order of $\mathbf{a}$ and $\mathbf{b}$, calling $\mathbf{a}$, $\mathbf{b}$ and $\mathbf{b}$, $\mathbf{a}$, we would have the sign of $\mathbf{c}$ reversed, because c_z would be $b_xa_y - b_ya_x$. Therefore the cross product is unlike ordinary multiplication, where $ab = ba$; for the cross product, $\mathbf{b} \times \mathbf{a} = -\mathbf{a} \times \mathbf{b}$. From this, we can prove at once that if $\mathbf{a} = \mathbf{b}$, the cross product is zero. Thus, $\mathbf{a} \times \mathbf{a} = 0$.

The cross product is very important for representing the features of rotation, and it is important that we understand the geometrical relationship of the three vectors $\mathbf{a}$, $\mathbf{b}$, and $\mathbf{c}$. Of course the relationship in components is given in Eq. (20.9) and from that one can determine what the relationship is in geometry. The answer is, first, that the vector $\mathbf{c}$ is perpendicular to both $\mathbf{a}$ and $\mathbf{b}$. (Try to calculate $\mathbf{c} \cdot \mathbf{a}$, and see if it does not reduce to zero.) Second, the magnitude of $\mathbf{c}$ turns out to be the magnitude of $\mathbf{a}$ times the magnitude of $\mathbf{b}$ times the sine of the angle between the two. In which direction does $\mathbf{c}$ point? Imagine that we turn $\mathbf{a}$ into $\mathbf{b}$ through an angle less than 180°; a screw with a right-hand thread turning in this way will advance in the direction of $\mathbf{c}$. The fact that we say a *right*-hand screw instead of a *left*-hand screw is a convention, and is a perpetual reminder that if $\mathbf{a}$ and $\mathbf{b}$ are "honest" vectors in the ordinary sense, the new kind of "vector" which we have created by $\mathbf{a} \times \mathbf{b}$ is artificial, or slightly different in its character from $\mathbf{a}$ and $\mathbf{b}$, because it was made up with a special rule. If $\mathbf{a}$ and $\mathbf{b}$ are called ordinary vectors, we have a special name for them, we call them *polar vectors*. Examples of such vectors are the coordinate $\mathbf{r}$, force $\mathbf{F}$, momentum $\mathbf{p}$, velocity $\mathbf{v}$, electric field $\mathbf{E}$, etc.; these are ordinary polar vectors. Vectors which involve just one cross product in their definition are called *axial vectors* or *pseudovectors*. Examples of pseudovectors are, of course, torque $\boldsymbol{\tau}$ and the angular momentum $\mathbf{L}$. It also turns out that the angular velocity $\boldsymbol{\omega}$ is a pseudovector, as is the magnetic field $\mathbf{B}$.

In order to complete the mathematical properties of vectors, we should know all the rules for their multiplication, using dot and cross products. In our applications at the moment, we will need very little of this, but for the sake of completeness we shall write down all of the rules for vector multiplication so that we can use the results later. These are

$$\begin{aligned} &\text{(a)} \quad \mathbf{a} \times (\mathbf{b} + \mathbf{c}) = \mathbf{a} \times \mathbf{b} + \mathbf{a} \times \mathbf{c}, \\ &\text{(b)} \quad (\alpha\mathbf{a}) \times \mathbf{b} = \alpha(\mathbf{a} \times \mathbf{b}), \\ &\text{(c)} \quad \mathbf{a} \cdot (\mathbf{b} \times \mathbf{c}) = (\mathbf{a} \times \mathbf{b}) \cdot \mathbf{c}, \\ &\text{(d)} \quad \mathbf{a} \times (\mathbf{b} \times \mathbf{c}) = \mathbf{b}(\mathbf{a} \cdot \mathbf{c}) - \mathbf{c}(\mathbf{a} \cdot \mathbf{b}), \\ &\text{(e)} \quad \mathbf{a} \times \mathbf{a} = 0, \\ &\text{(f)} \quad \mathbf{a} \cdot (\mathbf{a} \times \mathbf{b}) = 0. \end{aligned} \tag{20.10}$$

20–2 The rotation equations using cross products

Now let us ask whether any equations in physics can be written using the cross product. The answer, of course, is that a great many equations can be so written. For instance, we see immediately that the torque is equal to the position

vector cross the force:

$$\boldsymbol{\tau} = \mathbf{r} \times \mathbf{F}. \tag{20.11}$$

This is a vector summary of the three equations $\tau_x = yF_z - zF_y$, etc. By the same token, the angular momentum vector, if there is only one particle present, is the distance from the origin multiplied by the vector momentum:

$$\mathbf{L} = \mathbf{r} \times \mathbf{p}. \tag{20.12}$$

For three-dimensional space rotation, the dynamical law analogous to the law $\mathbf{F} = d\mathbf{p}/dt$ of Newton, is that the torque vector is the rate of change with time of the angular momentum vector:

$$\boldsymbol{\tau} = d\mathbf{L}/dt. \tag{20.13}$$

If we sum (20.13) over many particles, the external torque on a system is the rate of change of the total angular momentum:

$$\boldsymbol{\tau}_{\text{ext}} = d\mathbf{L}_{\text{tot}}/dt. \tag{20.14}$$

Another theorem: If the total external torque is zero, then the total vector angular momentum of the system is a constant. This is called the law of *conservation of angular momentum.* If there is no torque on a given system, its angular momentum cannot change.

What about angular velocity? Is it a vector? We have already discussed turning a solid object about a fixed axis, but for a moment suppose that we are turning it simultaneously about *two* axes. It might be turning about an axis inside a box, while the box is turning about some other axis. The net result of such combined motions is that the object simply turns about some new axis! The wonderful thing about this new axis is that it can be figured out this way. If the rate of turning in the xy-plane is written as a vector in the z-direction whose length is equal to the rate of rotation in the plane, and if another vector is drawn in the y-direction, say, which is the rate of rotation in the zx-plane, then if we add these together as a vector, the magnitude of the result tells us how fast the object is turning, and the direction tells us in what plane, by the rule of the parallelogram. That is to say, simply, angular velocity *is* a vector, where we draw the magnitudes of the rotations in the three planes as projections at right angles to those planes.*

As a simple application of the use of the angular velocity vector, we may evaluate the power being expended by the torque acting on a rigid body. The power, of course, is the rate of change of work with time; in three dimensions, the power turns out to be $P = \boldsymbol{\tau} \cdot \boldsymbol{\omega}$.

All the formulas that we wrote for plane rotation can be generalized to three dimensions. For example, if a rigid body is turning about a certain axis with angular velocity $\boldsymbol{\omega}$, we might ask, "What is the velocity of a point at a certain radial position $\mathbf{r}$?" We shall leave it as a problem for the student to show that the velocity of a particle in a rigid body is given by $\mathbf{v} = \boldsymbol{\omega} \times \mathbf{r}$, where $\boldsymbol{\omega}$ is the angular velocity and $\mathbf{r}$ is the position. Also, as another example of cross products, we had a formula for Coriolis force, which can also be written using cross products: $\mathbf{F}_c = 2m\mathbf{v} \times \boldsymbol{\omega}$. That is, if a particle is moving with velocity $\mathbf{v}$ in a coordinate system which is, in fact, rotating with angular velocity $\boldsymbol{\omega}$, and we want to think in terms of the rotating coordinate system, then we have to add the pseudoforce $\mathbf{F}_c$.

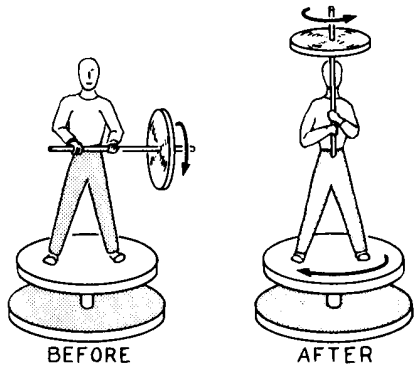


Fig. 20–1. *Before:* axis is horizontal; moment about vertical axis = 0. *After:* axis is vertical; momentum about vertical axis is still zero; man and chair spin in direction opposite to spin of the wheel.

20–3 The gyroscope

Let us now return to the law of conservation of angular momentum. This law may be demonstrated with a rapidly spinning wheel, or gyroscope, as follows (see Fig. 20–1). If we sit on a swivel chair and hold the spinning wheel with its axis horizontal, the wheel has an angular momentum about the horizontal axis.

* That this is true can be derived by compounding the displacements of the particles of the body during an infinitesimal time Δt. It is not self-evident, and is left to those who are interested to try to figure it out.

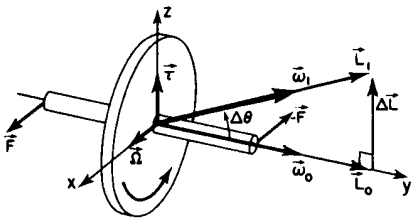


Fig. 20–2. A gyroscope.

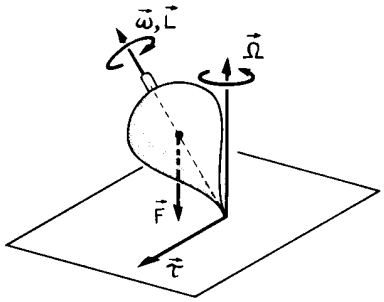


Fig. 20–3. A rapidly spinning top. Note that the direction of the torque vector is the direction of the precession.

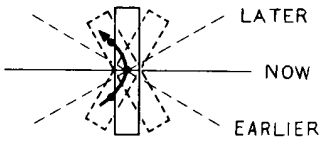


Fig. 20–4. The motion of particles in the spinning wheel of Fig. 20–2, whose axis is turning, is in curved lines.

Angular momentum around a *vertical* axis cannot change because of the (frictionless) pivot of the chair, so if we turn the axis of the wheel into the vertical, then the wheel would have angular momentum about the vertical axis, because it is now spinning about this axis. But the *system* (wheel, ourself, and chair) *cannot* have a vertical component, so we and the chair have to turn in the direction opposite to the spin of the wheel, to balance it.

First let us analyze in more detail the thing we have just described. What is surprising, and what we must understand, is the origin of the forces which turn us and the chair around as we turn the axis of the gyroscope toward the vertical. Figure 20–2 shows the wheel spinning rapidly about the y-axis. Therefore its angular velocity is about that axis and, it turns out, its angular momentum is likewise in that direction. Now suppose that we wish to rotate the wheel about the x-axis at a small angular velocity Ω; what forces are required? After a short time Δt, the axis has turned to a new position, tilted at an angle $\Delta\theta$ with the horizontal. Since the major part of the angular momentum is due to the spin on the axis (very little is contributed by the slow turning), we see that the angular momentum vector has changed. What is the change in angular momentum? The angular momentum does not change in *magnitude*, but it does change in *direction* by an amount $\Delta\theta$. The magnitude of the vector ΔL is thus $\Delta L = L_0\,\Delta\theta$, so that the torque, which is the time rate of change of the angular momentum, is $\tau = \Delta L/\Delta t = L_0\,\Delta\theta/\Delta t = L_0\Omega$. Taking the directions of the various quantities into account, we see that

$$\boldsymbol{\tau} = \boldsymbol{\Omega} \times \mathbf{L}_0. \tag{20.15}$$

Thus, if $\boldsymbol{\Omega}$ and $\mathbf{L}_0$ are both horizontal, as shown in the figure, τ is *vertical*. To produce such a torque, horizontal forces $\mathbf{F}$ and $-\mathbf{F}$ must be applied at the ends of the axle. How are these forces applied? By our hands, as we try to rotate the axis of the wheel into the vertical direction. But Newton's Third Law demands that equal and opposite forces (and equal and opposite *torques*) act on *us*. This causes us to rotate in the opposite sense about the vertical axis z.

This result can be generalized for a rapidly spinning top. In the familiar case of a spinning top, gravity acting on its center of mass furnishes a torque about the point of contact with the floor (see Fig. 20–3). This torque is in the horizontal direction, and causes the top to precess with its axis moving in a circular cone about the vertical. If $\boldsymbol{\Omega}$ is the (vertical) angular velocity of precession, we again find that

$$\boldsymbol{\tau} = d\mathbf{L}/dt = \boldsymbol{\Omega} \times \mathbf{L}_0.$$

Thus, when we apply a torque to a rapidly spinning top, the direction of the precessional motion is in the direction of the torque, or at right angles to the forces producing the torque.

We may now claim to understand the precession of gyroscopes, and indeed we do, mathematically. However, this is a mathematical thing which, in a sense, appears as a "miracle." It will turn out, as we go to more and more advanced physics, that many simple things can be deduced mathematically more rapidly than they can be really understood in a fundamental or simple sense. This is a strange characteristic, and as we get into more and more advanced work there are circumstances in which mathematics will produce results which *no one* has really been able to understand in any direct fashion. An example is the Dirac equation, which appears in a very simple and beautiful form, but whose consequences are hard to understand. In our particular case, the precession of a top looks like some kind of a miracle involving right angles and circles, and twists and right-hand screws. What we should try to do is to understand it in a more physical way.

How can we explain the torque in terms of the real forces and the accelerations? We note that when the wheel is precessing, the particles that are going around the wheel are not really moving in a plane because the wheel is precessing (see Fig. 20–4). As we explained previously (Fig. 19–4), the particles which are crossing through the precession axis are moving in *curved paths*, and this requires application of a lateral force. This is supplied by our pushing on the axle, which then com-

municates the force to the rim through the spokes. "Wait," someone says, "what about the particles that are going back on the other side?" It does not take long to decide that there must be a force in the *opposite direction* on that side. The net force that we have to apply is therefore zero. The *forces* balance out, but one of them must be applied at one side of the wheel, and the other must be applied at the other side of the wheel. We could apply these forces directly, but because the wheel is solid we are allowed to do it by pushing on the axle, since forces can be carried up through the spokes.

What we have so far proved is that if the wheel is precessing, it can balance the torque due to gravity or some other applied torque. But all we have shown is that this is *a* solution of an equation. That is, if the torque is given, and *if we get the spinning started right*, then the wheel will precess smoothly and uniformly. But we have not proved (and it is not true) that a uniform precession is the *most general* motion a spinning body can undergo as the result of a given torque. The general motion involves also a "wobbling" about the mean precession. This "wobbling" is called *nutation.*

Some people like to say that when one exerts a torque on a gyroscope, it turns and it precesses, and that the torque *produces* the precession. It is very strange that when one suddenly lets go of a gyroscope, it does not *fall* under the action of gravity, but moves sidewise instead! Why is it that the *downward* force of the gravity, which we *know* and *feel*, makes it go *sidewise*? All the formulas in the world like (20.15) are not going to tell us, because (20.15) is a special equation, valid only after the gyroscope is precessing nicely. What really happens, in detail, is the following. If we were to hold the axis absolutely fixed, so that it cannot precess in any manner (but the top is spinning) then there is no torque acting, not even a torque from gravity, because it is balanced by our fingers. But if we suddenly let go, then there will instantaneously be a torque from gravity. Anyone in his right mind would think that the top would fall, and that is what it starts to do, as can be seen if the top is not spinning too fast.

The gyro actually does fall, as we would expect. But as soon as it falls, it is then turning, and if this turning were to continue, a torque would be required. In the absence of a torque in this direction, the gyro begins to "fall" in the direction opposite that of the missing force. This gives the gyro a component of motion around the vertical axis, as it would have in steady precession. But the actual motion "overshoots" the steady precessional velocity, and the axis actually rises again to the level from which it started. The path followed by the end of the axle is a cycloid (the path followed by a pebble that is stuck in the tread of an automobile tire). Ordinarily, this motion is too quick for the eye to follow, and it damps out quickly because of the friction in the gimbal bearings, leaving only the steady precessional drift (Fig. 20–5). The slower the wheel spins, the more obvious the nutation is.

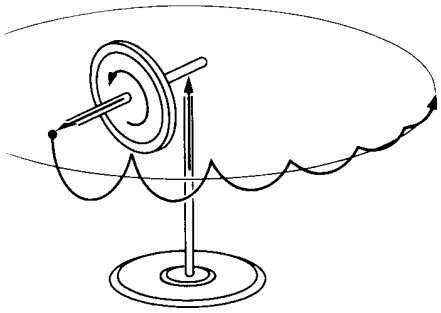


Fig. 20–5. Actual motion of tip of axis of gyroscope under gravity just after releasing axis previously held fixed.

When the motion settles down, the axis of the gyro is a little bit lower than it was at the start. Why? (These are the more complicated details, but we bring them in because we do not want the reader to get the idea that the gyroscope is an absolute miracle. It *is* a wonderful thing, but it is not a miracle.) If we were holding the axis absolutely horizontally, and suddenly let go, then the simple precession equation would tell us that it precesses, that it goes around in a horizontal plane. But that is impossible! Although we neglected it before, it is true that the wheel has *some* moment of inertia about the precession axis, and if it is moving about that axis, even slowly, it has a weak angular momentum about the axis. Where did it come from? If the pivots are perfect, there is no torque about the vertical axis. How then *does* it get to precess if there is no change in the angular momentum? The answer is that the cycloidal motion of the end of the axis damps down to the average, steady motion of the center of the equivalent rolling circle. That is, it settles down a little bit low. Because it is low, the spin angular momentum now has a small vertical component, which is exactly what is needed for the precession. So you see it has to go down a little, in order to go around. It has to yield a little bit to the gravity; by turning its axis down a little bit, it maintains the rotation about the vertical axis. That, then, is the way a gyroscope works.

20–4 Angular momentum of a solid body

Before we leave the subject of rotations in three dimensions, we shall discuss, at least qualitatively, a few effects that occur in three-dimensional rotations that are not self-evident. The main effect is that, in general, the angular momentum of a rigid body is *not necessarily* in the same direction as the angular velocity. Consider a wheel that is fastened onto a shaft in a lopsided fashion, but with the axis through the center of gravity, to be sure (Fig. 20–6). When we spin the wheel around the axis, anybody knows that there will be shaking at the bearings because of the lopsided way we have it mounted. Qualitatively, we know that in the rotating system there is centrifugal force acting on the wheel, trying to throw its mass as far as possible from the axis. This tends to line up the plane of the wheel so that it is perpendicular to the axis. To resist this tendency, a torque is exerted by the bearings. If there is a torque exerted by the bearings, there must be a rate of change of angular momentum. How can there be a rate of change of angular momentum when we are simply turning the wheel about the axis? Suppose we break the angular velocity $\boldsymbol{\omega}$ into components $\boldsymbol{\omega}_1$ and $\boldsymbol{\omega}_2$ perpendicular and parallel to the plane of the wheel. What is the angular momentum? The moments of inertia about these two axes are *different*, so the angular momentum components, which (in these particular, special axes only) are equal to the moments of inertia times the corresponding angular velocity components, are in a *different ratio* than are the angular velocity components. Therefore the angular momentum vector is in a direction in space *not* along the axis. When we turn the object, we have to turn the angular momentum vector in space, so we must exert torques on the shaft.

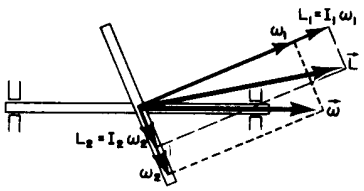


Fig. 20–6. The angular momentum of a rotating body is not necessarily parallel to the angular velocity.

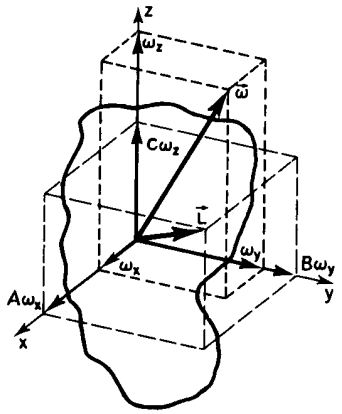


Fig. 20–7. The angular velocity and angular momentum of a rigid body ($A > B > C$).

Although it is much too complicated to prove here, there is a very important and interesting property of the moment of inertia which is easy to describe and to use, and which is the basis of our above analysis. This property is the following: Any rigid body, even an irregular one like a potato, possesses three mutually perpendicular axes through the CM, such that the moment of inertia about one of these axes has the greatest possible value for any axis through the CM, the moment of inertia about another of the axes has the *minimum* possible value, and the moment of inertia about the third axis is intermediate between these two (or equal to one of them). These axes are called the *principal axes* of the body, and they have the important property that if the body is rotating about one of them, its angular momentum is in the same direction as the angular velocity. For a body having axes of symmetry, the principal axes are along the symmetry axes.

If we take the x-, y-, and z-axes along the principal axes, and call the corresponding principal moments of inertia A, B, and C, we may easily evaluate the angular momentum and the kinetic energy of rotation of the body for any angular velocity $\boldsymbol{\omega}$. If we resolve $\boldsymbol{\omega}$ into components ω_x, ω_y, and ω_z along the x-, y-, z-axes, and use unit vectors $\mathbf{i}$, $\mathbf{j}$, $\mathbf{k}$, also along x, y, z, we may write the angular momentum as

$$\mathbf{L} = A\omega_x\mathbf{i} + B\omega_y\mathbf{j} + C\omega_z\mathbf{k}. \tag{20.16}$$

The kinetic energy of rotation is

$$\begin{aligned} \text{KE} &= \tfrac{1}{2}(A\omega_x^2 + B\omega_y^2 + C\omega_z^2) \\ &= \tfrac{1}{2}\mathbf{L}\cdot\boldsymbol{\omega}. \end{aligned} \tag{20.17}$$

21

The Harmonic Oscillator

21–1 Linear differential equations

21–2 The harmonic oscillator

21–3 Harmonic motion and circular motion

21–4 Initial conditions

21–5 Forced oscillations

21–1 Linear differential equations

In the study of physics, usually the course is divided into a series of subjects, such as mechanics, electricity, optics, etc., and one studies one subject after the other. For example, this course has so far dealt mostly with mechanics. But a strange thing occurs again and again: the equations which appear in different fields of physics, and even in other sciences, are often almost exactly the same, so that many phenomena have analogs in these different fields. To take the simplest example, the propagation of sound waves is in many ways analogous to the propagation of light waves. If we study acoustics in great detail we discover that much of the work is the same as it would be if we were studying optics in great detail. So the study of a phenomenon in one field may permit an extension of our knowledge in another field. It is best to realize from the first that such extensions are possible, for otherwise one might not understand the reason for spending a great deal of time and energy on what appears to be only a small part of mechanics.

The harmonic oscillator, which we are about to study, has close analogs in many other fields; although we start with a mechanical example of a weight on a spring, or a pendulum with a small swing, or certain other mechanical devices, we are really studying a certain *differential equation.* This equation appears again and again in physics and in other sciences, and in fact it is a part of so many phenomena that its close study is well worth our while. Some of the phenomena involving this equation are the oscillations of a mass on a spring; the oscillations of charge flowing back and forth in an electrical circuit; the vibrations of a tuning fork which is generating sound waves; the analogous vibrations of the electrons in an atom, which generate light waves; the equations for the operation of a servosystem, such as a thermostat trying to adjust a temperature; complicated interactions in chemical reactions; the growth of a colony of bacteria in interaction with the food supply and the poisons the bacteria produce; foxes eating rabbits eating grass, and so on; all these phenomena follow equations which are very similar to one another, and this is the reason why we study the mechanical oscillator in such detail. The equations are called *linear differential equations with constant coefficients.* A linear differential equation with constant coefficients is a differential equation consisting of a sum of several terms, each term being a derivative of the dependent variable with respect to the independent variable, and multiplied by some constant. Thus

$$a_n\, d^nx/dt^n + a_{n-1}\, d^{n-1}x/dt^{n-1} + \cdots + a_1\, dx/dt + a_0x = f(t) \qquad (21.1)$$

is called a linear differential equation of order n with constant coefficients (each a_i is constant).

21–2 The harmonic oscillator

Perhaps the simplest mechanical system whose motion follows a linear differential equation with constant coefficients is a mass on a spring: first the spring stretches to balance the gravity; once it is balanced, we then discuss the vertical displacement of the mass from its equilibrium position (Fig. 21–1). We shall call this upward displacement x, and we shall also suppose that the spring is perfectly linear, in which case the force pulling back when the spring is stretched is precisely proportional to the amount of stretch. That is, the force is $-kx$ (with a

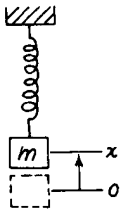


Fig. 21–1. A mass on a spring: a simple example of a harmonic oscillator.

minus sign to remind us that it pulls back). Thus the mass times the acceleration must equal $-kx$:

$$m\,d^2x/dt^2 = -kx. \tag{21.2}$$

For simplicity, suppose it happens (or we change our unit of time measurement) that the ratio $k/m = 1$. We shall first study the equation

$$d^2x/dt^2 = -x. \tag{21.3}$$

Later we shall come back to Eq. (21.2) with the k and m explicitly present.

We have already analyzed Eq. (21.3) in detail numerically; when we first introduced the subject of mechanics we solved this equation (see Eq. 9.12) to find the motion. By numerical integration we found a curve (Fig. 9–4) which showed that if m was initially displaced, but at rest, it would come down and go through zero; we did not then follow it any farther, but of course we know that it just keeps going up and down—it *oscillates*. When we calculated the motion numerically, we found that it went through the equilibrium point at $t = 1.570$. The length of the whole cycle is four times this long, or $t_0 = 6.28$ "sec." This was found numerically, before we knew much calculus. We assume that in the meantime the Mathematics Department has brought forth a function which, when differentiated twice, is equal to itself with a minus sign. (There are, of course, ways of getting at this function in a direct fashion, but they are more complicated than already knowing what the answer is.) The function is $x = \cos t$. If we differentiate this we find $dx/dt = -\sin t$ and $d^2x/dt^2 = -\cos t = -x$. The function $x = \cos t$ starts, at $t = 0$, with $x = 1$, and no initial velocity; that was the situation with which we started when we did our numerical work. Now that we know that $x = \cos t$, we can calculate a *precise* value for the time at which it should pass $x = 0$. The answer is $t = \pi/2$, or 1.57108. We were wrong in the last figure because of the errors of numerical analysis, but it was very close!

Now to go further with the original problem, we restore the time units to real seconds. What is the solution then? First of all, we might think that we can get the constants k and m in by multiplying cos t by something. So let us try the equation $x = A\cos t$; then we find $dx/dt = -A\sin t$, and $d^2x/dt^2 = -A\cos t = -x$. Thus we discover to our horror that we did not succeed in solving Eq. (21.2), but we got Eq. (21.3) again! That fact illustrates one of the most important properties of linear differential equations: *if we multiply a solution of the equation by any constant, it is again a solution.* The mathematical reason for this is clear. If x is a solution, and we multiply both sides of the equation, say by A, we see that all derivatives are also multiplied by A, and therefore Ax is just as good a solution of the original equation as x was. The physics of it is the following. If we have a weight on a spring, and pull it down twice as far, the force is twice as much, the resulting acceleration is twice as great, the velocity it acquires in a given time is twice as great, the distance covered in a given time is twice as great; but it *has* to cover twice as great a distance in order to get back to the origin because it is pulled down twice as far. So it takes the *same time* to get back to the origin, irrespective of the initial displacement. In other words, with a linear equation, the motion has the same *time pattern*, no matter how "strong" it is.

That was the wrong thing to do—it only taught us that we can multiply the solution by anything, and it satisfies the same equation, but not a different equation. After a little cut and try to get to an equation with a different constant multiplying x, we find that we must alter the scale of *time*. In other words, Eq. (21.2) has a solution of the form

$$x = \cos \omega_0 t. \tag{21.4}$$

(It is important to realize that in the present case, ω_0 is not an angular velocity of a spinning body, but we run out of letters if we are not allowed to use the same letter for more than one thing.) The reason we put a subscript "0" on ω is that we are going to have more omegas before long; let us remember that ω_0 refers to the natural motion of this oscillator. Now we try Eq. (21.4) and this time we are more successful, because $dx/dt = -\omega_0 \sin \omega_0 t$ and $d^2x/dt^2 = -\omega_0^2 \cos \omega_0 t = -\omega_0^2 x$.

So at last we have solved the equation that we really wanted to solve. The equation $d^2x/dt^2 = -\omega_0^2 x$ is the same as Eq. (21.2) if $\omega_0^2 = k/m$.

The next thing we must investigate is the physical significance of ω_0. We know that the cosine function repeats itself when the angle it refers to is 2π. So $x = \cos \omega_0 t$ will repeat its motion, it will go through a complete cycle, when the "angle" changes by 2π. The quantity $\omega_0 t$ is often called the *phase* of the motion. In order to change $\omega_0 t$ by 2π, the time must change by an amount t_0, called the *period* of one complete oscillation; of course t_0 must be such that $\omega_0 t_0 = 2\pi$. That is, $\omega_0 t_0$ must account for one cycle of the angle, and then everything will repeat itself—if we increase t by t_0, we add 2π to the phase. Thus

$$t_0 = 2\pi/\omega_0 = 2\pi\sqrt{m/k}. \tag{21.5}$$

Thus if we had a heavier mass, it would take longer to oscillate back and forth on a spring. That is because it has more inertia, and so, while the forces are the same, it takes longer to get the mass moving. Or, if the spring is stronger, it will move more quickly, and that is right: the period is less if the spring is stronger.

Note that the period of oscillation of a mass on a spring does not depend in any way on *how* it has been started, how far down we pull it. The *period* is determined, but the amplitude of the oscillation is *not* determined by the equation of motion (21.2). The amplitude *is* determined, in fact, by how we let go of it, by what we call the *initial conditions* or starting conditions.

Actually, we have not quite found the most general possible solution of Eq. (21.2). There are other solutions. It should be clear why: because all of the cases covered by $x = a \cos \omega_0 t$ start with an initial displacement and no initial velocity. But it is possible, for instance, for the mass to start at $x = 0$, and we may then give it an impulsive kick, so that it has some speed at $t = 0$. Such a motion is not represented by a cosine—it is represented by a sine. To put it another way, if $x = \cos \omega_0 t$ is a solution, then is it not obvious that if we were to happen to walk into the room at some time (which we would call "$t = 0$") and saw the mass as it was passing $x = 0$, it would keep on going just the same? Therefore, $x = \cos \omega_0 t$ cannot be the most general solution; it must be possible to shift the beginning of time, so to speak. As an example, we could write the solution this way: $x = a \cos \omega_0(t - t_1)$, where t_1 is some constant. This also corresponds to shifting the origin of time to some new instant. Furthermore, we may expand

$$\cos(\omega_0 t + \Delta) = \cos \omega_0 t \cos \Delta - \sin \omega_0 t \sin \Delta,$$

and write

$$x = A \cos \omega_0 t + B \sin \omega_0 t,$$

where $A = a \cos \Delta$ and $B = -a \sin \Delta$. Any one of these forms is a possible way to write the complete, general solution of (21.2): that is, every solution of the differential equation $d^2x/dt^2 = -\omega_0^2 x$ that exists in the world can be written as

$$\begin{aligned} &\text{(a)} \quad x = a \cos \omega_0(t - t_1), \\ &\text{or} \\ &\text{(b)} \quad x = a \cos(\omega_0 t + \Delta), \\ &\text{or} \\ &\text{(c)} \quad x = A \cos \omega_0 t + B \sin \omega_0 t. \end{aligned} \tag{21.6}$$

Some of the quantities in (21.6) have names: ω_0 is called the *angular frequency*; it is the number of radians by which the phase changes in a second. That is determined by the differential equation. The other constants are not determined by the equation, but by how the motion is started. Of these constants, a measures the maximum displacement attained by the mass, and is called the *amplitude* of oscillation. The constant Δ is sometimes called the *phase* of the oscillation, but that is a confusion, because other people call $\omega_0 t + \Delta$ the phase, and say the phase changes with time. We might say that Δ is a *phase shift* from some defined zero. Let us put it differently. Different Δ's correspond to motions in different phases. That is true, but whether we want to call Δ *the* phase, or not, is another question.

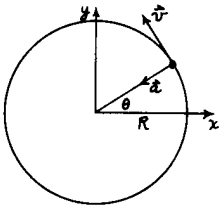


Fig. 21–2. A particle moving in a circular path at constant speed.

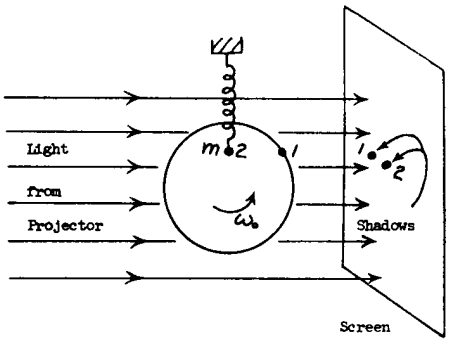


Fig. 21–3. Demonstration of the equivalence between simple harmonic motion and uniform circular motion.

21–3 Harmonic motion and circular motion

The fact that cosines are involved in the solution of Eq. (21.2) suggests that there might be some relationship to circles. This is artificial, of course, because there is no circle actually involved in the linear motion—it just goes up and down. We may point out that we have, in fact, already solved that differential equation when we were studying the mechanics of circular motion. If a particle moves in a circle with a constant speed v, the radius vector from the center of the circle to the particle turns through an angle whose size is proportional to the time. If we call this angle $\theta = vt/R$ (Fig. 21–2) then $d\theta/dt = \omega_0 = v/R$. We know that there is an acceleration $a = v^2/R = \omega_0^2 R$ toward the center. Now we also know that the position x, at a given moment, is the radius of the circle times cos θ, and that y is the radius times sin θ:

$$x = R\cos\theta, \qquad y = R\sin\theta.$$

Now what about the acceleration? What is the x-component of acceleration, d^2x/dt^2? We have already worked that out geometrically; it is the magnitude of the acceleration times the cosine of the projection angle, with a minus sign because it is toward the center.

$$a_x = -a\cos\theta = -\omega^2 R\cos\theta = -\omega^2 x. \tag{21.7}$$

In other words, when a particle is moving in a circle, the horizontal component of its motion has an acceleration which is proportional to the horizontal displacement from the center. Of course we also have the solution for motion in a circle: $x = R\cos\omega_0 t$. Equation (21.7) does not depend upon the radius of the circle, so for a circle of any radius, one finds the same equation for a given ω_0. Thus, for several reasons, we expect that the displacement of a mass on a spring will turn out to be proportional to cos $\omega_0 t$, and will, in fact, be exactly the same motion as we would see if we looked at the x-component of the position of an object rotating in a circle with angular velocity ω_0. As a check on this, one can devise an experiment to show that the up-and-down motion of a mass on a spring is the same as that of a point going around in a circle. In Fig. 21–3 an arc light projected on a screen casts shadows of a crank pin on a shaft and of a vertically oscillating mass, side by side. If we let go of the mass at the right time from the right place, and if the shaft speed is carefully adjusted so that the frequencies match, each should follow the other exactly. One can also check the numerical solution we obtained earlier with the cosine function, and see whether that agrees very well.

Here we may point out that because uniform motion in a circle is so closely related mathematically to oscillatory up-and-down motion, we can analyze oscillatory motion in a simpler way if we imagine it to be a projection of something going in a circle. In other words, although the distance y means nothing in the oscillator problem, we may still artificially supplement Eq. (21.2) with another equation using y, and put the two together. If we do this, we will be able to analyze our one-dimensional oscillator with *circular* motions, which is a lot easier than having to solve a differential equation. The trick in doing this is to use complex numbers, a procedure we shall introduce in the next chapter.

21–4 Initial conditions

Now let us consider what determines the constants A and B, or a and Δ. Of course these are determined by how we start the motion. If we start the motion with just a small displacement, that is one type of oscillation; if we start with an initial displacement and then push up when we let go, we get still a different motion. The constants A and B, or a and Δ, or any other way of putting it, are determined, of course, by the way the motion started, not by any other features of the situation. These are called the *initial conditions*. We would like to connect the initial conditions with the constants. Although this can be done using any one of the forms (21.6), it turns out to be easiest if we use Eq. (21.6c). Suppose that at $t = 0$ we have started with an initial displacement x_0 and a certain velocity v_0.

This is the most general way we can start the motion. (We cannot specify the *acceleration* with which it started, true, because that is determined by the spring, once we specify x_0.) Now let us calculate A and B. We start with the equation for x,

$$x = A\cos\omega_0 t + B\sin\omega_0 t.$$

Since we shall later need the velocity also, we differentiate x and obtain

$$v = -\omega_0 A\sin\omega_0 t + \omega_0 B\cos\omega_0 t.$$

These expressions are valid for all t, but we have special knowledge about x and v at $t = 0$. So if we put $t = 0$ into these equations, on the left we get x_0 and v_0, because that is what x and v are at $t = 0$; also, we know that the cosine of zero is unity, and the sine of zero is zero. Therefore we get

$$x_0 = A\cdot 1 + B\cdot 0 = A$$

and

$$v_0 = -\omega_0 A\cdot 0 + \omega_0 B\cdot 1 = \omega_0 B.$$

So for this particular case we find that

$$A = x_0, \qquad B = v_0/\omega_0.$$

From these values of A and B, we can get a and Δ if we wish.

That is the end of our solution, but there is one physically interesting thing to check, and that is the conservation of energy. Since there are no frictional losses, energy ought to be conserved. Let us use the formula

$$x = a\cos(\omega_0 t + \Delta);$$

then

$$v = -\omega_0 a\sin(\omega_0 t + \Delta).$$

Now let us find out what the kinetic energy T is, and what the potential energy U is. The potential energy at any moment is $\frac{1}{2}kx^2$, where x is the displacement and k is the constant of the spring. If we substitute for x, using our expression above, we get

$$U = \tfrac{1}{2}kx^2 = \tfrac{1}{2}ka^2\cos^2(\omega_0 t + \Delta).$$

Of course the potential energy is not constant; the potential never becomes negative, naturally—there is always some energy in the spring, but the amount of energy fluctuates with x. The kinetic energy, on the other hand, is $\frac{1}{2}mv^2$, and by substituting for v we get

$$T = \tfrac{1}{2}mv^2 = \tfrac{1}{2}m\omega_0^2 a^2\sin^2(\omega_0 t + \Delta).$$

Now the kinetic energy is zero when x is at the maximum, because then there is no velocity; on the other hand, it is maximal when x is passing through zero, because then it is moving fastest. This variation of the kinetic energy is just the opposite of that of the potential energy. But the total energy ought to be a constant. If we note that $k = m\omega_0^2$, we see that

$$T + U = \tfrac{1}{2}m\omega_0^2 a^2[\cos^2(\omega_0 t + \Delta) + \sin^2(\omega_0 t + \Delta)] = \tfrac{1}{2}m\omega_0^2 a^2.$$

The energy is dependent on the square of the amplitude; if we have twice the amplitude, we get an oscillation which has four times the energy. The *average* potential energy is half the maximum and, therefore, half the total, and the average kinetic energy is likewise half the total energy.

21–5 Forced oscillations

Next we shall discuss the *forced harmonic oscillator*, i.e., one in which there is an external driving force acting. The equation then is the following:

$$m\,d^2x/dt^2 = -kx + F(t). \tag{21.8}$$

We would like to find out what happens in these circumstances. The external driving force can have various kinds of functional dependence on the time; the first one that we shall analyze is very simple—we shall suppose that the force is oscillating:

$$F(t) = F_0 \cos \omega t. \tag{21.9}$$

Notice, however, that this ω is not necessarily ω_0: we have ω under our control; the forcing may be done at different frequencies. So we try to solve Eq. (21.8) with the special force (21.9). What is the solution of (21.8)? One special solution, (we shall discuss the more general cases later) is

$$x = C \cos \omega t, \tag{21.10}$$

where the constant is to be determined. In other words, we might suppose that if we kept pushing back and forth, the mass would follow back and forth in step with the force. We can try it anyway. So we put (21.10) into (21.9), and get

$$-m\omega^2 C \cos \omega t = -m\omega_0^2 C \cos \omega t + F_0 \cos \omega t. \tag{21.11}$$

We have also put in $k = m\omega_0^2$, so that we will understand the equation better at the end. Now because the cosine appears everywhere, we can divide it out, and that shows that (21.10) is, in fact, a solution, provided we pick C just right. The answer is that C must be

$$C = F_0/m(\omega_0^2 - \omega^2). \tag{21.12}$$

That is, m oscillates at the same frequency as the force, but with an amplitude which depends on the frequency of the force, and also upon the frequency of the natural motion of the oscillator. It means, first, that if ω is very small compared with ω_0, then the displacement and the force are in the same direction. On the other hand, if we shake it back and forth very fast, then (21.12) tells us that C is negative if ω is above the natural frequency ω_0 of the harmonic oscillator. (We will call ω_0 the natural frequency of the harmonic oscillator, and ω the applied frequency.) At very high frequency the denominator may become very large, and there is then not much amplitude.

Of course the solution we have found is the solution only if things are started just right, for otherwise there is a part which usually dies out after a while. This other part is called the *transient* response to $F(t)$, while (21.10) and (21.12) are called the *steady-state* response.

According to our formula (21.12), a very remarkable thing should also occur: if ω is almost exactly the same as ω_0, then C should approach infinity. So if we adjust the frequency of the force to be "in time" with the natural frequency, then we should get an enormous displacement. This is well known to anybody who has pushed a child on a swing. It does not work very well to close our eyes and push at a certain speed at random. If we happen to get the right timing, then the swing goes very high, but if we have the wrong timing, then sometimes we may be pushing when we should be pulling, and so on, and it does not work.

If we make ω exactly equal to ω_0, we find that it should oscillate at an *infinite* amplitude, which is, of course, impossible. The reason it does not is that something goes wrong with the equation, there are some other frictional terms, and other forces, which are not in (21.8) but which occur in the real world. So the amplitude does not reach infinity for some reason; it may be that the spring breaks!

22

Algebra

22–1 Addition and multiplication

22–2 The inverse operations

22–3 Abstraction and generalization

22–4 Approximating irrational numbers

22–5 Complex numbers

22–6 Imaginary exponents

22–1 Addition and multiplication

In our study of oscillating systems we shall have occasion to use one of the most remarkable, almost astounding, formulas in all of mathematics. From the physicist's point of view we could bring forth this formula in two minutes or so, and be done with it. But science is as much for intellectual enjoyment as for practical utility, so instead of just spending a few minutes on this amazing jewel, we shall surround the jewel by its proper setting in the grand design of that branch of mathematics which is called elementary algebra.

Now you may ask, "What is mathematics doing in a physics lecture?" We have several possible excuses: first, of course, mathematics is an important tool, but that would only excuse us for giving the formula in two minutes. On the other hand, in theoretical physics we discover that all our laws can be written in mathematical form; and that this has a certain simplicity and beauty about it. So, ultimately, in order to understand nature it may be necessary to have a deeper understanding of mathematical relationships. But the real reason is that the subject is enjoyable, and although we humans cut nature up in different ways, and we have different courses in different departments, such compartmentalization is really artificial, and we should take our intellectual pleasures where we find them.

Another reason for looking more carefully at algebra now, even though most of us studied algebra in high school, is that that was the first time we studied it; all the equations were unfamiliar, and it was hard work, just as physics is now. Every so often it is a great pleasure to look back to see what territory has been covered, and what the great map or plan of the whole thing is. Perhaps some day somebody in the Mathematics Department will present a lecture on mechanics in such a way as to show what it was we were trying to learn in the physics course!

The subject of algebra will not be developed from the point of view of a mathematician, exactly, because the mathematicians are mainly interested in how various mathematical facts are demonstrated, and how many assumptions are absolutely required, and what is not required. They are not so interested in the result of what they prove. For example, we may find the Pythagorean theorem quite interesting, that the sum of the squares of the sides of a right triangle is equal to the square of the hypotenuse; that is an interesting fact, a curiously simple thing, which may be appreciated without discussing the question of how to prove it, or what axioms are required. So, in the same spirit, we shall describe qualitatively, if we may put it that way, the system of elementary algebra. We say *elementary* algebra because there is a branch of mathematics called *modern* algebra in which some of the rules such as $ab = ba$, are abandoned, and it is still called algebra, but we shall not discuss that.

To discuss this subject we start in the middle. We suppose that we already know what integers are, what zero is, and what it means to increase a number by one unit. You may say, "That is not in the middle!" But it is the middle from a mathematical standpoint, because we could go even further back and describe the theory of sets in order to *derive* some of these properties of integers. But we are not going in that direction, the direction of mathematical philosophy and mathematical logic, but rather in the other direction, from the assumption that we know what integers are and we know how to count.

If we start with a certain number a, an integer, and we count successively one unit b times, the number we arrive at we call $a + b$, and that defines *addition* of integers.

Once we have defined addition, then we can consider this: if we start with nothing and add a to it, b times in succession, we call the result *multiplication* of integers; we call it *b times a.*

Now we can also have a *succession of multiplications:* if we start with 1 and multiply by a, b times in succession, we call that *raising to a power:* a^b.

Now as a consequence of these definitions it can be easily shown that all of the following relationships are true:

$$\begin{array}{ll} \text{(a)}\ a+b=b+a & \text{(b)}\ a+(b+c)=(a+b)+c \\ \text{(c)}\ ab=ba & \text{(d)}\ a(b+c)=ab+ac \\ \text{(e)}\ (ab)c=a(bc) & \text{(f)}\ (ab)^c=a^cb^c \\ \text{(g)}\ a^ba^c=a^{(b+c)} & \text{(h)}\ (a^b)^c=a^{(bc)} \\ \text{(i)}\ a+0=a & \text{(j)}\ a\cdot 1=a \\ \text{(k)}\ a^1=a & \end{array} \tag{22.1}$$

These results are well known and we shall not belabor the point, we merely list them. Of course, 1 and 0 have special properties; for example, $a+0$ is a, a times $1=a$, and a to the first power is a.

In this discussion we must also assume a few other properties like continuity and ordering, which are very hard to define; we will let the rigorous theory do it. Furthermore, it is definitely true that we have written down too many "rules"; some of them may be deducible from the others, but we shall not worry about such matters.

22–2 The inverse operations

In addition to the direct operations of addition, multiplication, and raising to a power, we have also the *inverse* operations, which are defined as follows. Let us assume that a and c are given, and that we wish to find what values of b satisfy such equations as $a+b=c$, $ab=c$, $b^a=c$. If $a+b=c$, b is defined as $c-a$, which is called *subtraction.* The operation called division is also clear: if $ab=c$, then $b=c/a$ defines division—a solution of the equation $ab=c$ "backwards." Now if we have a power $b^a=c$ and we ask ourselves, "What is b?," it is called the ath *root* of c: $b=\sqrt[a]{c}$. For instance, if we ask ourselves the following question, "What integer, raised to the third power, equals 8?," then the answer is called the *cube root* of 8; it is 2. Because b^a and a^b are not equal, there are *two* inverse problems associated with powers, and the other inverse problem would be, "To what power must we raise 2 to get 8?" This is called taking the *logarithm.* If $a^b=c$, we write $b=\log_a c$. The fact that it has a cumbersome notation relative to the others does not mean that it is any less elementary, at least applied to integers, than the other processes. Although logarithms come late in an algebra class, in practice they are, of course, just as simple as roots; they are just a different kind of solution of an algebraic equation. The direct and inverse operations are summarized as follows:

$$\begin{array}{ll} \text{(a) addition} & \text{(a′) subtraction} \\ a+b=c & b=c-a \\ \text{(b) multiplication} & \text{(b′) division} \\ ab=c & b=c/a \\ \text{(c) power} & \text{(c′) root} \\ b^a=c & b=\sqrt[a]{c} \\ \text{(d) power} & \text{(d′) logarithm} \\ a^b=c & b=\log_a c \end{array} \tag{22.2}$$

Now here is the idea. These relationships, or rules, are correct for integers, since they follow from the definitions of addition, multiplication, and raising to a power. *We are going to discuss whether or not we can broaden the class of objects*

which a, b, and c represent so that they will obey these same rules, although the processes for $a + b$, and so on, will not be definable in terms of the direct action of adding 1, for instance, or successive multiplications by integers.

22–3 Abstraction and generalization

When we try to solve simple algebraic equations using all these definitions, we soon discover some insoluble problems, such as the following. Suppose that we try to solve the equation $b = 3 - 5$. That means, according to our definition of subtraction, that we must find a number which, when added to 5, gives 3. And of course there *is* no such number, because we consider only positive integers; *this* is an insoluble problem. However, the plan, the great idea, is this: *abstraction and generalization.* From the whole structure of algebra, rules plus integers, we abstract the original definitions of addition and multiplication, but we leave the rules (22.1) and (22.2), and assume these to be true *in general* on a wider class of numbers, even though they are originally derived on a smaller class. Thus, rather than using integers symbolically to define the rules, we use the rules as the definition of the symbols, which then represent a more general kind of number. As an example, by working with the rules alone we can show that $3 - 5 = 0 - 2$. In fact we can show that one can make *all* subtractions, provided we define a whole set of new numbers: $0 - 1, 0 - 2, 0 - 3, 0 - 4$, and so on, called the *negative integers*. Then we may use all the other rules, like $a(b + c) = ab + ac$ and so forth, to find what the rules are for multiplying negative numbers, and we will discover, in fact, that all of the rules can be maintained with negative as well as positive integers.

So we have increased the range of objects over which the rules work, but the meaning of the symbols is different.

One cannot say, for instance, that -2 times 5 really means to add 5 together successively -2 times. That means nothing. But nevertheless everything will work out all right according to the rules.

An interesting problem comes up in taking powers. Suppose that we wish to discover what $a^{(3-5)}$ means. We know only that $3 - 5$ is a solution of the problem, $(3 - 5) + 5 = 3$. Knowing that, we know that $a^{(3-5)}a^5 = a^3$. Therefore $a^{(3-5)} = a^3/a^5$, by the definition of division. With a little more work, this can be reduced to $1/a^2$. So we find that the negative powers are the reciprocals of the positive powers, but $1/a^2$ is a meaningless symbol, because if a is a positive or negative integer, the square of it is greater than 1, and we do not yet know what we mean by 1 divided by a number greater than 1!

Onward! The great plan is to continue the process of generalization; whenever we find another problem that we cannot solve we extend our realm of numbers. Consider division: we cannot find a number which is an integer, even a negative integer, which is equal to the result of dividing 3 by 5. But if we suppose that all fractional numbers also satisfy the rules, then we can talk about multiplying and adding fractions, and everything works as well as it did before.

Take another example of powers: what is $a^{3/5}$? We know only that $(3/5)5 = 3$, since that was the definition of $3/5$. So we know also that $(a^{(3/5)})^5 = a^{(3/5)(5)} = a^3$, because this is one of the rules. Then by the definition of roots we find that $a^{(3/5)} = \sqrt[5]{a^3}$.

In this way, then, we can define what we mean by putting fractions in the various symbols, by using the rules themselves to help us determine the definition—it is not arbitrary. It is a remarkable fact that all the rules still work for positive and negative integers, as well as for fractions!

We go on in the process of generalization. Are there any other equations we cannot solve? Yes, there are. For example, it is impossible to solve this equation: $b = 2^{1/2} = \sqrt{2}$. It is impossible to find a number which is rational (a fraction) whose square is equal to 2. It is very easy for us in modern days to answer this question. We know the decimal system, and so we have no difficulty in appreciating the meaning of an unending decimal as a type of approximation to the square root of 2. Historically, this idea presented great difficulty to the Greeks. To really

define *precisely* what is meant here requires that we add some substance of continuity and ordering, and it is, in fact, quite the most difficult step in the processes of generalization just at this point. It was made, formally and rigorously, by Dedekind. However, without worrying about the mathematical rigor of the thing, it is quite easy to understand that what we mean is that we are going to find a whole sequence of approximate fractions, perfect fractions (because any decimal, when stopped somewhere, is of course rational), which just keeps on going, getting closer and closer to the desired result. That is good enough for what we wish to discuss, and it permits us to involve ourselves in irrational numbers, and to calculate things like the square root of 2 to any accuracy that we desire, with enough work.

22–4 Approximating irrational numbers

The next problem comes with what happens with the irrational powers. Suppose that we want to define, for instance, $10^{\sqrt{2}}$. In principle, the answer is simple enough. If we approximate the square root of 2 to a certain number of decimal places, then the power is rational, and we can take the approximate root, using the above method, and get an *approximation* to $10^{\sqrt{2}}$. Then we may run it up a few more decimal places (it is again rational), take the appropriate root, this time a much higher root because there is a much bigger denominator in the fraction, and get a better approximation. Of course we are going to get some enormously high roots involved here, and the work is quite difficult. How can we cope with this problem?

In the computations of square roots, cube roots, and other small roots, there is an arithmetical process available by which we can get one decimal place after another. But the amount of labor needed to calculate irrational powers and the logarithms that go with them (the inverse problem) is so great that there is no simple arithmetical process we can use. Therefore tables have been built up which permit us to calculate these powers, and these are called the tables of logarithms, or the tables of powers, depending on which way the table is set up. It is merely a question of saving time; if we must raise some number to an irrational power, we can look it up rather than having to compute it. Of course, such a computation is just a technical problem, but it is an interesting one, and of great historical value. In the first place, not only do we have the problem of solving $x = 10^{\sqrt{2}}$, but we also have the problem of solving $10^x = 2$, or $x = \log_{10} 2$. This is not a problem where we have to define a new kind of number for the result, it is merely a computational problem. The answer is simply an irrational number, an unending decimal, not a new kind of a number.

Let us now discuss the problem of calculating solutions of such equations. The general idea is really very simple. If we could calculate 10^1, and $10^{4/10}$, and $10^{1/100}$, and $10^{4/1000}$ and so on, and multiply them all together, we would get $10^{1.414\cdots}$ or $10^{\sqrt{2}}$, and that is the general idea on which things work. But instead of calculating $10^{1/10}$ and so on, we shall calculate $10^{1/2}$, $10^{1/4}$, and so on. Before we start, we should explain why we make so much work with 10, instead of some other number. Of course, we realize that logarithm tables are of great practical utility, quite aside from the mathematical problem of taking roots, since with any base at all,

$$\log_b (ac) = \log_b a + \log_b c. \tag{22.3}$$

We are all familiar with the fact that one can use this fact in a practical way to multiply numbers if we have a table of logarithms. The only question is, with what base b shall we compute? It makes no difference what base is used; we can use the same principle all the time, and if we are using logarithms to any particular base, we can find logarithms to any other base merely by a change in scale, a multiplying factor. If we multiply Eq. (22.3) by 61, it is just as true, and if we had a table of logs with a base b, and somebody else multiplied all of our table by 61, there would be no essential difference. Suppose that we know the logarithms of all the numbers to the base b. In other words, we can solve the equation $b^a = c$ for any c because

we have a table. The problem is to find the logarithm of the same number c to some other base, let us say the base x. We would like to solve $x^{a'} = c$. It is easy to do, because we can always write $x = b^t$, which defines t, knowing x and b. As a matter of fact, $t = \log_b x$. Then if we put that in and solve for a', we see that $(b^t)^{a'} = b^{ta'} = c$. In other words, ta' is the logarithm of c in base b. Thus $a' = a/t$. Thus logs to base x are just $1/t$, which is a constant, times the logs to the base, b. Therefore any log table is equivalent to any other log table if we multiply by a constant, and the constant is $1/\log_b x$. This permits us to choose a particular base, and for convenience we take the base 10. (The question may arise as to whether there is any natural base, any base in which things are somehow simpler, and we shall try to find an answer to that later. At the moment we shall just use the base 10.)

Table 22–1

Successive Square Roots of Ten

Power s	$1024\,s$	10^s	$(10^s - 1)/s$
1	1024	10.00000	9.00
1/2	512	3.16228	4.32
1/4	256	1.77828	3.113
1/8	128	1.33352	2.668
1/16	64	1.15478	2.476
1/32	32	1.074607	2.3874
1/64	16	1.036633	2.3445
1/128	8	1.018152	2.3234^{211}
1/256	4	1.0090350	2.3130^{104}
1/512	2	1.0045073	2.3077^{53}
1/1024	1	1.0022511	2.3051^{26}
			^{26} ↓
$\Delta/1024$ ($\Delta \to 0$)	Δ	$1 + .0022486\Delta$ ←	2.3025

Now let us see how to calculate logarithms. We begin by computing successive square roots of 10, by cut and try. The results are shown in Table 22–1. The powers of 10 are given in the first column, and the result, 10^s, is given in the third column. Thus $10^1 = 10$. The one-half power of 10 we can easily work out, because that is the square root of 10, and there is a known, simple process for taking square roots of any number.* Using this process, we find the first square root to be 3.16228. What good is that? It already tells us something, it tells us how to take $10^{0.5}$, so we now know at least *one* logarithm, if we happen to need the logarithm of 3.16228, we know the answer is close to 0.50000. But we must do a little bit better than that; we clearly need more information. So we take the square root again, and find $10^{1/4}$, which is 1.77828. Now we have the logarithm of more numbers than we had before, 1.250 is the logarithm of 17.78 and, incidentally, if it happens that somebody asks for $10^{0.75}$, we can get it, because that is $10^{(0.5+0.25)}$; it is therefore the product of the second and third numbers. If we can get enough numbers in column s to be able to make up almost any number, then by multiplying the proper things in column 3, we can get 10 to any power; that is the plan. So we evaluate ten successive square roots of 10, and that is the main work which is involved in the calculations.

Why don't we keep on going for more and more accuracy? Because we begin to notice something. When we raise 10 to a very small power, we get 1 plus a small amount. The reason for this is clear, because we are going to have to take the

* There is a definite arithmetic procedure, but the easiest way to find the square root of any number N is to choose some a fairly close, find N/a, average $a' = \frac{1}{2}[a + (N/a)]$, and use this average a' for the next choice for a. The convergence is very rapid—the number of significant figures doubles each time.

1000th power of $10^{1/1000}$ to get back to 10, so we had better not start with too big a number; it has to be close to 1. What we notice is that the small numbers that are added to 1 begin to look as though we are merely dividing by 2 each time; we see 1815 becomes 903, then 450, 225; so it is clear that, to an excellent approximation, if we take another root, we shall get 1.00112 something, and rather than actually *take* all the square roots, we *guess* at the ultimate limit. When we take a small fraction Δ of 1024 as Δ approaches zero, what will the answer be? Of course it will be some number close to $0.0022511\,\Delta$. Not exactly $0.0022511\,\Delta$, however—we can get a better value by the following trick: we subtract the 1, and then divide by the power s. This ought to correct all the excesses to the same value. We see that they are very closely equal. At the top of the table they are not equal, but as they come down, they get closer and closer to a constant value. What is the value? Again we look to see how the series is going, how it has changed with s. It changed by 211, by 104, by 53, by 26. These changes are obviously half of each other, very closely, as we go down. Therefore, if we kept going, the changes would be 13, 7, 3, 2 and 1, more or less, or a total of 26. Thus we have only 26 more to go, and so we find that the true number is 2.3025. (Actually, we shall later see that the *exact* number should be 2.3026, but to keep it realistic, we shall not alter anything in the arithmetic.) From this table we can now calculate any power of 10, by compounding the power out of 1024ths.

Let us now actually calculate a logarithm, because the process we shall use is where logarithm tables actually come from. The procedure is shown in Table 22–2, and the numerical values are shown in Table 22–1 (columns 2 and 3).

Table 22–2

Calculation of a logarithm: $\log_{10} 2$

$$2 \div 1.77828 = 1.124682$$

$$1.124682 \div 1.074607 = 1.046598, \text{ etc.}$$

$$\therefore 2 = (1.77828)(1.074607)(1.036633)(1.090350)(1.000573)$$

$$= 10\left[\frac{1}{1024}(256 + 32 + 16 + 4 + 0.254)\right] = 10\left[\frac{308.254}{1024}\right]$$

$$= 10^{0.30103} \qquad \left(\frac{573}{2249} = 0.254\right)$$

$$\therefore \log_{10} 2 = 0.30103$$

Suppose we want the logarithm of 2. That is, we want to know to what power we must raise 10 to get 2. Can we raise 10 to the 1/2 power? No; that is too big. In other words, we can see that the answer is going to be bigger than 1/4, and less than 1/2. Let us take the factor $10^{1/4}$ out; we divide 2 by 1.778 . . . , and get 1.124 . . . , and so on, and now we know that we have taken away 0.250000 from the logarithm. The number 1.124 . . . is now the number whose logarithm we need. When we are finished we shall add back the 1/4, or 256/1024. Now we look in the table for the next number just below 1.124 . . . , and that is 1.074607. We therefore divide by 1.074607 and get 1.046598. From that we discover that 2 can be made up of a product of numbers that are in Table 22–1, as follows:

$$2 = (1.77828)(1.074607)(1.036633)(1.0090350)(1.000573).$$

There was one factor (1.000573) left over, naturally, which is beyond the range of our table. To get the logarithm of this factor, we use our result that $10^{\Delta/1024} \approx 1 + 2.3025\,\Delta/1024$. We find $\Delta = 0.254$. Therefore our answer is 10 to the following power: $(256 + 32 + 16 + 4 + 0.254)/1024$. Adding those together, we get 308.254/1024. Dividing, we get 0.30103, so we know that the $\log_{10} 2 = 0.30103$, which happens to be right to 5 figures!

This is how logarithms were originally computed by Mr. Briggs of Halifax, in 1620. He said, "I computed successively 54 square roots of 10." We know he

really computed only the first 27, because the rest of them can be obtained by this trick with Δ. His work involved calculating the square root of 10 twenty-seven times, which is not much more than the ten times we did; however, it was more work because he calculated to sixteen decimal places, and then reduced his answer to fourteen when he published it, so that there were no rounding errors. He made tables of logarithms to fourteen decimal places by this method, which is quite tedious. But all logarithm tables for three hundred years were borrowed from Mr. Briggs' tables by reducing the number of decimal places. Only in modern times, with the WPA and computing machines, have new tables been independently computed. There are much more efficient methods of computing logarithms today, using certain series expansions.

In the above process, we discovered something rather interesting, and that is that for very small powers ϵ we can calculate 10^ϵ easily; we have discovered that $10^\epsilon = 1 + 2.3025\epsilon$, by sheer numerical analysis. Of course this also means that $10^{n/2.3025} = 1 + n$ if n is very small. Now logarithms to any other base are merely multiples of logarithms to the base 10. The base 10 was used only because we have 10 fingers, and the arithmetic of it is easy, but if we ask for a mathematically natural base, one that has nothing to do with the number of fingers on human beings, we might try to change our *scale* of logarithms in some convenient and natural manner, and the method which people have chosen is to redefine the logarithms by multiplying all the logarithms to the base 10 by 2.3025 . . . This then corresponds to using some other base, and this is called the *natural* base, or base e. Note that $\log_e (1 + n) \approx n$, or $e^n \approx 1 + n$ as $n \to 0$.

It is easy enough to find out what e is: $e = 10^{1/2.3025}$ or $10^{0.434294\ldots}$, an irrational power. Our table of the successive square roots of 10 can be used to compute, not just logarithms, but also 10 to any power, so let us use it to calculate this natural base e. For convenience we transform 0.434294 . . . into 444.73/1024. Now, 444.73 is 256 + 128 + 32 + 16 + 2 + 0.73. Therefore e, since it is an exponent of a sum, will be a product of the numbers

$$(1.77828)(1.33352)(1.074607)(1.036633)(1.018152)(1.009035)(1.001643) = 2.7184.$$

(The only problem is the last one, which is 0.73, and which is not in the table, but we know that if Δ is small enough, the answer is $1 + 2.3025\,\Delta$.) When we multiply all these together, we get 2.7184 (it should be 2.7183, but it is good enough). The use of such tables, then, is the way in which irrational powers and the logarithms of irrational numbers are all calculated. That takes care of the irrationals.

22–5 Complex numbers

Now it turns out that after all that work we *still* cannot solve every equation! For instance, what is the square root of -1? Suppose we have to find $x^2 = -1$. The square of no rational, of no irrational, of *nothing* that we have discovered so far, is equal to -1. So we again have to generalize our numbers to a still wider class. Let us suppose that a specific solution of $x^2 = -1$ is called something, we shall call it i; i has the property, by definition, that its square is -1. That is about all we are going to say about it; of course, there is more than one root of the equation $x^2 = -1$. Someone could write i, but another could say, "No, I prefer $-i$. My i is minus your i." It *is* just as good a solution, and since the only definition that i has is that $i^2 = -1$, it must be true that any equation we can write is equally true if the sign of i is changed everywhere. This is called taking the *complex conjugate*. Now we are going to make up numbers by adding successive i's, and multiplying i's by numbers, and adding other numbers, and so on, according to all of our rules. In this way we find that numbers can all be written in the form $p + iq$, where p and q are what we call *real* numbers, i.e., the numbers we have been defining up until now. The number i is called the *unit imaginary* number. Any real multiple of i is called *pure imaginary*. The most general number, a, is of the form $p + iq$ and is called a *complex number*. Things do not get any worse if, for instance, we multiply two such numbers, let us say $(r + is)(p + iq)$. Then,

using the rules, we get

$$\begin{aligned}(r + is)(p + iq) &= rp + r(iq) + (is)p + (is)(iq)\\ &= rp + i(rq) + i(sp) + (ii)(sq)\\ &= (rp - sq) + i(rq + sp),\end{aligned} \tag{22.4}$$

since $ii = i^2 = -1$. Therefore all the numbers that now belong in the rules (22.1) have this mathematical form.

Now you say, "This can go on forever! We have defined powers of imaginaries and all the rest, and when we are all finished, somebody else will come along with another equation which cannot be solved, like $x^6 + 3x^2 = -2$. Then we have to generalize all over again!" But it turns out that *with this one more invention*, just the square root of -1, *every algebraic equation can be solved!* This is a fantastic fact, which we must leave to the Mathematics Department to prove. The proofs are very beautiful and very interesting, but certainly not self-evident. In fact, the most obvious supposition is that we are going to have to invent again and again and again. But the greatest miracle of all is that we do not. This is the last invention. After this invention of complex numbers, we find that the rules still work with complex numbers, and we are finished inventing new things. We can find the complex power of any complex number, we can solve any equation that is written algebraically, in terms of a finite number of those symbols. We do not find any new numbers. The square root of i, for instance, has a definite result, it is not something new; and i^i is something. We will discuss that now.

We have already discussed multiplication, and addition is also easy; if we add two complex numbers, $(p + iq) + (r + is)$, the answer is $(p + r) + i(q + s)$. Now we can add and multiply complex numbers. But the real problem, of course, is to compute *complex powers of complex numbers.* It turns out that the problem is actually no more difficult than computing complex powers of real numbers. So let us concentrate now on the problem of calculating 10 to a complex power, not just an irrational power, but $10^{(r+is)}$. Of course, we must at all times use our rules (22.1) and (22.2). Thus

$$10^{(r+is)} = 10^r 10^{is}. \tag{22.5}$$

But 10^r we already know how to compute, and we can always multiply anything by anything else; therefore the problem is to compute only 10^{is}. Let us call it some complex number, $x + iy$. Problem: given s, find x, find y. Now if

$$10^{is} = x + iy,$$

then the complex conjugate of this equation must also be true, so that

$$10^{-is} = x - iy.$$

(Thus we see that we can deduce a number of things without actually computing anything, by using our rules.) We deduce another interesting thing by multiplying these together:

$$10^{is}10^{-is} = 10^0 = 1 = (x + iy)(x - iy) = x^2 + y^2. \tag{22.6}$$

Thus if we find x, we have y also.

Now the problem is *how* to compute 10 to an imaginary power. What guide is there? We may work over our rules until we can go no further, but here is a reasonable guide: if we can compute it for any particular s, we can get it for all the rest. If we know 10^{is} for any one s and then we want it for twice that s, we can square the number, and so on. But how can we find 10^{is} for even one special value of s? To do so we shall make one additional assumption, which is not quite in the category of all the other rules, but which leads to reasonable results and permits us to make progress: when the power is small, we shall suppose that the "law" $10^\epsilon = 1 + 2.3025\epsilon$ is right, as ϵ gets very small, not only for real ϵ, *but for complex ϵ as well.* Therefore, we begin with the supposition that this law is true in general, and that tells us that $10^{is} = 1 + 2.3025 \cdot is$, for $s \to 0$. So we assume that if s is very small, say one part in 1024, we have a rather good approximation to 10^{is}.

Now we make a table by which we can compute *all* the imaginary powers of 10, that is, compute x and y. It is done as follows. The first power we start with is the 1/1024 power, which we presume is very nearly $1 + 2.3025i/1024$. Thus we start with

$$10^{i/1024} = 1.00000 + 0.0022486i, \tag{22.7}$$

and if we keep multiplying the number by itself, we can get to a higher imaginary power. In fact, we may just reverse the procedure we used in making our logarithm table, and calculate the square, 4th power, 8th power, etc., of (22.7), and thus build up the values shown in Table 22–3. We notice an interesting thing, that the x numbers are positive at first, but then swing negative. We shall look into that a little bit more in a moment. But first we may be curious to find for what number s the real part of 10^{is} is *zero*. The y-value would be i, and so we would have $10^{is} = i$, or $is = \log_{10} i$. As an example of how to use this table, just as we calculated $\log_{10} 2$ before, let us now use Table 22–3 to find $\log_{10} i$.

Which of the numbers in Table 22–3 do we have to multiply together to get a pure imaginary result? After a little trial and error, we discover that to reduce x the most, it is best to multiply "512" by "128." This gives $0.13056 + 0.99144i$. Then we discover that we should multiply this by a number whose imaginary part is about equal to the size of the real part we are trying to remove. Thus we choose "64" whose y-value is 0.14349, since that is closest to 0.13056. This then gives $-0.01350 + 0.99993i$. Now we have overshot, and must *divide* by $0.99996 + 0.00900i$. How do we do that? By changing the sign of i and multiplying by $0.99996 - 0.00900i$ (which works if $x^2 + y^2 = 1$). Continuing in this way, we find that the entire power to which 10 must be raised to give i is $i(512 + 128 + 64 - 4 - 2 + 0.20)/1024$, or $698.20i/1024$. If we raise 10 to that power, we can get i. Therefore $\log_{10} i = 0.68226i$.

Table 22–3

Successive Squares of

$10^{i/1024} = 1 + 0.0022486i$

Power is	$1024s$	10^{is}
$i/1024$	1	$1.00000 + 0.00225i$*
$i/512$	2	$1.00000 + 0.00450i$
$i/256$	4	$0.99996 + 0.00900i$
$i/128$	8	$0.99984 + 0.01800i$
$i/64$	16	$0.99936 + 0.03599i$
$i/32$	32	$0.99742 + 0.07193i$
$i/16$	64	$0.98967 + 0.14349i$
$i/8$	128	$0.95885 + 0.28402i$
$i/4$	256	$0.83872 + 0.54467i$
$i/2$	512	$0.40679 + 0.91365i$
$i/1$	1024	$-0.66928 + 0.74332i$

* Should be $0.0022486i$

Table 22–4

Successive Powers of $10^{i/8}$

p = power · $8i$	$10^{ip/8}$
0	$1.00000 + 0.00000i$
1	$0.95882 + 0.28402i$
2	$0.83867 + 0.54465i$
3	$0.64944 + 0.76042i$
4	$0.40672 + 0.91356i$
5	$0.13050 + 0.99146i$
6	$-0.15647 + 0.98770i$
7	$-0.43055 + 0.90260i$
8	$-0.66917 + 0.74315i$
9	$-0.85268 + 0.52249i$
10	$-0.96596 + 0.25880i$
11	$-0.99969 - 0.02620i$
12	$-0.95104 - 0.30905i$
14	$-0.62928 - 0.77717i$
16	$-0.10447 - 0.99453i$
18	$+0.45454 - 0.89098i$
20	$+0.86648 - 0.49967i$
22	$+0.99884 + 0.05287i$
24	$+0.80890 + 0.58836i$

22–6 Imaginary exponents

To further investigate the subject of taking complex imaginary powers, let us look at the powers of 10 taking *successive powers*, not doubling the power each time, in order to follow Table 22–3 further and to see what happens to those minus signs. This is shown in Table 22–4, in which we take $10^{i/8}$, and just keep multiplying it. We see that x decreases, passes through zero, swings almost to -1 (if we could get in between $p = 10$ and $p = 11$ it would obviously swing to -1), and swings back. The y-value is going back and forth too.

In Fig. 22–1 the dots represent the numbers that appear in Table 22–4, and the lines are just drawn to help you visually. So we see that the numbers x and y oscillate; 10^{is} *repeats itself*, it is a *periodic* thing, and as such, it is easy enough to explain, because if a certain power is i, then the fourth power of that would be i^2 *squared*. It would be $+1$ again, and therefore, since $10^{0.68i}$ is equal to i, by taking the fourth power we discover that $10^{2.72i}$ is equal to $+1$. Therefore, if we wanted $10^{3.00i}$, for instance, we could write it as $10^{2.72i}$ times $10^{0.28i}$. In other words, it has a period, it repeats. Of course, we recognize what the curves look like! They look like the sine and cosine, and we shall call them, for a while, the algebraic sine and algebraic cosine. However, instead of using the base 10, we shall put them into our natural base, which only changes the horizontal scale; so we denote $2.3025s$ by t, and write $10^{is} = e^{it}$, where t is a real number. Now $e^{it} = x + iy$, and we shall write this as the algebraic cosine of t plus i times the algebraic sine of t. Thus

$$e^{it} = \underline{\cos}\, t + i\, \underline{\sin}\, t. \tag{22.8}$$

What are the properties of $\underline{\cos}\, t$ and $\underline{\sin}\, t$? First, we know, for instance, that $x^2 + y^2$ must be 1; we have proved that before, and it is just as true for base e as for base 10. Therefore $\underline{\cos}^2 t + \underline{\sin}^2 t = 1$. We also know that, for small t, $e^{it} = 1 + it$, and therefore $\underline{\cos}\, t$ is nearly 1, and $\underline{\sin}\, t$ is nearly t, and so it goes, that *all of the various properties of these remarkable functions*, which come from taking imaginary powers, *are the same as the sine and cosine of trigonometry.*

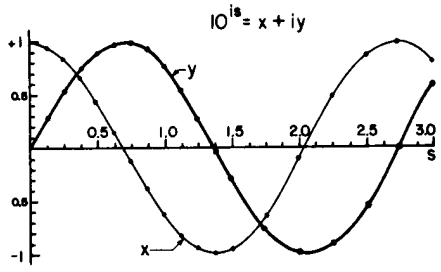


Figure 22–1

Is the period the same? Let us find out. e to what power is equal to i? What is the logarithm of i to the base e? We worked it out before, in the base 10 it was $0.68226i$, but when we change our logarithmic scale to e, we have to multiply by 2.3025, and if we do that it comes out 1.5709. So this will be called "algebraic $\pi/2$." But, we see, it differs from the regular $\pi/2$ by only one place in the last point, and that, of course, is the result of errors in our arithmetic! So we have created two new functions in a purely algebraic manner, the cosine and the sine, which belong to algebra, and only to algebra. We wake up at the end to discover the very functions that are natural to geometry. So there is a connection, ultimately, between algebra and geometry.

We summarize with this, the most remarkable formula in mathematics:

$$e^{i\theta} = \cos\theta + i\sin\theta. \tag{22.9}$$

This is our jewel.

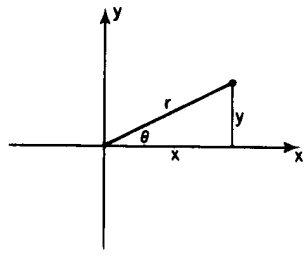


Fig. 22–2. $x + iy = re^{i\theta}$.

We may relate the geometry to the algebra by representing complex numbers in a plane; the horizontal position of a point is x, the vertical position of a point is y (Fig. 22–2). We represent every complex number, $x + iy$. Then if the radial distance to this point is called r and the angle is called θ, the algebraic law is that $x + iy$ is written in the form $re^{i\theta}$, where the geometrical relationships between x, y, r, and θ are as shown. This, then, is the unification of algebra and geometry.

When we began this chapter, armed only with the basic notions of integers and counting, we had little idea of the power of the processes of abstraction and generalization. Using the set of algebraic "laws," or properties of numbers, Eq. (22.1), and the definitions of inverse operations (22.2), we have been able here, ourselves, to manufacture not only numbers but useful things like tables of logarithms, powers, and trigonometric functions (for these are what the imaginary powers of real numbers are), all merely by extracting ten successive square roots of ten!

23

Resonance

23–1 Complex numbers and harmonic motion

23–2 The forced oscillator with damping

23–3 Electrical resonance

23–4 Resonance in nature

23–1 Complex numbers and harmonic motion

In the present chapter we shall continue our discussion of the harmonic oscillator and, in particular, the forced harmonic oscillator, using a new technique in the analysis. In the preceding chapter we introduced the idea of complex numbers, which have real and imaginary parts and which can be represented on a diagram in which the ordinate represents the imaginary part and the abscissa represents the real part. If a is a complex number, we may write it as $a = a_r + ia_i$, where the subscript r means the real part of a, and the subscript i means the imaginary part of a. Referring to Fig. 23–1, we see that we may also write a complex number $a = x + iy$ in the form $x + iy = re^{i\theta}$, where $r^2 = x^2 + y^2 = (x + iy)(x - iy) = aa^*$. (The complex conjugate of a, written a^*, is obtained by reversing the sign of i in a.) So we shall represent a complex number in either of two forms, a real plus an imaginary part, or a magnitude r and a phase angle θ, so-called. Given r and θ, x and y are clearly $r\cos\theta$ and $r\sin\theta$ and, in reverse, given a complex number $x + iy$, $r = \sqrt{x^2 + y^2}$ and $\tan\theta = y/x$, the ratio of the imaginary to the real part.

We are going to apply complex numbers to our analysis of physical phenomena by the following trick. We have examples of things that oscillate; the oscillation may have a driving force which is a certain constant times $\cos\omega t$. Now such a force, $F = F_0\cos\omega t$, can be written as the real part of a complex number $F = F_0e^{i\omega t}$ because $e^{i\omega t} = \cos\omega t + i\sin\omega t$. The reason we do this is that it is easier to work with an exponential function than with a cosine. So the whole trick is to represent our oscillatory functions as the real parts of certain complex functions. The complex number F that we have so defined is not a real physical force, because no force in physics is really complex; actual forces have no imaginary part, only a real part. We shall, however, speak of the "force" $F_0e^{i\omega t}$, but of course the actual force is the *real part* of that expression.

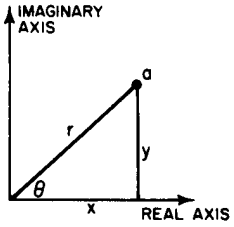


Fig. 23–1. A complex number may be represented by a point in the "complex plane."

Let us take another example. Suppose we want to represent a force which is a cosine wave that is out of phase with a delayed phase Δ. This, of course, would be the real part of $F_0e^{i(\omega t-\Delta)}$, but exponentials being what they are, we may write $e^{i(\omega t-\Delta)} = e^{i\omega t}e^{-i\Delta}$. Thus we see that the algebra of exponentials is much easier than that of sines and cosines; this is the reason we choose to use complex numbers. We shall often write

$$F = F_0e^{-i\Delta}e^{i\omega t} = \hat{F}e^{i\omega t}. \tag{23.1}$$

We write a little caret ($\hat{}$) over the F to remind ourselves that this quantity is a complex number: here the number is

$$\hat{F} = F_0e^{-i\Delta}.$$

Now let us solve an equation, using complex numbers, to see whether we can work out a problem for some real case. For example, let us try to solve

$$\frac{d^2x}{dt^2} + \frac{kx}{m} = \frac{F}{m} = \frac{F_0}{m}\cos\omega t, \tag{23.2}$$

where F is the force which drives the oscillator and x is the displacement. Now, absurd though it may seem, let us suppose that x and F are actually complex numbers, for a mathematical purpose only. That is to say, x has a real part and an imaginary part times i, and F has a real part and an imaginary part times i.

Now if we had a solution of (23.2) with complex numbers, and substituted the complex numbers in the equation, we would get

$$\frac{d^2(x_r + ix_i)}{dt^2} + \frac{k(x_r + ix_i)}{m} = \frac{F_r + iF_i}{m}$$

or

$$\frac{d^2x_r}{dt^2} + \frac{kx_r}{m} + i\left(\frac{d^2x_i}{dt^2} + \frac{kx_i}{m}\right) = \frac{F_r}{m} + \frac{iF_i}{m}.$$

Now, since if two complex numbers are equal, their real parts must be equal *and* their complex parts must be equal, we deduce that *the real part of x satisfies the equation with the real part of the force.* We must emphasize, however, that this separation into a real part and an imaginary part is *not* valid in general, but is valid only for equations which are *linear*, that is, for equations in which x appears in every term only in the first power or the zeroth power. For instance, if there were in the equation a term λx^2, then when we substitute $x_r + ix_i$, we would get $\lambda(x_r + ix_i)^2$, but when separated into real and imaginary parts this would yield $\lambda(x_r^2 - x_i^2)$ as the real part and $2i\lambda x_r x_i$ as the imaginary part. So we see that the real part of the equation would not involve just λx_r^2, but also $-\lambda x_i^2$. In this case we get a different equation than the one we wanted to solve, with x_i, the completely artificial thing we introduced in our analysis, mixed in.

Let us now try our new method for the problem of the forced oscillator, that we already know how to solve. We want to solve Eq. (23.2) as before, but we say that we are going to try to solve

$$\frac{d^2x}{dt^2} + \frac{kx}{m} = \frac{\hat{F}e^{i\omega t}}{m}, \tag{23.3}$$

where $\hat{F}e^{i\omega t}$ is a complex number. Of course x will also be complex, but remember the rule: take the real part to find out what is really going on. So we try to solve (23.3) for the forced solution; we shall discuss other solutions later. The forced solution has the same frequency as the applied force, and has some amplitude of oscillation and some phase, and so it can be represented also by some complex number $\hat{x}$ whose magnitude represents the swing of x and whose phase represents the time delay in the same way as for the force. Now a wonderful feature of an exponential function is that $d(\hat{x}e^{i\omega t})/dt = i\omega\hat{x}e^{i\omega t}$. When we differentiate an exponential function, we bring down the exponent as a simple multiplier. The second derivative does the same thing, it brings down another $i\omega$, and so it is very simple to write immediately, by inspection, what the equation is for $\hat{x}$: every time we see a differentiation, we simply multiply by $i\omega$. (Differentiation is now as easy as multiplication! This idea of using exponentials in linear differential equations is almost as great as the invention of logarithms, in which multiplication is replaced by addition. Here differentiation is replaced by multiplication.) Thus our equation becomes

$$(i\omega)^2\hat{x} + (k\hat{x}/m) = \hat{F}/m. \tag{23.4}$$

(We have cancelled the common factor $e^{i\omega t}$.) See how simple it is! Differential equations are immediately converted, by sight, into mere algebraic equations; we virtually have the solution by sight, that

$$\hat{x} = \frac{\hat{F}/m}{(k/m) - \omega^2},$$

since $(i\omega)^2 = -\omega^2$. This may be slightly simplified by substituting $k/m = \omega_0^2$, which gives

$$\hat{x} = \hat{F}/m(\omega_0^2 - \omega^2). \tag{23.5}$$

This, of course, is the solution we had before; for since $m(\omega_0^2 - \omega^2)$ is a real number, the phase angles of $\hat{F}$ and of $\hat{x}$ are the same (or perhaps 180° apart, if $\omega^2 > \omega_0^2$), as advertised previously, The magnitude of $\hat{x}$, which measures how far it oscillates, is related to the size of the $\hat{F}$ by the factor $1/m(\omega_0^2 - \omega^2)$, and this factor becomes

enormous when ω is nearly equal to ω_0. So we get a very strong response when we apply the right frequency ω (if we hold a pendulum on the end of a string and shake it at just the right frequency, we can make it swing very high).

23–2 The forced oscillator with damping

That, then, is how we analyze oscillatory motion with the more elegant mathematical technique. But the elegance of the technique is not at all exhibited in such a problem that can be solved easily by other methods. It is only exhibited when one applies it to more difficult problems. Let us therefore solve another, more difficult problem, which furthermore adds a relatively realistic feature to the previous one. Equation (23.5) tells us that if the frequency ω were exactly equal to ω_0, we would have an infinite response. Actually, of course, no such infinite response occurs because some other things, like friction, which we have so far ignored, limits the response. Let us therefore add to Eq. (23.2) a friction term.

Ordinarily such a problem is very difficult because of the character and complexity of the frictional term. There are, however, many circumstances in which the frictional force is *proportional to the speed* with which the object moves. An example of such friction is the friction for slow motion of an object in oil or a thick liquid. There is no force when it is just standing still, but the faster it moves the faster the oil has to go past the object, and the greater is the resistance. So we shall assume that there is, in addition to the terms in (23.2), another term, a resistance force proportional to the velocity: $F_f = -c\,dx/dt$. It will be convenient, in our mathematical analysis, to write the constant c as m times γ to simplify the equation a little. This is just the same trick we use with k when we replace it by $m\omega_0^2$, just to simplify the algebra. Thus our equation will be

$$m(d^2x/dt^2) + c(dx/dt) + kx = F \tag{23.6}$$

or, writing $c = m\gamma$ and $k = m\omega_0^2$ and dividing out the mass m,

$$(d^2x/dt^2) + \gamma(dx/dt) + \omega_0^2 x = F/m. \tag{23.6a}$$

Now we have the equation in the most convenient form to solve. If γ is very small, that represents very little friction; if γ is very large, there is a tremendous amount of friction. How do we solve this new linear differential equation? Suppose that the driving force is equal to $F_0 \cos(\omega t + \Delta)$; we could put this into (23.6a) and try to solve it, but we shall instead solve it by our new method. Thus we write F as the real part of $\hat{F}e^{i\omega t}$ and x as the real part of $\hat{x}e^{i\omega t}$, and substitute these into Eq. (23.6a). It is not even necessary to do the actual substituting, for we can see by inspection that the equation would become

$$[(i\omega)^2\hat{x} + \gamma(i\omega)\hat{x} + \omega_0^2\hat{x}]e^{i\omega t} = (\hat{F}/m)e^{i\omega t}. \tag{23.7}$$

[As a matter of fact, if we tried to solve Eq. (23.6a) by our old straightforward way, we would really appreciate the magic of the "complex" method.] If we divide by $e^{i\omega t}$ on both sides, then we can obtain the response $\hat{x}$ to the given force $\hat{F}$; it is

$$\hat{x} = \hat{F}/m(\omega_0^2 - \omega^2 + i\gamma\omega). \tag{23.8}$$

Thus again $\hat{x}$ is given by $\hat{F}$ times a certain factor. There is no technical name for this factor, no particular letter for it, but we may call it R for discussion purposes:

$$R = \frac{1}{m(\omega_0^2 - \omega^2 + i\gamma\omega)}$$

and

$$\hat{x} = \hat{F}R. \tag{23.9}$$

(Although the letters γ and ω_0 are in very common use, this R has no particular name.) This factor R can either be written as $p + iq$, or as a certain magnitude ρ times $e^{i\theta}$. If it is written as a certain magnitude times $e^{i\theta}$, let us see what it means.

Now $\hat{F} = F_0 e^{i\Delta}$, and the actual force F is the real part of $F_0 e^{i\Delta} e^{i\omega t}$, that is, $F_0 \cos(\omega t + \Delta)$. Next, Eq. (23.9) tells us that $\hat{x}$ is equal to $\hat{F}R$. So, writing $R = \rho e^{i\theta}$ as another name for R, we get

$$\hat{x} = R\hat{F} = \rho e^{i\theta} F_0 e^{i\Delta} = \rho F_0 e^{i(\theta + \Delta)}.$$

Finally, going even further back, we see that the physical x, which is the real part of the complex $\hat{x}$, is equal to the real part of $\rho F_0 e^{i(\theta+\Delta)} e^{i\omega t}$. But ρ and F_0 are real, and the real part of $e^{i(\theta+\Delta+\omega t)}$ is simply $\cos(\omega t + \Delta + \theta)$. Thus

$$x = \rho F_0 \cos(\omega t + \Delta + \theta). \tag{23.10}$$

This tells us that the amplitude of the response is the magnitude of the force F multiplied by a certain magnification factor, ρ; this gives us the "amount" of oscillation. It also tells us, however, that x is not oscillating in phase with the force, which has the phase Δ, but is shifted by an extra amount θ. Therefore ρ and θ represent the size of the response and the phase shift of the response.

Now let us work out what ρ is. If we have a complex number, the square of the magnitude is equal to the number times its complex conjugate; thus

$$\begin{aligned} \rho^2 &= \frac{1}{m^2(\omega_0^2 - \omega^2 + i\gamma\omega)(\omega_0^2 - \omega^2 - i\gamma\omega)} \\ &= \frac{1}{m^2[(\omega^2 - \omega_0^2)^2 + \gamma^2\omega^2]}. \end{aligned} \tag{23.11}$$

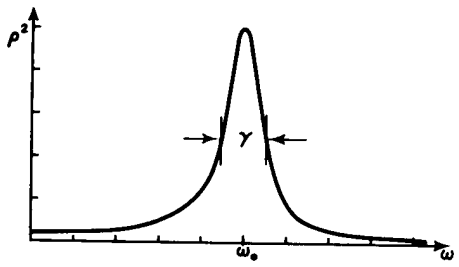


Fig. 23–2. Plot of ρ^2 versus ω.

In addition, the phase angle θ is easy to find, for if we write

$$1/R = 1/\rho e^{i\theta} = (1/\rho)e^{-i\theta} = m(\omega_0^2 - \omega^2 + i\gamma\omega),$$

we see that

$$\tan\theta = -\gamma\omega/(\omega_0^2 - \omega^2). \tag{23.12}$$

It is minus because $\tan(-\theta) = -\tan\theta$. A negative value for θ results for all ω, and this corresponds to the displacement x lagging the force F.

Figure 23–2 shows how ρ^2 varies as a function of frequency (ρ^2 is physically more interesting than ρ, because ρ^2 is proportional to the square of the amplitude, or more or less to the *energy* that is developed in the oscillator by the force). We see that if γ is very small, then $1/(\omega_0^2 - \omega^2)^2$ is the most important term, and the response tries to go up toward infinity when ω equals ω_0. Now the "infinity" is not actually infinite because if $\omega = \omega_0$, then $1/\gamma^2\omega^2$ is still there. The phase shift varies as shown in Fig. 23–3.

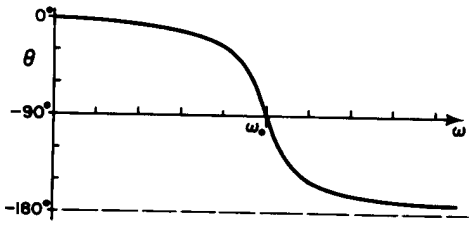


Fig. 23–3. Plot of θ versus ω.

In certain circumstances we get a slightly different formula than (23.8), also called a "resonance" formula, and one might think that it represents a different phenomenon, but it does not. The reason is that if γ is very small the most interesting part of the curve is near $\omega = \omega_0$, and we may replace (23.8) by an approximate formula which is very accurate if γ is small and ω is near ω_0. Since $\omega_0^2 - \omega^2 = (\omega_0 - \omega)(\omega_0 + \omega)$, if ω is near ω_0 this is nearly the same as $2\omega_0(\omega_0 - \omega)$ and $\gamma\omega$ is nearly the same as $\gamma\omega_0$. Using these in (23.8), we see that $\omega_0^2 - \omega^2 + i\gamma\omega \approx 2\omega_0(\omega_0 - \omega + i\gamma/2)$, so that

$$\hat{x} \approx \hat{F}/2m\omega_0(\omega_0 - \omega + i\gamma/2) \quad \text{if} \quad \gamma \ll \omega_0 \quad \text{and} \quad \omega \approx \omega_0. \tag{23.13}$$

It is easy to find the corresponding formula for ρ^2. It is

$$\rho^2 \approx 1/4m^2\omega_0^2[(\omega_0 - \omega)^2 + \gamma^2/4].$$

We shall leave it to the student to show the following: if we call the maximum height of the curve of ρ^2 vs. ω one unit, and we ask for the width $\Delta\omega$ of the curve, at one half the maximum height, the full width at half the maximum height of the curve is $\Delta\omega = \gamma$, supposing that γ is small. The resonance is sharper and sharper as the frictional effects are made smaller and smaller.

As another measure of the width, some people use a quantity Q which is defined as $Q = \omega_0/\gamma$. The narrower the resonance, the higher the Q: $Q = 1000$ means a resonance whose width is only 1000th of the frequency scale. The Q of the resonance curve shown in Fig. 23–2 is 5.

The importance of the resonance phenomenon is that it occurs in many other circumstances, and so the rest of this chapter will describe some of these other circumstances.

23–3 Electrical resonance

The simplest and broadest technical applications of resonance are in electricity. In the electrical world there are a number of objects which can be connected to make electric circuits. These *passive circuit elements*, as they are often called, are of three main types, although each one has a little bit of the other two mixed in. Before describing them in greater detail, let us note that the whole idea of our mechanical oscillator being a mass on the end of a spring is only an approximation. All the mass is not actually at the "mass"; some of the mass is in the inertia of the spring. Similarly, all of the spring is not at the "spring"; the mass itself has a little elasticity, and although it may appear so, it is not *absolutely* rigid, and as it goes up and down, it flexes ever so slightly under the action of the spring pulling it. The same thing is true in electricity. There is an approximation in which we can lump things into "circuit elements" which are assumed to have pure, ideal characteristics. It is not the proper time to discuss that approximation here, we shall simply assume that it is true in the circumstances.

The three main kinds of circuit elements are the following. The first is called a *capacitor* (Fig. 23–4); an example is two plane metallic plates spaced a very small distance apart by an insulating material. When the plates are charged there is a certain voltage difference, that is, a certain difference in potential, between them. The same difference of potential appears between the terminals A and B, because if there were any difference along the connecting wire, electricity would flow right away. So there is a certain voltage difference V between the plates if there is a certain electric charge $+q$ and $-q$ on them, respectively. Between the plates there will be a certain electric field; we have even found a formula for it (Chapters 13 and 14):

$$V = \sigma d/\epsilon_0 = qd/\epsilon_0 A, \tag{23.14}$$

where d is the spacing and A is the area of the plates. Note that the potential difference is a linear function of the charge. If we do not have parallel plates, but insulated electrodes which are of any other shape, the difference in potential is still precisely proportional to the charge, but the constant of proportionality may not be so easy to compute. However, all we need to know is that the potential difference across a capacitor *is proportional to the charge:* $V = q/C$; the proportionality constant is $1/C$, where C is the *capacitance* of the object.

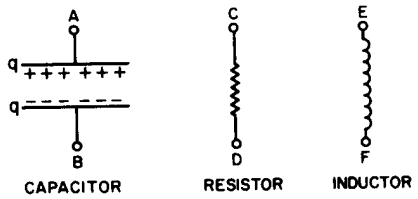


Fig. 23–4. The three passive circuit elements.

The second kind of circuit element is called a *resistor*; it offers resistance to the flow of electrical current. It turns out that metallic wires and many other substances resist the flow of electricity in this manner: if there is a voltage difference across a piece of some substance, there exists an electric current $I = dq/dt$ that is proportional to the electric voltage difference:

$$V = RI = R\,dq/dt. \tag{23.15}$$

The proportionality coefficient is called the *resistance* R. This relationship may already be familiar to you; it is Ohm's law.

If we think of the charge q on a capacitor as being analogous to the displacement x of a mechanical system, we see that the current, $I = dq/dt$, is analogous to velocity, $1/C$ is analogous to a spring constant k, and R is analogous to the resistive coefficient γ. Now it is very interesting that there exists another circuit element which is the analog of *mass*! This is a coil which builds up a magnetic field within itself when there is a current in it. A *changing* magnetic field develops in the coil a voltage that is proportional to dI/dt (this is how a transformer works, in

fact). The magnetic field is proportional to a current, and the induced voltage (so-called) in such a coil is proportional to the rate of change of the current:

$$V = L\,dI/dt = L\,d^2q/dt^2. \tag{23.16}$$

The coefficient L is the *self-inductance*, and is analogous to the mass in a mechanical oscillating circuit.

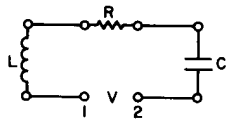


Fig. 23–5. An oscillatory electrical circuit with resistance, inductance, and capacitance.

Suppose we make a circuit in which we have connected the three circuit elements in series (Fig. 23–5); then the voltage across the whole thing from 1 to 2 is the work done in carrying a charge through, and it consists of the sum of several pieces: across the inductor, $V_L = L\,d^2q/dt^2$; across the resistance, $V_R = R\,dq/dt$; across the capacitor, $V_c = q/C$. The sum of these is equal to the applied voltage, V:

$$L\,d^2q/dt^2 + R\,dq/dt + q/C = V(t). \tag{23.17}$$

Now we see that this equation is exactly the same as the mechanical equation (23.6), and of course it can be solved in exactly the same manner. We suppose that $V(t)$ is oscillatory: we are driving the circuit with a generator with a pure sine wave oscillation. Then we can write our $V(t)$ as a complex $\hat{V}$ with the understanding that it must be ultimately multiplied by $e^{i\omega t}$, and the real part taken in order to find the true V. Likewise, the charge q can thus be analyzed, and then in exactly the same manner as in Eq. (23.8) we write the corresponding equation: the second derivative of $\hat{q}$ is $(i\omega)^2\hat{q}$; the first derivative is $(i\omega)q$. Thus Eq. (23.17) translates to

$$\left[L(i\omega)^2 + R(i\omega) + \frac{1}{C}\right]\hat{q} = \hat{V}$$

or

$$q = \frac{\hat{V}}{L(i\omega)^2 + R(i\omega) + \dfrac{1}{C}}$$

which we can write in the form

$$\hat{q} = \hat{V}/L(\omega_0^2 - \omega^2 + i\gamma\omega), \tag{23.18}$$

where $\omega_0^2 = 1/LC$ and $\gamma = R/L$. It is exactly the same denominator as we had in the mechanical case, with exactly the same resonance properties! The correspondence between the electrical and mechanical cases is outlined in Table 23–1.

Table 23–1

General characteristic	Mechanical property	Electrical property
indep. variable	time (t)	time (t)
dep. variable	position (x)	charge (q)
inertia	mass (m)	inductance (L)
resistance	drag coeff. ($c = \gamma m$)	resistance ($R = \gamma L$)
stiffness	stiffness (k)	(capacitance)$^{-1}$ ($1/C$)
resonant frequency	$\omega_0^2 = k/m$	$\omega_0^2 = 1/LC$
period	$t_0 = 2\pi\sqrt{m/k}$	$t_0 = 2\pi\sqrt{LC}$
figure of merit	$Q = \omega_0/\gamma$	$Q = \omega_0 L/R$

We must mention a small technical point. In the electrical literature, a different notation is used. (From one field to another, the subject is not really any different, but the way of writing the notations is often different.) First, j is commonly used instead of i in electrical engineering, to denote $\sqrt{-1}$. (After all, i must be the current!) Also, the engineers would rather have a relationship between $\hat{V}$ and $\hat{I}$ than between $\hat{V}$ and $\hat{q}$, just because they are more used to it that way. Thus, since $\hat{I} = d\hat{q}/dt = i\omega\hat{q}$, we can just substitute $\hat{I}/i\omega$ for $\hat{q}$ and get

$$\hat{V} = (i\omega L + R + 1/i\omega C)\hat{I} = \hat{Z}\hat{I}. \tag{23.19}$$

Another way is to rewrite Eq. (23.17), so that it looks more familiar; one often sees it written this way:

$$L\,dI/dt + RI + (1/C)\int^t I\,dt = V(t). \tag{23.20}$$

At any rate, we find the relation (23.19) between voltage $\hat{V}$ and current $\hat{I}$ which is just the same as (23.18) except divided by $i\omega$, and that produces Eq. (23.19). The quantity $R + i\omega L + 1/i\omega C$ is a complex number, and is used so much in electrical engineering that it has a name: it is called the *complex impedance*, $\hat{Z}$. Thus we can write $\hat{V} = \hat{Z}\hat{I}$. The reason that the engineers like to do this is that they learned something when they were young: $V = RI$ for resistances, when they only knew about resistances and DC. Now they have become more educated and have AC circuits, so they want the equation to look the same. Thus they write $\hat{V} = \hat{Z}\hat{I}$, the only difference being that the resistance is replaced by a more complicated thing, a complex quantity. So they insist that they cannot use what everyone else in the world uses for imaginary numbers, they have to use a j for that; it is a miracle that they did not insist also that the letter Z be an R! (Then they get into trouble when they talk about current densities, for which they also use j. The difficulties of science are to a large extent the difficulties of notations, the units, and all the other artificialities which are invented by man, not by nature.)

23–4 Resonance in nature

Although we have discussed the electrical case in detail, we could also bring up case after case in many fields, and show exactly how the resonance equation is the same. There are many circumstances in nature in which something is "oscillating" and in which the resonance phenomenon occurs. We said that in an earlier chapter; let us now demonstrate it. If we walk around our study, pulling books off the shelves and simply looking through them to find an example of a curve that corresponds to Fig. 23–2 and comes from the same equation, what do we find? Just to demonstrate the wide range obtained by taking the smallest possible sample, it takes only five or six books to produce quite a series of phenomena which show resonances.

The first two are from mechanics, the first on a large scale: the atmosphere of the whole earth. If the atmosphere, which we suppose surrounds the earth evenly on all sides, is pulled to one side by the moon or, rather, squashed prolate into a double tide, and if we could then let it go, it would go sloshing up and down; it is an oscillator. This oscillator is *driven* by the moon, which is effectively revolving about the earth; any one component of the force, say in the x-direction, has a cosine component, and so the response of the earth's atmosphere to the tidal pull of the moon is that of an oscillator. The expected response of the atmosphere is shown in Fig. 23–6, curve b (curve a is another theoretical curve under discussion in the book from which this is taken out of context). Now one might think that we only have one point on this resonance curve, since we only have the one frequency, corresponding to the rotation of the earth under the moon, which occurs at a period of 12.42 hours—12 hours for the earth (the tide is a double bump), plus a little more because the moon is going around. But from the *size* of the atmospheric tides, and from the *phase*, the amount of delay, we can get both ρ and θ. From those we can get ω_0 and γ, and thus draw the entire curve! This is an example of very poor science. From two numbers we obtain two numbers, and from those two numbers we draw a beautiful curve, which of course goes through the very point that determined the curve! It is of no use *unless we can measure something else*, and in the case of geophysics that is often very difficult. But in this particular case there is another thing which we can show theoretically must have the same timing as the natural frequency ω_0: that is, if someone disturbed the atmosphere, it would oscillate with the frequency ω_0. Now there *was* such a sharp disturbance in 1883; the Krakatoa volcano exploded and half the island blew off, and it made such a terrific explosion in the atmosphere that the period of oscillation of the atmosphere could be measured. It came out to $10\frac{1}{2}$ hours. The ω_0 obtained from

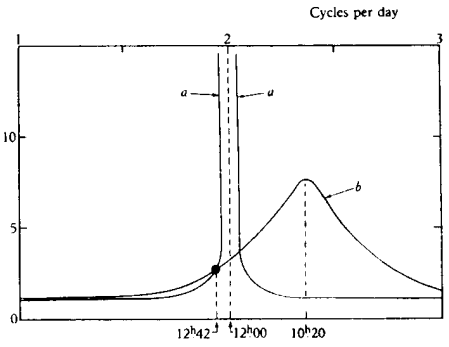


Fig. 23–6. Response of the atmosphere to external excitation. a is the required response if the atmospheric S_2-tide is of gravitational origin; peak amplification is 100:1. b is derived from observed magnification and phase of M_2-tide. [Munk and MacDonald, "Rotation of the Earth," Cambridge University Press (1960)]

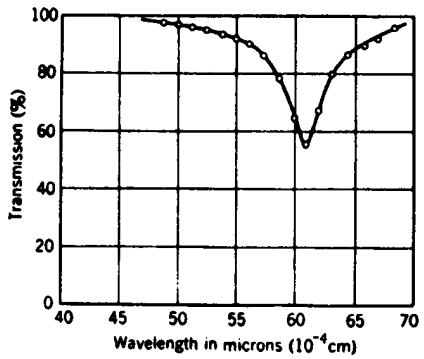


Fig. 23–7. Transmission of infrared radiation through a thin (0.17 μ) sodium chloride film. [After R. B. Barnes, *Z. Physik* **75**, 723 (1932). Kittel, *Introduction to Solid State Physics*, Wiley, 1956.]

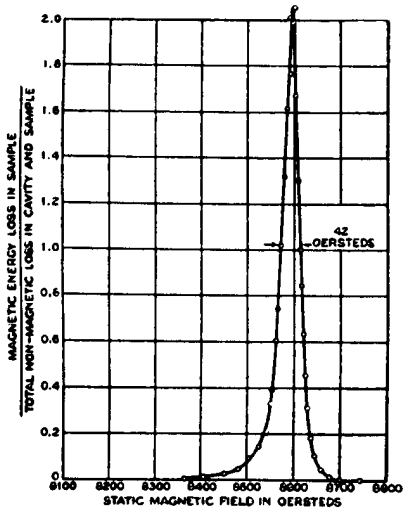


Fig. 23–8. Magnetic energy loss in paramagnetic organic compound as function of applied magnetic field intensity. [Holden *et al.*, *Phys. Rev.* **75**, 1614 (1949)]

Fig. 23–6 comes out 10 hours and 20 minutes, so there we have at least one check on the reality of our understanding of the atmospheric tides.

Next we go to the small scale of mechanical oscillation. This time we take a sodium chloride crystal, which has sodium ions and chlorine ions next to each other, as we described in an early chapter. These ions are electrically charged, alternately plus and minus. Now there is an interesting oscillation possible. Suppose that we could drive all the plus charges to the right and all the negative charges to the left, and let go; they would then oscillate back and forth, the sodium lattice against the chlorine lattice. How can we ever drive such a thing? That is easy, for if we apply an electric field on the crystal, it will push the plus charge one way and the minus charge the other way! So, by having an external electric field we can perhaps get the crystal to oscillate. The frequency of the electric field needed is so high, however, that it corresponds to *infrared radiation!* So we try to find a resonance curve by measuring the absorption of infrared light by sodium chloride. Such a curve is shown in Fig. 23–7. The abscissa is not frequency, but is given in terms of wavelength, but that is just a technical matter, of course, since for a wave there is a definite relation between frequency and wavelength; so it is really a frequency scale, and a certain frequency corresponds to the resonant frequency.

But what about the width? What determines the width? There are many cases in which the width that is seen on the curve is not really the natural width γ that one would have theoretically. There are two reasons why there can be a wider curve than the theoretical curve. If the objects do not all have the same frequency, as might happen if the crystal were strained in certain regions, so that in those regions the oscillation frequency were slightly different than in other regions, then what we have is many resonance curves on top of each other; so we apparently get a wider curve. The other kind of width is simply this: perhaps we cannot measure the frequency precisely enough—if we open the slit of the spectrometer fairly wide, so although we thought we had only one frequency, we actually had a certain range $\Delta\omega$, then we may not have the resolving power needed to see a narrow curve. Offhand, we cannot say whether the width in Fig. 23–7 is natural, or whether it is due to inhomogeneities in the crystal or the finite width of the slit of the spectrometer.

Now we turn to a more esoteric example, and that is the swinging of a magnet. If we have a magnet, with north and south poles, in a constant magnetic field, the N end of the magnet will be pulled one way and the S end the other way, and there will in general be a torque on it, so it will vibrate about its equilibrium position, like a compass needle. However, the magnets we are talking about are *atoms*. These atoms have an angular momentum, the torque does not produce a simple motion in the direction of the field, but instead, of course, a *precession*. Now, looked at from the side, any one component is "swinging," and we can disturb or drive that swinging and measure an absorption. The curve in Fig. 23–8 represents a typical such resonance curve. What has been done here is slightly different technically. The frequency of the lateral field that is used to drive this swinging is always kept the same, while we would have expected that the investigators would vary that and plot the curve. They could have done it that way, but technically it was easier for them to leave the frequency ω fixed, and change the strength of the constant magnetic field, which corresponds to changing ω_0 in our formula. They have plotted the resonance curve against ω_0. Anyway, this is a typical resonance with a certain ω_0 and γ.

Now we go still further. Our next example has to do with atomic nuclei. The motions of protons and neutrons in nuclei are oscillatory in certain ways, and we can demonstrate this by the following experiment. We bombard a lithium atom with protons, and we discover that a certain reaction, producing γ-rays, actually has a very sharp maximum typical of resonance. We note in Fig. 23–9, however, one difference from other cases: the horizontal scale is not a frequency, it is an *energy!* The reason is that in quantum mechanics what we think of classically as the energy will turn out to be really related to a frequency of a wave amplitude. When we analyze something which in simple large-scale physics has to do with a frequency, we find that when we do quantum-mechanical experiments with atomic

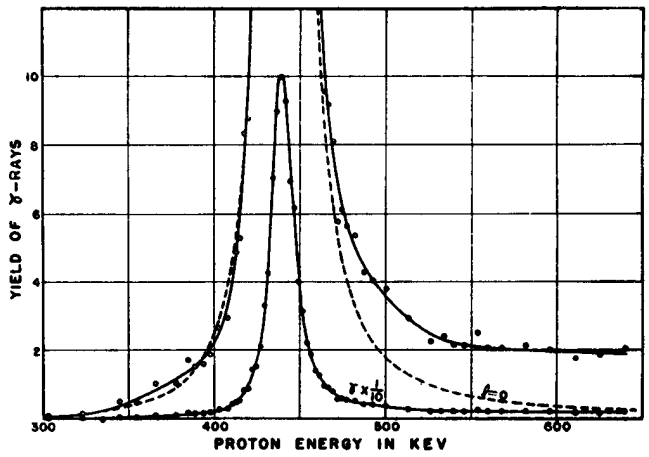


Fig. 23–9. The intensity of gamma-radiation from lithium as a function of the energy of the bombarding protons. The dashed curve is a theoretical one calculated for protons with an angular momentum $\ell = 0$. [Bonner and Evans, *Phys. Rev.* **73**, 666 (1948)]

matter, we get the corresponding curve as a function of energy. In fact, this curve is a demonstration of this relationship, in a sense. It shows that frequency and energy have some deep interrelationship, which of course they do.

Now we turn to another example which also involves a nuclear energy level, but now a much, much narrower one. The ω_0 in Fig. 23–10 corresponds to an energy of 100,000 electron volts, while the width γ is approximately 10^{-5} electron volt; in other words, this has a Q of 10^{10}! When this curve was measured it was the largest Q of any oscillator that had ever been measured. It was measured by Dr. Möessbauer, and it was the basis of his Nobel prize. The horizontal scale here is velocity, because the technique for obtaining the slightly different frequencies was to use the Doppler effect, by moving the source relative to the absorber. One can see how delicate the experiment is when we realize that the speed involved is a few centimeters per second! On the actual scale of the figure, zero frequency would correspond to a point about 10^{10} cm to the left—slightly off the paper!

Finally, if we look in an issue of the *Physical Review*, say that of January 1, 1962, will we find a resonance curve? Every issue has a resonance curve, and Fig. 23–11 is the resonance curve for this one. This resonance curve turns out be be very interesting. It is the resonance found in a certain reaction among strange particles, a reaction in which a K^- and a proton interact. The resonance is detected by seeing how many of some kinds of particles come out, and depending on what and how many come out, one gets different curves, but of the same shape and with the peak at the same energy. We thus determine that there is a resonance at a certain energy for the K^- meson. That presumably means that there is some kind of a state, or condition, corresponding to this resonance, which can be attained by putting together a K^- and a proton. This is a new particle, or resonance. Today we do not know whether to call a bump like this a "particle" or simply a resonance. When there is a very *sharp* resonance, it corresponds to a very *definite energy*, just as though there were a particle of that energy present in nature. When the resonance gets wider, then we do not know whether to say there is a particle which does not last very long, or simply a resonance in the reaction probability. In the second chapter, this point is made about the particles, but when the second chapter was written this resonance was not known, so our chart should now have still another particle in it!

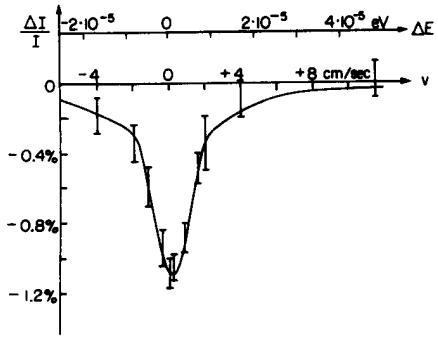


Fig. 23–10. [Courtesy of Dr. R. Mössbauer]

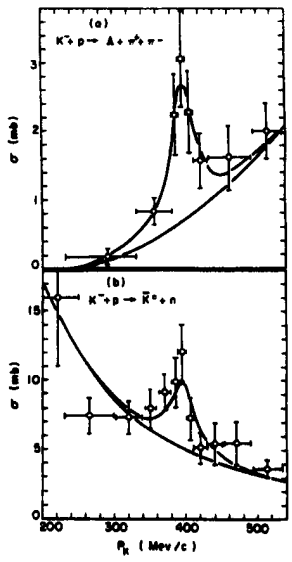


Fig. 23–11. Momentum dependence of the cross section for the reactions (a) $K^- + p \rightarrow \Lambda + \pi^+ + \pi^-$ and (b) $K^- + p \rightarrow K^0 + n$. The lower curves in (a) and (b) represent the presumed nonresonant backgrounds, while the upper curves contain in addition the superposed resonance. [Ferro-Luzzi *et al.*, *Phys. Rev., Lett.* **8**, 28 (1962)]

24

Transients

24–1 The energy of an oscillator

24–2 Damped oscillations

24–3 Electrical transients

24–1 The energy of an oscillator

Although this chapter is entitled "transients," certain parts of it are, in a way, part of the last chapter on forced oscillation. One of the features of a forced oscillation which we have not yet discussed is the *energy* in the oscillation. Let us now consider that energy.

In a mechanical oscillator, how much kinetic energy is there? It is proportional to the square of the velocity. Now we come to an important point. Consider an arbitrary quantity A, which may be the velocity or something else that we want to discuss. When we write $A = \hat{A}e^{i\omega t}$, a complex number, the true and honest A, in the physical world, is only the *real part*; therefore if, for some reason, we want to use the *square* of A, it is not right to square the complex number and then take the real part, because the real part of the square of a complex number is not just the square of the real part, but also involves the *imaginary* part. So when we wish to find the energy we have to get away from the complex notation for a while to see what the inner workings are.

Now the true physical A is the real part of $A_0e^{i(\omega t+\Delta)}$, that is, $A = A_0 \cos(\omega t + \Delta)$, where $\hat{A}$, the complex number, is written as $A_0e^{i\Delta}$. Now the square of this real physical quantity is $A^2 = A_0^2 \cos^2(\omega t + \Delta)$. The square of the quantity, then, goes up and down from a maximum to zero, like the square of the cosine. The square of the cosine has a maximum of 1 and a minimum of 0, and its average value is 1/2.

In many circumstances we are not interested in the energy at any specific moment during the oscillation; for a large number of applications we merely want the average of A^2, the *mean* of the square of A over a period of time large compared with the period of oscillation. In those circumstances, the average of the cosine squared may be used, so we have the following theorem: if A is represented by a complex number, then the mean of A^2 is equal to $\frac{1}{2}A_0^2$. Now A_0^2 is the square of the magnitude of the complex $\hat{A}$. (This can be written in many ways—some people like to write $|\hat{A}|^2$; others write, $\hat{A}\hat{A}^*$, $\hat{A}$ times its complex conjugate.) We shall use this theorem several times.

Now let us consider the energy in a forced oscillator. The equation for the forced oscillator is

$$m\,d^2x/dt^2 + \gamma m\,dx/dt + m\omega_0^2 x = F(t). \tag{24.1}$$

In our problem, of course, $F(t)$ is a cosine function of t. Now let us analyze the situation: how much work is done by the outside force F? The work done by the force per second, i.e., the power, is the force times the velocity. (We know that the differential work in a time dt is $F\,dx$, and the power is $F\,dx/dt$.) Thus

$$P = F\frac{dx}{dt} = m\left[\left(\frac{dx}{dt}\right)\left(\frac{d^2x}{dt^2}\right) + \omega_0^2 x\left(\frac{dx}{dt}\right)\right] + \gamma m\left(\frac{dx}{dt}\right)^2. \tag{24.2}$$

But the first two terms on the right can also be written as $d/dt[\frac{1}{2}m(dx/dt)^2 + \frac{1}{2}m\omega_0^2x^2]$, as is immediately verified by differentiating. That is to say, the term in brackets is a pure derivative of two terms that are easy to understand—one is the kinetic energy of motion, and the other is the potential energy of the spring. Let us call this quantity the *stored energy*, that is, the energy stored in the oscillation. Suppose that we want the average power over many cycles when the oscillator is being forced and has been running for a long time. In the long run, the stored

energy does not change—its derivative gives zero average effect. In other words, if we average the power in the long run, *all the energy ultimately ends up in the resistive term* $\gamma m(dx/dt)^2$. There is *some* energy stored in the oscillation, but that does not change with time, if we average over many cycles. Therefore the mean power $\langle P\rangle$ is

$$\langle P\rangle = \langle \gamma m(dx/dt)^2\rangle. \tag{24.3}$$

Using our method of writing complex numbers, and our theorem that $\langle A^2\rangle = \frac{1}{2}A_0^2$, we may find this mean power. Thus if $x = \hat{x}e^{i\omega t}$, then $dx/dt = i\omega\hat{x}e^{i\omega t}$. Therefore, in these circumstances, the average power could be written as

$$\langle P\rangle = \tfrac{1}{2}\gamma m\omega^2 x_0^2. \tag{24.4}$$

In the notation for electrical circuits, dx/dt is replaced by the current I (I is dq/dt, where q corresponds to x), and $m\gamma$ corresponds to the resistance R. Thus the rate of the energy loss—the power used up by the forcing function—is the resistance in the circuit times the average square of the current:

$$\langle P\rangle = R\langle I^2\rangle = R\cdot\tfrac{1}{2}I_0^2. \tag{24.5}$$

This energy, of course, goes into heating the resistor; it is sometimes called the heating loss or the Joule heating.

Another interesting feature to discuss is how much energy is *stored.* That is not the same as the power, because although power was at first used to store up some energy, after that the system keeps on absorbing power, insofar as there are any heating (resistive) losses. At any moment there is a certain amount of stored energy, so we would like to calculate the mean stored energy $\langle E\rangle$ also. We have already calculated what the average of $(dx/dt)^2$ is, so we find

$$\begin{aligned}\langle E\rangle &= \tfrac{1}{2}m\langle(dx/dt)^2\rangle + \tfrac{1}{2}m\omega_0^2\langle x^2\rangle\\ &= \tfrac{1}{2}m(\omega^2+\omega_0^2)\tfrac{1}{2}x_0^2.\end{aligned} \tag{24.6}$$

Now, when an oscillator is very efficient, and if ω is near ω_0, so that $|\hat{x}|$ is large, the stored energy is very high—we can get a large stored energy from a relatively small force. The force does a great deal of work in getting the oscillation going, but then to keep it steady, all it has to do is to fight the friction. The oscillator can have a great deal of energy if the friction is very low, and even though it is oscillating strongly, not much energy is being lost. The efficiency of an oscillator can be measured by how much energy is stored, compared with how much work the force does per oscillation.

How does the stored energy compare with the amount of work that is done in one cycle? This is called the Q of the system, and Q is defined as 2π times the mean stored energy, divided by the work done per cycle. (If we were to say the work done per *radian* instead of per cycle, then the 2π disappears.)

$$Q = 2\pi\,\frac{\tfrac{1}{2}m(\omega^2+\omega_0^2)\cdot\langle x^2\rangle}{\gamma m\omega^2\langle x^2\rangle\cdot 2\pi/\omega} = \frac{\omega^2+\omega_0^2}{2\gamma\omega}. \tag{24.7}$$

Q is not a very useful number unless it is very large. When it is relatively large, it gives a measure of how good the oscillator is. People have tried to define Q in the simplest and most useful way; various definitions differ a bit from one another, but if Q is very large, all definitions are in agreement. The most generally accepted definition is Eq. (24.7), which depends on ω. For a good oscillator, close to resonance, we can simplify (24.7) a little by setting $\omega = \omega_0$, and we then have $Q = \omega_0/\gamma$, which is the definition of Q that we used before.

What is Q for an electrical circuit? To find out, we merely have to translate L for m, R for $m\gamma$, and $1/C$ for $m\omega_0^2$ (see Table 23–1). The Q at resonance is $L\omega/R$, where ω is the resonance frequency. If we consider a circuit with a high Q, that means that the amount of energy stored in the oscillation is very large compared with the amount of work done per cycle by the machinery that drives the oscillations.

24–2 Damped oscillations

We now turn to our main topic of discussion: transients. By a *transient* is meant a solution of the differential equation when there is no force present, but when the system is not simply at rest. (Of course, if it is standing still at the origin with no force acting, that is a nice problem—it stays there!) Suppose the oscillation starts another way: say it was driven by a force for a while, and then we turn off the force. What happens then? Let us first get a rough idea of what will happen for a very high Q system. So long as a force is acting, the stored energy stays the same, and there is a certain amount of work done to maintain it. Now suppose we turn off the force, and no more work is being done; then the losses which are eating up the energy of the supply are no longer eating up its energy—there *is* no more driver. The losses will have to consume, so to speak, the energy that is stored. Let us suppose that $Q/2\pi = 1000$. Then the work done per cycle is 1/1000 of the stored energy. Is it not reasonable, since it is oscillating with no driving force, that in one cycle the system will still lose a thousandth of its energy E, which ordinarily would have been supplied from the outside, and that it will continue oscillating, always losing 1/1000 of its energy per cycle? So, as a guess, for a relatively high Q system, we would suppose that the following equation might be roughly right (we will later do it exactly, and it will turn out that it *was* right!):

$$dE/dt = -\omega E/Q. \tag{24.8}$$

This is rough because it is true only for large Q. In each radian the system loses a fraction $1/Q$ of the stored energy E. Thus in a given amount of time dt the energy will change by an amount $\omega\, dt/Q$, since the number of radians associated with the time dt is $\omega\, dt$. What is the frequency? Let us suppose that the system moves so nicely, with hardly any force, that if we let go it will oscillate at essentially the same frequency all by itself. So we will guess that ω is the resonant frequency ω_0. Then we deduce from Eq. (24.8) that the stored energy will vary as

$$E = E_0 e^{-\omega_0 t/Q} = E_0 e^{-\gamma t}. \tag{24.9}$$

This would be the measure of the *energy* at any moment. What would the formula be, roughly, for the amplitude of the oscillation as a function of the time? The same? No! The amount of energy in a spring, say, goes as the *square* of the displacement; the kinetic energy goes as the *square* of the velocity; so the total energy goes as the *square* of the displacement. Thus the displacement, the amplitude of oscillation, will decrease half as fast because of the square. In other words, we guess that the solution for the damped transient motion will be an oscillation of frequency close to the resonance frequency ω_0, in which the amplitude of the sine-wave motion will diminish as $e^{-\gamma t/2}$:

$$x = A_0 e^{-\gamma t/2} \cos \omega_0 t. \tag{24.10}$$

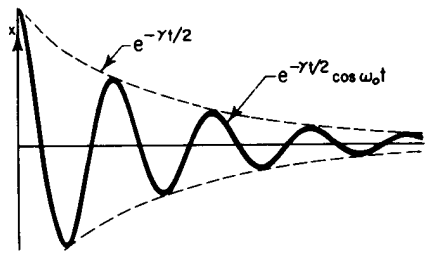


Fig. 24–1. A damped cosine oscillation.

This equation and Fig. 24–1 give us an idea of what we should expect; now let us try to analyze the motion *precisely* by solving the differential equation of the motion itself.

So, starting with Eq. (24.1), with no outside force, how do we solve it? Being physicists, we do not have to worry about the *method* as much as we do about what the solution *is*. Armed with our previous experience, let us try as a solution an exponential curve, $x = Ae^{i\alpha t}$. (Why do we try this? It is the easiest thing to differentiate!) We put this into (24.1) (with $F(t) = 0$), using the rule that each time we differentiate x with respect to time, we multiply by $i\alpha$. So it is really quite simple to substitute. Thus our equation looks like this:

$$(-\alpha^2 + i\gamma\alpha + \omega_0^2)Ae^{i\alpha t} = 0. \tag{24.11}$$

The net result must be zero for *all times*, which is impossible unless (a) $A = 0$, which is no solution at all—it stands still, or (b)

$$-\alpha^2 + i\alpha\gamma + \omega_0^2 = 0. \tag{24.12}$$

If we can solve this and find an α, then we will have a solution in which A need not be zero!

$$\alpha = i\gamma/2 \pm \sqrt{\omega_0^2 - \gamma^2/4}. \tag{24.13}$$

For a while we shall assume that γ is fairly small compared with ω_0, so that $\omega_0^2 - \gamma^2/4$ is definitely positive, and there is nothing the matter with taking the square root. The only bothersome thing is that we get *two* solutions! Thus

$$\alpha_1 = i\gamma/2 + \sqrt{\omega_0^2 - \gamma^2/4} = i\gamma/2 + \omega_\gamma \tag{24.14}$$

and

$$\alpha_2 = i\gamma/2 - \sqrt{\omega_0^2 - \gamma^2/4} = i\gamma/2 - \omega_\gamma. \tag{24.15}$$

Let us consider the first one, supposing that we had not noticed that the square root has two possible values. Then we know that a solution for x is $x_1 = Ae^{i\alpha_1 t}$, where A is any constant whatever. Now, in substituting α_1, because it is going to come so many times and it takes so long to write, we shall call $\sqrt{\omega_0^2 - \gamma^2/4} = \omega_\gamma$. Thus $i\alpha_1 = -\gamma/2 + i\omega_\gamma$, and we get $x = Ae^{(-\gamma/2+i\omega_\gamma)t}$, or what is the same, because of the wonderful properties of an exponential,

$$x_1 = Ae^{-\gamma t/2}e^{i\omega_\gamma t}. \tag{24.16}$$

First, we recognize this as an oscillation, an oscillation at a frequency ω_γ which is not *exactly* the frequency ω_0, but is rather close to ω_0 if it is a good system. Second, the amplitude of the oscillation is decreasing exponentially! If we take, for instance, the real part of (24.16), we get

$$x_1 = Ae^{-\gamma t/2}\cos\omega_\gamma t. \tag{24.17}$$

This is very much like our guessed-at solution (24.10), except that the frequency really is ω_γ. This is the only error, so it is the same thing—we have the right idea. But everything is *not* all right! What is not all right is that *there is another solution.*

The other solution is α_2, and we see that the difference is only that the sign of ω_γ is reversed:

$$x_2 = Be^{-\gamma t/2}e^{-i\omega_\gamma t}. \tag{24.18}$$

What does this mean? We shall soon prove that if x_1 and x_2 are each a possible solution of Eq. (24.1) with $F = 0$, then $x_1 + x_2$ is also a solution of the same equation! So the general solution x is of the mathematical form

$$x = e^{-\gamma t/2}(Ae^{i\omega_\gamma t} + Be^{-i\omega_\gamma t}). \tag{24.19}$$

Now we may wonder why we bother to give this other solution, since we were so happy with the first one all by itself. What is the extra one for, because of course we know we should only take the real part? *We* know that we must take the real part, but how did the *mathematics* know that we only wanted the real part? When we had a nonzero driving force $F(t)$, we put in an *artificial* force to go with it, and the *imaginary* part of the equation, so to speak, was driven in a definite way. But when we put $F(t) \equiv 0$, our convention that x should be only the real part of whatever we write down is purely our own, and the mathematical equations do not know it yet. The physical world *has* a real solution, but the answer that we were so happy with before is not real, it is *complex.* The *equation* does not know that we are arbitrarily going to take the real part, so it has to present us, so to speak, with a complex conjugate type of solution, so that by putting them together we can *make a truly real* solution; that is what α_2 is doing for us. In order for x to be real, $Be^{-i\omega_\gamma t}$ will have to be the complex conjugate of $Ae^{i\omega_\gamma t}$, so that the imaginary parts disappear. So it turns out that B is the complex conjugate of A, and our real solution is

$$x = e^{-\gamma t/2}(Ae^{i\omega_\gamma t} + A^*e^{-i\omega_\gamma t}). \tag{24.20}$$

So our real solution is an oscillation with a *phase shift* and a damping—just as advertised.

24–3 Electrical transients

Now let us see if the above really works. We construct the electrical circuit shown in Fig. 24–2, in which we apply to an oscilloscope the voltage across the inductance L after we suddenly turn on a voltage by closing the switch S. It is an oscillatory circuit, and it generates a transient of some kind. It corresponds to a circumstance in which we suddenly apply a force and the system starts to oscillate. It is the electrical analog of a damped mechanical oscillator, and we watch the oscillation on an oscilloscope, where we should see the curves that we were trying to analyze. (The horizontal motion of the oscilloscope is driven at a uniform speed, while the vertical motion is the voltage across the inductor. The rest of the circuit is only a technical detail. We would like to repeat the experiment many, many times, since the persistence of vision is not good enough to see only one trace on the screen. So we do the experiment again and again by closing the switch 60 times a second; each time we close the switch, we also start the oscilloscope horizontal sweep, and it draws the curve over and over.) In Figs. 24–3 to 24–6 we see examples of damped oscillations, actually photographed on an oscilloscope screen. Figure 24–3 shows a damped oscillation in a circuit which has a high Q, a small γ. It does not die out very fast; it oscillates many times on the way down.

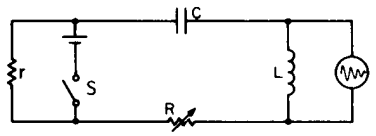


Fig. 24–2. An electrical circuit for demonstrating transients.

But let us see what happens as we decrease Q, so that the oscillation dies out more rapidly. We can decrease Q by increasing the resistance R in the circuit. When we increase the resistance in the circuit, it dies out faster (Fig. 24–4). Then if we increase the resistance in the circuit still more, it dies out faster still (Fig. 24–5). But when we put in more than a certain amount, we cannot see any oscillation at all! The question is, is this because our eyes are not good enough? If we increase the resistance still more, we get a curve like that of Fig. 24–6, which does not appear to have any oscillations, except perhaps one. Now, how can we explain that by mathematics?

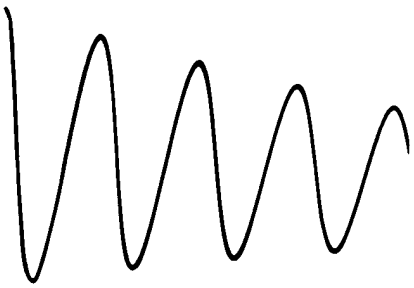


Figure 24–3

The resistance is, of course, proportional to the γ term in the mechanical device. Specifically, γ is R/L. Now if we increase the γ in the solutions (24.14) and (24.15) that we were so happy with before, chaos sets in when $\gamma/2$ exceeds ω_0; we must write it a different way, as

$$i\gamma/2 + i\sqrt{\gamma^2/4 - \omega_0^2} \quad \text{and} \quad i\gamma/2 - i\sqrt{\gamma^2/4 - \omega_0^2}.$$

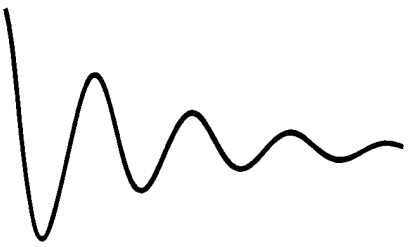


Figure 24–4

Those are now the two solutions and, following the same line of mathematical reasoning as previously, we again find two solutions: $e^{i\alpha_1 t}$ and $e^{i\alpha_2 t}$. If we now substitute for α_1, we get

$$x = Ae^{-(\gamma/2+\sqrt{\gamma^2/4-\omega_0^2})t},$$

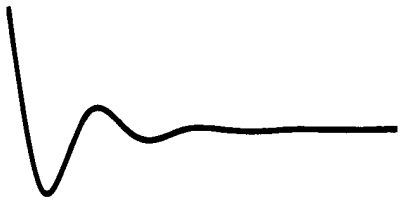


Figure 24–5

a nice exponential decay with no oscillations. Likewise, the other solution is

$$x = Be^{-(\gamma/2-\sqrt{\gamma^2/4-\omega_0^2})t}.$$

Note that the square root cannot exceed $\gamma/2$, because even if $\omega_0 = 0$, one term just equals the other. But ω_0^2 is taken away from $\gamma^2/4$, so the square root is less than $\gamma/2$, and the term in parentheses is, therefore, always a positive number. Thank goodness! Why? Because if it were negative, we would find e raised to a *positive* factor times t, which would mean it was exploding! In putting more and more resistance into the circuit, we know it is not going to explode—quite the contrary. So now we have two solutions, each one by itself a dying exponential, but one having a much faster "dying rate" than the other. The general solution is of course a combination of the two; the coefficients in the combination depending upon how the motion starts—what the initial conditions of the problem are. In the particular way this circuit happens to be starting, the A is negative and the B is positive, so we get the difference of two exponential curves.

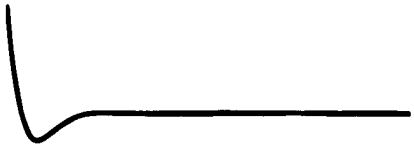


Figure 24–6

Now let us discuss how we can find the two coefficients A and B (or A and A^*), if we know how the motion was started.

Suppose that at $t = 0$ we know that $x = x_0$, and $dx/dt = v_0$. If we put $t = 0$, $x = x_0$, and $dx/dt = v_0$ into the expressions

$$x = e^{-\gamma t/2}(Ae^{i\omega_\gamma t} + A^*e^{-i\omega_\gamma t}),$$
$$dx/dt = e^{-\gamma t/2}[(-\gamma/2 + i\omega_\gamma)Ae^{i\omega_\gamma t} + (-\gamma/2 - i\omega_\gamma)A^*e^{-i\omega_\gamma t}],$$

we find, since $e^0 = e^{i0} = 1$,

$$\begin{aligned} x_0 &= A + A^* = 2A_R, \\ v_0 &= (-\gamma/2)(A + A^*) + i\omega_\gamma(A - A^*) \\ &= -\gamma x_0/2 + i\omega_\gamma(2iA_I), \end{aligned}$$

where $A = A_R + iA_I$, and $A^* = A_R - iA_I$. Thus we find

$$A_R = x_0/2$$

and

$$A_I = (v_0 + \gamma x_0/2)/2\omega_\gamma. \tag{24.21}$$

This completely determines A and A^*, and therefore the complete curve of the transient solution, in terms of how it begins. Incidentally, we can write the solution another way if we note that

$$e^{i\theta} + e^{-i\theta} = 2\cos\theta \quad \text{and} \quad e^{i\theta} - e^{-i\theta} = 2i\sin\theta.$$

We may then write the complete solution as

$$x = e^{-\gamma t/2}\left[x_0\cos\omega_\gamma t + \frac{v_0 + \gamma x_0/2}{\omega_\gamma}\sin\omega_\gamma t\right], \tag{24.22}$$

where $\omega_\gamma = +\sqrt{\omega_0^2 - \gamma^2/4}$. This is the mathematical expression for the way an oscillation dies out. We shall not make direct use of it, but there are a number of points we should like to emphasize that are true in more general cases.

First of all the behavior of such a system with no external force is expressed by a sum, or superposition, of pure exponentials in time (which we wrote as $e^{i\alpha t}$). This is a good solution to try in such circumstances. The values of α may be complex in general, the imaginary parts representing damping. Finally the intimate mathematical relation of the sinusoidal and exponential function discussed in Chapter 22 often appears physically as a change from oscillatory to exponential behavior when some physical parameter (in this case resistance, γ) exceeds some critical value.

25

Linear Systems and Review

25–1 Linear differential equations

25–2 Superposition of solutions

25–3 Oscillations in linear systems

25–4 Analogs in physics

25–5 Series and parallel impedances

25–1 Linear differential equations

In this chapter we shall discuss certain aspects of oscillating systems that are found somewhat more generally than just in the particular systems we have been discussing. For our particular system, the differential equation that we have been solving is

$$m\frac{d^2x}{dt^2} + \gamma m\frac{dx}{dt} + m\omega_0^2 x = F(t). \tag{25.1}$$

Now this particular combination of "operations" on the variable x has the interesting property that if we substitute $(x + y)$ for x, then we get the sum of the same operations on x and y; or, if we multiply x by a, then we get just a times the same combination. This is easy to prove. Just as a "shorthand" notation, because we get tired of writing down all those letters in (25.1), we shall use the symbol $\underline{L}(x)$ instead. When we see this, it means the left-hand side of (25.1), with x substituted in. With this system of writing, $\underline{L}(x + y)$ would mean the following:

$$\underline{L}(x + y) = m\frac{d^2(x + y)}{dt^2} + \gamma m\frac{d(x + y)}{dt} + m\omega_0^2(x + y). \tag{25.2}$$

(We underline the $\underline{L}$ so as to remind ourselves that it is not an ordinary function.) We sometimes call this an *operator notation*, but it makes no difference what we call it, it is just "shorthand."

Our first statement was that

$$\underline{L}(x + y) = \underline{L}(x) + \underline{L}(y), \tag{25.3}$$

which of course follows from the fact that $a(x + y) = ax + ay$, $d(x + y)/dt = dx/dt + dy/dt$, etc.

Our second statement was, for constant a,

$$\underline{L}(ax) = a\underline{L}(x). \tag{25.4}$$

[Actually, (25.3) and (25.4) are very closely related, because if we put $x + x$ into (25.3), this is the same as setting $a = 2$ in (25.4), and so on.]

In more complicated problems, there may be more derivatives, and more terms in $\underline{L}$; the question of interest is whether the two equations (25.3) and (25.4) are maintained or not. If they are, we call such a problem a *linear* problem. In this chapter we shall discuss some of the properties that exist because the system is linear, to appreciate the generality of some of the results that we have obtained in our special analysis of a special equation.

Now let us study some of the properties of linear differential equations, having illustrated them already with the specific equation (25.1) that we have studied so closely. The first property of interest is this: suppose that we have to solve the differential equation for a transient, the free oscillation with no driving force. That is, we want to solve

$$\underline{L}(x) = 0. \tag{25.5}$$

Suppose that, by some hook or crook, we have found a particular solution, which we shall call x_1. That is, we have an x_1 for which $\underline{L}(x_1) = 0$. Now we notice that ax_1 is also a solution to the same equation; we can multiply this special solution by any constant whatever, and get a new solution. In other words, if we had a

motion of a certain "size," then a motion twice as "big" is again a solution. *Proof:* $\underline{L}(ax_1) = a\underline{L}(x_1) = a \cdot 0 = 0$.

Next, suppose that, by hook or by crook, we have not only found *one* solution x_1, but also another solution, x_2. (Remember that when we substituted $x = e^{i\alpha t}$ for finding the transients, we found *two* values for α, that is, two solutions, x_1 and x_2.) Now let us show that the combination $(x_1 + x_2)$ is also a solution. In other words, if we put $x = x_1 + x_2$, x is again a solution of the equation. Why? Because, if $\underline{L}(x_1) = 0$ and $\underline{L}(x_2) = 0$, then $\underline{L}(x_1 + x_2) = \underline{L}(x_1) + \underline{L}(x_2) = 0 + 0 = 0$. So if we have found a number of solutions for the motion of a linear system we can add them together.

Combining these two ideas, we see, of course, that we can also add six of one and two of the other: if x_1 is a solution, so is αx_1. Therefore any sum of these two solutions, such as $(\alpha x_1 + \beta x_2)$, is also a solution. If we happen to be able to find three solutions, then we find that any combination of the three solutions is again a solution, and so on. It turns out that the number of what we call *independent solutions** that we have obtained for our oscillator problem is only *two*. The number of independent solutions that one finds in the general case depends upon what is called the number of *degrees of freedom*. We shall not discuss this in detail now, but if we have a second-order differential equation, there are only two independent solutions, and we have found both of them; so we have the most general solution.

Now let us go on to another proposition, which applies to the situation in which the system is subjected to an outside force. Suppose the equation is

$$\underline{L}(x) = F(t), \tag{25.6}$$

and suppose that we have found a special solution of it. Let us say that Joe's solution is x_J, and that $\underline{L}(x_J) = F(t)$. Suppose we want to find yet another solution; suppose we add to Joe's solution one of those that was a solution of the free equation (25.5), say x_1. Then we see by (25.3) that

$$\underline{L}(x_J + x_1) = \underline{L}(x_J) + \underline{L}(x_1) = F(t) + 0 = F(t). \tag{25.7}$$

Therefore, to the "forced" solution we can add any "free" solution, and we still have a solution. The free solution is called a *transient* solution.

When we have no force acting, and suddenly turn one on, we do not immediately get the steady solution that we solved for with the sine wave solution, but for a while there is a transient which sooner or later dies out, if we wait long enough. The "forced" solution does not die out, since it keeps on being driven by the force. Ultimately, for long periods of time, the solution is unique, but initially the motions are different for different circumstances, depending on how the system was started.

25–2 Superposition of solutions

Now we come to another interesting proposition. Suppose that we have a certain particular driving force F_a (let us say an oscillatory one with a certain $\omega = \omega_a$, but our conclusions will be true for any functional form of F_a) and we have solved for the forced motion (with or without the transients; it makes no difference). Now suppose some other force is acting, let us say F_b, and we solve the same problem, but for this different force. Then suppose someone comes along and says, "I have a new problem for you to solve; I have the force $F_a + F_b$." Can we do it? Of course we can do it, because the solution is the sum of the two solutions x_a and x_b for the forces taken separately—a most remarkable circumstance indeed. If we use (25.3), we see that

$$\underline{L}(x_a + x_b) = \underline{L}(x_a) + \underline{L}(x_b) = F_a(t) + F_b(t). \tag{25.8}$$

---

* Solutions which cannot be expressed as linear combinations of each other are called independent.

This is an example of what is called the *principle of superposition* for linear systems, and it is very important. It means the following: if we have a complicated force which can be broken up in any convenient manner into a sum of separate pieces, each of which is in some way simple, in the sense that for each special piece into which we have divided the force we can solve the equation, then the answer is available for the *whole* force, because we may simply add the pieces of the *solution* back together, in the same manner as the total *force* is compounded out of pieces (Fig. 25–1).

Let us give another example of the principle of superposition. In Chapter 12 we said that it was one of the great facts of the laws of electricity that if we have a certain distribution of charges q_a and calculate the electric field $\mathbf{E}_a$ arising from these charges at a certain place P, and if, on the other hand, we have another set of charges q_b and we calculate the field $\mathbf{E}_b$ due to these at the corresponding place, then if both charge distributions are present at the same time, the field $\mathbf{E}$ at P is the *sum* of $\mathbf{E}_a$ due to one set plus $\mathbf{E}_b$ due to the other. In other words, if we know the field due to a certain charge, then the field due to many charges is merely the vector sum of the fields of these charges taken individually. This is exactly analogous to the above proposition that if we know the result of two given forces taken at one time, then if the force is considered as a sum of them, the response is a sum of the corresponding individual responses.

The reason why this is true in electricity is that the great laws of electricity, Maxwell's equations, which determine the electric field, turn out to be differential equations which are *linear*, i.e., which have the property (25.3). What corresponds to the force is the *charge* generating the electric field, and the equation which determines the electric field in terms of the charge is linear.

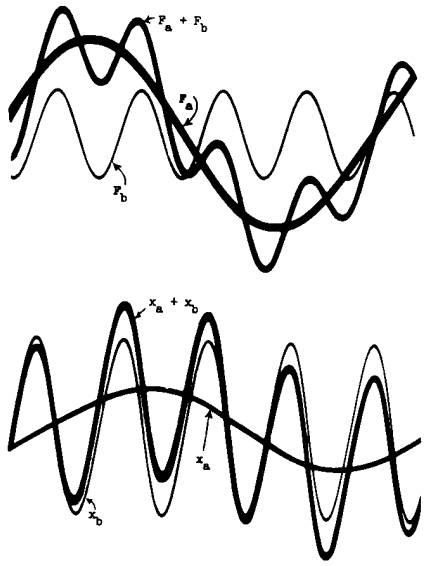


Fig. 25–1. An example of the principle of superposition for linear systems.

As another interesting example of this proposition, let us ask how it is possible to "tune in" to a particular radio station at the same time as all the radio stations are broadcasting. The radio station transmits, fundamentally, an oscillating electric field of very high frequency which acts on our radio antenna. It is true that the amplitude of the oscillation of the field is changed, modulated, to carry the signal of the voice, but that is very slow, and we are not going to worry about it. When one hears "This station is broadcasting at a frequency of 780 kilocycles," this indicates that 780,000 oscillations per second is the frequency of the electric field of the station antenna, and this drives the electrons up and down at that frequency in our antenna. Now at the same time we may have another radio station in the same town radiating at a different frequency, say 550 kilocycles per second; then the electrons in our antenna are also being driven by that frequency. Now the question is, how is it that we can separate the signals coming into the one radio at 780 kilocycles from those coming in at 550 kilocycles? We certainly do not hear both stations at the same time.

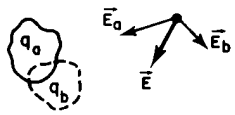


Fig. 25–2. The principle of superposition in electrostatics.

By the principle of superposition, the response of the electric circuit in the radio, the first part of which is a linear circuit, to the forces that are acting due to the electric field $F_a + F_b$, is $x_a + x_b$. It therefore looks as though we will never disentangle them. In fact, the very proposition of superposition seems to insist that we cannot *avoid* having both of them in our system. But remember, for a *resonant* circuit, the response curve, the amount of x per unit F, as a function of the frequency, looks like Fig. 25–3. If it were a very high Q circuit, the response would show a very sharp maximum. Now suppose that the two stations are comparable in strength, that is, the two *forces* are of the same magnitude. The *response* that we get is the sum of x_a and x_b. But, in Fig. 25–3, x_a is tremendous, while x_b is small. So, in spite of the fact that the two signals are equal in strength, when they go through the sharp resonant circuit of the radio tuned for ω_a, the frequency of the transmission of one station, then the response to this station is much greater than to the other. Therefore the complete response, with both signals acting, is almost all made up of ω_a, and we have selected the station we want.

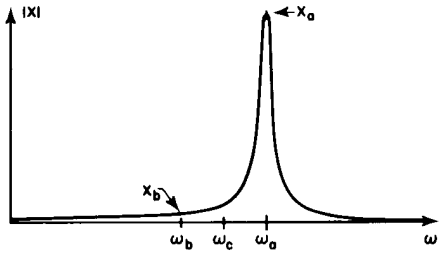


Fig. 25–3. A sharply tuned resonance curve.

Now what about the tuning? How do we tune it? We change ω_0 by changing the L or the C of the circuit, because the frequency of the circuit has to do with the combination of L and C. In particular, most radios are built so that one can change the capacitance. When we retune the radio, we can make a new setting of

the dial, so that the natural frequency of the circuit is shifted, say, to ω_c. In those circumstances we hear neither one station nor the other; we get silence, provided there is no other station at frequency ω_c. If we keep on changing the capacitance until the resonance curve is at ω_b, then of course we hear the other station. That is how radio tuning works; it is again the principle of superposition, combined with a resonant response.*

To conclude this discussion, let us describe qualitatively what happens if we proceed further in analyzing a linear problem with a given force, when the force is quite complicated. Out of the many possible procedures, there are two especially useful general ways that we can solve the problem. One is this: suppose that we can solve it for special known forces, such as sine waves of different frequencies. We know it is child's play to solve it for sine waves. So we have the so-called "child's play" cases. Now the question is whether our very complicated force can be represented as the sum of two or more "child's play" forces. In Fig. 25–1 we already had a fairly complicated curve, and of course we can make it more complicated still if we add in more sine waves. So it is certainly possible to obtain very complicated curves. And, in fact, the reverse is also true: practically every curve can be obtained by adding together *infinite numbers* of sine waves of different wavelengths (or frequencies) for each one of which we know the answer. We just have to know how much of each sine wave to put in to make the given F, and then our answer, x, is the corresponding sum of the F sine waves, each multiplied by its effective ratio of x to F. This method of solution is called the method of *Fourier transforms* or *Fourier analysis*. We are not going to actually carry out such an analysis just now; we only wish to describe the idea involved.

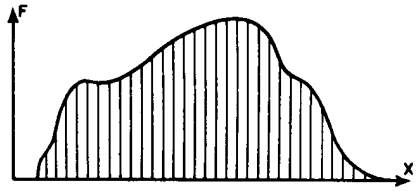


Fig. 25–4. A complicated force may be treated as a succession of sharp impulses.

Another way in which our complicated problem can be solved is the following very interesting one. Suppose that, by some tremendous mental effort, it were possible to solve our problem for a special force, namely an *impulse*. The force is quickly turned on and then off; it is all over. Actually we need only solve for an impulse of some unit strength, any other strength can be gotten by multiplication by an appropriate factor. We know that the response x for an impulse is a damped oscillation. Now what can we say about some other force, for instance a force like that of Fig. 25–4?

Such a force can be likened to a succession of blows with a hammer. First there is no force, and all of a sudden there is a steady force—impulse, impulse, impulse, impulse, . . . and then it stops. In other words, we imagine the continuous force to be a series of impulses, very close together. Now, we know the result for an impulse, so the result for a whole series of impulses will be a whole series of damped oscillations: it will be the curve for the first impulse, and then (slightly later) we add to that the curve for the second impulse, and the curve for the third impulse, and so on. Thus we can represent, mathematically, the complete solution for arbitrary functions if we know the answer for an impulse. We get the answer for any other force simply by integrating. This method is called the *Green's function method.* A Green's function is a response to an impulse, and the method of analyzing any force by putting together the response of impulses is called the Green's function method.

The physical principles involved in both of these schemes are so simple, involving just the linear equation, that they can be readily understood, but the *mathematical* problems that are involved, the complicated integrations and so on, are a little too advanced for us to attack right now. You will most likely return to this some day when you have had more practice in mathematics. But the *idea* is very simple indeed.

Finally, we make some remarks on why *linear* systems are so important. The answer is simple: because we can solve them! So most of the time we solve linear

---

* In modern superheterodyne receivers the actual operation is more complex. The amplifiers are all tuned to a fixed frequency (called IF frequency) and an oscillator of variable tunable frequency is combined with the input signal in a *nonlinear* circuit to produce a new frequency (the difference of signal and oscillator frequency) equal to the IF frequency, which is then amplified. This will be discussed in Chapter 50.

problems. Second (and most important), it turns out that *the fundamental laws of physics are often linear.* The Maxwell equations for the laws of electricity are linear, for example. The great laws of quantum mechanics turn out, so far as we know, to be linear equations. *That* is why we spend so much time on linear equations: because if we understand linear equations, we are ready, in principle, to understand a lot of things.

We mention another situation where linear equations are found. When displacements are small, many functions can be *approximated* linearly. For example, if we have a simple pendulum, the correct equation for its motion is

$$d^2\theta/dt^2 = -(g/L)\sin\theta. \tag{25.9}$$

This equation can be solved by elliptic functions, but the easiest way to solve it is numerically, as was shown in Chapter 9 on Newton's Laws of Motion. A nonlinear equation cannot be solved, ordinarily, any other way *but* numerically. Now for small θ, $\sin\theta$ is practically equal to θ, and we have a linear equation. It turns out that there are many circumstances where small effects are linear: for the example here the swing of a pendulum through small arcs. As another example, if we pull a little bit on a spring, the force is proportional to the extension. If we pull hard, we break the spring, and the force is a completely different function of the distance! Linear equations are important. In fact they are *so* important that perhaps fifty percent of the time we are solving linear equations in physics and in engineering.

25–3 Oscillations in linear systems

Let us now review the things we have been talking about in the past few chapters. It is very easy for the physics of oscillators to become obscured by the mathematics. The physics is actually very simple, and if we may forget the mathematics for a moment we shall see that we can understand almost everything that happens in an oscillating system. First, if we have only the spring and the weight, it is easy to understand why the system oscillates—it is a consequence of inertia. We pull the mass down and the force pulls it back up; as it passes zero, which is the place it likes to be, it cannot just suddenly stop; because of its momentum it keeps on going and swings to the other side, and back and forth. So, if there were no friction, we would surely *expect* an oscillatory motion, and indeed we get one. But if there is even a little bit of friction, then on the return cycle, the swing will not be quite as high as it was the first time.

Now what happens, cycle by cycle? That depends on the kind and amount of friction. Suppose that we could concoct a kind of friction force that always remains in the same proportion to the other forces, of inertia and in the spring, as the amplitude of oscillation varies. In other words, for smaller oscillations the friction should be weaker than for big oscillations. Ordinary friction does not have this property, so a special kind of friction must be carefully invented for the very purpose of creating a friction that is directly proportional to the velocity—so that for big oscillations it is stronger and for small oscillations it is weaker. If we happen to have that kind of friction, then at the end of each successive cycle the system is in the same condition as it was at the start, except a little bit smaller. All the forces are smaller in the same proportion: the spring force is reduced, the inertial effects are lower because the accelerations are now weaker, and the friction is less too, by our careful design. When we actually have that kind of friction, we find that each oscillation is exactly the same as the first one, except reduced in amplitude. If the first cycle dropped the amplitude, say, to 90 percent of what it was at the start, the next will drop it to 90 percent of 90 percent, and so on: *the sizes of the oscillations are reduced by the same fraction of themselves in every cycle.* An exponential function is a curve which does just that. It changes by the same factor in each equal interval of time. That is to say, if the amplitude of one cycle, relative to the preceding one, is called a, then the amplitude of the next is a^2, and of the next, a^3. So the amplitude is some constant raised to a power equal to the number of cycles traversed:

$$A = A_0 a^n. \tag{25.10}$$

But of course $n \sim t$, so it is perfectly clear that the general solution will be some kind of an oscillation, sine or cosine ωt, times an amplitude which goes as b^t more or less. But b can be written as e^{-c}, if b is positive and less than 1. So this is why the solution looks like $e^{-ct} \cos \omega t$. It is very simple.

What happens if the friction is not so artificial; for example, ordinary rubbing on a table, so that the friction force is a certain constant amount, and is independent of the size of the oscillation that reverses its direction each half-cycle? Then the equation is no longer linear, it becomes hard to solve, and must be solved by the numerical method given in Chapter 2, or by considering each half-cycle separately. The numerical method is the most powerful method of all, and can solve any equation. It is only when we have a simple problem that we can use mathematical analysis.

Mathematical analysis is not the grand thing it is said to be; it solves only the simplest possible equations. As soon as the equations get a little more complicated, just a shade—they cannot be solved analytically. But the numerical method, which was advertised at the beginning of the course, can take care of any equation of physical interest.

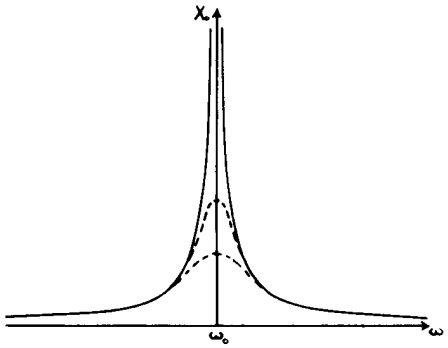


Fig. 25–5. Resonance curves with various amounts of friction present.

Next, what about the resonance curve? Why is there a resonance? First, imagine for a moment that there is no friction, and we have something which could oscillate by itself. If we tapped the pendulum just right each time it went by, of course we could make it go like mad. But if we close our eyes and do not watch it, and tap at arbitrary equal intervals, what is going to happen? Sometimes we will find ourselves tapping when it is going the wrong way. When we happen to have the timing just right, of course, each tap is given at just the right time, and so it goes higher and higher and higher. So without friction we get a curve which looks like the solid curve in Fig. 25–5 for different frequencies. Qualitatively, we understand the resonance curve; in order to get the exact shape of the curve it is probably just as well to do the mathematics. The curve goes toward infinity as $\omega \to \omega_0$, where ω_0 is the natural frequency of the oscillator.

Now suppose there is a little bit of friction; then when the displacement of the oscillator is small, the friction does not affect it much; the resonance curve is the same, except when we are near resonance. Instead of becoming infinite near resonance, the curve is only going to get so high that the work done by our tapping each time is enough to compensate for the loss of energy by friction during the cycle. So the top of the curve is rounded off—it does not go to infinity. If there is more friction, the top of the curve is rounded off still more. Now someone might say, "I thought the widths of the curves depended on the friction." That is because the curve is usually plotted so that the top of the curve is called one unit. However, the mathematical expression is even simpler to understand if we just plot all the curves on the same scale; then all that happens is that the friction cuts down the top! If there is less friction, we can go farther up into that little pinnacle before the friction cuts it off, so it looks relatively narrow. That is, the higher the peak of the curve, the narrower the width at half the maximum height.

Finally, we take the case where there is an enormous amount of friction. It turns out that if there is too much friction, the system does not oscillate at all. The energy in the spring is barely able to move it against the frictional force, and so it slowly oozes down to the equilibrium point.

25–4 Analogs in physics

The next aspect of this review is to note that masses and springs are not the only linear systems; there are others. In particular, there are electrical systems called linear circuits, in which we find a complete analog to mechanical systems. We did not learn exactly *why* each of the objects in an electrical circuit works in the way it does—that is not to be understood at the present moment; we may assert it as an experimentally verifiable fact that they behave as stated.

For example, let us take the simplest possible circumstance. We have a piece of wire, which is just a resistance, and we have applied to it a difference in potential, V. Now the V means this: if we carry a charge q through the wire from one terminal

to another terminal, the work done is qV. The higher the voltage difference, the more work was done when the charge, as we say, "falls" from the high potential end of the terminal to the low potential end. So charges release energy in going from one end to the other. Now the charges do not simply fly from one end straight to the other end; the atoms in the wire offer some resistance to the current, and this resistance obeys the following law for almost all ordinary substances: if there is a current I, that is, so and so many charges per second tumbling down, the number per second that comes tumbling through the wire is proportional to how hard we push them—in other words, proportional to how much voltage there is:

$$V = IR = R(dq/dt). \tag{25.11}$$

The coefficient R is called the *resistance*, and the equation is called Ohm's Law. The unit of resistance is the ohm; it is equal to one volt per ampere. In mechanical situations, to get such a frictional force in proportion to the velocity is difficult; in an electrical system it is very easy, and this law is extremely accurate for most metals.

We are often interested in how much work is done per second, the power loss, or the energy liberated by the charges as they tumble down the wire. When we carry a charge q through a voltage V, the work is qV, so the work done per second would be $V(dq/dt)$, which is the same as VI, or also $IR \cdot I = I^2R$. This is called the *heating loss*—this is how much heat is generated in the resistance per second, by the conservation of energy. It is this heat that makes an ordinary incandescent light bulb work.

Of course, there are other interesting properties of mechanical systems, such as the mass (inertia), and it turns out that there is an electrical analog to inertia also. It is possible to make something called an *inductor*, having a property called *inductance*, such that a current, once started through the inductance, *does not want to stop*. It requires a voltage in order to change the current! If the current is constant, there is no voltage across an inductance. DC circuits do not know anything about inductance; it is only when we *change* the current that the effects of inductance show up. The equation is

$$V = L(dI/dt) = L(d^2q/dt^2), \tag{25.12}$$

and the unit of inductance, called the *henry*, is such that one volt applied to an inductance of one henry produces a change of one ampere per second in the current. Equation (25.12) is the analog of Newton's law for electricity, if we wish: V corresponds to F, L corresponds to m, and I corresponds to velocity! All of the consequent equations for the two kinds of systems will have the same derivations because, in all the equations, we can change any letter to its corresponding analog letter and we get the same equation; everything we deduce will have a correspondence in the two systems.

Now what electrical thing corresponds to the mechanical spring, in which there was a force proportional to the stretch? If we start with $F = kx$ and replace $F \to V$ and $x \to q$, we get $V = \alpha q$. It turns out that there *is* such a thing, in fact it is the only one of the three circuit elements we can really understand, because we did study a pair of parallel plates, and we found that if there were a charge of certain equal, opposite amounts on each plate, the electric field between them would be proportional to the size of the charge. So the work done in moving a unit charge across the gap from one plate to the other is precisely proportional to the charge. This work is the *definition* of the voltage difference, and it is the line integral of the electric field from one plate to another. It turns out, for historical reasons, that the constant of proportionality is not called C, but $1/C$. It could have been called C, but it was not. So we have

$$V = q/C. \tag{25.13}$$

The unit of capacitance, C, is the farad; a charge of one coulomb on each plate of a one-farad capacitor yields a voltage difference of one volt.

There are our analogies, and the equation corresponding to the oscillating circuit becomes the following, by direct substitution of L for m, q for x, etc:

$$m(d^2x/dt^2) + \gamma m(dx/dt) + kx = F, \tag{25.14}$$

$$L(d^2q/dt^2) + R(dq/dt) + q/C = V. \tag{25.15}$$

Now everything we learned about (25.14) can be transformed to apply to (25.15). Every *consequence* is the same; so much the same that there is a brilliant thing we can do.

Suppose we have a mechanical system which is quite complicated, not just one mass on a spring, but several masses on several springs, all hooked together. What do we do? Solve it? Perhaps; but look, we can make an *electrical* circuit which will have the same equations as the thing we are trying to analyze! For instance, if we wanted to analyze a mass on a spring, why can we not build an electrical circuit in which we use an inductance proportional to the mass, a resistance proportional to the corresponding $m\gamma$, $1/C$ proportional to k, all in the same ratio? Then, of course, this electrical circuit will be the exact analog of our mechanical one, in the sense that whatever q does, in response to V (V also is made to correspond to the forces that are acting), so the x would do in response to the force! So if we have a complicated thing with a whole lot of interconnecting elements, we can interconnect a whole lot of resistances, inductances, and capacitances, to *imitate* the mechanically complicated system. What is the advantage to that? One problem is just as hard (or as easy) as the other, because they are exactly equivalent. The advantage is not that it is any easier to solve the *mathematical equations* after we discover that we have an electrical circuit (although that *is* the method used by electrical engineers!), but instead, the real reason for looking at the analog is that it is easier to *make* the electrical circuit, and to *change* something in the system.

Suppose we have designed an automobile, and want to know how much it is going to shake when it goes over a certain kind of bumpy road. We build an electrical circuit with inductances to represent the inertia of the wheels, spring constants as capacitances to represent the springs of the wheels, and resistors to represent the shock absorbers, and so on for the other parts of the automobile. Then we need a bumpy road. All right, we apply a *voltage* from a generator, which represents such and such a kind of bump, and then look at how the left wheel jiggles by measuring the charge on some capacitor. Having measured it (it is easy to do), we find that it is bumping too much. Do we need more shock absorber, or less shock absorber? With a complicated thing like an automobile, do we actually change the shock absorber, and solve it all over again? No!, we simply turn a dial; dial number ten is shock absorber number three, so we put in more shock absorber. The bumps are worse—all right, we try less. The bumps are still worse; we change the stiffness of the spring (dial 17), and we adjust all these things *electrically*, with merely the turn of a knob.

This is called an *analog computer*. It is a device which imitates the problem that we want to solve by making another problem, which has the same equation, but in another circumstance of nature, and which is easier to build, to measure, to adjust, and to destroy!

25–5 Series and parallel impedances

Finally, there is an important item which is not quite in the nature of review. This has to do with an electrical circuit in which there is more than one circuit element. For example, when we have an inductor, a resistor, and a capacitor connected as in Fig. 24–2, we note that all the charge went through every one of the three, so that the current in such a singly connected thing is the same at all points along the wire. Since the current is the same in each one, the voltage across R is IR, the voltage across L is $L(dI/dt)$, and so on. So, the total voltage drop is the sum of these, and this leads to Eq. (25.15). Using complex numbers, we found that we could solve the equation for the steady-state motion in response to a

sinusoidal force. We thus found that $\hat{V} = \hat{Z}\hat{I}$. Now $\hat{Z}$ is called the *impedance* of this particular circuit. It tells us that if we apply a sinusoidal voltage, $\hat{V}$, we get a current $\hat{I}$.

Now suppose we have a more complicated circuit which has two pieces, which by themselves have certain impedances, $\hat{Z}_1$ and $\hat{Z}_2$, and we put them in *series* (Fig. 25–6a) and apply a voltage. What happens? It is now a little more complicated, but if $\hat{I}$ is the current through $\hat{Z}_1$, the voltage difference across $\hat{Z}_1$ is $\hat{V}_1 = \hat{I}\hat{Z}_1$; similarly, the voltage across $\hat{Z}_2$ is $\hat{V}_2 = \hat{I}\hat{Z}_2$. *The same current goes through both.* Therefore the total voltage is the sum of the voltages across the two sections and is equal to $\hat{V} = \hat{V}_1 + \hat{V}_2 = (\hat{Z}_1 + \hat{Z}_2)\hat{I}$. This means that the voltage on the complete circuit can be written $\hat{V} = \hat{I}\hat{Z}_s$ where the $\hat{Z}_s$ of the combined system in series is the sum of the two $\hat{Z}$'s of the separate pieces:

$$\hat{Z}_s = \hat{Z}_1 + \hat{Z}_2. \tag{25.16}$$

This is not the only way things may be connected. We may also connect them in another way, called a *parallel* connection (Fig. 25–6b). Now we see that a given voltage across the terminals, if the connecting wires are perfect conductors, is effectively applied to both of the impedances, and will cause currents in each independently. Therefore the current through $\hat{Z}_1$ is equal to $\hat{I}_1 = \hat{V}/\hat{Z}_1$. The current in $\hat{Z}_2$ is $\hat{I}_2 = \hat{V}/\hat{Z}_2$. It is the *same voltage*. Now the total current which is supplied to the terminals is the *sum* of the currents in the two sections: $\hat{I} = \hat{V}/\hat{Z}_1 + \hat{V}/\hat{Z}_2$. This can be written as

$$\hat{V} = \frac{\hat{I}}{(\hat{1}/\hat{Z}_1) + (\hat{1}/\hat{Z}_2)} = \hat{I}\hat{Z}_p.$$

Thus

$$1/\hat{Z}_p = 1/\hat{Z}_1 + 1/\hat{Z}_2. \tag{25.17}$$

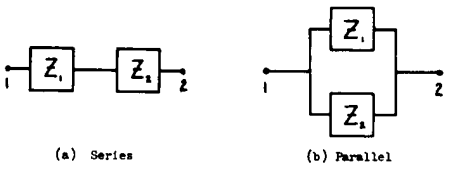


Fig. 25–6. Two impedances, connected in series and in parallel.

More complicated circuits can sometimes be simplified by taking pieces of them, working out the succession of impedances of the pieces, and combining the circuit together step by step, using the above rules. If we have any kind of circuit with many impedances connected in all kinds of ways, and if we include the voltages in the form of little generators having no impedance (when we pass charge through it, the generator adds a voltage V), then the following principles apply: (1) At any junction, the sum of the currents into a junction is zero. That is, all the current which comes in must come back out. (2) If we carry a charge around any loop, and back to where it started, the net work done is zero. These rules are called *Kirchhoff's laws* for electrical circuits. Their systematic application to complicated circuits often simplifies the analysis of such circuits. We mention them here in conjunction with Eqs. (25.16) and (25.17), in case you have already come across such circuits that you need to analyze in laboratory work. They will be discussed again in more detail next year.

26

Optics: The Principle of Least Time

26–1 Light

26–2 Reflection and refraction

26–3 Fermat's principle of least time

26–4 Applications of Fermat's principle

26–5 A more precise statement of Fermat's principle

26–6 How it works

26–1 Light

This is the first of a number of chapters on the subject of *electromagnetic radiation*. Light, with which we see, is only one small part of a vast spectrum of the same kind of thing, the various parts of this spectrum being distinguished by different values of a certain quantity which varies. This variable quantity could be called the "wavelength." As it varies in the visible spectrum, the light apparently changes color from red to violet. If we explore the spectrum systematically, from long wavelengths toward shorter ones, we would begin with what are usually called *radiowaves*. Radiowaves are technically available in a wide range of wavelengths, some even longer than those used in regular broadcasts; regular broadcasts have wavelengths corresponding to about 500 meters. Then there are the so-called "short waves," i.e., radar waves, millimeter waves, and so on. There are no actual boundaries between one range of wavelengths and another, because nature did not present us with sharp edges. The number associated with a given name for the waves are only approximate and, of course, so are the names we give to the different ranges.

Then, a long way down through the millimeter waves, we come to what we call the *infrared*, and thence to the visible spectrum. Then going in the other direction, we get into a region which is called the *ultraviolet*. Where the ultraviolet stops, the x-rays begin, but we cannot define precisely where this is; it is roughly at 10^{-8} m, or $10^{-2}\,\mu$. These are "soft" x-rays; then there are ordinary x-rays and very hard x-rays; then γ-rays, and so on, for smaller and smaller values of this dimension called the wavelength.

Within this vast range of wavelengths, there are three or more regions of approximation which are especially interesting. In one of these, a condition exists in which the wavelengths involved are very small compared with the dimensions of the equipment available for their study; furthermore, the photon energies, using the quantum theory, are small compared with the energy sensitivity of the equipment. Under these conditions we can make a rough first approximation by a method called *geometrical optics*. If, on the other hand, the wavelengths are comparable to the dimensions of the equipment, which is difficult to arrange with visible light but easier with radiowaves, and if the photon energies are still negligibly small, then a very useful approximation can be made by studying the behavior of the waves, still disregarding the quantum mechanics. This method is based on the *classical theory of electromagnetic radiation*, which will be discussed in a later chapter. Next, if we go to very short wavelengths, where we can disregard the wave character but the photons have a very *large* energy compared with the sensitivity of our equipment, things get simple again. This is the simple *photon* picture, which we will describe only very roughly. The complete picture, which unifies the whole thing into one model, will not be available to us for a long time.

In this chapter our discussion is limited to the geometrical optics region, in which we forget about the wavelength and the photon character of the light, which will all be explained in due time. We do not even bother to say what the light *is*, but just find out *how it behaves* on a large scale compared with the dimensions of interest. All this must be said in order to emphasize the fact that what we are going to talk about is only a very crude approximation; this is one of the chapters that we shall have to "unlearn" again. But we shall very quickly unlearn it, because we shall almost immediately go on to a more accurate method.

Although geometrical optics is just an approximation, it is of very great importance technically and of great interest historically. We shall present this subject more historically than some of the others in order to give some idea of the development of a physical theory or physical idea.

First, light is, of course, familiar to everybody, and has been familiar since time immemorial. Now one problem is, by what process do we *see* light? There have been many theories, but it finally settled down to one, which is that there is something which enters the eye—which bounces off objects into the eye. We have heard that idea so long that we accept it, and it is almost impossible for us to realize that very intelligent men have proposed contrary theories—that something comes out of the eye and feels for the object, for example. Some other important observations are that, as light goes from one place to another, it goes in *straight lines*, if there is nothing in the way, and that the rays do not seem to interfere with one another. That is, light is crisscrossing in all directions in the room, but the light that is passing across our line of vision does not affect the light that comes to us from some object. This was once a most powerful argument against the corpuscular theory; it was used by Huygens. If light were like a lot of arrows shooting along, how could other arrows go through them so easily? Such philosophical arguments are not of much weight. One could always say that light is made up of arrows which go through each other!

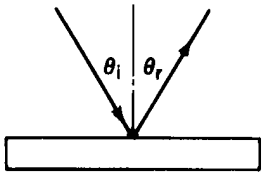


Fig. 26–1. The angle of incidence is equal to the angle of reflection.

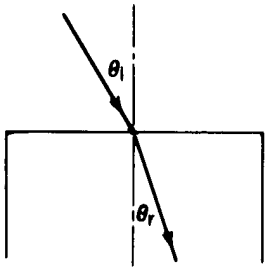


Fig. 26–2. A light ray is refracted when it passes from one medium into another.

Table 26–1

Angle in air	Angle in water
10°	8°
20°	15-1/2°
30°	22-1/2°
40°	29°
50°	35°
60°	40-1/2°
70°	45-1/2°
80°	50°

26–2 Reflection and refraction

The discussion above gives enough of the basic *idea* of geometrical optics—now we have to go a little further into the quantitative features. Thus far we have light going only in straight lines between two points; now let us study the behavior of light when it hits various materials. The simplest object is a mirror, and the law for a mirror is that when the light hits the mirror, it does not continue in a straight line, but bounces off the mirror into a new straight line, which changes when we change the inclination of the mirror. The question for the ancients was, what is the relation between the two angles involved? This is a very simple relation, discovered long ago. The light striking a mirror travels in such a way that the two angles, between each beam and the mirror, are equal. For some reason it is customary to measure the angles from the normal to the mirror surface. Thus the so-called law of reflection is

$$\theta_i = \theta_r. \tag{26.1}$$

That is a simple enough proposition, but a more difficult problem is encountered when light goes from one medium into another, for example from air into water; here also, we see that it does not go in a straight line. In the water the ray is at an inclination to its path in the air; if we change the angle θ_i so that it comes down more nearly vertically, then the angle of "breakage" is not as great. But if we tilt the beam of light at quite an angle, then the deviation angle is very large. The question is, what is the relation of one angle to the other? This also puzzled the ancients for a long time, and here they never found the answer! It is, however, one of the few places in all of Greek physics that one may find any experimental results listed. Claudius Ptolemy made a list of the angle in water for each of a number of different angles in air. Table 26–1 shows the angles in the air, in degrees, and the corresponding angle as measured in the water. (Ordinarily it is said that Greek scientists never did any experiments. But it would be impossible to obtain this table of values without knowing the right law, except by experiment. It should be noted, however, that these do not represent independent careful measurements for each angle but only some numbers interpolated from a few measurements, for they all fit perfectly on a parabola.)

This, then, is one of the important steps in the development of physical law: first we observe an effect, then we measure it and list it in a table; then we try to find the *rule* by which one thing can be connected with another. The above numerical table was made in 140 A.D., but it was not until 1621 that someone finally found the rule connecting the two angles! The rule, found by Willebrord

Snell, a Dutch mathematician, is as follows: if θ_i is the angle in air and θ_r is the angle in the water, then it turns out that the sine of θ_i is equal to some constant multiple of the sine of θ_r:

$$\sin \theta_i = n \sin \theta_r. \tag{26.2}$$

For water the number n is approximately 1.33. Equation (26.2) is called *Snell's law;* it permits us to *predict* how the light is going to bend when it goes from air into water. Table 26–2 shows the angles in air and in water according to Snell's law. Note the remarkable agreement with Ptolemy's list.

Table 26–2

Angle in air	Angle in water
10°	7-1/2°
20°	15°
30°	22°
40°	29°
50°	35°
60°	40-1/2°
70°	45°
80°	48°

26–3 Fermat's principle of least time

Now in the further development of science, we want more than just a formula. First we have an observation, then we have numbers that we measure, then we have a law which summarizes all the numbers. But the real *glory* of science is that *we can find a way of thinking* such that the law is *evident.*

The first way of thinking that made the law about the behavior of light evident was discovered by Fermat in about 1650, and it is called *the principle of least time,* or *Fermat's principle.* His idea is this: that out of all possible paths that it might take to get from one point to another, light takes the path which requires the *shortest time.*

Let us first show that this is true for the case of the mirror, that this simple principle contains both the law of straight-line propagation and the law for the mirror. So, we are growing in our understanding! Let us try to find the solution to the following problem. In Fig. 26–3 are shown two points, A and B, and a plane mirror, MM'. What is the way to get from A to B in the shortest time? The answer is to go straight from A to B! But if we add the extra rule that the light has to *strike the mirror* and come back in the shortest time, the answer is not so easy. One way would be to go as quickly as possible to the mirror and then go to B, on the path ADB. Of course, we then have a long path DB. If we move over a little to the right, to E, we slightly increase the first distance, but we greatly *decrease* the second one, and so the total path length, and therefore the travel time, is less. How can we find the point C for which the time is the shortest? We can find it very nicely by a geometrical trick.

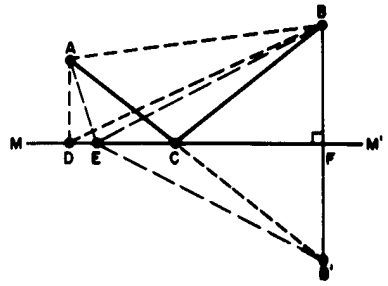


Fig. 26–3. Illustration of the principle of least time.

We construct on the other side of MM' an artificial point B', which is the same distance below the plane MM' as the point B is above the plane. Then we draw the line EB'. Now because BFM is a right angle and $BF = FB'$, EB is equal to EB'. Therefore the sum of the two distances, $AE + EB$, which is proportional to the time it will take if the light travels with constant velocity, is also the sum of the two lengths $AE + EB'$. Therefore the problem becomes, when is the sum of these two lengths the least? The answer is easy: when the line goes through point C as a *straight line* from A to B'! In other words, we have to find the point where we go toward the artificial point, and that will be the correct one. Now if ACB' is a straight line, then angle BCF is equal to angle $B'CF$ and thence to angle ACM. Thus the statement that the angle of incidence equals the angle of reflection is equivalent to the statement that the light goes to the mirror in such a way that it comes back to the point B' in the *least possible time.* Originally, the statement was made by Hero of Alexandria that the light travels in such a way that it goes to the mirror and to the other point in the shortest possible *distance,* so it is not a modern theory. It was this that inspired Fermat to suggest to himself that perhaps refraction operated on a similar basis. But for refraction, light obviously does not use the path of shortest *distance,* so Fermat tried the idea that it takes the shortest *time.*

Before we go on to analyze refraction, we should make one more remark about the mirror. If we have a source of light at the point B and it sends light toward the mirror, then we see that the light which goes to A from the point B comes to A in exactly the same manner as it would have come to A if there were an object at B', and *no* mirror. Now of course the eye detects only the light which enters it physically, so if we have an object at B and a mirror which makes the light come

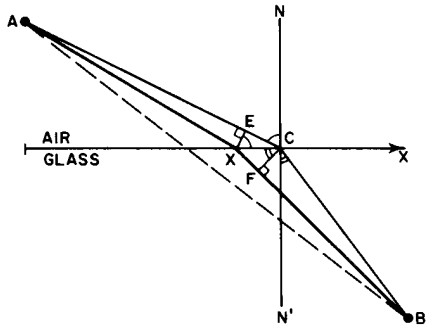


Fig. 26–4. Illustration of Fermat's principle for refraction.

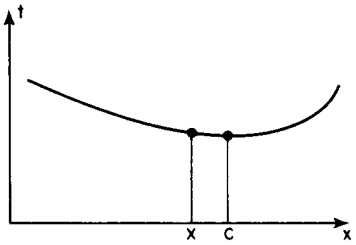


Fig. 26–5. The minimum time corresponds to point C, but nearby points correspond to nearly the same time.

into the eye in exactly the same manner as it would have come into the eye if the object were at B', then the eye-brain system interprets that, assuming it does not know too much, as *being* an object at B'. So the illusion that there is an object behind the mirror is merely due to the fact that the light which is entering the eye is entering in exactly the same manner, physically, as it would have entered had there *been* an object back there (except for the dirt on the mirror, and our knowledge of the existence of the mirror, and so on, which is corrected in the brain).

Now let us demonstrate that the principle of least time will give Snell's law of refraction. We must, however, make an assumption about the speed of light in water. We shall assume that the speed of light in water is lower than the speed of light in air by a certain factor, n.

In Fig. 26–4, our problem is again to go from A to B in *the shortest time.* To illustrate that the best thing to do is not just to go in a straight line, let us imagine that a beautiful girl has fallen out of a boat, and she is screaming for help in the water at point B. The line marked X is the shoreline. We are at point A on land, and we see the accident, and we can run and can also swim. But we can run faster than we can swim. What do we do? Do we go in a straight line? (Yes, no doubt!) However, by using a little more intelligence we would realize that it would be advantageous to travel a little greater distance on land in order to decrease the distance in the water, because we go so much slower in the water. (Following this line of reasoning out, we would say the right thing to do is to *compute* very carefully what should be done!) At any rate, let us try to show that the final solution to the problem is the path ACB, and that this path takes the shortest time of all possible ones. If it is the shortest path, that means that if we take any other, it will be longer. So, if we were to plot the time it takes against the position of point X, we would get a curve something like that shown in Fig. 26–5, where point C corresponds to the shortest of all possible times. This means that if we move the point X to points *near* C, in the first approximation there is essentially *no change* in time because the slope is zero at the bottom of the curve. So our way of finding the law will be to consider that we move the place by a very small amount, and to demand that there be essentially no change in time. (Of course there is an infinitesimal change of a *second* order; we ought to have a positive increase for displacements in either direction from C.) So we consider a nearby point X and we calculate how long it would take to go from A to B by the two paths, and compare the new path with the old path. It is very easy to do. We want the difference, of course, to be nearly zero if the distance XC is short. First, look at the path on land. If we draw a perpendicular XE, we see that this path is shortened by the amount EC. Let us say we gain by not having to go that extra distance. On the other hand, in the water, by drawing a corresponding perpendicular, CF, we find that we have to go the extra distance XF, and that is what we lose. Or, in *time*, we gain the time it would have taken to go the distance EC, but we lose the time it would have taken to go the distance XF. Those times must be equal since, in the first approximation, there is to be no change in time. But supposing that in the water the speed is $1/n$ times as fast as in air, then we must have

$$EC = n \cdot XF. \tag{26.3}$$

Therefore we see that when we have the right point, $XC \sin EXC = n \cdot XC \sin XCF$ or, cancelling the common hypotenuse length XC and noting that

$$EXC = ECN = \theta_i \quad \text{and} \quad XCF = BCN' = \theta_r,$$

we have

$$\sin \theta_i = n \sin \theta_r. \tag{26.4}$$

So we see that to get from one point to another in the least time when the ratio of speeds is n, the light should enter at such an angle that the ratio of the sines of the angles θ_i and θ_r is the ratio of the speeds in the two media.

26–4 Applications of Fermat's principle

Now let us consider some of the interesting consequences of the principle of least time. First is the principle of reciprocity. If to go from A to B we have found the path of the least time, then to go in the opposite direction (assuming that light goes at the same speed in any direction), the shortest time will be the same path, and therefore, if light can be sent one way, it can be sent the other way.

An example of interest is a glass block with plane parallel faces, set at an angle to a light beam. Light, in going through the block from a point A to a point B (Fig. 26–6) does not go through in a straight line, but instead it decreases the time in the block by making the angle in the block less inclined, although it loses a little bit in the air. The beam is simply displaced parallel to itself because the angles in and out are the same.

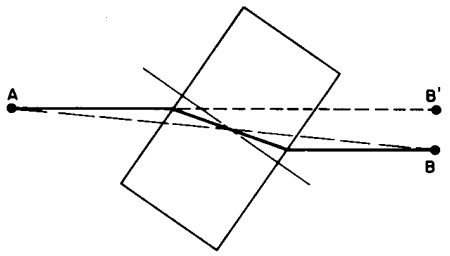


Fig. 26–6. A beam of light is offset as it passes through a transparent block.

A third interesting phenomenon is the fact that when we see the sun setting, it is already below the horizon! It does not *look* as though it is below the horizon, but it is (Fig. 26–7). The earth's atmosphere is thin at the top and dense at the bottom. Light travels more slowly in air than it does in a vacuum, and so the light of the sun can get to point S beyond the horizon more quickly if, instead of just going in a straight line, it avoids the dense regions where it goes slowly by getting through them at a steeper tilt. When it appears to go below the horizon, it is actually already well below the horizon. Another example of this phenomenon is the mirage that one often sees while driving on hot roads. One sees "water" on the road, but when he gets there, it is as dry as the desert! The phenomenon is the following. What we are really seeing is the sky light "reflected" on the road: light from the sky, heading for the road, can end up in the eye, as shown in Fig. 26–8. Why? The air is very hot just above the road but it is cooler up higher. Hotter air is more expanded than cooler air and is thinner, and this decreases the speed of light less. That is to say, light goes faster in the hot region than in the cool region. Therefore, instead of the light deciding to come in the straightforward way, it also has a least-time path by which it goes into the region where it goes faster for awhile, in order to save time. So, it can go in a curve.

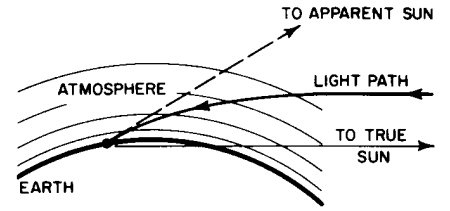


Fig. 26–7. Near the horizon, the apparent sun is higher than the true sun by about 1/2 degree.

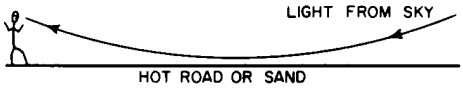


Fig. 26–8. A mirage.

As another important example of the principle of least time, suppose that we would like to arrange a situation where we have all the light that comes out of one point, P, collected back together at another point, P' (Fig. 26–9). That means, of course, that the light can go in a straight line from P to P'. That is all right. But how can we arrange that not only does it go straight, but also so that the light starting out from P toward Q also ends up at P'? We want to bring all the light back to what we call a *focus*. How? If the light always takes the path of least time, then certainly it should not want to go over all these other paths. The only way that the light can be perfectly satisfied to take several adjacent paths is to make those times *exactly equal!* Otherwise, it would select the one of least time. Therefore the problem of making a focusing system is merely to arrange a device so that it takes the same time for the light to go on *all* the different paths!

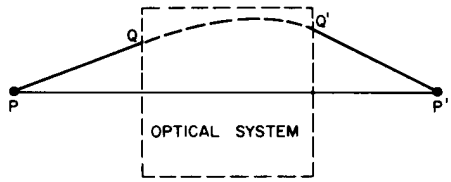


Fig. 26–9. An optical "black box."

This is easy to do. Suppose that we had a piece of glass in which light goes slower than it does in the air (Fig. 26–10). Now consider a ray which goes in air in the path PQP'. That is a longer path than from P directly to P' and no doubt takes a longer time. But if we were to insert a piece of glass of just the right thickness (we shall later figure out how thick) it might exactly compensate the excess time that it would take the light to go at an angle! In those circumstances we can arrange that the time the light takes to go straight through is the same as the time it takes to go in the path PQP'. Likewise, if we take a ray $PRR'P'$ which is partly inclined, it is not quite as long as PQP', and we do not have to compensate as much as for the straight one, but we do have to compensate somewhat. We end up with a piece of glass that looks like Fig. 26–10. With this shape, all the light which comes from P will go to P'. This, of course, is well known to us, and we call such a device a converging *lens*. In the next chapter we shall actually calculate what shape the lens has to have to make a perfect focus.

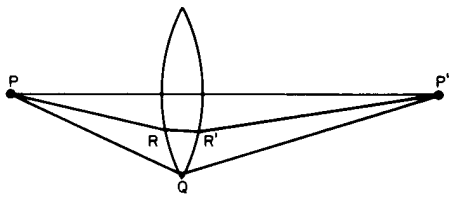


Fig. 26–10. A focusing optical system.

Take another example: suppose we wish to arrange some mirrors so that the light from P always goes to P' (Fig. 26–11). On any path, it goes to some mirror

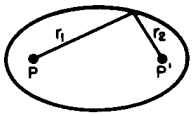


Fig. 26–11. An ellipsoidal mirror.

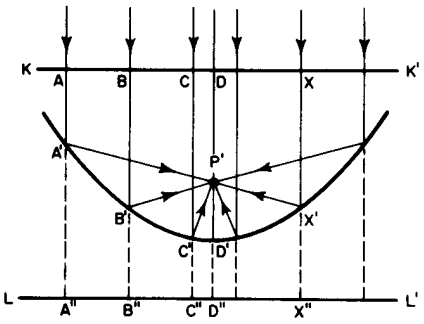


Fig. 26–12. A paraboloidal mirror.

and comes back, and all times must be equal. Here the light always travels in air, so the time and the distance are proportional. Therefore the statement that all the times are the same is the same as the statement that the total distance is the same. Thus the sum of the two distances r_1 and r_2 must be a constant. An *ellipse* is that curve which has the property that the sum of the distances from two points is a constant for every point on the ellipse; thus we can be sure that the light from one focus will come to the other.

The same principle works for gathering the light of a star. The great 200-inch Palomar telescope is built on the following principle. Imagine a star billions of miles away; we would like to cause all the light that comes in to come to a focus. Of course we cannot draw the rays that go all the way up to the star, but we still want to check whether the times are equal. Of course we know that when the various rays have arrived at some plane KK', perpendicular to the rays, all the times in this plane are equal (Fig. 26–12). The rays must then come down to the mirror and proceed toward P' in equal times. That is, we must find a curve which has the property that the sum of the distances $XX' + X'P'$ is a constant, no matter where X is chosen. An easy way to find it is to extend the length of the line XX' down to a plane LL'. Now if we arrange our curve so that $A'A'' = A'P'$, $B'B'' = B'P'$, $C'C'' = C'P'$, and so on, we will have our curve, because then of course, $AA' + A'P' = AA' + A'A''$ will be constant. Thus our curve is the locus of all points equidistant from a line and a point. Such a curve is called a *parabola*; the mirror is made in the shape of a parabola.

The above examples illustrate the principle upon which such optical devices can be designed. The exact curves can be calculated using the principle that, to focus perfectly, the travel times must be exactly equal for all light rays, as well as being less than for any other nearby path.

We shall discuss these focusing optical devices further in the next chapter; let us now discuss the further development of the theory. When a new theoretical principle is developed, such as the principle of least time, our first inclination might be to say, "Well, that is very pretty; it is delightful; but the question is, does it help at all in understanding the physics?" Someone may say, "Yes, look at how many things we can now understand!" Another says, "Very well, but I can understand mirrors, too. I need a curve such that every tangent plane makes equal angles with the two rays. I can figure out a lens, too, because every ray that comes to it is bent through an angle given by Snell's law." Evidently the statement of least time and the statement that angles are equal on reflection, and that the sines of the angles are proportional on refraction, are the same. So is it merely a philosophical question, or one of beauty? There can be arguments on both sides.

However, the importance of a powerful principle is that *it predicts new things.*

It is easy to show that there are a number of new things predicted by Fermat's principle. First, suppose that there are *three* media, glass, water, and air, and we perform a refraction experiment and measure the index n for one medium against another. Let us call n_{12} the index of air (1) against water (2); n_{13} the index of air (1) against glass (3). If we measured water against glass, we should find another index, which we shall call n_{23}. But there is no *a priori* reason why there should be any connection between n_{12}, n_{13}, and n_{23}. On the other hand, according to the idea of least time, there *is* a definite relationship. The index n_{12} is the ratio of two things, the speed in air to the speed in water; n_{13} is the ratio of the speed in air to the speed in glass; n_{23} is the ratio of the speed in water to the speed in glass. Therefore we cancel out the air, and get

$$n_{23} = \frac{v_2}{v_3} = \frac{v_1/v_3}{v_1/v_2} = \frac{n_{13}}{n_{12}}. \tag{26.5}$$

In other words, we *predict* that the index for a new pair of materials can be obtained from the indexes of the individual materials, both against air or against vacuum. So if we measure the speed of light in all materials, and from this get a single number for each material, namely its index relative to vacuum, called n_i

(n_1 is the speed in air relative to the speed in vacuum, etc.), then our formula is easy. The index for any two materials i and j is

$$n_{ij} = \frac{v_i}{v_j} = \frac{n_j}{n_i}. \tag{26.6}$$

Using only Snell's law, there is no basis for a prediction of this kind.* But of course this prediction works. The relation (26.5) was known very early, and was a very strong argument for the principle of least time.

Another argument for the principle of least time, another prediction, is that if we *measure* the speed of light in water, it will be lower than in air. This is a prediction of a completely different type. It is a brilliant prediction, because all we have so far measured are *angles*; here we have a theoretical prediction which is quite different from the observations from which Fermat deduced the idea of least time. It turns out, in fact, that the speed in water *is* slower than the speed in air, by just the proportion that is needed to get the right index!

26–5 A more precise statement of Fermat's principle

Actually, we must make the statement of the principle of least time a little more accurately. It was not stated correctly above. It is *incorrectly* called the principle of least time and we have gone along with the incorrect description for convenience, but we must now see what the correct statement is. Suppose we had a mirror as in Fig. 26–3. What makes the light think it has to go to the mirror? The path of *least* time is clearly AB. So some people might say, "Sometimes it is a maximum time." It is *not* a maximum time, because certainly a curved path would take a still longer time! The correct statement is the following: a ray going in a certain particular path has the property that if we make a small change (say a one percent shift) in the ray in any manner whatever, say in the location at which it comes to the mirror, or the shape of the curve, or anything, there will be *no* first-order change in the time; there will be only a *second*-order change in the time. In other words, the principle is that light takes a path such that there are many other paths nearby which take almost exactly the *same* time.

The following is another difficulty with the principle of least time, and one which people who do not like this kind of a theory could never stomach. With Snell's theory we can "understand" light. Light goes along, it sees a surface, it bends because it does something at the surface. The idea of causality, that it goes from one point to another, and another, and so on, is easy to understand. But the principle of least time is a completely different philosophical principle about the way nature works. Instead of saying it is a causal thing, that when we do one thing, something else happens, and so on, it says this: we set up the situation, and *light* decides which is the shortest time, or the extreme one, and chooses that path. But *what* does it do, *how* does it find out? Does it *smell* the nearby paths, and check them against each other? The answer is, yes, it does, in a way. That is the feature which is, of course, not known in geometrical optics, and which is involved in the idea of *wavelength*; the wavelength tells us approximately how far away the light must "smell" the path in order to check it. It is hard to demonstrate this fact on a large scale with light, because the wavelengths are so terribly short. But with radiowaves, say 3-cm waves, the distances over which the radiowaves are checking are larger. If we have a source of radiowaves, a detector, and a slit, as in Fig. 26–13, the rays of course go from S to D because it is a straight line, and if we close down the slit it is all right—they still go. But now if we move the detector aside to D', the waves will not go through the wide slit from S to D', because they check several paths nearby, and say, "No, my friend, those all correspond to different times." On the other hand, if we *prevent* the radiation from checking the paths by closing the slit down to a very narrow crack, then there is but one path available, and the

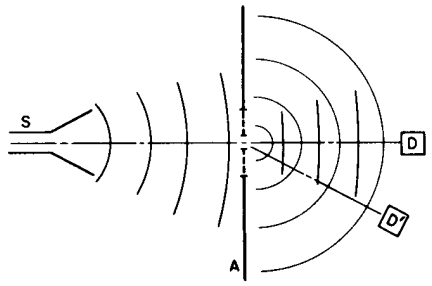


Fig. 26–13. The passage of radiowaves through a narrow slit.

* Although it can be deduced if the additional assumption is made that adding a layer of one substance to the surface of another does not change the eventual angle of refraction in the latter material.

radiation takes it! With a narrow slit, more radiation reaches D' than reaches it with a wide slit!

One can do the same thing with light, but it is hard to demonstrate on a large scale. The effect can be seen under the following simple conditions. Find a small, bright light, say an unfrosted bulb in a street light far away or the reflection of the sun in a curved automobile bumper. Then put two fingers in front of one eye, so as to look through the crack, and squeeze the light to zero very gently. You will see that the image of the light, which was a little dot before, becomes quite elongated, and even stretches into a long line. The reason is that the fingers are very close together, and the light which is supposed to come in a straight line is spread out at an angle, so that when it comes into the eye it comes in from several directions. Also you will notice, if you are very careful, side maxima, a lot of fringes along the edges too, Furthermore, the whole thing is colored. All of this will be explained in due time, but for the present it is a demonstration that light does not always go in straight lines, and it is one that is very easily performed.

26–6 How it works

Finally, we give a very crude view of what actually happens, how the whole thing really works, from what we now believe is the correct, quantum-dynamically accurate viewpoint, but of course only qualitatively described. In following the light from A to B in Fig. 26–3, we find that the light does not seem to be in the form of waves at all. Instead the rays seem to be made up of photons, and they actually produce clicks in a photon counter, if we are using one. The brightness of the light is proportional to the average number of photons that come in per second, and what we calculate is the *chance* that a photon gets from A to B, say by hitting the mirror. The *law* for that chance is the following very strange one. Take any path and find the time for that path; then make a complex number, or draw a little complex vector, $pe^{i\theta}$, *whose angle* θ *is proportional to the time.* The number of turns per second is the frequency of the light. Now take another path; it has, for instance, a different time, so the vector for it is turned through a different angle—the angle being always proportional to the time. Take *all* the available paths and add on a little vector for each one; then the answer is that the chance of arrival of the photon is proportional to the square of the length of the final vector, from the beginning to the end!

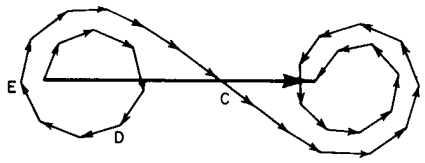


Fig. 26–14. The summation of probability amplitudes for many neighboring paths.

Now let us show how this implies the principle of least time for a mirror. We consider all rays, all possible paths ADB, AEB, ACB, etc., in Fig. 26–3. The path ADB makes a certain small contribution, but the next path, AEB, takes a quite different time, so its angle θ is quite different. Let us say that point C corresponds to minimum time, where if we change the paths the times do not change. So for awhile the times do change, and then they begin to change less and less as we get near point C (Fig. 26–14). So the arrows which we have to add are coming almost exactly at the same angle for awhile near C, and then gradually the time begins to increase again, and the phases go around the other way, and so on. Eventually, we have quite a tight knot. The total probability is the distance from one end to the other, squared. *Almost all of that accumulated probability occurs in the region where all the arrows are in the same direction* (or in the same phase). All the contributions from the paths which have very *different* times as we change the path, cancel themselves out by pointing in different directions. That is why, if we hide the extreme parts of the mirror, it still reflects almost exactly the same, because all we did was to take out a piece of the diagram inside the spiral ends, and that makes only a very small change in the light. So this is the relationship between the ultimate picture of photons with a probability of arrival depending on an accumulation of arrows, and the principle of least time.

27

Geometrical Optics

27–1 Introduction

27–2 The focal length of a spherical surface

27–3 The focal length of a lens

27–4 Magnification

27–5 Compound lenses

27–6 Aberrations

27–7 Resolving power

27–1 Introduction

In this chapter we shall discuss some elementary applications of the ideas of the previous chapter to a number of practical devices, using the approximation called *geometrical optics*. This is a most useful approximation in the practical design of many optical systems and instruments. Geometrical optics is either very simple or else it is very complicated. By that we mean that we can either study it only superficially, so that we can design instruments roughly, using rules that are so simple that we hardly need deal with them here at all, since they are practically of high school level, or else, if we want to know about the small errors in lenses and similar details, the subject gets so complicated that it is too advanced to discuss here! If one has an actual, detailed problem in lens design, including analysis of aberrations, then he is advised to read about the subject or else simply to trace the rays through the various surfaces (which is what the book tells how to do), using the law of refraction from one side to the other, and to find out where they come out and see if they form a satisfactory image. People have said that this is too tedious, but today, with computing machines, it is the right way to do it. One can set up the problem and make the calculation for one ray after another very easily. So the subject is really ultimately quite simple, and involves no new principles. Furthermore, it turns out that the rules of either elementary or advanced optics are seldom characteristic of other fields, so that there is no special reason to follow the subject very far, with one important exception.

The most advanced and abstract theory of geometrical optics was worked out by Hamilton, and it turns out that this has very important applications in mechanics. It is actually even more important in mechanics than it is in optics, and so we leave Hamilton's theory for the subject of advanced analytical mechanics, which is studied in the senior year or in graduate school. So, appreciating that geometrical optics contributes very little, except for its own sake, we now go on to discuss the elementary properties of simple optical systems on the basis of the principles outlined in the last chapter.

In order to go on, we must have one geometrical formula, which is the following: if we have a triangle with a small altitude h and a long base d, then the diagonal s (we are going to need it to find the difference in time between two different routes) is longer than the base (Fig. 27–1). How much longer? The difference $\Delta = s - d$ can be found in a number of ways. One way is this. We see that $s^2 - d^2 = h^2$, or $(s - d)(s + d) = h^2$. But $s - d = \Delta$, and $s + d \sim 2s$. Thus

$$\Delta \sim h^2/2s. \tag{27.1}$$

This is all the geometry we need to discuss the formation of images by curved surfaces!

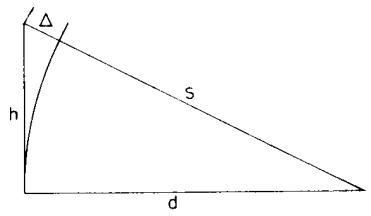


Figure 27–1

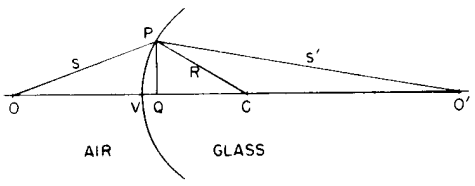


Fig. 27–2. Focusing by a single refracting surface.

27–2 The focal length of a spherical surface

The first and simplest situation to discuss is a single refracting surface, separating two media with different indices of refraction (Fig. 27–2). We leave the case of arbitrary indices of refraction to the student, because *ideas* are always the most important thing, not the specific situation, and the problem is easy enough to do in any case. So we shall suppose that, on the left, the speed is 1 and on the right it is $1/n$, where n is the index of refraction. The light travels more slowly in the glass by a factor n.

Now suppose that we have a point at O, at a distance s from the front surface of the glass, and another point O' at a distance s' inside the glass, and we desire to arrange the curved surface in such a manner that every ray from O which hits the surface, at any point P, will be bent so as to proceed toward the point O'. For that to be true, we have to shape the surface in such a way that the time it takes for the light to go from O to P, that is, the distance OP divided by the speed of light (the speed here is unity), plus $n \cdot O'P$, which is the time it takes to go from P to O', is equal to a constant independent of the point P. This condition supplies us with an equation for determining the surface. The answer is that the surface is a very complicated fourth-degree curve, and the student may entertain himself by trying to calculate it by analytic geometry. It is simpler to try a special case that corresponds to $s \to \infty$, because then the curve is a second-degree curve and is more recognizable. It is interesting to compare this curve with the parabolic curve we found for a focusing mirror when the light is coming from infinity.

So the proper surface cannot easily be made—to focus the light from one point to another requires a rather complicated surface. It turns out in practice that we do not try to make such complicated surfaces ordinarily, but instead we make a compromise. Instead of trying to get *all* the rays to come to a focus, we arrange it so that only the rays fairly close to the axis OO' come to a focus. The farther ones may deviate if they want to, unfortunately, because the ideal surface is complicated, and we use instead a spherical surface with the right curvature at the axis. It is so much easier to fabricate a sphere than other surfaces that it is profitable for us to find out what happens to rays striking a spherical surface, supposing that only the rays near the axis are going to be focused perfectly. Those rays which are near the axis are sometimes called *paraxial rays*, and what we are analyzing are the conditions for the focusing of paraxial rays. We shall discuss later the errors that are introduced by the fact that all rays are not always close to the axis.

Thus, supposing P is close to the axis, we drop a perpendicular PQ such that the height PQ is h. For a moment, we imagine that the surface is a plane passing through P. In that case, the time needed to go from O to P would exceed the time from O to Q, and also, the time from P to O' would exceed the time from Q to O'. But that is why the glass must be curved, because the total excess time must be compensated by the delay in passing from V to Q! Now the *excess* time along route OP is $h^2/2s$, and the excess time on the other route is $nh^2/2s'$. This excess time, which must be matched by the delay in going along VQ, differs from what it would have been in a vacuum, because there is a medium present. In other words, the time to go from V to Q is not as if it were straight in the air, but it is slower by the factor n, so that the excess delay in this distance is then $(n - 1)VQ$. And now, how large is VQ? If the point C is the center of the sphere and if its radius is R, we see by the same formula that the distance VQ is equal to $h^2/2R$. Therefore we discover that the law that connects the distances s and s', and that gives us the radius of curvature R of the surface that we need, is

$$(h^2/2s) + (nh^2/2s') = (n - 1)h^2/2R \tag{27.2}$$

or

$$(1/s) + (n/s') = (n - 1)/R. \tag{27.3}$$

If we have a position O and another position O', and want to focus light from O to O', then we can calculate the required radius of curvature R of the surface by this formula.

Now it turns out, interestingly, that the same lens, with the same curvature R, will focus for other distances, namely, for any pair of distances such that the sum of the two reciprocals, one multiplied by n, is a constant. Thus a given lens will (so long as we limit ourselves to paraxial rays) focus not only from O to O', but between an infinite number of other pairs of points, so long as those pairs of points bear the relationship that $1/s + n/s'$ is a constant, characteristic of the lens.

In particular, an interesting case is that in which $s \to \infty$. We can see from the formula that as one s increases, the other decreases. In other words, if point O

goes out, point O' comes in, and vice versa. As point O goes toward infinity, point O' keeps moving in until it reaches a certain distance, called the *focal length* f', inside the material. If parallel rays come in, they will meet the axis at a distance f'. Likewise, we could imagine it the other way. (Remember the reciprocity rule: if light will go from O to O', of course it will also go from O' to O.) Therefore, if we had a light source inside the glass, we might want to know where the focus is. In particular, if the light in the glass were at infinity (same problem) where would it come to a focus outside? This distance is called f. Of course, we can also put it the other way. If we had a light source at f and the light went through the surface, then it would go out as a parallel beam. We can easily find out what f and f' are:

$$n/f' = (n - 1)/R \quad \text{or} \quad f' = Rn/(n - 1), \tag{27.4}$$

$$1/f = (n - 1)/R \quad \text{or} \quad f = R/(n - 1). \tag{27.5}$$

We see an interesting thing: if we divide each focal length by the corresponding index of refraction we get the same result! This theorem, in fact, is general. It is true of any system of lenses, no matter how complicated, so it is interesting to remember. We did not prove here that it is general—we merely noted it for a single surface, but it happens to be true in general that the two focal lengths of a system are related in this way. Sometimes Eq. (27.3) is written in the form

$$1/s + n/s' = 1/f. \tag{27.6}$$

This is more useful than (27.3) because we can measure f more easily than we can measure the curvature and index of refraction of the lens: if we are not interested in designing a lens or in knowing how it got that way, but simply lift it off a shelf, the interesting quantity is f, not the n and the 1 and the R!

Now an interesting situation occurs if s becomes less than f. What happens then? If $s < f$, then $(1/s) > (1/f)$, and therefore s' is negative; our equation says that the light will focus only with a negative value of s', whatever that means! It does mean something very interesting and very definite. It is still a useful formula, in other words, even when the numbers are negative. What it means is shown in Fig. 27–3. If we draw the rays which are diverging from O, they will be bent, it is true, at the surface, and they will not come to a focus, because O is so close in that they are "beyond parallel." However, they diverge as if they had come from a point O' *outside* the glass. This is an apparent image, sometimes called a *virtual image*. The image O' in Fig. 27–2 is called a *real image*. If the light really comes to a point, it is a real image. But if the light *appears* to be coming *from* a point, a fictitious point different from the original point, it is a virtual image. So when s' comes out negative, it means that O' is on the other side of the surface, and everything is all right.

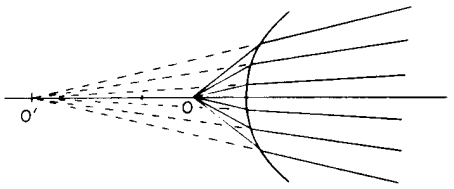


Fig. 27–3. A virtual image.

Now consider the interesting case where R is equal to infinity; then we have $(1/s) + (n/s') = 0$. In other words, $s' = -ns$, which means that if we look from a dense medium into a rare medium and see a point in the rare medium, it appears to be deeper by a factor n. Likewise, we can use the same equation backwards, so that if we look into a plane surface at an object that is at a certain distance inside the dense medium, it will appear as though the light is coming from not as far back (Fig. 27–4). When we look at the bottom of a swimming pool from above, it does not look as deep as it really is, by a factor 3/4, which is the reciprocal of the index of refraction of water.

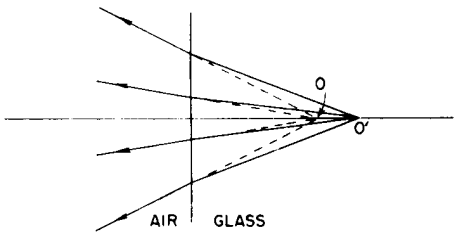


Fig. 27–4. A plane surface re-images the light from O′ to O.

We could go on, of course, to discuss the spherical mirror. But if one appreciates the ideas involved, he should be able to work it out for himself. Therefore we leave it to the student to work out the formula for the spherical mirror, but we mention that it is well to adopt certain conventions concerning the distances involved:

(1) The object distance s is positive if the point O is to the left of the surface.

(2) The image distance s' is positive if the point O' is to the right of the surface.

(3) The radius of curvature of the surface is positive if the center is to the right of the surface.

In Fig. 27–2, for example, s, s', and R are all positive; in Fig. 27–3, s and R are positive, but s' is negative. If we had used a concave surface, our formula (27.3) would still give the correct result if we merely make R a negative quantity.

In working out the corresponding formula for a mirror, using the above conventions, you will find that if you put $n = -1$ throughout the formula (27.3) (as though the material behind the mirror had an index -1), the right formula for a mirror results!

Although the derivation of formula (27.3) is simple and elegant, using least time, one can of course work out the same formula using Snell's law, remembering that the angles are so small that the sines of angles can be replaced by the angles themselves.

27–3 The focal length of a lens

Now we go on to consider another situation, a very practical one. Most of the lenses that we use have two surfaces, not just one. How does this affect matters? Suppose that we have two surfaces of different curvature, with glass filling the space between them (Fig. 27–5). We want to study the problem of focusing from a point O to an alternate point O'. How can we do that? The answer is this: First, use formula (27.3) for the first surface, forgetting about the second surface. This will tell us that the light which was diverging from O will appear to be converging or diverging, depending on the sign, from some other point, say O'. Now we consider a new problem. We have a different surface, between glass and air, in which rays are converging toward a certain point O'. Where will they actually converge? We use the same formula again! We find that they converge at O''. Thus, if necessary, we can go through 75 surfaces by just using the same formula in succession, from one to the next!

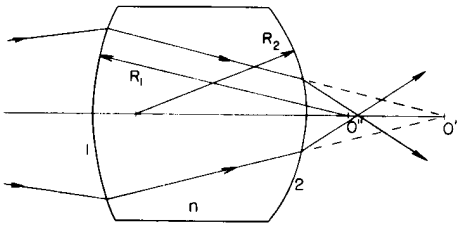


Fig. 27–5. Image formation by a two-surface lens.

There are some rather high-class formulas that would save us considerable energy in the few times in our lives that we might have to chase the light through five surfaces, but it is easier just to chase it through five surfaces when the problem arises than it is to memorize a lot of formulas, because it may be we will never have to chase it through any surfaces at all!

In any case, the principle is that when we go through one surface we find a new position, a new focal point, and then take that point as the starting point for the next surface, and so on. In order to actually do this, since on the second surface we are going from n to 1 rather than from 1 to n, and since in many systems there is more than one kind of glass, so that there are indices $n_1, n_2, \ldots$, we really need a generalization of formula (27.3) for a case where there are two different indices, n_1 and n_2, rather than only n. Then it is not difficult to prove that the general form of (27.3) is

$$(n_1/s) + (n_2/s') = (n_2 - n_1)/R. \tag{27.7}$$

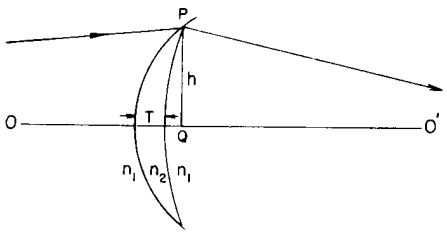


Fig. 27–6. A thin lens with two positive radii.

Particularly simple is the special case in which the two surfaces are very close together—so close that we may ignore small errors due to the thickness. If we draw the lens as shown in Fig. 27–6, we may ask this question: How must the lens be built so as to focus light from O to O'? Suppose the light comes exactly to the edge of the lens, at point P. Then the excess time in going from O to O' is $(n_1h^2/2s) + (n_1h^2/2s')$, ignoring for a moment the presence of the thickness T of glass of index n_2. Now, to make the time for the direct path equal to that for the path OPO', we have to use a piece of glass whose thickness T at the center is such that the delay introduced in going through this thickness is enough to compensate for the excess time above. Therefore the thickness of the lens at the center must be given by the relationship

$$(n_1h^2/2s) + (n_1h^2/2s') = (n_2 - n_1)T. \tag{27.8}$$

We can also express T in terms of the radii R_1 and R_2 of the two surfaces. Paying attention to our convention (3), we thus find, for $R_1 < R_2$ (a convex lens),

$$T = (h^2/2R_1) - (h^2/2R_2). \tag{27.9}$$

Therefore, we finally get

$$(n_1/s) + (n_1/s') = (n_2 - n_1)(1/R_1 - 1/R_2). \tag{27.10}$$

Now we note again that if one of the points is at infinity, the other will be at a point which we will call the focal length f. The focal length f is given by

$$1/f = (n - 1)(1/R_1 - 1/R_2), \tag{27.11}$$

where $n = n_2/n_1$.

Now, if we take the opposite case, where s goes to infinity, we see that s' is at the focal length f'. This time the focal lengths are equal. (This is another special case of the general rule that the ratio of the two focal lengths is the ratio of the indices of refraction in the two media in which the rays focus. In this particular optical system, the initial and final indices are the same, so the two focal lengths are equal.)

Forgetting for a moment about the actual formula for the focal length, if we bought a lens that somebody designed with certain radii of curvature and a certain index, we could measure the focal length, say, by seeing where a point at infinity focuses. Once we had the focal length, it would be better to write our equation in terms of the focal length directly, and the formula then is

$$(1/s) + (1/s') = 1/f. \tag{27.12}$$

Now let us see how the formula works and what it implies in different circumstances. First, it implies that if s or s' is infinite the other one is f. That means that parallel light focuses at a distance f, and this in effect *defines* f. Another interesting thing it says is that both points move in the same direction. If one moves to the right, the other does also. Another thing it says is that s and s' are equal if they are both equal to $2f$. In other words, if we want a symmetrical situation, we find that they will both focus at a distance $2f$.

27–4 Magnification

So far we have discussed the focusing action only for points on the axis. Now let us discuss also the imaging of objects not exactly on the axis, but a little bit off, so that we can understand the properties of *magnification*. When we set up a lens so as to focus light from a small filament onto a "point" on a screen, we notice that on the screen we get a "picture" of the same filament, except of a larger or smaller size than the true filament. This must mean that the light comes to a focus from *each point* of the filament. In order to understand this a little better, let us analyze the thin lens system shown schematically in Fig. 27–7. We know the following facts:

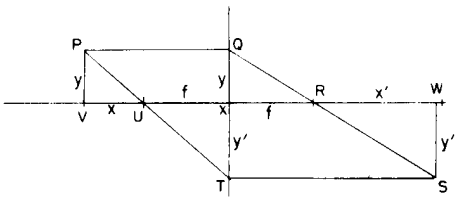


Fig. 27–7. The geometry of imaging by a thin lens.

(1) Any ray that comes in parallel on one side proceeds toward a certain particular point called the focus on the other side, at a distance f from the lens.

(2) Any ray that arrives at the lens from the focus on one side comes out parallel to the axis on the other side.

This is all we need to establish formula (27.12) by geometry, as follows: Suppose we have an object at some distance x from the focus; let the height of the object be y. Then we know that one of the rays, namely PQ, will be bent so as to pass through the focus R on the other side. Now if the lens will focus point P at all, we can find out where if we find out where just one other ray goes, because the new focus will be where the two intersect again. We need only use our ingenuity to find the exact direction of *one* other ray. But we remember that a parallel ray goes through the focus and *vice versa*: a ray which goes through the focus will come out parallel! So we draw ray PT through U. (It is true that the actual rays which are doing the focusing may be much more limited than the two we have drawn, but they are harder to figure, so we make believe that we can make this ray.) Since it would come out parallel, we draw TS parallel to XW. The intersection S is the point we need. This will determine the correct place and the correct height. Let

us call the height y' and the distance from the focus, x'. Now we may derive a lens formula. Using the similar triangles PVU and TXU, we find

$$\frac{y'}{f} = \frac{y}{x}. \tag{27.13}$$

Similarly, from triangles SWR and QXR, we get

$$\frac{y'}{x'} = \frac{y}{f}. \tag{27.14}$$

Solving each for y'/y, we find that

$$\frac{y'}{y} = \frac{x'}{f} = \frac{f}{x}. \tag{27.15}$$

Equation (27.15) is the famous lens formula; in it is everything we need to know about lenses: It tells us the magnification, y'/y, in terms of the distances and the focal lengths. It also connects the two distances x and x' with f:

$$xx' = f^2, \tag{27.16}$$

which is a much neater form to work with than Eq. (27.12). We leave it to the student to demonstrate that if we call $s = x + f$ and $s' = x' + f$, Eq. (27.12) is the same as Eq. (27.16).

27–5 Compound lenses

Without actually deriving it, we shall briefly describe the general result when we have a number of lenses. If we have a system of several lenses, how can we possibly analyze it? That is easy. We start with some object and calculate where its image is for the first lens, using formula (27.16) or (27.12) or any other equivalent formula, or by drawing diagrams. So we find an image. Then we treat this image as the source for the next lens, and use the second lens with whatever its focal length is to again find an image. We simply chase the thing through the succession of lenses. That is all there is to it. It involves nothing new in principle, so we shall not go into it. However, there is a very interesting net result of the effects of any sequence of lenses on light that starts and ends up in the same medium, say air. Any optical instrument—a telescope or a microscope with any number of lenses and mirrors—has the following property: There exist two planes, called the *principal planes* of the system (these planes are often fairly close to the first surface of the first lens and the last surface of the last lens), which have the following properties: (1) If light comes into the system parallel from the first side, it comes out at a certain focus, at a distance from the *second* principal plane equal to the focal length, just as though the system were a thin lens situated at this plane. (2) If parallel light comes in the other way, it comes to a focus at the same distance f from the *first* principal plane, again as if a thin lens where situated there. (See Fig. 27–8.)

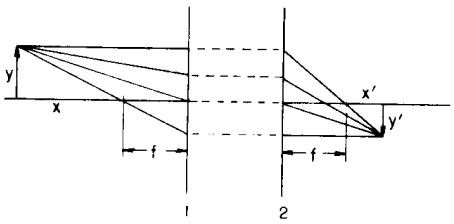


Fig. 27–8. Illustration of the principal planes of an optical system.

Of course, if we measure the distances x and x', and y and y' as before, the formula (27.16) that we have written for the thin lens is absolutely general, provided that we measure the focal length from the principal planes and not from the center of the lens. It so happens that for a thin lens the principal planes are coincident. It is just as though we could take a thin lens, slice it down the middle, and separate it, and not notice that it was separated. Every ray that comes in pops out immediately on the other side of the second plane from the same point as it went into the first plane! The principal planes and the focal length may be found either by experiment or by calculation, and then the whole set of properties of the optical system are described. It is very interesting that the result is not complicated when we are all finished with such a big, complicated optical system.

27–6 Aberrations

Before we get too excited about how marvelous lenses are, we must hasten to add that there are also serious limitations, because of the fact that we have limited ourselves, strictly speaking, to paraxial rays, the rays near the axis. A real lens having a finite size will, in general, exhibit *aberrations*. For example, a ray that is on the axis, of course, goes through the focus; a ray that is very close to the axis will still come to the focus very well. But as we go farther out, the ray begins to deviate from the focus, perhaps by falling short, and a ray striking near the top edge comes down and misses the focus by quite a wide margin. So, instead of getting a point image, we get a smear. This effect is called *spherical aberration*, because it is a property of the spherical surfaces we use in place of the right shape. This could be remedied, for any specific object distance, by re-forming the shape of the lens surface, or perhaps by using several lenses arranged so that the aberrations of the individual lenses tend to cancel each other.

Lenses have another fault: light of different colors has different speeds, or different indices of refraction, in the glass, and therefore the focal length of a given lens is different for different colors. So if we image a white spot, the image will have colors, because when we focus for the red, the blue is out of focus, or vice versa. This property is called *chromatic aberration*.

There are still other faults. If the object is off the axis, then the focus really isn't perfect any more, when it gets far enough off the axis. The easiest way to verify this is to focus a lens and then tilt it so that the rays are coming in at a large angle from the axis. Then the image that is formed will usually be quite crude, and there may be no place where it focuses well. There are thus several kinds of errors in lenses that the optical designer tries to remedy by using many lenses to compensate each other's errors.

How careful do we have to be to eliminate aberrations? Is it possible to make an absolutely perfect optical system? Suppose we had built an optical system that is supposed to bring light exactly to a point. Now, arguing from the point of view of least time, can we find a condition on how perfect the system has to be? The system will have some kind of an entrance opening for the light. If we take the farthest ray from the axis that can come to the focus (if the system is perfect, of course), the times for all rays are exactly equal. But nothing is perfect, so the question is, how wrong can the time be for this ray and not be worth correcting any further? That depends on how perfect we want to make the image. But suppose we want to make the image as perfect as it possibly can be made. Then, of course, our impression is that we have to arrange that every ray takes as nearly the same time as possible. But it turns out that this is not true, that beyond a certain point we are trying to do something that is too fine, because the theory of geometrical optics does not work!

Remember that the principle of least time is not an accurate formulation, unlike the principle of conservation of energy or the principle of conservation of momentum. The principle of least time is only an *approximation*, and it is interesting to know how much error can be allowed and still not make any apparent difference. The answer is that if we have arranged that between the maximal ray—the worst ray, the ray that is farthest out—and the central ray, the difference in time is less than about the period that corresponds to one oscillation of the light, then there is no use improving it any further. Light is an oscillatory thing with a definite frequency that is related to the wavelength, and if we have arranged that the time difference for different rays is less than about a period, there is no use going any further.

27–7 Resolving power

Another interesting question—a very important technical question with all optical instruments—is how much *resolving power* they have. If we build a microscope, we want to see the objects that we are looking at. That means, for instance, that if we are looking at a bacterium with a spot on each end, we want to *see* that

there are two dots when we magnify them. One might think that all we have to do is to get enough magnification—we can always add another lens, and we can always magnify again and again, and with the cleverness of designers, all the spherical aberrations and chromatic aberrations can be cancelled out, and there is no reason why we cannot keep on magnifying the image. So the limitations of a microscope are not that it is impossible to build a lens that magnifies more than 2000 diameters. We can build a system of lenses that magnifies 10,000 diameters, but we *still* could not see two points that are too close together because of the limitations of geometrical optics, because of the fact that least time is not precise.

To discover the rule that determines how far apart two points have to be so that at the image they appear as separate points can be stated in a very beautiful way associated with the time it takes for different rays. Suppose that we disregard the aberrations now, and imagine that for a particular point P (Fig. 27–9) all the rays from object to image T take exactly the same time. (It is not true, because it is not a perfect system, but that is another problem.) Now take another nearby point, P', and ask whether its image will be distinct from T. In other words, whether we can make out the difference between them. Of course, according to geometrical optics, there should be two point images, but what we see may be rather smeared and we may not be able to make out that there are two points. The condition that the second point is focused in a distinctly different place from the first one is that the two times for the extreme rays $P'ST$ and $P'RT$ on each side of the big opening of the lenses to go from one end to the other, must *not* be equal from the two possible object points to a given image point. Why? Because, if the times were equal, of course both would *focus* at the same point. So the times are not going to be equal. But by how much do they have to differ so that we can say that both do *not* come to a common focus, so that we can distinguish the two image points? The general rule for the resolution of any optical instrument is this: two different point sources can be resolved only if one source is focused at such a point that the times for the maximal rays from the other source to reach that point, as compared with its own true image point, differ by more than one period. It is necessary that the difference in time between the top ray and the bottom ray to the *wrong* focus shall exceed a certain amount, namely, approximately the period of oscillation of the light:



Fig. 27–9. The resolving power of an optical system.

$$t_2 - t_1 > 1/\nu, \tag{27.17}$$

where ν is the frequency of the light (number of oscillations per second; also speed divided by wavelength). If the distance of separation of the two points is called D, and if the opening angle of the lens is called θ, then one can demonstrate that (27.17) is exactly equivalent to the statement that D must exceed $\lambda/n \sin\theta$, where n is the index of refraction at P and λ is the wavelength. The smallest things that we can see are therefore approximately the wavelength of light. A corresponding formula exists for telescopes, which tells us the smallest difference in angle between two stars that can just be distinguished.*

* The angle is about λ/D, where D is the lens diameter. Can you see why?

28

Electromagnetic Radiation

28–1 Electromagnetism
28–2 Radiation
28–3 The dipole radiator
28–4 Interference

28–1 Electromagnetism

The most dramatic moments in the development of physics are those in which great syntheses take place, where phenomena which previously had appeared to be different are suddenly discovered to be but different aspects of the same thing. The history of physics is the history of such syntheses, and the basis of the success of physical science is mainly that we are *able* to synthesize.

Perhaps the most dramatic moment in the development of physics during the 19th century occurred to J. C. Maxwell one day in the 1860's, when he combined the laws of electricity and magnetism with the laws of the behavior of light. As a result, the properties of light were partly unravelled—that old and subtle stuff that is so important and mysterious that it was felt necessary to arrange a special creation for it when writing Genesis. Maxwell could say, when he was finished with his discovery, "Let there be electricity and magnetism, and there is light!"

For this culminating moment there was a long preparation in the gradual discovery and unfolding of the laws of electricity and magnetism. This story we shall reserve for detailed study next year. However, the story is, briefly, as follows. The gradually discovered properties of electricity and magnetism, of electric forces of attraction and repulsion, and of magnetic forces, showed that although these forces were rather complex, they all fell off inversely as the square of the distance. We know, for example, that the simple coulomb law for stationary charges is that the electric force field varies inversely as the square of the distance. As a consequence, for sufficiently great distances there is very little influence of one system of charges on another. Maxwell noted that the equations or the laws that had been discovered up to this time were mutually inconsistent when he tried to put them all together, and in order for the whole system to be consistent, he had to add another term to his equations. With this new term there came an amazing prediction, which was that a part of the electric and magnetic fields would fall off much more slowly with the distance than the inverse square, namely, inversely as the first power of the distance! And so he realized that electric currents in one place can affect other charges far away, and he predicted the basic effects with which we are familiar today—radio transmission, radar, and so on.

It seems a miracle that someone talking in Europe can, with mere electrical influences, be heard thousands of miles away in Los Angeles. How is it possible? It is because the fields do not vary as the inverse square, but only inversely as the first power of the distance. Finally, then, even light itself was recognized to be electric and magnetic influences extending over vast distances, generated by an almost incredibly rapid oscillation of the electrons in the atoms. All these phenomena we summarize by the word *radiation* or, more specifically, *electromagnetic radiation*, there being one or two other kinds of radiation also. Almost always, radiation means electromagnetic radiation.

And thus is the universe knit together. The atomic motions of a distant star still have sufficient influence at this great distance to set the electrons in our eye in motion, and so we know about the stars. If this law did not exist, we would all be literally in the dark about the exterior world! And the electric surgings in a galaxy five billion light years away—which is the farthest object we have found so far—can still influence in a significant and detectable way the currents in the great "dish" in front of a radio telescope. And so it is that we see the stars and the galaxies.

This remarkable phenomenon is what we shall discuss in the present chapter. At the beginning of this course in physics we outlined a broad picture of the world, but we are now better prepared to understand some aspects of it, and so we shall now go over some parts of it again in greater detail. We begin by describing the position of physics at the end of the 19th century. All that was then known about the fundamental laws can be summarized as follows.

First, there were laws of forces: one force was the law of gravitation, which we have written down several times; the force on an object of mass m, due to another of mass M, is given by

$$\mathbf{F} = GmM\mathbf{e}_r/r^2, \tag{28.1}$$

where $\mathbf{e}_r$ is a unit vector directed from m to M, and r is the distance between them.

Next, the laws of electricity and magnetism, as known at the end of the 19th century, are these: the electrical forces acting on a charge q can be described by two fields, called $\mathbf{E}$ and $\mathbf{B}$, and the velocity $\mathbf{v}$ of the charge q, by the equation

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}). \tag{28.2}$$

To complete this law, we have to say what the formulas for $\mathbf{E}$ and $\mathbf{B}$ are in a given circumstance: if a number of charges are present, $\mathbf{E}$ and the $\mathbf{B}$ are each the sum of contributions, one from each individual charge. So if we can find the $\mathbf{E}$ and $\mathbf{B}$ produced by a single charge, we need only to add all the effects from all the charges in the universe to get the total $\mathbf{E}$ and $\mathbf{B}$! This is the principle of superposition.

What is the formula for the electric and magnetic field produced by one individual charge? It turns out that this is very complicated, and it takes a great deal of study and sophistication to appreciate it. But that is not the point. We write down the law now only to impress the reader with the beauty of nature, so to speak, i.e., that it is *possible* to summarize all the fundamental knowledge on one page, with notations that he is now familiar with. This law for the fields of an individual charge is complete and accurate, so far as we know (except for quantum mechanics) but it looks rather complicated. We shall not study all the pieces now; we only write it down to give an impression, to show that it can be written, and so that we can see ahead of time roughly what it looks like. As a matter of fact, the most *useful* way to write the correct laws of electricity and magnetism is not the way we shall now write them, but involves what are called *field equations*, which we shall learn about next year. But the mathematical notations for these are different and new, and so we write the law in an inconvenient form for calculation, but in notations that we now know.

The electric field, $\mathbf{E}$, is given by

$$\mathbf{E} = \frac{-q}{4\pi\epsilon_0}\left[\frac{\mathbf{e}_{r'}}{r'^2} + \frac{r'}{c}\frac{d}{dt}\left(\frac{\mathbf{e}_{r'}}{r'^2}\right) + \frac{1}{c^2}\frac{d^2}{dt^2}\,\mathbf{e}_{r'}\right]. \tag{28.3}$$

What do the various terms tell us? Take the first term, $\mathbf{E} = -q\mathbf{e}_{r'}/4\pi\epsilon_0 r'^2$. That, of course, is Coulomb's law, which we already know: q is the charge that is producing the field; $\mathbf{e}_{r'}$ is the unit vector in the direction from the point P where $\mathbf{E}$ is measured, r is the distance from P to q. But, Coulomb's law is wrong. The discoveries of the 19th century showed that influences cannot travel faster than a certain fundamental speed c, which we now call the speed of light. It is not correct that the first term is Coulomb's law, not only because it is not possible to know where the charge is *now* and at what distance it is *now*, but also because the only thing that can affect the field at a given place and time is the behavior of the charges in the *past*. How *far* in the past? The time delay, or *retarded time*, so-called, is the time it takes, at speed c, to get from the charge to the field point P. The delay is r'/c.

So to allow for this time delay, we put a little prime on r, meaning how far away it *was* when the information now arriving at P left q. Just for a moment suppose that the charge carried a light, and that the light could only come to P at the speed c. Then when we look at q, we would not see where it is now, of course, but where it *was* at some earlier time. What appears in our formula is the *apparent*

direction $\theta_{r'}$ —the direction it used to be—the so-called *retarded* direction—and at the *retarded* distance r'. That would be easy enough to understand, too, but it is also wrong. The whole thing is much more complicated.

There are several more terms. The next term is as though nature were trying to allow for the fact that the effect is retarded, if we might put it very crudely. It suggests that we should calculate the delayed coulomb field and add a correction to it, which is its rate of change times the time delay that we use. Nature seems to be attempting to guess what the field at the present time is going to be, by taking the rate of change and multiplying by the time that is delayed. But we are not yet through. There is a third term—the second derivative, with respect to t, of the unit vector in the direction of the charge. Now the formula *is* finished, and that is all there is to the electric field from an arbitrarily moving charge.

The magnetic field is given by

$$\mathbf{B} = -\mathbf{e}_{r'} \times \mathbf{E}/c. \tag{28.4}$$

We have written these down only for the purpose of showing the beauty of nature or, in a way, the power of mathematics. We do not pretend to understand *why* it is possible to write so much in such a small space, but (28.3) and (28.4) contain the machinery by which electric generators work, how light operates, all the phenomena of electricity and magnetism. Of course, to complete the story we also need to know something about the behavior of the materials involved—the properties of matter—which are not described properly by (28.3).

To finish with our description of the world of the 19th century we must mention one other great synthesis which occurred in that century, one with which Maxwell had a great deal to do also, and that was the synthesis of the phenomena of heat and mechanics. We shall study that subject soon.

What had to be added in the 20th century was that the dynamical laws of Newton were found to be all wrong, and quantum mechanics had to be introduced to correct them. Newton's laws are approximately valid when the scale of things is sufficiently large. These quantum-mechanical laws, combined with the laws of electricity, have only recently been combined to form a set of laws called *quantum electrodynamics*. In addition, there were discovered a number of new phenomena, of which the first was radioactivity, discovered by Becquerel in 1898—he just sneaked it in under the 19th century. This phenomenon of radioactivity was followed up to produce our knowledge of nuclei and new kinds of forces that are not gravitational and not electrical, but new particles with different interactions, a subject which has still not been unravelled.

For those purists who know more (the professors who happen to be reading this), we should add that when we say that (28.3) is a complete expression of the knowledge of electrodynamics, we are not being entirely accurate. There was a problem that was not quite solved at the end of the 19th century. When we try to calculate the field from all the charges *including the charge itself that we want the field to act on*, we get into trouble trying to find the distance, for example, of a charge from itself, and dividing something by that distance, which is zero. The problem of how to handle the part of this field which is generated by the very charge on which we want the field to act is not yet solved today. So we leave it there; we do not have a complete solution to that puzzle yet, and so we shall avoid the puzzle for as long as we can.

28–2 Radiation

That, then, is a summary of the world picture. Now let us use it to discuss the phenomena called radiation. To discuss these phenomena, we must select from Eq. (28.3) only that piece which varies inversely as the distance and not as the square of the distance. It turns out that when we finally do find that piece, it is so simple in its form that it is legitimate to study optics and electrodynamics in an elementary way by taking it as "the law" of the electric field produced by a moving charge far away. We shall take it temporarily as a given law which we will learn about in detail next year.

Of the terms appearing in (28.3), the first one evidently goes inversely as the square of the distance, and the second is only a correction for delay, so it is easy to show that both of them vary inversely as the square of the distance. All of the effects we are interested in come from the third term, which is not very complicated, after all. What this term says is: look at the charge and note the direction of the unit vector (we can project the end of it onto the surface of a unit sphere). As the charge moves around, the unit vector wiggles, and *the acceleration of that unit vector is what we are looking for.* That is all. Thus

$$\mathbf{E} = \frac{-q}{4\pi\epsilon_0 c^2}\frac{d^2\mathbf{e}_{r'}}{dt^2} \tag{28.5}$$

is a statement of the laws of radiation, because that is the only important term when we get far enough away that the fields are varying inversely as the distance. (The parts that go as the square have fallen off so much that we are not interested in them.)

Now we can go a little bit further in studying (28.5) to see what it means. Suppose a charge is moving in any manner whatsoever, and we are observing it from a distance. We imagine for a moment that in a sense it is "lit up" (although it is light that we are trying to explain); we imagine it as a little white dot. Then we would see this white dot running around. But we don't see *exactly* how it is running around right *now*, because of the delay that we have been talking about. What counts is how it was moving *earlier*. The unit vector $\mathbf{e}_{r'}$ is pointed toward the apparent position of the charge. Of course, the end of $\mathbf{e}_{r'}$ goes on a slight curve, so that its acceleration has two components. One is the transverse piece, because the end of it goes up and down, and the other is a radial piece because it stays on a sphere. It is easy to demonstrate that the latter is much smaller and varies as the inverse square of r when r is very great. This is easy to see, for when we imagine that we move a given source farther and farther away, then the wigglings of $\mathbf{e}_{r'}$ look smaller and smaller, inversely as the distance, but the radial component of acceleration is varying much more rapidly than inversely as the distance. So for practical purposes all we have to do is project the motion on a plane at unit distance. Therefore we find the following rule: Imagine that we look at the moving charge and that everything we see is delayed—like a painter trying to paint a scene on a screen at a unit distance. A real painter, of course, does not take into account the fact that light is going at a certain speed, but paints the world as he sees it. We want to see what his picture would look like. So we see a dot, representing the charge, moving about in the picture. The acceleration of that dot is proportional to the electric field. That is all—all we need.

Thus Eq. (28.5) is the complete and correct formula for radiation; even relativity effects are all contained in it. However, we often want to apply it to a still simpler circumstance in which the charges are moving only a small distance at a relatively slow rate. Since they are moving slowly, they do not move an appreciable distance from where they start, so that the delay time is practically constant. Then the law is still simpler, because the delay time is fixed. Thus we imagine that the charge is executing a very tiny motion at an effectively constant distance. The delay at the distance r is r/c. Then our rule becomes the following: If the charged object is moving in a very small motion and it is laterally displaced by the distance $x(t)$, then the angle that the unit vector $\mathbf{e}_{r'}$ is displaced is x/r, and since r is practically constant, the x-component of $d^2\mathbf{e}_{r'}/dt^2$ is simply the acceleration of x itself at an earlier time, and so finally we get the law we want, which is

$$E_x(t) = \frac{-q}{4\pi\epsilon_0 c^2 r}\, a_x\left(t - \frac{r}{c}\right). \tag{28.6}$$

Only the component of a_x perpendicular to the line of sight is important. Let us see why that is. Evidently, if the charge is moving in and out straight at us, the unit vector in that direction does not wiggle at all, and it has no acceleration. So it is only the sidewise motion which is important, only the acceleration that we see projected on the screen.

28–3 The dipole radiator

As our fundamental "law" of electromagnetic radiation, we are going to assume that (28.6) is true, i.e., that the electric field produced by an accelerating charge which is moving nonrelativistically at a very large distance r approaches that form. The electric field varies inversely as r and is proportional to the acceleration of the charge, projected onto the "plane of sight," and this acceleration is not today's acceleration, but the acceleration that it had at an earlier time, the amount of delay being a time, r/c. In the remainder of this chapter we shall discuss this law so that we can understand it better physically, because we are going to use it to understand all of the phenomena of light and radio propagation, such as reflection, refraction, interference, diffraction, and scattering. It is the central law, and is all we need. All the rest of Eq. (28.3) was written down only to set the stage, so that we could appreciate where (28.6) fits and how it comes about.

We shall discuss (28.3) further next year. In the meantime, we shall accept it as true, but not just on a theoretical basis. We may devise a number of experiments which illustrate the character of the law. In order to do so, we need an accelerating charge. It should be a single charge, but if we can make a great many charges move together, all the same way, we know that the field will be the sum of the effects of each of the individual charges; we just add them together. As an example, consider two pieces of wire connected to a generator, as shown in Fig. 28–1. The idea is that the generator makes a potential difference, or a field, which pulls electrons away from piece A and pushes them into B at one moment, and then, an infinitesimal time later, it reverses the effect and pulls the electrons out of B and pumps them back into A! So in these two wires charges, let us say, are accelerating upward in wire A and upward in wire B for one moment, and a moment later they are accelerating downward in wire A and downward in wire B. The fact that we need two wires and a generator is merely that this is a way of doing it. The net result is that we merely have a charge accelerating up and down as though A and B were one single wire. A wire that is very short compared with the distance light travels in one oscillation period is called an *electric dipole oscillator*. Thus we have the circumstance that we need to apply our law, which tells us that this charge makes an electric field, and so we need an instrument to detect an electric field, and the instrument we use is the same thing—a pair of wires like A and B! If an electric field is applied to such a device, it will produce a force which will pull the electrons up on both wires or down on both wires. This signal is detected by means of a rectifier mounted between A and B, and a tiny, fine wire carries the information into an amplifier, where it is amplified so we can hear the audio-frequency tone with which the radiofrequency is modulated. When this probe feels an electric field, there will be a loud noise coming out of the loudspeaker, and when there is no electric field driving it, there will be no noise.

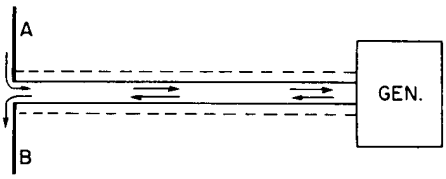


Fig. 28–1. A high-frequency signal generator drives charges up and down on two wires.

Because the room in which the waves we are measuring has other objects in it, our electric field will shake electrons in these other objects; the electric field makes these other charges go up and down, and in going up and down, these also produce an effect on our probe. Thus for a successful experiment we must hold things fairly close together, so that the influences from the walls and from ourselves—the reflected waves—are relatively small. So the phenomena will not turn out to appear to be precisely and perfectly in accord with Eq. (28.6), but will be close enough that we shall be able to appreciate the law.

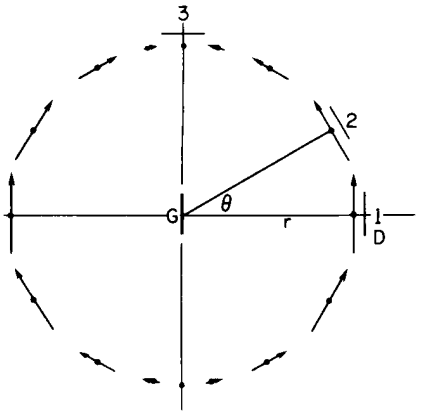


Fig. 28–2. The instantaneous electric field on a sphere centered at a localized, linearly oscillating charge.

Now we turn the generator on and hear the audio signal. We find a strong field when the detector D is parallel to the generator G at point 1 (Fig. 28–2). We find the same amount of field also at any other azimuth angle about the axis of G, because it has no directional effects. On the other hand, when the detector is at 3 the field is zero. That is all right, because our formula said that the field should be the acceleration of the charge *projected perpendicular* to the line of sight. Therefore when we look down on G, the charge is moving toward and away from D, and there is no effect. So that checks the first rule, that there is no effect when the charge is moving directly toward us. Secondly, the formula says that the electric field should be perpendicular to r and in the plane of G and r; so if we put D at 1 but rotate it 90°, we should get no signal. And this is just what we find, the

electric field is indeed vertical, and not horizontal. When we move D to some intermediate angle, we see that the strongest signal occurs when it is oriented as shown, because although G is vertical, it does not produce a field that is simply parallel to itself—it is the *projection of the acceleration perpendicular to the line of sight* that counts. The signal is weaker at 2 than it is at 1, because of the projection effect.

28–4 Interference

Next, we may test what happens when we have two sources side by side a few centimeters apart (Fig. 28–3). The law is that the two sources should add their effects at point 1 when both of the sources are connected to the same generator and are both moving up and down the same way, so that the total electric field is the sum of the two and is twice as strong as it was before.

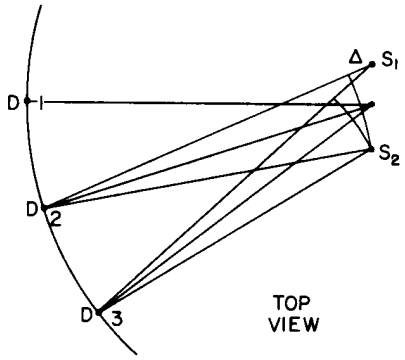


Fig. 28–3. Illustration of interference of sources.

Now comes an interesting possibility. Suppose we make the charges in S_1 and S_2 both accelerate up and down, but delay the timing of S_2 so that they are 180° out of phase. Then the field produced by S_1 will be in one direction and the field produced by S_2 will be in the opposite direction at any instant, and therefore we should get *no* effect at point 1. The phase of oscillation is neatly adjustable by means of a pipe which is carrying the signal to S_2. By changing the length of this pipe we change the time it takes the signal to arrive at S_2 and thus we change the phase of that oscillation. By adjusting this length, we can indeed find a place where there is no more signal left, in spite of the fact that both S_1 and S_2 are moving! The fact that they are both moving can be checked, because if we cut one out, we can see the motion of the other. So the two of them together can produce *zero* if everything is adjusted correctly.

Now, it is very interesting to show that the addition of the two fields is in fact a *vector* addition. We have just checked it for up and down motion, but let us check two nonparallel directions. First, we restore S_1 and S_2 to the same phase; that is, they are again moving together. But now we turn S_1 through 90°, as shown in Fig. 28–4. Now we should have at point 1 the sum of two effects, one of which is vertical and the other horizontal. The electric field is the vector sum of these two in-phase signals—they are both strong at the same time and go through zero together; the total field should be a signal R at 45°. If we turn D to get the maximum noise, it should be at about 45°, and not vertical. And if we turn it at right angles to that direction, we should get zero, which is easy to measure. Indeed, we observe just such behavior!

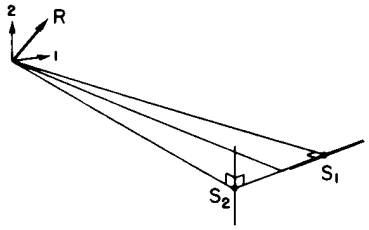


Fig. 28–4. Illustration of the vector character of the combination of sources.

Now, how about the retardation? How can we demonstrate that the signal is retarded? We could, with a great deal of equipment, measure the time at which it arrives, but there is another, very simple way. Referring again to Fig. 28–3, suppose that S_1 and S_2 are in phase. They are both shaking together, and they produce equal electric fields at point 1. But suppose we go to a certain place 2 which is closer to S_2 and farther from S_1. Then, in accordance with the principle that the acceleration should be retarded by an amount equal to r/c, if the retardations are not equal, the signals are no longer in phase. Thus it should be possible to find a position at which the distances of D from S_1 and S_2 differ by some amount Δ, in such a manner that there is no net signal. That is, the distance Δ is to be the distance light goes in one-half an oscillation of the generator. We may go still further, and find a point where the difference is greater by a whole cycle; that is to say, the signal from the first antenna reaches point 3 with a delay in time that is greater than that of the second antenna by just the length of time it takes for the electric current to oscillate once, and therefore the two electric fields produced at 3 are in phase again. At point 3 the signal is strong again.

This completes our discussion of the experimental verification of some of the important features of Eq. (28.6). Of course we have not really checked the $1/r$ variation of the electric field strength, or the fact that there is also a magnetic field that goes along with the electric field. To do so would require rather sophisticated techniques and would hardly add to our understanding at this point. In any case, we have checked those features that are of the greatest importance for our later applications, and we shall come back to study some of the other properties of electromagnetic waves next year.

29

Interference

29–1 Electromagnetic waves

29–2 Energy of radiation

29–3 Sinusoidal waves

29–4 Two dipole radiators

29–5 The mathematics of interference

29–1 Electromagnetic waves

In this chapter we shall discuss the subject of the preceding chapter more mathematically. We have qualitatively demonstrated that there are maxima and minima in the radiation field from two sources, and our problem now is to describe the field in mathematical detail, not just qualitatively.

We have already physically analyzed the meaning of formula (28.6) quite satisfactorily, but there are a few points to be made about it mathematically. In the first place, if a charge is accelerating up and down along a line, in a motion of very small amplitude, the field at some angle θ from the axis of the motion is in a direction at right angles to the line of sight and in the plane containing both the acceleration and the line of sight (Fig. 29–1). If the distance is called r, then at time t the electric field has the magnitude

$$E(t) = \frac{-qa(t - r/c)\sin\theta}{4\pi\epsilon_0 c^2 r}, \tag{29.1}$$

where $a(t - r/c)$ is the acceleration at the time $(t - r/c)$, called the *retarded* acceleration.

Now it would be interesting to draw a picture of the field under different conditions. The thing that is interesting, of course, is the factor $a(t - r/c)$, and to understand it we can take the simplest case, $\theta = 90°$, and plot the field graphically. What we had been thinking of before is that we stand in one position and ask how the field there changes with time. But instead of that, we are now going to see what the field looks like at different positions in space at a given instant. So what we want is a "snapshot" picture which tells us what the field is in different places. Of course it depends upon the acceleration of the charge. Suppose that the charge at first had some particular motion: it was initially standing still, and it suddenly accelerated in some manner, as shown in Fig. 29–2, and then stopped. Then, a little bit later, we measure the field at a different place. Then we may assert that the field will appear as shown in Fig. 29–3. At each point the field is determined by the acceleration of the charge at an earlier time, the amount earlier being the delay r/c. The field at farther and farther points is determined by the acceleration at earlier and earlier times. So the curve in Fig. 29–3 is really, in a sense, a "reversed" plot of the acceleration as a function of time; the distance is related to time by a constant scale factor c, which we often take as unity. This is easily seen by considering the mathematical behavior of $a(t - r/c)$. Evidently, if we add a little time Δt, we get the same value for $a(t - r/c)$ as we would have if we had subtracted a little distance: $\Delta r = -c\,\Delta t$.

Stated another way: if we add a little time Δt, we can restore $a(t - r/c)$ to its former value by *adding* a little distance $\Delta r = c\,\Delta t$. That is, as time goes on *the field moves as a wave outward from the source.* That is the reason why we sometimes say light is propagated as waves. It is equivalent to saying that the field is delayed, or to saying that the electric field is moving outward as time goes on.

An interesting special case is that where the charge q is moving up and down in an oscillatory manner. The case which we studied experimentally in the last chapter was one in which the displacement x at any time t was equal to a certain constant x_0, the magnitude of the oscillation, times $\cos\omega t$. Then the acceleration is

$$a = -\omega^2 x_0 \cos\omega t = a_0 \cos\omega t, \tag{29.2}$$

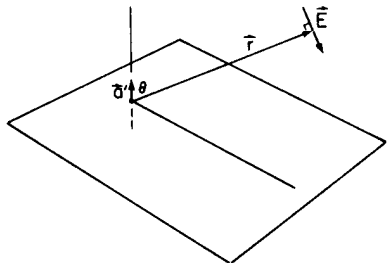


Fig. 29–1. The electric field **E** due to a positive charge whose retarded acceleration is a′.

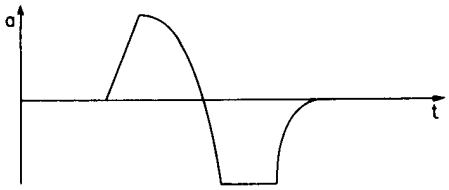


Fig. 29–2. The acceleration of a certain charge as a function of time.

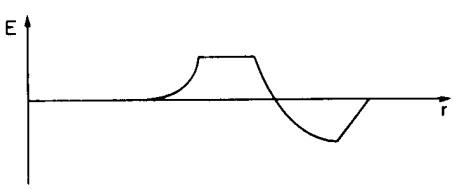


Fig. 29–3. The electric field as a function of position at a later time. (The $1/r$ variation is ignored.)

where a_0 is the maximum acceleration, $-\omega^2 x_0$. Putting this formula into (29.1), we find

$$E = -q \sin\theta \frac{a_0 \cos\omega(t - r/c)}{4\pi\epsilon_0 r c^2}. \tag{29.3}$$

Now, ignoring the angle θ and the constant factors, let us see what that looks like as a function of position or as a function of time.

29–2 Energy of radiation

First of all, at any particular moment or in any particular place, the strength of the field varies inversely as the distance r, as we mentioned previously. Now we must point out that the *energy* content of a wave, or the energy effects that such an electric field can have, are proportional to the *square* of the field, because if, for instance, we have some kind of a charge or an oscillator in the electric field, then if we let the field act on the oscillator, it makes it move. If this is a linear oscillator, the acceleration, velocity, and displacement produced by the electric field acting on the charge are all proportional to the field. So the kinetic energy which is developed in the charge is proportional to the *square* of the field. So we shall take it that the energy that a field can deliver to a system is proportional somehow to the square of the field.

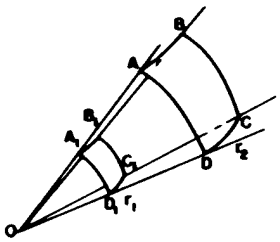


Fig. 29–4. The energy flowing within the cone *OABCD* is independent of the distance r at which it is measured.

This means that the energy that the source can deliver decreases as we get farther away; in fact, it varies *inversely as the square of the distance*. But that has a very simple interpretation: if we wanted to pick up all the energy we could from the wave in a certain cone at a distance r_1 (Fig. 29–4), and we do the same at another distance r_2, we find that the amount of energy per unit area at any one place goes inversely as the square of r, but the area of the surface intercepted by the cone goes *directly* as the square of r. So the energy that we can take out of the wave within a given conical angle is the same, no matter how far away we are! In particular, the total energy that we could take out of the whole wave by putting absorbing oscillators all around is a certain fixed amount. So the fact that the amplitude of E varies as $1/r$ is the same as saying that there is an energy flux which is never lost, an energy which goes on and on, spreading over a greater and greater effective area. Thus we see that after a charge has oscillated, it has lost some energy which it can never recover; the energy keeps going farther and farther away without diminution. So if we are far enough away that our basic approximation is good enough, the charge cannot recover the energy which has been, as we say, radiated away. Of course the energy still exists somewhere, and is available to be picked up by other systems. We shall study this energy "loss" further in Chapter 32.

Let us now consider more carefully how the wave (29.3) varies as a function of time at a given place, and as a function of position at a given time. Again we ignore the $1/r$ variation and the constants.

29–3 Sinusoidal waves

First let us fix the position r, and watch the field as a function of time. It is oscillatory at the angular frequency ω. The angular frequency ω can be defined as the *rate of change of phase with time* (radians per second). We have already studied such a thing, so it should be quite familiar to us by now. The *period* is the time needed for one oscillation, one complete cycle, and we have worked that out too; it is $2\pi/\omega$, because ω times the period is one cycle of the cosine.

Now we introduce a new quantity which is used a great deal in physics. This has to do with the opposite situation, in which we fix t and look at the wave as a function of distance r. Of course we notice that, as a function of r, the wave (29.3) is also oscillatory. That is, aside from $1/r$, which we are ignoring, we see that E oscillates as we change the position. So, in analogy with ω, we can define a quantity called the *wave number*, symbolized as k. This is defined as *the rate of change of phase with distance* (radians per meter). That is, as we move in space at a fixed time, the phase changes.

There is another quantity that corresponds to the period, and we might call it the period in space, but it is usually called the wavelength, symbolized λ. The wavelength is the distance occupied by one complete cycle. It is easy to see, then, that the wavelength is $2\pi/k$, because k times the wavelength would be the number of radians that the whole thing changes, being the product of the rate of change of the radians per meter, times the number of meters, and we must make a 2π change for one cycle. So $k\lambda = 2\pi$ is exactly analogous to $\omega t_0 = 2\pi$.

Now in our particular wave there is a definite relationship between the frequency and the wavelength, but the above definitions of k and ω are actually quite general. That is, the wavelength and the frequency may not be related in the same way in other physical circumstances. However, in our circumstance the rate of change of phase with distance is easily determined, because if we call $\phi = \omega(t - r/c)$ the phase, and differentiate (partially) with respect to distance r, the rate of change, $\partial\phi/\partial r$, is

$$\left|\frac{\partial\phi}{\partial r}\right| = k = \frac{\omega}{c}. \tag{29.4}$$

There are many ways to represent the same thing, such as

$$\lambda = ct_0 \qquad (29.5) \qquad\qquad \lambda\nu = c \qquad (29.7)$$

$$\omega = ck \qquad (29.6) \qquad\qquad \omega\lambda = 2\pi c \qquad (29.8)$$

Why is the wavelength equal to c times the period? That's very easy, of course, because if we sit still and wait for one period to elapse, the waves, travelling at the speed c, will move a distance ct_0, and will of course have moved over just one wavelength.

In a physical situation other than that of light, k is not necessarily related to ω in this simple way. If we call the distance along an axis x, then the formula for a cosine wave moving in a direction x with a wave number k and an angular frequency ω will be written in general as $\cos(\omega t - kx)$.

Now that we have introduced the idea of wavelength, we may say something more about the circumstances in which (29.1) is a legitimate formula. We recall that the field is made up of several pieces, one of which varies inversely as r, another part which varies inversely as r^2, and others which vary even faster. It would be worth while to know in what circumstances the $1/r$ part of the field is the most important part, and the other parts are relatively small. Naturally, the answer is "if we go 'far enough' away," because terms which vary inversely as the square ultimately become negligible compared with the $1/r$ term. How far is "far enough"? The answer is, qualitatively, that the other terms are of order λ/r smaller than the $1/r$ term. Thus, so long as we are beyond a few wavelengths, (29.1) is an excellent approximation to the field. Sometimes the region beyond a few wavelengths is called the "wave zone."

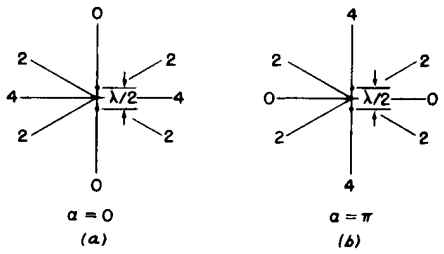


Fig. 29–5. The intensities in various directions from two dipole oscillators one-half wavelength apart. Left: in phase ($\alpha = 0$). Right: one-half period out of phase ($\alpha = \pi$).

29–4 Two dipole radiators

Next let us discuss the mathematics involved in combining the effects of two oscillators to find the net field at a given point. This is very easy in the few cases that we considered in the previous chapter. We shall first describe the effects qualitatively, and then more quantitatively. Let us take the simple case, where the oscillators are situated with their centers in the same horizontal plane as the detector, and the line of vibration is vertical.

Figure 29–5(a) represents the top view of two such oscillators, and in this particular example they are half a wavelength apart in a N–S direction, and are oscillating together in the same phase, which we call zero phase. Now we would like to know the intensity of the radiation in various directions. By the intensity we mean the amount of energy that the field carries past us per second, which is proportional to the square of the field, averaged in time. So the thing to look at, when we want to know how bright the light is, is the square of the electric field, not the electric field itself. (The electric field tells the strength of the force felt by a

stationary charge, but the amount of energy that is going past, in watts per square meter, is proportional to the square of the electric field. We shall derive the constant of proportionality in the next chapter.) If we look at the array from the W side, both oscillators contribute equally and in phase, so the electric field is twice as strong as it would be from a single oscillator. Therefore *the intensity is four times as strong as it would be if there were only one oscillator.* (The numbers in Fig. 29–5 represent how strong the intensity would be in this case, compared with what it would be if there were only a single oscillator of unit strength.) Now, in either the N or S direction along the line of the oscillators, since they are half a wavelength apart, the effect of one oscillator turns out to be out of phase by exactly half an oscillation from the other, and therefore the fields add to zero. At a certain particular intermediate angle (in fact, at 30°) the intensity is 2, and it falls off, 4, 2, 0, and so forth. We have to learn how to find these numbers at other angles. It is a question of adding two oscillations with different phases.

Let us quickly look at some other cases of interest. Suppose the oscillators are again one-half a wavelength apart, but the phase α of one is set half a period behind the other in its oscillation (Fig. 29–5b). In the W direction the intensity is now zero, because one oscillator is "pushing" when the other one is "pulling." But in the N direction the signal from the near one comes at a certain time, and that of the other comes half a period later. But the latter was *originally* half a period behind in timing, and therefore it is now exactly *in time* with the first one, and so the intensity in this direction is 4 units. The intensity in the direction at 30° is still 2, as we can prove later.

Now we come to an interesting case which shows up a possibly useful feature. Let us remark that one of the reasons that phase relations of oscillators are interesting is for beaming radio transmitters. For instance, if we build an antenna system and want to send a radio signal, say, to Hawaii, we set the antennas up as in Fig. 29–5(a) and we broadcast with our two antennas in phase, because Hawaii is to the west of us. Then we decide that tomorrow we are going to broadcast toward Alberta, Canada. Since that is north, not west, all we have to do is to reverse the phase of one of our antennas, and we can broadcast to the north. So we can build antenna systems with various arrangements. Ours is one of the simplest possible ones; we can make them much more complicated, and by changing the phases in the various antennas we can send the beams in various directions and send most of the power in the direction in which we wish to transmit, without ever moving the antenna! In both of the preceding cases, however, while we are broadcasting toward Alberta we are wasting a lot of power on Easter Island, and it would be interesting to ask whether it is possible to send it in only *one* direction. At first sight we might think that with a pair of antennas of this nature the result is always going to be symmetrical. So let us consider a case that comes out unsymmetrical, to show the possible variety.

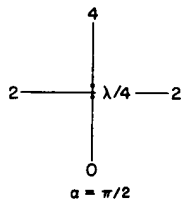


Fig. 29–6. A pair of dipole antennas giving maximum power in one direction.

If the antennas are separated by one-quarter wavelength, and if the N one is one-fourth period behind the S one in time, then what happens (Fig. 29–6)? In the W direction we get 2, as we will see later. In the S direction we get *zero*, because the signal from S comes at a certain time; that from N comes 90° later in *time*, but it is already 90° behind in its built-in phase, therefore it arrives, altogether, 180° out of phase, and there is no effect. On the other hand, in the N direction, the N signal arrives earlier than the S signal by 90° in time, because it is a quarter wavelength closer. But its phase is set so that it is oscillating 90° *behind* in time, which just compensates the delay difference, and therefore the two signals appear *together* in phase, making the field strength twice as large, and the energy four times as great.

Thus, by using some cleverness in spacing and phasing our antennas, we can send the power all in one direction. But still it is distributed over a great range of angles. Can we arrange it so that it is focused still more sharply in a particular direction? Let us consider the case of Hawaii again, where we are sending the beam east and west but it is spread over quite an angle, because even at 30° we are still getting half the intensity—we are wasting the power. Can we do better than that? Let us take a situation in which the separation is ten wavelengths (Fig.

29–7), which is more nearly comparable to the situation in which we experimented in the previous chapter, with separations of several wavelengths rather than a small fraction of a wavelength. Here the picture is quite different.

If the oscillators are ten wavelengths apart (we take the in-phase case to make it easy), we see that in the E–W direction, they are in phase, and we get a strong intensity, four times what we would get if one of them were there alone. On the other hand, at a very small angle away, the arrival times differ by 180° and the intensity is zero. To be precise, if we draw a line from each oscillator to a distant point and the difference Δ in the two distances is $\lambda/2$, half an oscillation, then they will be out of phase. So this first null occurs when that happens. (The figure is not drawn to scale; it is only a rough sketch.) This means that we do indeed have a very sharp beam in the direction we want, because if we just move over a little bit we lose all our intensity. Unfortunately for practical purposes, if we were thinking of making a radio broadcasting array and we doubled the distance Δ, then we would be a whole cycle out of phase, which is the same as being exactly *in* phase again! Thus we get many successive maxima and minima, just as we found with the $2\frac{1}{2}\lambda$ spacing in Chapter 28.

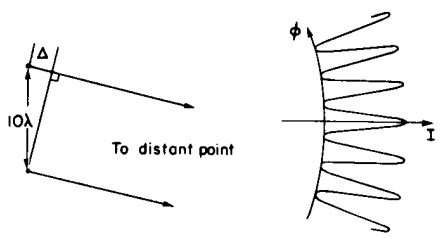


Fig. 29–7. The intensity pattern for two dipoles separated by 10λ.

Now how can we arrange to get rid of all these extra maxima, or "lobes," as they are called? We could get rid of the unwanted lobes in a rather interesting way. Suppose that we were to place another set of antennas between the two that we already have. That is, the outside ones are still 10λ apart, but between them, say every 2λ, we have put another antenna, and we drive them all in phase. There are now six antennas, and if we looked at the intensity in the E–W direction, it would, of course, be much higher with six antennas than with one. The field would be six times and the intensity thirty-six times as great (the square of the field). We get 36 units of intensity in that direction. Now if we look at neighboring points, we find a zero as before, roughly, but if we go farther, to where we used to get a big "bump," we get a much smaller "bump" now. Let us try to see why.

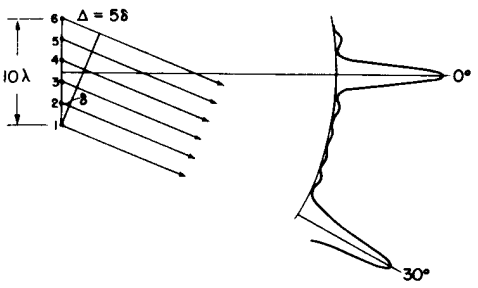


Fig. 29–8. A six-dipole antenna array and part of its intensity pattern.

The reason is that although we might expect to get a big bump when the distance Δ is exactly equal to the wavelength, it is true that dipoles 1 and 6 are then in phase and are cooperating in trying to get some strength in that direction. But numbers 3 and 4 are roughly $\frac{1}{2}$ a wavelength out of phase with 1 and 6, and although 1 and 6 push together, 3 and 4 push together too, but in opposite phase. Therefore there is very little intensity in this direction—but there is something; it does not balance exactly. This kind of thing keeps on happening; we get very little bumps, and we have the strong beam in the direction where we want it. But in this particular example, something else will happen: namely, since the distance between successive dipoles is 2λ, it is possible to find an angle where the distance δ between *successive dipoles* is exactly one wavelength, so that the effects from all of them are in phase again. Each one is delayed relative to the next one by 360°, so they all come back in phase, and we have another strong beam in that direction! It is easy to avoid this in practice because it is possible to put the dipoles closer than one wavelength apart. If we put in more antennas, closer than one wavelength apart, then this cannot happen. But the fact that this *can* happen at certain angles, if the spacing is bigger than one wavelength, is a very interesting and useful phenomenon in other applications—not radio broadcasting, but in *diffraction gratings*.

29–5 The mathematics of interference

Now we have finished our analysis of the phenomena of dipole radiators qualitatively, and we must learn how to analyze them quantitatively. To find the effect of two sources at some particular angle in the most general case, where the two oscillators have some intrinsic relative phase α from one another and the strengths A_1 and A_2 are not equal, we find that we have to add two cosines having the same frequency, but with different phases. It is very easy to find this phase difference; it is made up of a delay due to the difference in distance, and the intrinsic, built-in phase of the oscillation. Mathematically, we have to find the sum R of two waves: $R = A_1 \cos(\omega t + \phi_1) + A_2 \cos(\omega t + \phi_2)$. How do we do it?

It is really very easy, and we presume that we already know how to do it. However, we shall outline the procedure in some detail. First, we can, if we are clever with mathematics and know enough about cosines and sines, simply work it out. The easiest such case is the one where A_1 and A_2 are equal, let us say they are both equal to A. In those circumstances, for example (we could call this the trigonometric method of solving the problem), we have

$$R = A[\cos(\omega t + \phi_1) + \cos(\omega t + \phi_2)]. \tag{29.9}$$

Once, in our trigonometry class, we may have learned the rule that

$$\cos A + \cos B = 2\cos\tfrac{1}{2}(A + B)\cos\tfrac{1}{2}(A - B). \tag{29.10}$$

If we know that, then we can immediately write R as

$$R = 2A\cos\tfrac{1}{2}(\phi_1 - \phi_2)\cos(\omega t + \tfrac{1}{2}\phi_1 + \tfrac{1}{2}\phi_2). \tag{29.11}$$

So we find that we have an oscillatory wave with a new phase and a new amplitude. In general, the result *will* be an oscillatory wave with a new amplitude A_R, which we may call the resultant amplitude, oscillating at the same frequency but with a phase difference ϕ_R, called the resultant phase. In view of this, our particular case has the following result: that the resultant amplitude is

$$A_R = 2A\cos\tfrac{1}{2}(\phi_1 - \phi_2), \tag{29.12}$$

and the resultant phase is the average of the two phases, and we have completely solved our problem.

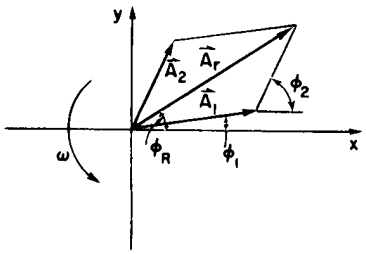


Fig. 29–9. A geometrical method for combining two cosine waves. The entire diagram is thought of as rotating counterclockwise with angular frequency ω.

Now suppose that we cannot remember that the sum of two cosines is twice the cosine of half the sum times the cosine of half the difference. Then we may use another method of analysis which is more geometrical. Any cosine function of ωt can be considered as the horizontal projection of a *rotating vector*. Suppose there were a vector $\mathbf{A}_1$ of length A_1 rotating with time, so that its angle with the horizontal axis is $\omega t + \phi_1$. (We shall leave out the ωt in a minute, and see that it makes no difference.) Suppose that we take a snapshot at the time $t = 0$, although, in fact, the picture is rotating with angular velocity ω (Fig. 29–9). The projection of A_1 along the horizontal axis is precisely $A_1\cos(\omega t + \phi_1)$. Now at $t = 0$ the second wave could be represented by another vector, $\mathbf{A}_2$, of length A_2 and at an angle ϕ_2, and also rotating. They are both rotating with the same angular velocity ω, and therefore the *relative* positions of the two are fixed. The system goes around like a rigid body. The horizontal projection of $\mathbf{A}_2$ is $A_2\cos(\omega t + \phi_2)$. But we know from the theory of vectors that if we add the two vectors in the ordinary way, by the parallelogram rule, and draw the resultant vector $\mathbf{A}_R$, the x-component of the resultant is the sum of the x-components of the other two vectors. That solves our problem. It is easy to check that this gives the correct result for the special case we treated above, where $A_1 = A_2 = A$. In this case, we see from Fig. 29–9 that $\mathbf{A}_R$ lies midway between $\mathbf{A}_1$ and $\mathbf{A}_2$ and makes an angle $\frac{1}{2}(\phi_2 - \phi_1)$ with each. Therefore we see that $A_R = 2A\cos\frac{1}{2}(\phi_2 - \phi_1)$, as before. Also, as we see from the triangle, the phase of $\mathbf{A}_R$, as it goes around, is the average angle of $\mathbf{A}_1$ and $\mathbf{A}_2$ when the two amplitudes are equal. Clearly, we can also solve for the case where the amplitudes are not equal, just as easily. We can call that the *geometrical* way of solving the problem.

There is still another way of solving the problem, and that is the *analytical* way. That is, instead of having actually to draw a picture like Fig. 29–9, we can write something down which says the same thing as the picture: instead of drawing the vectors, we write a *complex number* to represent each of the vectors. The real parts of the complex numbers are the actual physical quantities. So in our particular case the waves could be written in this way: $A_1e^{i(\omega t+\phi_1)}$ [the real part of this is $A_1\cos(\omega t + \phi_1)$] and $A_2e^{i(\omega t+\phi_2)}$. Now we can add the two:

$$R = A_1e^{i(\omega t+\phi_1)} + A_2e^{i(\omega t+\phi_2)} = (A_1e^{i\phi_1} + A_2e^{i\phi_2})e^{i\omega t} \tag{29.13}$$

or

$$\hat{R} = A_1e^{i\phi_1} + A_2e^{i\phi_2} = A_Re^{i\phi_R}. \tag{29.14}$$

This solves the problem that we wanted to solve, because it represents the result as a complex number of magnitude A_R and phase ϕ_R.

To see how this method works, let us find the amplitude A_R which is the "length" of $\hat{R}$. To get the "length" of a complex quantity, we always multiply the quantity by its complex conjugate, which gives the length squared. The complex conjugate is the same expression, but with the sign of the i's reversed. Thus we have

$$A_R^2 = (A_1e^{i\phi_1} + A_2e^{i\phi_2})(A_1e^{-i\phi_1} + A_2e^{-i\phi_2}). \tag{29.15}$$

In multiplying this out, we get $A_1^2 + A_2^2$ (here the e's cancel), and for the cross terms we have

$$A_1A_2(e^{i(\phi_1-\phi_2)} + e^{i(\phi_2-\phi_1)}).$$

Now

$$e^{i\theta} + e^{-i\theta} = \cos\theta + i\sin\theta + \cos\theta - i\sin\theta.$$

That is to say, $e^{i\theta} + e^{-i\theta} = 2\cos\theta$. Our final result is therefore

$$A_R^2 = A_1^2 + A_2^2 + 2A_1A_2\cos(\phi_2 - \phi_1). \tag{29.16}$$

As we see, this agrees with the length of A_R in Fig. 29–9, using the rules of trigonometry.

Thus the sum of the two effects has the intensity A_1^2 we would get with one of them alone, plus the intensity A_2^2 we would get with the other one alone, plus a correction. This correction we call the *interference effect*. It is really only the difference between what we get simply by adding the intensities, and what actually happens. We call it interference whether it is positive or negative. (Interference in ordinary language usually suggests opposition or hindrance, but in physics we often do not use language the way it was originally designed!) If the interference term is positive, we call that case *constructive* interference, horrible though it may sound to anybody other than a physicist! The opposite case is called *destructive* interference.

Now let us see how to apply our general formula (29.16) for the case of two oscillators to the special situations which we have discussed qualitatively. To apply this general formula, it is only necessary to find what phase difference, $\phi_1 - \phi_2$, exists between the signals arriving at a given point. (It depends only on the phase difference, of course, and not on the phase itself.) So let us consider the case where the two oscillators, of equal amplitude, are separated by some distance d and have an intrinsic relative phase α. (When one is at phase zero, the phase of the other is α.) Then we ask what the intensity will be in some azimuth direction θ from the E–W line. [Note that this is *not* the same θ as appears in (29.1). We are torn between using an unconventional symbol like $\not\theta$, or the conventional symbol θ (Fig. 29–10).] The phase relationship is found by noting that the difference in distance from P to the two oscillators is $d\sin\theta$, so that the phase difference contribution from this is the number of wavelengths in $d\sin\theta$, multiplied by 2π. (Those who are more sophisticated might want to multiply the wave number k, which is the rate of change of phase with distance, by $d\sin\theta$; it is exactly the same.) The phase difference due to the distance difference is thus $2\pi d\sin\theta/\lambda$, but, due to the timing of the oscillators, there is an additional phase α. So the phase difference at arrival would be

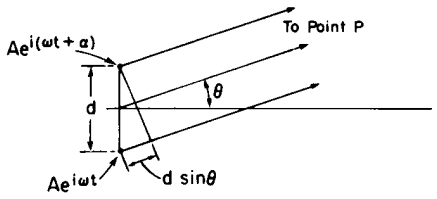


Fig. 29–10. Two oscillators of equal amplitude, with a phase difference α between them.

$$\phi_2 - \phi_1 = \alpha + 2\pi d\sin\theta/\lambda. \tag{29.17}$$

This takes care of all the cases. Thus all we have to do is substitute this expression into (29.16) for the case $A_1 = A_2$, and we can calculate all the various results for two antennas of equal intensity.

Now let us see what happens in our various cases. The reason we know, for example, that the intensity is 2 at 30° in Fig. 29–5 is the following: the two oscillators are $\frac{1}{2}\lambda$ apart, so at 30°, $d\sin\theta = \lambda/4$. Thus $\phi_2 - \phi_1 = 2\pi\lambda/4\lambda = \pi/2$, and so the interference term is zero. (We are adding two vectors at 90°.) The result is the hypotenuse of a 45° right-angle triangle, which is $\sqrt{2}$ times the unit amplitude; squaring it, we get twice the intensity of one oscillator alone. All the other cases can be worked out in this same way.

30

Diffraction

30–1 The resultant amplitude due to *n* equal oscillators

30–2 The diffraction grating

30–3 Resolving power of a grating

30–4 The parabolic antenna

30–5 Colored films; crystals

30–6 Diffraction by opaque screens

30–7 The field of a plane of oscillating charges

30–1 The resultant amplitude due to *n* equal oscillators

This chapter is a direct continuation of the previous one, although the name has been changed from *Interference* to *Diffraction.* No one has ever been able to define the difference between interference and diffraction satisfactorily. It is just a question of usage, and there is no specific, important physical difference between them. The best we can do, roughly speaking, is to say that when there are only a few sources, say two, interfering, then the result is usually called interference, but if there is a large number of them, it seems that the word diffraction is more often used. So, we shall not worry about whether it is interference or diffraction, but continue directly from where we left off in the middle of the subject in the last chapter.

Thus we shall now discuss the situation where there are n equally spaced oscillators, all of equal amplitude but different from one another in phase, either because they are driven differently in phase, or because we are looking at them at an angle such that there is a difference in time delay. For one reason or another, we have to add something like this:

$$R = A[\cos\omega t + \cos(\omega t + \phi) + \cos(\omega t + 2\phi) + \cdots + \cos(\omega t + (n-1)\phi)], \tag{30.1}$$

where ϕ is the phase difference between one oscillator and the next one, as seen in a particular direction. Specifically, $\phi = \alpha + 2\pi d \sin\theta/\lambda$. Now we must add all the terms together. We shall do this geometrically. The first one is of length A, and it has zero phase. The next is also of length A and it has a phase equal to ϕ. The next one is again of length A and it has a phase equal to 2ϕ, and so on. So we are evidently going around an equiangular polygon with n sides (Fig. 30–1).

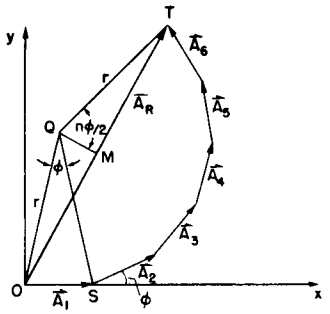


Fig. 30–1. The resultant amplitude of $n = 6$ equally spaced sources with net successive phase differences ϕ.

Now the vertices, of course, all lie on a circle, and we can find the net amplitude most easily if we find the radius of that circle. Suppose that Q is the center of the circle. Then we know that the angle OQS is just a phase angle ϕ. (This is because the radius QS bears the same geometrical relation to $\mathbf{A}_2$ as QO bears to $\mathbf{A}_1$, so they form an angle ϕ between them.) Therefore the radius r must be such that $A = 2r\sin\phi/2$, which fixes r. But the large angle OQT is equal to $n\phi$, and we thus find that $A_R = 2r\sin n\phi/2$. Combining these two results to eliminate r, we get

$$A_R = A\frac{\sin n\phi/2}{\sin\phi/2}. \tag{30.2}$$

The resultant intensity is thus

$$I = I_0\frac{\sin^2 n\phi/2}{\sin^2\phi/2}. \tag{30.3}$$

Now let us analyze this expression and study some of its consequences. In the first place, we can check it for $n = 1$. It checks: $I = I_0$. Next, we check it for $n = 2$: writing $\sin\phi = 2\sin\phi/2\cos\phi/2$, we find that $A_R = 2A\cos\phi/2$, which agrees with (29.12).

Now the idea that led us to consider the addition of several sources was that we might get a much stronger intensity in one direction than in another; that the nearby maxima which would have been present if there were only two sources will have gone down in strength. In order to see this effect, we plot the curve that comes from (30.3), taking n to be enormously large and plotting the region near

$\phi = 0$. In the first place, if ϕ is exactly 0, we have 0/0, but if ϕ is infinitesimal, the ratio of the two sines squared is simply n^2, since the sine and the angle are approximately equal. Thus the intensity of the maximum of the curve is equal to n^2 times the intensity of one oscillator. That is easy to see, because if they are all in phase, then the little vectors have no relative angle and all n of them add up so the amplitude is n times, and the intensity n^2 times, stronger.

As the phase ϕ increases, the ratio of the two sines begins to fall off, and the first time it reaches zero is when $n\phi/2 = \pi$, because $\sin \pi = 0$. In other words, $\phi = 2\pi/n$ corresponds to the first minimum in the curve (Fig. 30–2). In terms of what is happening with the arrows in Fig. 30–1, the first minimum occurs when all the arrows come back to the starting point; that means that the total accumulated angle in all the arrows, the total phase difference between the first and last oscillator, must be 2π to complete the circle.

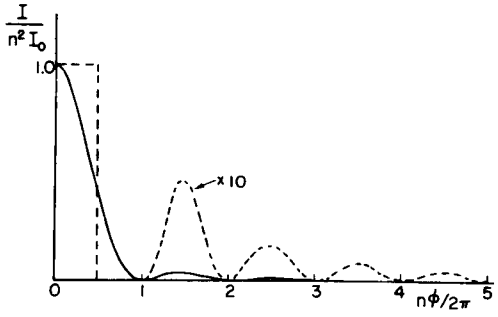


Fig. 30–2. The intensity as a function of phase angle for a large number of oscillators of equal strength.

Now we go to the next maximum, and we want to see that it is really much smaller than the first one, as we had hoped. We shall not go precisely to the maximum position, because both the numerator and the denominator of (30.3) are variant, but $\sin \phi/2$ varies quite slowly compared with $\sin n\phi/2$ when n is large, so when $\sin n\phi/2 = 1$ we are very close to the maximum. The next maximum of $\sin^2 n\phi/2$ comes at $n\phi/2 = 3\pi/2$, or $\phi = 3\pi/n$. This corresponds to the arrows having traversed the circle one and a half times. On putting $\phi = 3\pi/n$ into the formula to find the size of the maximum, we find that $\sin^2 3\pi/2 = 1$ in the numerator (because that is why we picked this angle), and in the denominator we have $\sin^2 3\pi/2n$. Now if n is sufficiently large, then this angle is very small and the sine is equal to the angle; so for all practical purposes, we can put $\sin 3\pi/2n = 3\pi/2n$. Thus we find that the intensity at this maximum is $I = I_0(4n^2/9\pi^2)$. But $n^2 I_0$ was the maximum intensity, and so we have $4/9\pi^2$ times the maximum intensity, which is about 0.047, less than 5 percent, of the maximum intensity! Of course there are decreasing intensities farther out. So we have a very sharp central maximum with very weak subsidiary maxima on the sides.

It is possible to prove that the area of the whole curve, including all the little bumps, is equal to $2\pi n I_0$, or twice the area of the dotted rectangle in Fig. 30–2.

Now let us consider further how we may apply Eq. (30.3) in different circumstances, and try to understand what is happening. Let us consider our sources to be all on a line, as drawn in Fig. 30–3. There are n of them, all spaced by a distance d, and we shall suppose that the intrinsic relative phase, one to the next, is α. Then if we are observing in a given direction θ from the normal, there is an additional phase $2\pi d \sin\theta/\lambda$ because of the time delay between each successive two, which we talked about before. Thus

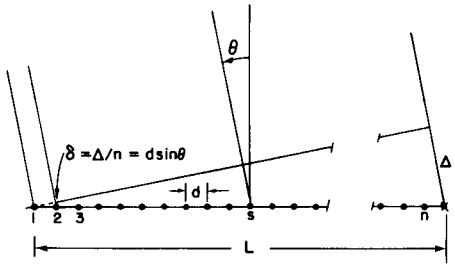


Fig. 30–3. A linear array of n equal oscillators, driven with phases $\alpha_s = s\alpha$.

$$\begin{aligned} \phi &= \alpha + 2\pi d \sin\theta/\lambda \\ &= \alpha + kd \sin\theta. \end{aligned} \tag{30.4}$$

First, we shall take the case $\alpha = 0$. That is, all oscillators are in phase, and we want to know what the intensity is as a function of the angle θ. In order to find out, we merely have to put $\phi = kd \sin\theta$ into formula (30.3) and see what happens. In the first place, there is a maximum when $\phi = 0$. That means that when all the oscillators are in phase there is a strong intensity in the direction $\theta = 0$. On the other hand, an interesting question is, where is the first minimum? That occurs when $\phi = 2\pi/n$. In other words, when $2\pi d \sin\theta/\lambda = 2\pi/n$, we get the first minumum of the curve. If we get rid of the 2π's so we can look at it a little better, it says that

$$nd \sin\theta = \lambda. \tag{30.5}$$

Now let us understand physically why we get a minimum at that position. nd is the total length L of the array. Referring to Fig. 30–3, we see that $nd \sin\theta = L \sin\theta = \Delta$. What (30.5) says is that when Δ is equal to *one wavelength*, we get a minimum. Now why do we get a minimum when $\Delta = \lambda$? Because the contributions of the various oscillators are then uniformly distributed in phase from 0° to

360°. The arrows (Fig. 30–1) are going around a whole circle—we are adding equal vectors in all directions, and such a sum is zero. So when we have an angle such that $\Delta = \lambda$, we get a minimum. That is the first minimum.

There is another important feature about formula (30.3), which is that if the angle ϕ is increased by any multiple of 2π, it makes no difference to the formula. So we will get other strong maxima at $\phi = 2\pi, 4\pi, 6\pi$, and so forth. Near each of these great maxima the pattern of Fig. 30–2 is repeated. We may ask ourselves, what is the geometrical circumstance that leads to these other great maxima? The condition is that $\phi = 2\pi m$, where m is any integer. That is, $2\pi d \sin\theta/\lambda = 2\pi m$. Dividing by 2π, we see that

$$d \sin\theta = m\lambda. \tag{30.6}$$

This looks like the other formula, (30.5). No, that formula was $nd \sin\theta = \lambda$. The difference is that here we have to look at the *individual sources*, and when we say $d \sin\theta = m\lambda$, that means that we have an angle θ such that $\delta = m\lambda$. In other words, each source is now contributing a certain amount, and successive ones are out of phase by a whole multiple of 360°, and therefore are contributing *in phase*, because out of phase by 360° is the same as being in phase. So they all contribute in phase and produce just as good a maximum as the one for $m = 0$ that we discussed before. The subsidiary bumps, the whole shape of the pattern, is just like the one near $\phi = 0$, with exactly the same minima on each side, etc. Thus such an array will send beams in various directions—each beam having a strong central maximum and a certain number of weak "side lobes." The various strong beams are referred to as the zero-order beam, the first-order beam, etc., according to the value of m. m is called the *order* of the beam.

We call attention to the fact that if d is less than λ, Eq. (30.6) can have no solution except $m = 0$, so that if the spacing is too small there is only one possible beam, the zero-order one centered at $\theta = 0$. (Of course, there is also a beam in the opposite direction.) In order to get subsidiary great maxima, we must have the spacing d of the array greater than one wavelength.

30–2 The diffraction grating

In technical work with antennas and wires it is possible to arrange that all the phases of the little oscillators, or antennas, are equal. The question is whether and how we can do a similar thing with light. We cannot at the present time literally make little optical-frequency radio stations and hook them up with infinitesimal wires and drive them all with a given phase. But there is a very easy way to do what amounts to the same thing.

Suppose that we had a lot of parallel wires, equally spaced at a spacing d, and a radiofrequency source very far away, practically at infinity, which is generating an electric field which arrives at each one of the wires at the same phase (it is so far away that the time delay is the same for all of the wires). (One can work out cases with curved arrays, but let us take a plane one.) Then the external electric field will drive the electrons up and down in each wire. That is, the field which is coming from the original source will shake the electrons up and down, and in moving, these represent *new generators*. This phenomenon is called scattering: a light wave from some source can induce a motion of the electrons in a piece of material, and these motions generate their own waves. Therefore all that is necessary is to set up a lot of wires, equally spaced, drive them with a radiofrequency source far away, and we have the situation that we want, without a whole lot of special wiring. If the incidence is normal, the phases will be equal, and we will get exactly the circumstance we have been discussing. Therefore, if the wire spacing is greater than the wavelength, we will get a strong intensity of scattering in the normal direction, and in certain other directions given by (30.6).

This can also be done with light! Instead of wires, we use a flat piece of glass and make notches in it such that each of the notches scatters a little differently than the rest of the glass. If we then shine light on the glass, each one of the notches

will represent a source, and if we space the lines very finely, but not closer than a wavelength (which is technically almost impossible anyway), then we would expect a miraculous phenomenon: the light not only will pass straight through, but there will also be a strong beam at a finite angle, depending on the spacing of the notches! Such objects have actually been made and are in common use—they are called *diffraction gratings.*

In one of its forms, a diffraction grating consists of nothing but a plane glass sheet, transparent and colorless, with scratches on it. There are often several hundred scratches to the millimeter, *very* carefully arranged so as to be equally spaced. The effect of such a grating can be seen by arranging a projector so as to throw a narrow, vertical line of light (the image of a slit) onto a screen. When we put the grating into the beam, with its scratches vertical, we see that the line is still there but, in addition, on each side we have *another* strong patch of light which is *colored.* This, of course, is the slit image spread out over a wide angular range, because the angle θ in (30.6) depends upon λ, and lights of different colors, as we know, correspond to different frequencies, and therefore different wavelengths. The longest visible wavelength is red, and since $d \sin \theta = \lambda$, that requires a larger θ. And we do, in fact, find that red is at a greater angle out from the central image! There should also be a beam on the other side, and indeed we see one on the screen. Then, there might be another solution of (30.6) when $m = 2$. We do see that there is something vaguely there—very weak—and there are even other beams beyond.

We have just argued that all these beams ought to be of the same strength, but we see that they actually are not and, in fact, not even the first ones on the right and left are equal! The reason is that the grating has been carefully built to do just this. How? If the grating consists of very fine notches, infinitesimally wide, spaced evenly, then all the intensities would indeed be equal. But, as a matter of fact, although we have taken the simplest case, we could also have considered an array of *pairs* of antennas, in which each member of the pair has a certain strength and some relative phase. In this case, it is possible to get intensities which are different in the different orders. A grating is often made with little "sawtooth" cuts instead of little symmetrical notches. By carefully arranging the "sawteeth," more light may be sent into one particular order of spectrum than into the others. In a practical grating, we would like to have as much light as possible in one of the orders. This may seem a complicated point to bring in, but it is a very clever thing to do, because it makes the grating more useful.

So far, we have taken the case where all the phases of the sources are equal. But we also have a formula for ϕ when the phases differ from one to the next by an angle α. That requires wiring up our antennas with a slight phase shift between each one. Can we do that with light? Yes, we can do it very easily, for suppose that there were a source of light at infinity, *at an angle* such that the light is coming in at an angle θ_{in}, and let us say that we wish to discuss the scattered beam, which is leaving at an angle θ_{out}. The θ_{out} is the same θ as we have had before, but the θ_{in} is merely a means for arranging that the phase of each source is different: the light coming from the distant driving source first hits one scratch, then the next, then the next, and so on, with a phase shift from one to the other, which, as we see, is $\alpha = -d \sin \theta_{\text{in}}/\lambda$. Therefore we have the formula for a grating in which light both comes in and goes out at an angle:

$$\phi = 2\pi d \sin \theta_{\text{out}}/\lambda - 2\pi d \sin \theta_{\text{in}}/\lambda. \tag{30.7}$$

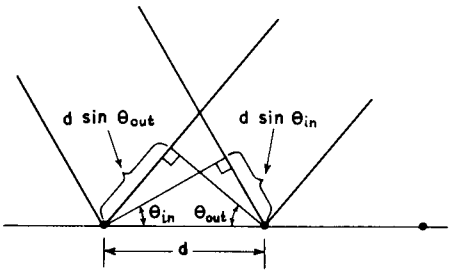


Fig. 30–4. The path difference for rays scattered from adjacent rulings of a grating is $d \sin \theta_{\text{out}} - d \sin \theta_{\text{in}}$.

Let us try to find out where we get strong intensity in these circumstances. The condition for strong intensities is, of course, that ϕ should be a multiple of 2π. There are several interesting points to be noted.

One case of rather great interest is that which corresponds to $m = 0$, where d is less than λ; in fact, this is the only solution. In this case we see that $\sin \theta_{\text{out}} = \sin \theta_{\text{in}}$, which means that the light comes out in the *same direction* as the light which was exciting the grating. We might think that the light "goes right through." No, it is *different light* that we are talking about. The light that goes right through is from the original source; what we are talking about is the new light *which is*

generated by scattering. It turns out that the scattered light is going in the same direction as the original light, in fact it can interfere with it—a feature which we will study later.

There is another solution for this same case. For a given θ_{in}, θ_{out} may be the *supplement* of θ_{in}. So not only do we get a beam in the same direction as the incoming beam but also one in another direction, which, if we consider it carefully, is such that the *angle of incidence is equal to the angle of scattering*. This we call the *reflected beam.*

So we begin to understand the basic machinery of reflection: the light that comes in generates motions of the atoms in the reflector, and the reflector then regenerates a *new wave*, and one of the solutions for the direction of scattering, the *only* solution if the spacing of the scatterers is small compared with one wavelength, is that the angle at which the light comes out is equal to the angle at which it comes in!

Next, we discuss the special case when $d \to 0$. That is, we have just a solid piece of material, so to speak, but of finite length. In addition, we want the phase shift from one scatterer to the next to go to zero. In other words, we put more and more antennas between the other ones, so that each of the phase differences is getting smaller, but the number of antennas is increasing in such a way that the total phase difference, between one end of the line and the other, is constant. Let us see what happens to (30.3) if we keep the difference in phase $n\phi$ from one end to the other constant (say $n\phi = \Phi$), letting the number go to infinity and the phase shift ϕ of each one go to zero. But now ϕ is so small that $\sin \phi = \phi$, and if we also recognize $n^2 I_0$ as I_m, the maximum intensity at the center of the beam, we find

$$I = 4I_m \sin^2 \tfrac{1}{2}\Phi/\Phi^2. \tag{30.8}$$

This limiting case is what is shown in Fig. 30–2.

In such circumstances we find the same general kind of a picture as for finite spacing with $d > \lambda$; all the side lobes are practically the same as before, but there are no higher-order maxima. If the scatterers are all in phase, we get a maximum in the direction $\theta_{out} = 0$, and a minimum when the distance Δ is equal to λ, just as for finite d and n. So we can even analyze a *continuous* distribution of scatterers or oscillators, by using integrals instead of summing.

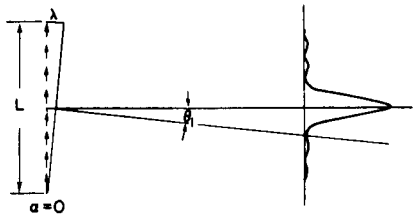


Fig. 30–5. The intensity pattern of a continuous line of oscillators has a single strong maximum and many weak "side lobes."

As an example, suppose there were a long line of oscillators, with the charge oscillating along the direction of the line (Fig. 30–5). From such an array the greatest intensity is perpendicular to the line. There is a little bit of intensity up and down from the equatorial plane, but it is very slight. With this result, we can handle a more complicated situation. Suppose we have a set of such lines, each producing a beam only in a plane perpendicular to the line. To find the intensity in various directions from a series of long wires, instead of infinitesimal wires, is the same problem as it was for infinitesimal wires, so long as we are in the central plane perpendicular to the wires; we just add the contribution from each of the long wires. That is why, although we actually analyzed only tiny antennas, we might as well have used a grating with long, narrow slots. Each of the long slots produces an effect only in its own direction, not up and down, but they are all set next to each other horizontally, so they produce interference that way.

Thus we can build up more complicated situations by having various distributions of scatterers in lines, planes, or in space. The first thing we did was to consider scatterers in a line, and we have just extended the analysis to strips; we can work it out by just doing the necessary summations, adding the contributions from the individual scatterers. The principle is always the same.

30–3 Resolving power of a grating

We are now in a position to understand a number of interesting phenomena. For example, consider the use of a grating for separating wavelengths. We noticed that the whole spectrum was spread out on the screen, so a grating can be used as an instrument for separating light into its different wavelengths. One of the

interesting questions is: supposing that there were two sources of slightly different frequency, or slightly different wavelength, how close together in wavelength could they be such that the grating would be unable to tell that there were really two different wavelengths there? The red and the blue were clearly separated. But when one wave is red and the other is slightly redder, very close, how close can they be? This is called the *resolving power* of the grating, and one way of analyzing the problem is as follows. Suppose that for light of a certain color we happen to have the maximum of the diffracted beam occurring at a certain angle. If we vary the wavelength the phase $2\pi d \sin\theta/\lambda$ is different, so of course the maximum occurs at a different angle. That is why the red and blue are spread out. How different in angle must it be in order for us to be able to see it? If the two maxima are exactly on top of each other, of course we cannot see them. If the maximum of one is far enough away from the other, then we can see that there is a double bump in the distribution of light. In order to be able to just make out the double bump, the following simple criterion, called *Rayleigh's criterion*, is usually used (Fig. 30–6). It is that the first minimum from one bump should sit at the maximum of the other. Now it is very easy to calculate, when one minimum sits on the other maximum, how much the difference in wavelength is. The best way to do it is geometrically.

In order to have a maximum for wavelength λ', the distance Δ (Fig. 30–3) must be $n\lambda'$, and if we are looking at the mth-order beam, it is $mn\lambda'$. In other words, $2\pi d \sin\theta/\lambda' = 2\pi m$, so $nd \sin\theta$, which is Δ, is λ' times n, or $mn\lambda'$. For the other beam, of wavelength λ, we want to have a *minimum* at this angle. That is, we want Δ to be exactly one wavelength λ more than $mn\lambda$. That is, $\Delta = mn\lambda + \lambda = mn\lambda'$. Thus if $\lambda' = \lambda + \Delta\lambda$, we find

$$\Delta\lambda/\lambda = 1/mn. \tag{30.9}$$



Fig. 30–6. Illustration of the Rayleigh criterion. The maximum of one pattern falls on the first minimum of the other.

The ratio $\lambda/\Delta\lambda$ is called the *resolving power* of a grating; we see that it is equal to the total number of lines in the grating, times the order. It is not hard to prove that this formula is equivalent to the formula that the error in frequency is equal to the reciprocal time difference between extreme paths that are allowed to interfere:*

$$\Delta\nu = 1/T.$$

In fact, that is the best way to remember it, because the general formula works not only for gratings, but for any other instrument whatsoever, while the special formula (30.9) depends on the fact that we are using a grating.

30–4 The parabolic antenna

Now let us consider another problem in resolving power. This has to do with the antenna of a radio telescope, used for determining the position of radio sources in the sky, i.e., how large they are in angle. Of course if we use any old antenna and find signals, we would not know from what direction they came. We are very interested to know whether the source is in one place or another. One way we can find out is to lay out a whole series of equally spaced dipole wires on the Australian landscape. Then we take all the wires from these antennas and feed them into the same receiver, in such a way that all the delays in the feed lines are equal. Thus the receiver receives signals from all of the dipoles in phase. That is, it adds all the waves from every one of the dipoles in the same phase. Now what happens? If the source is directly above the array, at infinity or nearly so, then its radiowaves will excite all the antennas in the same phase, so they all feed the receiver together.

Now suppose that the radio source is at a slight angle θ from the vertical. Then the various antennas are receiving signals a little out of phase. The receiver adds all these out-of-phase signals together, and so we get nothing, if the angle

* In our case $T = \Delta/c = mn\lambda/c$, where c is the speed of light. The frequency $\nu = c/\lambda$, so $\Delta\nu = c\,\Delta\lambda/\lambda^2$.

θ is too big. How big may the angle be? *Answer:* we get zero if the angle $\Delta/L = \theta$ (Fig. 30–3) corresponds to a 360° phase shift, that is, if Δ is the wavelength λ. This is because the vector contributions form together a complete polygon with zero resultant. The smallest angle that can be resolved by an antenna array of length L is $\theta = \lambda/L$. Notice that the receiving pattern of an antenna such as this is exactly the same as the intensity distribution we would get if we turned the receiver around and made it into a transmitter. This is an example of what is called a *reciprocity principle*. It turns out, in fact, to be generally true for any arrangement of antennas, angles, and so on, that if we first work out what the relative intensities would be in various directions if the receiver were a transmitter instead, then the relative directional sensitivity of a receiver with the same external wiring, the same array of antennas, is the same as the relative intensity of emission would be if it were a transmitter.

Some radio antennas are made in a different way. Instead of having a whole lot of dipoles in a long line, with a lot of feed wires, we may arrange them not in a line but in a curve, and put the receiver at a certain point where it can detect the scattered waves. This curve is cleverly designed so that if the radiowaves are coming down from above, and the wires scatter, making a new wave, the wires are so arranged that the scattered waves reach the receiver all at the same time (Fig. 26–12). In other words, the curve is a *parabola*, and when the source is exactly on its axis, we get a very strong intensity at the focus. In this case we understand very clearly what the resolving power of such an instrument is. The arranging of the antennas on a parabolic curve is not an essential point. It is only a convenient way to get all the signals to the same point with no relative delay and without feed wires. The angle such an instrument can resolve is still $\theta = \lambda/L$, where L is the separation of the first and last antennas. It does not depend on the spacing of the antennas and they may be very close together or in fact be all one piece of metal. Now we are describing a telescope mirror, of course. We have found the resolving power of a telescope! (Sometimes the resolving power is written $\theta = 1.22\lambda/L$, where L is the diameter of the telescope. The reason that it is not exactly λ/L is this: when we worked out that $\theta = \lambda/L$, we assumed that all the lines of dipoles were equal in strength, but when we have a circular telescope, which is the way we usually arrange a telescope, not as much signal comes from the outside edges, because it is not like a square, where we get the same intensity all along a side. We get somewhat less because we are using only part of the telescope there; thus we can appreciate that the effective diameter is a little shorter than the true diameter, and that is what the 1.22 factor tells us. In any case, it seems a little pedantic to put such precision into the resolving power formula.*)

30–5 Colored films; crystals

The above, then, are some of the effects of interference obtained by adding the various waves. But there are a number of other examples, and even though we do not understand the fundamental mechanism yet, we will some day, and we can understand even now how the interference occurs. For example, when a light wave hits a surface of a material with an index n, let us say at normal incidence, some of the light is reflected. The *reason* for the reflection we are not in a position to understand right now; we shall discuss it later. But suppose we know that some of the light is reflected both on entering and leaving a refracting medium. Then, if we look at the reflection of a light source in a thin film, we see the sum of two waves; if the thicknesses are small enough, these two waves will produce an interference, either constructive or destructive, depending on the signs of the phases. It might be, for instance, that for red light, we get an enhanced reflection, but for

* This is because Rayleigh's criterion is a rough idea in the first place, It tells you where it begins to get very hard to tell whether the image was made by one or by two stars. Actually, if sufficiently careful measurements of the exact intensity distribution over the diffracted image spot can be made, the fact that two sources make the spot can be proved even if θ is less than λ/L.

blue light, which has a different wavelength, perhaps we get a destructively interfering reflection, so that we see a bright red reflection. If we change the thickness, i.e., if we look at another place where the film is thicker, it may be reversed, the red interfering and the blue not, so it is bright blue, or green, or yellow, or whatnot. So we see *colors* when we look at thin films and the colors change if we look at different angles, because we can appreciate that the timings are different at different angles. Thus we suddenly appreciate another hundred thousand situations involving the colors that we see on oil fims, soap bubbles, etc. at different angles. But the principle is all the same: we are only adding waves at different phases.

As another important application of diffraction, we may mention the following. We used a grating and we saw the diffracted image on the screen. If we had used monochromatic light, it would have been at a certain specific place. Then there were various higher-order images also. From the positions of the images, we could tell how far apart the lines on the grating were, if we knew the wavelength of the light. From the difference in intensity of the various images, we could find out the shape of the grating scratches, whether the grating was made of wires, sawtooth notches, or whatever, *without being able to see them.* This principle is used to discover the positions of the *atoms in a crystal.* The only complication is that a crystal is three-dimensional; it is a repeating three-dimensional array of atoms. We cannot use ordinary light, because we must use something whose wavelength is less than the space between the atoms or we get no effect; so we must use radiation of very short wavelength, i.e., x-rays. So, by shining x-rays into a crystal and by noticing how intense is the reflection in the various orders, we can determine the arrangement of the atoms inside without ever being able to see them with the eye! It is in this way that we know the arrangement of the atoms in various substances, which permitted us to draw those pictures in the first chapter, showing the arrangement of atoms in salt, and so on. We shall later come back to this subject and discuss it in more detail, and therefore we say no more about this most remarkable idea at present.

30–6 Diffraction by opaque screens

Now we come to a very interesting situation. Suppose that we have an opaque sheet with holes in it, and a light on one side of it. We want to know what the intensity is on the other side. What most people say is that the light shines through the holes, and produces an effect on the other side. It will turn out that one gets the right answer, to an excellent approximation, if he assumes that there are sources distributed with uniform density across the open holes, and that the phases of these sources are the same as they would have been if the opaque material were absent. Of course, actually there are *no* sources at the holes, in fact that is the only place that there are *certainly* no sources. Nevertheless, we get the correct diffraction patterns by considering the holes to be the only places that there *are* sources; that is a rather peculiar fact. We shall explain later why this is true, but for now let us just suppose that it is.

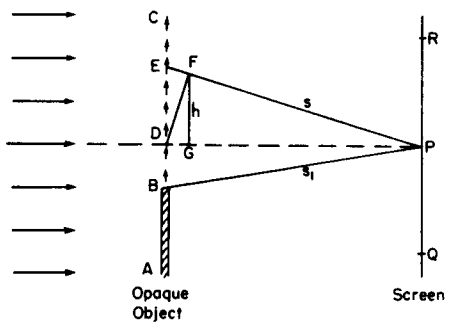


Fig. 30–7. A distant light source casts a shadow of an opaque object on a screen.

In the theory of diffraction there is another kind of diffraction that we should briefly discuss. It is usually not discussed in an elementary course as early as this, only because the mathematical formulas involved in adding these little vectors are a little elaborate. Otherwise it is exactly the same as we have been doing all along. All the interference phenomena are the same; there is nothing very much more advanced involved, only the circumstances are more complicated and it is harder to add the vectors together, that is all.

Suppose that we have light coming in from infinity, casting a shadow of an object. Figure 30–7 shows a screen on which the shadow of an object AB is made by a light source very far away compared with one wavelength. Now we would expect that outside the shadow, the intensity is all bright, and inside it, it is all dark. As a matter of fact, if we plot the intensity as a function of position near the shadow edge, the intensity rises and then overshoots, and wobbles, and oscillates about in a very peculiar manner near this edge (Fig. 30–8). We now shall discuss the reason for this. If we use the theorem that we have not yet proved, then we can

replace the actual problem by a set of effective sources uniformly distributed over the open space beyond the object.

We imagine a large number of very closely spaced antennas, and we want the intensity at some point P. That looks just like what we have been doing. Not quite; because our screen is not at infinity. We do not want the intensity at infinity, but at a finite point. To calculate the intensity at some particular place, we have to add the contributions from all the antennas. First there is an antenna at D, exactly opposite P; if we go up a little bit in angle, let us say a height h, then there is an increase in delay (there is also a change in amplitude because of the change in distance, but this is a very small effect if we are at all far away, and is much less important than the difference in the phases). Now the path difference $EP - DP$ is $h^2/2s$, so that the phase difference is proportional to the *square* of how far we go from D, while in our previous work s was infinite, and the phase difference was *linearly* proportional to h. When the phases are linearly proportional, each vector adds at a constant angle to the next vector. What we now need is a curve which is made by adding a lot of infinitesimal vectors with the requirement that the angle they make shall increase, not linearly, but as the *square* of the length of the curve. To construct that curve involves slightly advanced mathematics, but we can always construct it by actually drawing the arrows and measuring the angles. In any case, we get the marvelous curve (called Cornu's spiral) shown in Fig. 30–8. Now how do we use this curve?

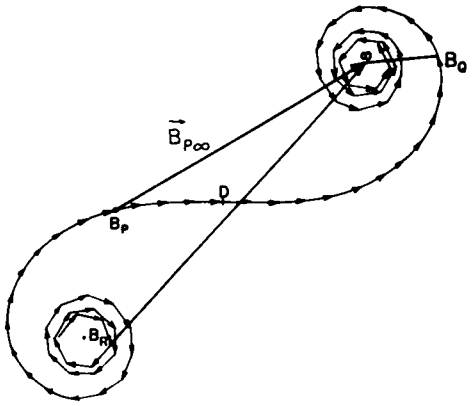


Fig. 30–8. The addition of amplitudes for many in-phase oscillators whose phase delays vary as the square of the distance from point D of the previous figure.

If we want the intensity, let us say, at point P, we add a lot of contributions of different phases from point D on up to infinity, and from D down only to point B_P. So we start at B_P in Fig. 30–8, and draw a series of arrows of ever-increasing angle. Therefore the total contribution above point B_P all goes along the spiraling curve. If we were to stop integrating at some place, then the total amplitude would be a vector from B to that point; in this particular problem we are going to infinity, so the total answer is the vector $\mathbf{B}_{P\infty}$. Now the position on the curve which corresponds to point B_P on the object depends upon where point P is located, since point D, the inflection point, always corresponds to the position of point P. Thus, depending upon where P is located above B, the beginning point will fall at various positions on the lower left part of the curve, and the resultant vector $\mathbf{B}_{P\infty}$ will have many maxima and minima (Fig. 30–9).

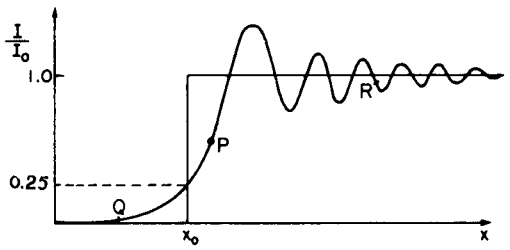


Fig. 30–9. The intensity near the edge of a shadow. The geometrical shadow edge is at x_0.

On the other hand, if we are at Q, on the other side of P, then we are using only one end of the spiral curve, and not the other end. In other words, we do not even start at D, but at B_Q, so on this side we get an intensity which continuously falls off as Q goes farther into the shadow.

One point that we can immediately calculate with ease, to show that we really understand it, is the intensity exactly opposite the edge. The intensity here is 1/4 that of the incident light. Reason: Exactly at the edge (so the endpoint B of the arrow is at D in Fig. 30–8) we have half the curve that we would have had if we were far into the bright region. If our point R is far into the light we go from one end of the curve to the other, that is, one full unit vector; but if we are at the edge of the shadow, we have only half the amplitude—1/4 the intensity.

In this chapter we have been finding the intensity produced in various directions from various distributions of sources. As a final example we shall derive a formula which we shall need for the next chapter on the theory of the index of refraction. Up to this point relative intensities have been sufficient for our purpose, but this time we shall find the complete formula for the field in the following situation.

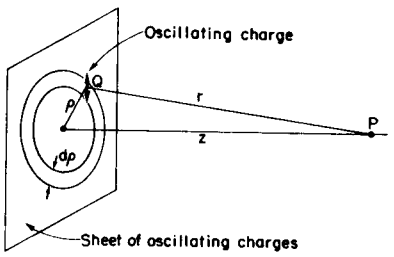


Fig. 30–10. Radiation field of a sheet of oscillating charges.

30–7 The field of a plane of oscillating charges

Suppose that we have a plane full of sources, all oscillating together, with their motion in the plane and all having the same amplitude and phase. What is the field at a finite, but very large, distance away from the plane? (We cannot get very close, of course, because we do not have the right formulas for the field close to the sources.) If we let the plane of the charges be the XY-plane, then we want to find the field at the point P far out on the Z-axis (Fig. 30–10). We suppose that there are η charges per unit area of the plane, and that each one of them has a charge q. All of the charges move with simple harmonic motion, with the same direction, amplitude, and phase. We let the motion of each charge, *with respect to its own average position*, be $x_0 \cos \omega t$. Or, using the complex notation and remembering that the real part represents the actual motion, the motion can be described by $x_0 e^{i\omega t}$.

Now we find the field at the point P from all of the charges by finding the field there from each charge q, and then adding the contributions from all the charges. We know that the radiation field is proportional to the acceleration of the charge, which is $-\omega^2 x_0 e^{i\omega t}$ (and is the same for every charge). The electric field that we want at the point P due to a charge at the point Q is proportional to the acceleration of the charge q, but we have to remember that the field at the point P at the instant t is given by the acceleration of the charge at the earlier time $t' = t - r/c$, where r/c is the time it takes the waves to travel the distance r from Q to P. Therefore the field at P is proportional to

$$-\omega^2 x_0 e^{i\omega(t-r/c)}. \tag{30.10}$$

Using this value for the acceleration as seen from P in our formula for the electric field at large distances from a radiating charge, we get

$$\begin{pmatrix}\text{Electric field at } P\\ \text{from charge at } Q\end{pmatrix} = \frac{q}{4\pi\epsilon_0 c^2}\,\frac{\omega^2 x_0 e^{i\omega(t-r/c)}}{r}\ \text{(approx.)}. \tag{30.11}$$

Now this formula is not quite right, because we should have used *not* the acceleration of the charge but *its component* perpendicular to the line QP. We shall suppose, however, that the point P is so far away, compared with the distance of the point Q from the axis (the distance ρ in Fig. 30–9), for those changes that we need to take into account, that we can leave out the cosine factor (which would be nearly equal to 1 anyway).

To get the total field at P, we now add the effects of all the charges in the plane. We should, of course, make a *vector* sum. But since the direction of the electric field is nearly the same for all the charges, we may, in keeping with the approximation we have already made, just add the magnitudes of the fields. To our approximation the field at P depends only on the distance r, so all charges at the same r produce equal fields. So we add, first, the fields of those charges in a ring of width $d\rho$ and radius ρ. Then, by taking the integral over all ρ, we will obtain the total field.

The number of charges in the ring is the product of the surface area of the ring, $2\pi\rho\,d\rho$, and η, the number of charges per unit area. We have, then,

$$\text{Total field at } P = \int \frac{q}{4\pi\epsilon_0 c^2}\,\frac{\omega^2 x_0 e^{i\omega(t-r/c)}}{r}\cdot\eta\cdot 2\pi\rho\,d\rho. \tag{30.12}$$

We wish to evaluate this integral from $\rho = 0$ to $\rho = \infty$. The variable t, of course, is to be held fixed while we do the integral, so the only varying quantities are ρ and r. Leaving out all the constant factors, *including the factor* $e^{i\omega t}$, for the moment, the integral we wish is

$$\int_{\rho=0}^{\rho=\infty} \frac{e^{-i\omega r/c}}{r}\,\rho\,d\rho. \tag{30.13}$$

To do this integral we need to use the relation between r and ρ:

$$r^2 = \rho^2 + z^2. \tag{30.14}$$

Since z is independent of ρ, when we take the differential of this equation, we get

$$2r\,dr = 2\rho\,d\rho,$$

which is lucky, since in our integral we can replace $\rho\,d\rho$ by $r\,dr$ and the r will cancel the one in the denominator. The integral we want is then the simpler one

$$\int_{r=z}^{r=\infty} e^{-i\omega r/c}\,dr. \tag{30.15}$$

To integrate an exponential is very easy. We divide by the coefficient of r in the exponent and evaluate the exponential at the limits. But the limits of r are not the same as the limits of ρ. When $\rho = 0$, we have $r = z$, so the limits of r are z to infinity. We get for the integral

$$-\frac{c}{i\omega}\left[e^{-i\infty} - e^{-(i\omega/c)z}\right], \tag{30.16}$$

where we have written ∞ for $(r/c)\,\infty$, since they both just mean a very large number!

Now $e^{-i\infty}$ is a mysterious quantity. Its real part, for example, is cos $(-\infty)$, which, mathematically speaking, is completely indefinite (although we would expect it to be somewhere—or everywhere (?)—between $+1$ and -1!). But in a *physical* situation, it can mean something quite reasonable, and usually can just be taken to be zero. To see that this is so in our case, we go back to consider again the original integral (30.15).

We can understand (30.15) as a sum of many small complex numbers, each of magnitude Δr, and with the angle $\theta = -\omega r/c$ in the complex plane. We can try to evaluate the sum by a graphical method. In Fig. 30–11 we have drawn the first five pieces of the sum. Each segment of the curve has the length Δr and is placed at the angle $\Delta\theta = -\omega\,\Delta r/c$ with respect to the preceding piece. The sum for these first five pieces is represented by the arrow from the starting point to the end of the fifth segment. As we continue to add pieces we shall trace out a polygon until we get back to the starting point (approximately) and then start around once more. Adding more pieces, we just go round and round, staying close to a circle whose radius is easily shown to be c/ω. We can see now why the integral does not give a definite answer!

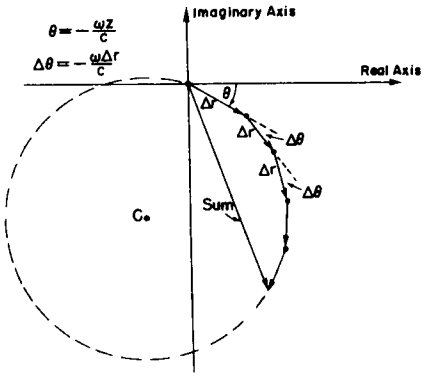


Fig. 30–11. Graphical solution of $\int_z^\infty e^{-i\omega r/c}dr$.

But now we have to go back to the *physics* of the situation. In any real situation the plane of charges *cannot* be infinite in extent, but must sometime stop. If it stopped suddenly, and was exactly circular in shape, our integral would have some value on the circle in Fig. 30–11. If, however, we let the number of charges in the plane gradually taper off at some large distance from the center (or else stop suddenly but in an irregular shape so for larger ρ the entire ring of width $d\rho$ no longer contributes), then the coefficient η in the exact integral would decrease toward zero. Since we are adding smaller pieces but still turning through the same angle, the graph of our integral would then become a curve which is a spiral. The spiral would eventually end up at the center of our original circle, as drawn in Fig. 30–12. The *physically* correct integral is the complex number A in the figure represented by the interval from the starting point to the center of the circle, which is just equal to

$$\frac{c}{i\omega}\,e^{-i\omega z/c}, \tag{30.17}$$

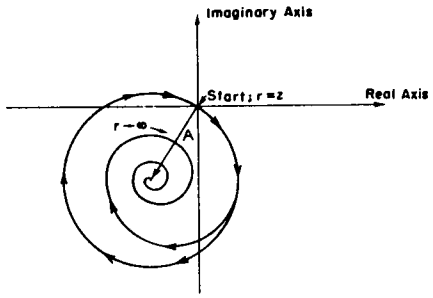


Fig. 30–12. Graphical solution of $\int_z^\infty \eta e^{-i\omega r/c}dr$.

as you can work out for yourself. This is the same result we would get from Eq. (30.16) if we set $e^{-i\infty} = 0$.

(There is also another reason why the contribution to the integral tapers off for large values of r, and that is the factor we have omitted for the projection of the acceleration on the plane perpendicular to the line PQ.)

We are, of course, interested only in physical situations, so we will take $e^{-i\infty}$ equal to zero. Returning to our original formula (30.12) for the field and putting

back all of the factors that go with the integral, we have the result

$$\text{Total field at } P = -\frac{\eta q}{2\epsilon_0 c} i\omega x_0 e^{i\omega(t-z/c)} \tag{30.18}$$

(remembering that $1/i = -i$).

It is interesting to note that $(i\omega x_0 e^{i\omega t})$ is just equal to the *velocity* of the charges, so that we can also write the equation for the field as

$$\text{Total field at } P = -\frac{\eta q}{2\epsilon_0 c} [\text{velocity of charges}]_{\text{at } t-z/c}, \tag{30.19}$$

which is a little strange, because the retardation is just by the distance z, which is the shortest distance from P to the plane of charges. But that is the way it comes out—fortunately a rather simple formula. (We may add, by the way, that although our derivation is valid only for distances far from the plane of oscillatory charges, it turns out that the formula (30.18) or (30.19) is correct at any distance z, even for $z < \lambda$.)

31

The Origin of the Refractive Index

31–1 The index of refraction

31–2 The field due to the material

31–3 Dispersion

31–4 Absorption

31–5 The energy carried by an electric wave

31–6 Diffraction of light by a screen

31–1 The index of refraction

We have said before that light goes slower in water than in air, and slower, slightly, in air than in vacuum. This effect is described by the index of refraction n. Now we would like to understand how such a slower velocity could come about. In particular, we should try to see what the relation is to some physical assumptions, or statements, we made earlier, which were the following:

(a) That the total electric field in any physical circumstance can always be represented by the sum of the fields from all the charges in the universe.

(b) That the field from a single charge is given by its acceleration evaluated with a retardation at the speed c, *always* (for the *radiation* field).

But, for a piece of glass, you might think: "Oh, no, you should modify all this. You should say it is retarded at the speed c/n." That, however, is not right, and we have to understand why it is not.

It *is* approximately true that light or any electrical wave *does appear* to travel at the speed c/n through a material whose index of refraction is n, but the fields are still produced by the motions of *all* the charges—including the charges moving in the material—and with these basic contributions of the field travelling at the ultimate velocity c. Our problem is to understand how the *apparently* slower velocity comes about.

We shall try to understand the effect in a very simple case. A source which we shall call "the *external* source" is placed a large distance away from a thin plate of transparent material, say glass. We inquire about the field at a large distance on the opposite side of the plate. The situation is illustrated by the diagram of Fig. 31–1, where S and P are imagined to be very far away from the plate. According to the principles we have stated earlier, an electric field anywhere that is far from all moving charges is the (vector) sum of the fields produced by the external source (at S) *and* the fields produced by *each* of the charges in the plate of glass, *every one with its proper retardation at the velocity c*. Remember that the contribution of each charge is not changed by the presence of the other charges. These are our basic principles. The field at P can be written thus:

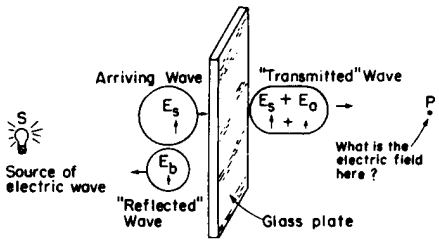


Fig. 31–1. Electric waves passing through a layer of transparent material.

$$\mathbf{E} = \sum_{\text{all charges}} \mathbf{E}_{\text{each charge}} \tag{31.1}$$

or

$$\mathbf{E} = \mathbf{E}_s + \sum_{\text{all other charges}} \mathbf{E}_{\text{each charge}}, \tag{31.2}$$

where $\mathbf{E}_s$ is the field due to the source alone and would be precisely the field at P *if there were no material present*. We expect the field at P to be different from $\mathbf{E}_s$ if there are any other moving charges.

Why should there be charges moving in the glass? We know that all material consists of atoms which contain electrons. When the electric field *of the source* acts on these atoms it drives the electrons up and down, because it exerts a force on the electrons. And moving electrons generate a field—they constitute new radiators. These new radiators are related to the source S, because they are driven by the field of the source. The total field is not just the field of the source S, but it is modified by the additional contribution from the other moving charges. This means that the field is not the same as the one which was there before the glass was there, but is modified, and it turns out that it is modified in such a way that

the field inside the glass appears to be moving at a different speed. That is the idea which we would like to work out quantitatively.

Now this is, in the exact case, pretty complicated, because although we have said that all the other moving charges are driven by the source field, that is not quite true. If we think of a particular charge, it feels not only the source, but like anything else in the world, it feels *all* of the charges that are moving. It feels, in particular, the charges that are moving somewhere else in the glass. So the total field which is acting on a *particular charge* is a combination of the fields from the other charges, *whose motions depend on what this particular charge is doing!* You can see that it would take a complicated set of equations to get the complete and exact formula. It is so complicated that we postpone this problem until next year.

Instead we shall work out a very simple case in order to understand all the physical principles very clearly. We take a circumstance in which the effects from the other atoms is very small relative to the effects from the source. In other words, we take a material in which the total field is not modified very much by the motion of the other charges. That corresponds to a material in which the index of refraction is very close to 1, which will happen, for example, if the density of the atoms is very low. Our calculation will be valid for any case in which the index is for any reason very close to 1. In this way we shall avoid the complications of the most general, complete solution.

Incidentally, you should notice that there is another effect caused by the motion of the charges in the plate. These charges will also radiate waves back toward the source S. This backward-going field is the light we see reflected from the surfaces of transparent materials. It does not come from just the surface. The backward radiation comes from everywhere in the interior, but it turns out that the total effect is equivalent to a reflection from the surfaces. These reflection effects are beyond our approximation at the moment because we shall be limited to a calculation for a material with an index so close to 1 that very little light is reflected.

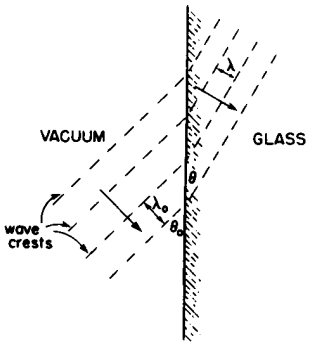


Fig. 31–2. Relation between refraction and velocity change.

Before we proceed with our study of how the index of refraction comes about, we should understand that all that is required to understand refraction is to understand why the apparent wave *velocity* is different in different materials. The *bending* of light rays comes about just *because* the effective speed of the waves is different in the materials. To remind you how that comes about we have drawn in Fig. 31–2 several successive crests of an electric wave which arrives from a vacuum onto the surface of a block of glass. The arrow perpendicular to the wave crests indicates the direction of travel of the wave. Now all oscillations in the wave must have the same *frequency*. (We have seen that driven oscillations have the same frequency as the driving source.) This means, also, that the wave crests for the waves on both sides of the surface must have the *same spacing along the surface* because they must travel together, so that a charge sitting at the boundary will feel only one frequency. The *shortest* distance between crests of the wave, however, is the wavelength which is the velocity divided by the frequency. On the vacuum side it is $\lambda_0 = 2\pi c/\omega$, and on the other side it is $\lambda = 2\pi v/\omega$ or $2\pi c/\omega n$, if $v = c/n$ is the velocity of the wave. From the figure we can see that the only way for the waves to "fit" properly at the boundary is for the waves in the material to be travelling at a different angle with respect to the surface. From the geometry of the figure you can see that for a "fit" we must have $\lambda_0/\sin\theta_0 = \lambda/\sin\theta$, or $\sin\theta_0/\sin\theta = n$, which is Snell's law. We shall, for the rest of our discussion, consider only why light has an effective speed of c/n in material of index n, and no longer worry, in this chapter, about the bending of the light direction.

We go back now to the situation shown in Fig. 31–1. We see that what we have to do is to calculate the field produced at P by all the oscillating charges in the glass plate. We shall call this part of the field E_a, and it is just the sum written as the second term in Eq. (31.2). When we add it to the term E_s, due to the source, we will have the total field at P.

This is probably the most complicated thing that we are going to do this year, but it is complicated only in that there are many pieces that have to be put together; each piece, however, is very simple. Unlike other derivations where we say, "Forget the derivation, just look at the answer!," in this case we do not need the answer so much as the derivation. In other words, the thing to understand now is the physical machinery for the production of the index.

To see where we are going, let us first find out what the "correction field" E_a would have to be if the total field at P is going to look like radiation from the source that is slowed down while passing through the thin plate. If the plate had no effect on it, the field of a wave travelling to the right (along the z-axis) would be

$$E_s = E_0 \cos \omega(t - z/c) \tag{31.3}$$

or, using the exponential notation,

$$E_s = E_0 e^{i\omega(t-z/c)}. \tag{31.4}$$

Now what would happen if the wave travelled more slowly in going through the plate? Let us call the thickness of the plate Δz. If the plate were not there the wave would travel the distance Δz in the time $\Delta z/c$. But if it appears to travel at the speed c/n then it should take the longer time $n\,\Delta z/c$ or the *additional* time $\Delta t = (n - 1)\,\Delta z/c$. After that it would continue to travel at the speed c again. We can take into account the extra delay in getting through the plate by replacing t in Eq. (31.4) by $(t - \Delta t)$ or by $[t - (n - 1)\,\Delta z/c]$. So the wave after insertion of the plate should be written

$$E_{\text{after plate}} = E_0 e^{i\omega[t-(n-1)\Delta z/c-z/c]}. \tag{31.5}$$

We can also write this equation as

$$E_{\text{after plate}} = e^{-i\omega(n-1)\Delta z/c} E_0 e^{i\omega(t-z/c)}, \tag{31.6}$$

which says that the wave after the plate is obtained from the wave which could exist without the plate, i, e., from E_s, by multiplying by the factor $e^{-i\omega(n-1)\Delta z/c}$. Now we know that multiplying an oscillating function like $e^{i\omega t}$ by a factor $e^{i\theta}$ just says that we change the phase of the oscillation by the angle θ, which is, of course, what the extra delay in passing through the thickness Δz has done. It has retarded the phase by the amount $\omega(n - 1)\,\Delta z/c$ (retarded, because of the minus sign in the exponent).

We have said earlier that the plate should *add* a field E_a to the original field $E_s = E_0 e^{i\omega(t-z/c)}$, but we have found instead that the effect of the plate is to *multiply* the field by a factor which shifts its phase. However, that is really all right because we can get the same result by adding a suitable complex number. It is particularly easy to find the right number to add in the case that Δz is small, for you will remember that if x is a small number then e^x is nearly equal to $(1 + x)$. We can write, therefore,

$$e^{-i\omega(n-1)\Delta z/c} = 1 - i\omega(n - 1)\,\Delta z/c. \tag{31.7}$$

Using this equality in Eq. (31.6), we have

$$E_{\text{after plate}} = \underbrace{E_0 e^{i\omega(t-z/c)}}_{E_s} - \underbrace{\frac{i\omega(n-1)\,\Delta z}{c} E_0 e^{i\omega(t-z/c)}}_{E_a}. \tag{31.8}$$

The first term is just the field from the source, and the second term must just be equal to E_a, the field produced to the right of the plate by the oscillating charges of the plate—expressed here in terms of the index of refraction n, and depending, of course, on the strength of the wave from the source.

What we have been doing is easily visualized if we look at the complex number diagram in Fig. 31–3. We first draw the number E_s (we chose some values for z and t so that E_s comes out horizontal, but this is not necessary). The delay due to

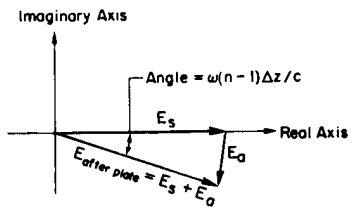


Fig. 31–3. Diagram for the transmitted wave at a particular t and z.

slowing down in the plate would delay the phase of this number, that is, it would rotate E_s through a negative angle. But this is equivalent to adding the small vector E_a at roughly right angles to E_s. But that is just what the factor $-i$ means in the second term of Eq. (31.8). It says that if E_s is real, then E_a is negative imaginary or that, in general, E_s and E_a make a right angle.

31–2 The field due to the material

We now have to ask: Is the field E_a obtained in the second term of Eq. (31.8) the kind we would expect from oscillating charges in the plate? If we can show that it is, we will then have calculated what the index n should be! [Since n is the only nonfundamental number in Eq. (31.8).] We turn now to calculating what field E_a the charges in the material will produce. (To help you keep track of the many symbols we have used up to now, and will be using in the rest of our calculation, we have put them all together in Table 31–1.)

Table 31–1

Symbols used in the calculations

E_s	= field from the source
E_a	= field produced by charges in the plate
Δz	= thickness of the plate
z	= perpendicular distance from the plate
n	= index of refraction
ω	= frequency (angular) of the radiation
N	= number of charges per unit volume in the plate
η	= number of charges per unit area of the plate
q_e	= charge on an electron
m	= mass of an electron
ω_0	= resonant frequency of an electron bound in an atom

If the source S (of Fig. 31–1) is far off to the left, then the field E_s will have the same phase everywhere on the plate, so we can write that in the neighborhood of the plate

$$E_s = E_0 e^{i(\omega t - z/c)} \tag{31.9}$$

Right at the plate, where $z = 0$, we will have

$$E_s = E_0 e^{i\omega t} \text{ (at the plate)} \tag{31.10}$$

Each of the electrons in the atoms of the plate will feel this electric field and will be driven up and down (we assume the direction of E_0 is vertical) by the electric force qE. To find what motion we expect for the electrons, we will assume that the atoms are little oscillators, that is, that the electrons are fastened elastically to the atoms, which means that if a force is applied to an electron its displacement from its normal position will be proportional to the force.

You may think that this is a funny model of an atom if you have heard about electrons whirling around in orbits. But that is just an oversimplified picture. The correct picture of an atom, which is given by the theory of wave mechanics, says that, *so far as problems involving light are concerned*, the electrons behave as though they were held by springs. So we shall suppose that the electrons have a linear restoring force which, together with their mass m, makes them behave like little oscillators, with a resonant frequency ω_0. We have already studied such oscillators, and we know that the equation of their motion is written this way:

$$m\left(\frac{d^2x}{dt^2} + \omega_0^2 x\right) = F, \tag{31.11}$$

where F is the driving force.

For our problem, the driving force comes from the electric field of the wave from the source, so we should use

$$F = q_e E_s = q_e E_0 e^{i\omega t}, \tag{31.12}$$

where q_e is the electric charge on the electron and for E_s we use the expression $E_s = E_0 e^{i\omega t}$ from (31.10). Our equation of motion for the electron is then

$$m\left(\frac{d^2x}{dt^2} + \omega_0^2 x\right) = q_e E_0 e^{i\omega t}. \tag{31.13}$$

We have solved this equation before, and we know that the solution is

$$x = x_0 e^{i\omega t}, \tag{31.14}$$

where, by substituting in (31.13), we find that

$$x_0 = \frac{q_e E_0}{m(\omega_0^2 - \omega^2)}, \tag{31.15}$$

so that

$$x = \frac{q_e E_0}{m(\omega_0^2 - \omega^2)} e^{i\omega t}. \tag{31.16}$$

We have what we needed to know—the motion of the electrons in the plate. And it is the same for every electron, except that the mean position (the "zero" of the motion) is, of course, different for each electron.

Now we are ready to find the field E_a that these atoms produce at the point P, because we have already worked out (at the end of Chapter 30) what field is produced by a sheet of charges that all move together. Referring back to Eq. (30.19), we see that the field E_a at P is just a negative constant times the velocity of the charges retarded in time the amount z/c. Differentiating x in Eq. (31.16) to get the velocity, and sticking in the retardation [or just putting x_0 from (31.15) into (30.18)] yields

$$E_a = -\frac{\eta q_e}{2\epsilon_0 c}\left[i\omega \frac{q_e E_0}{m(\omega_0^2 - \omega^2)} e^{i\omega(t - z/c)}\right]. \tag{31.17}$$

Just as we expected, the driven motion of the electrons produced an extra wave which travels to the right (that is what the factor $e^{i\omega(t-z/c)}$ says), and the amplitude of this wave is proportional to the number of atoms per unit area in the plate (the factor η) and also proportional to the strength of the source field (the factor E_0). Then there are some factors which depend on the atomic properties (q_e, m, and ω_0), as we should expect.

The most important thing, however, is that this formula (31.17) for E_a looks very much like the expression for E_a that we got in Eq. (31.8) by saying that the original wave was delayed in passing through a material with an index of refraction n. The two expressions will, in fact, be identical if

$$(n - 1)\,\Delta z = \frac{\eta q_e^2}{2\epsilon_0 m(\omega_0^2 - \omega^2)}. \tag{31.18}$$

Notice that both sides are proportional to Δz, since η, which is the number of atoms *per unit area*, is equal to $N\,\Delta Z$, where N is the number of atoms *per unit volume* of the plate. Substituting $N\,\Delta Z$ for η and cancelling the Δz, we get our main result, a formula for the index of refraction in terms of the properties of the atoms of the material—and of the frequency of the light:

$$n = 1 + \frac{N q_e^2}{2\epsilon_0 m(\omega_0^2 - \omega^2)}. \tag{31.19}$$

This equation gives the "explanation" of the index of refraction that we wished to obtain.

Notice that in the above process we have obtained something very interesting. For we have not only a number for the index of refraction which can be computed from the basic atomic quantities, but we have also learned how the index of refraction should vary with the frequency ω of the light. This is something we would never understand from the simple statement that "light travels slower in a transparent material." We still have the problem, of course, of knowing how many atoms per unit volume there are, and what is their natural frequency ω_0. We do not know this just yet, because it is different for every different material, and we cannot get a general theory of that now. Formulation of a general theory of the properties of different substances—their natural frequencies, and so on—is possible only with quantum atomic mechanics. Also, different materials have different properties and different indexes, so we cannot expect, anyway, to get a general formula for the index which will apply to all substances.

However, we shall discuss the formula we have obtained, in various possible circumstances. First of all, for most ordinary gases (for instance, for air, most colorless gases, hydrogen, helium, and so on) the natural frequencies of the electron oscillators correspond to ultraviolet light. These frequencies are higher than the frequencies of visible light, that is, ω_0 is much larger than ω of visible light, and to a first approximation, we can disregard ω^2 in comparison with ω_0^2. Then we find that the index is nearly constant. So for a gas, the index is nearly constant. This is also true for most other transparent substances, like glass. If we look at our expression a little more closely, however, we notice that as ω rises, taking a little bit more away from the denominator, the index also rises. So n rises slowly with frequency. The index is higher for blue light than for red light. That is the reason why a prism bends the light more in the blue than in the red.

The phenomenon that the index depends upon the frequency is called the phenomenon of *dispersion*, because it is the basis of the fact that light is "dispersed" by a prism into a spectrum. The equation for the index of refraction as a function of frequency is called a *dispersion equation*. So we have obtained a dispersion equation. (In the past few years "dispersion equations" have been finding a new use in the theory of elementary particles.)

Our dispersion equation suggests other interesting effects. If we have a natural frequency ω_0 which lies in the visible region, or if we measure the index of refraction of a material like glass in the ultraviolet, where ω gets near ω_0, we see that at frequencies very close to the natural frequency the index can get enormously large, because the denominator can go to zero. Next, suppose that ω is greater than ω_0. This would occur, for example, if we take a material like glass, say, and shine x-ray radiation on it. In fact, since many materials which are opaque to visible light, like graphite for instance, are transparent to x-rays, we can also talk about the index of refraction of carbon for x-rays. All the natural frequencies of the carbon atoms would be much lower than the frequency we are using in the x-rays, since x-ray radiation has a very high frequency. The index of refraction is that given by our dispersion equation if we set ω_0 equal to zero (we neglect ω_0^2 in comparison with ω^2).

A similar situation would occur if we beam radiowaves (or light) on a gas of free electrons. In the upper atmosphere electrons are liberated from their atoms by ultraviolet light from the sun and they sit up there as free electrons. For free electrons $\omega_0 = 0$ (there is no elastic restoring force). Setting $\omega_0 = 0$ in our dispersion equation yields the correct formula for the index of refraction for radiowaves in the stratosphere, where N is now to represent the density of free electrons (number per unit volume) in the stratosphere. But let us look again at the equation, if we beam x-rays on matter, or radiowaves (or any electric waves) on free electrons the term $(\omega_0^2 - \omega^2)$ becomes *negative*, and we obtain the result that n is *less than one*. That means that the effective speed of the waves in the substance is *faster* than c! Can that be correct?

It is correct. In spite of the fact that it is said that you cannot send signals any faster than the speed of light, it is nevertheless true that the index of refraction of materials at a particular frequency can be either greater or less than 1. This

just means that the *phase shift* which is produced by the scattered light can be either positive or negative. It can be shown, however, that the speed at which you can send a *signal* is not determined by the index at one frequency, but depends on what the index is at *many* frequencies. What the index tells us is the speed at which the *nodes* (or crests) of the wave travel. The *node* of a wave is not a signal by itself. In a perfect wave, which has no modulations of any kind, i.e., which is a steady oscillation, you cannot really say when it "starts," so you cannot use it for a timing signal. In order to send a *signal* you have to change the wave somehow, make a notch in it, make it a little bit fatter or thinner. That means that you have to have more than one frequency in the wave, and it can be shown that the speed at which *signals* travel is not dependent upon the index alone, but upon the way that the index changes with the frequency. This subject we must also delay (until Chapter 48). Then we will calculate for you the actual speed of *signals* through such a piece of glass, and you will see that it will not be faster than the speed of light, although the nodes, which are mathematical points, do travel faster than the speed of light.

Just to give a slight hint as to how that happens, you will note that the real difficulty has to do with the fact that the responses of the charges are opposite to the field, i.e., the sign has gotten reversed. Thus in our expression for x (Eq. 31.16) the displacement of the charge is in the direction opposite to the driving field, because $(\omega_0^2 - \omega^2)$ is negative for small ω_0. The formula says that when the electric field is pulling in one direction, the charge is moving in the opposite direction.

How does the charge happen to be going in the opposite direction? It certainly does not start off in the opposite direction when the field is first turned on. When the motion first starts there is a transient, which settles down after awhile, and only *then* is the phase of the oscillation of the charge opposite to the driving field. And it is then that the *phase* of the transmitted field can appear to be *advanced* with respect to the source wave. It is this *advance in phase* which is meant when we say that the "phase velocity" or velocity of the nodes is greater than c. In Fig. 31–4 we give a schematic idea of how the waves might look for a case where the wave is suddenly turned on (to make a signal). You will see from the diagram that the *signal* (i.e., the *start* of the wave) is *not* earlier for the wave which ends up with an advance in phase.

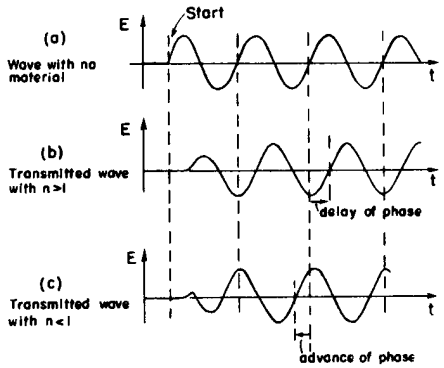


Fig. 31–4. Wave "signals."

Let us now look again at our dispersion equation. We should remark that our analysis of the refractive index gives a result that is somewhat simpler than you would actually find in nature. To be completely accurate we must add some refinements. First, we should expect that our model of the atomic oscillator should have some damping force (otherwise once started it would oscillate forever, and we do not expect that to happen). We have worked out before (Eq. 23.8) the motion of a damped oscillator and the result is that the denominator in Eq. (31.16), and therefore in (31.19), is changed from $(\omega_0^2 - \omega^2)$ to $(\omega_0^2 - \omega^2 + i\gamma\omega)$, where γ is the damping coefficient.

We need a second modification to take into account the fact that there are several resonant frequencies for a particular kind of atom. It is easy to fix up our

dispersion equation by imagining that there are several different kinds of oscillators, but that each oscillator acts separately, and so we simply add the contributions of all the oscillators. Let us say that there are N_k electrons per unit of volume, whose natural frequency is ω_k and whose damping factor is γ_k. We would then have for our dispersion equation

$$n = 1 + \frac{q_e^2}{2\epsilon_0 m} \sum_k \frac{N_k}{\omega_k^2 - \omega^2 + i\gamma_k\omega}. \tag{31.20}$$

We have, finally, a complete expression which describes the index of refraction that is observed for many substances.* The index described by this formula varies with frequency roughly like the curve shown in Fig. 31–5.

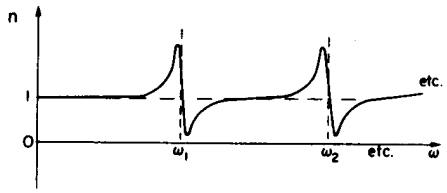


Fig. 31–5. The index of refraction as a function of frequency.

You will note that so long as ω is not too close to one of the resonant frequencies, the slope of the curve is positive. Such a positive slope is called "normal" dispersion (because it is clearly the most common occurrence). Very near the resonant frequencies, however, there is a small range of ω's for which the slope is negative. Such a negative slope is often referred to as "anomalous" (meaning abnormal) dispersion, because it seemed unusual when it was first observed, long before anyone even knew there were such things as electrons. From our point of view both slopes are quite "normal"!

31–4 Absorption

Perhaps you have noticed something a little strange about the last form (Eq. 31.20) we obtained for our dispersion equation. Because of the term $i\gamma$ we put in to take account of damping, the index of refraction is now a *complex number!* What does *that* mean? By working out what the real and imaginary parts of n are we could write

$$n = n' - in'', \tag{31.21}$$

where n' and n'' are real numbers. (We use the minus sign in front of the in'' because then n'' will turn out to be a positive number, as you can show for yourself.)

We can see what such a complex index means by going back to Eq. (31.6), which is the equation of the wave after it goes through a plate of material with an index n. If we put our complex n into this equation, and do some rearranging, we get

$$E_{\text{after plate}} = \underbrace{e^{-\omega n''\Delta z/c}}_{\text{A}} \underbrace{e^{-i\omega(n'-1)\Delta z/c} E_0 e^{i\omega(t-z/c)}}_{\text{B}} \tag{31.22}$$

The last factors, marked B in Eq. (31.22), are just the form we had before, and again describe a wave whose phase has been delayed by the angle $\omega(n' - 1)\,\Delta z/c$ in traversing the material. The first term (A) is new and is an exponential factor with a *real* exponent, because there were two i's that cancelled. Also, the exponent is negative, so the factor is a real number less than one. It describes a *decrease* in the magnitude of the field and, as we should expect, by an amount which is more the larger Δz is. As the wave goes through the material, it is weakened. The material is "absorbing" part of the wave. The wave comes out the other side with less energy. We should not be surprised at this, because the damping we put in for the oscillators is indeed a friction force and must be expected to cause a loss of energy. We see that the imaginary part n'' of a complex index of refraction represents an absorption (or "attenuation") of the wave. In fact, n'' is sometimes referred to as the "absorption index."

We may also point out that an imaginary part to the index n corresponds to bending the arrow E_a in Fig. 31–3 toward the origin. It is clear why the transmitted field is then decreased.

---

* Actually, although in quantum mechanics Eq. (31.20) is still valid, its interpretation is somewhat different. In quantum mechanics even an atom with one electron, like hydrogen, has several resonant frequencies. Therefore N_k is not really the number of electrons having the frequency ω_k, but is replaced instead by Nf_k, where N is the number of atoms per unit volume and f_k (called the oscillator strength) is a factor that tells how strongly the atom exhibits each of its resonant frequencies ω_k.

Normally, for instance as in glass, the absorption of light is very small. This is to be expected from our Eq. (31.20), because the imaginary part of the denominator, $i\gamma_k\omega$, is much smaller than the term $(\omega_k^2 - \omega^2)$. But if the light frequency ω is very close to ω_k then the resonance term $(\omega_k^2 - \omega^2)$ can become small compared with $i\gamma_k\omega$ and the index becomes almost completely imaginary. The absorption of the light becomes the dominant effect. It is just this effect that gives the dark lines in the spectrum of light which we receive from the sun. The light from the solar surface has passed through the sun's atmosphere (as well as the earth's), and the light has been strongly absorbed at the resonant frequencies of the atoms in the solar atmosphere.

The observation of such spectral lines in the sunlight allows us to tell the resonant frequencies of the atoms and hence the chemical composition of the sun's atmosphere. The same kind of observations tell us about the materials in the stars. From such measurements we know that the chemical elements in the sun and in the stars are the same as those we find on the earth.

31–5 The energy carried by an electric wave

We have seen that the imaginary part of the index means absorption. We shall now use this knowledge to find out how much energy is carried by a light wave. We have given earlier an argument that the energy carried by light is proportional to $\overline{E^2}$, the time average of the square of the electric field in the wave. The decrease in E due to absorption must mean a loss of energy, which would go into some friction of the electrons and, we might guess, would end up as heat in the material.

If we consider the light arriving on a unit area, say one square centimeter, of our plate in Fig. 31–1, then we can write the following energy equation (if we assume that energy is conserved, as we *do!*):

$$\text{Energy in per sec} = \text{energy out per sec} + \text{work done per sec.} \tag{31.23}$$

For the first term we can write $\alpha\overline{E_s^2}$, where α is the as yet unknown constant of proportionality which relates the average value of E^2 to the energy being carried. For the second term we must include the part from the radiating atoms of the material, so we should use $\alpha\overline{(E_s + E_a)^2}$, or (evaluating the square) $\alpha(\overline{E_s^2} + 2\overline{E_sE_a} + \overline{E_a^2})$.

All of our calculations have been made for a thin layer of material whose index is not too far from 1, so that E_a would always be much less than E_s (just to make the calculations easier). In keeping with our approximations, we should, therefore, leave out the term $\overline{E_a^2}$, because it is much smaller than $\overline{E_sE_a}$. You may say: "Then you should leave out $\overline{E_sE_a}$ also, because *it* is much smaller than $\overline{E_s^2}$." It is true that $\overline{E_sE_a}$ is much smaller than $\overline{E_s^2}$, but we must keep $\overline{E_sE_a}$ or our approximation will be the one that would apply if we neglected the presence of the material completely! One way of checking that our calculations are consistent is to see that we always keep terms which are proportional to $N\,\Delta z$, the area density of atoms in the material, but we leave out terms which are proportional to $(N\,\Delta z)^2$ or any higher power of $N\,\Delta z$. Ours is what should be called a "low-density approximation."

In the same spirit, we might remark that our energy equation has neglected the energy in the reflected wave. But that is OK because this term, too, is proportional to $(N\,\Delta z)^2$, since the amplitude of the reflected wave is proportional to $N\,\Delta z$.

For the last term in Eq. (31.23) we wish to compute the rate at which the incoming wave is doing work on the electrons. We know that work is force times distance, so the *rate* of doing work (also called power) is the force times the velocity. It is really $\mathbf{F}\cdot\mathbf{V}$, but we do not need to worry about the dot product when the velocity and force are along the same direction as they are here (except for a possible minus sign). So for each atom we take $\overline{q_eE_sv}$ for the average rate of

doing work. Since there are $N\,\Delta z$ atoms in a unit area, the last term in Eq. (31.23) should be $N\,\Delta z q_e\,\overline{E_s v}$. Our energy equation now looks like

$$\alpha\overline{E_s^2} = \alpha\overline{E_s^2} + 2\alpha\overline{E_s E_a} + N\,\Delta z\, q_e\,\overline{E_s v}. \tag{31.24}$$

The $\overline{E_s^2}$ terms cancel, and we have

$$2\alpha\overline{E_s E_a} = N\,\Delta z\, q_e\,\overline{E_s v}. \tag{31.25}$$

We now go back to Eq. (30.19), which tells us that for large z

$$E_a = \frac{N\,\Delta z q_e}{2\epsilon_0 c}\, v(\text{ret by } z/c) \tag{31.26}$$

(recalling that $\eta = N\,\Delta z$). Putting Eq. (31.26) into the left-hand side of (31.25), we get

$$2\alpha\,\frac{N\,\Delta z q_e}{2\epsilon_0 c}\,\overline{E_s(\text{at } z)\cdot v(\text{ret by } z/c)}.$$

However, E_s(at z) is E_s (at atoms) retarded by z/c. Since the average is independent of time, it is the same now as retarded by z/c, or is $\overline{E_s\ (\text{at atom})\cdot v}$, the same average that appears on the right-hand side of (31.25). The two sides are therefore equal if

$$\frac{\alpha}{\epsilon_0 c} = 1, \qquad \text{or} \qquad \alpha = \epsilon_0 c. \tag{31.27}$$

We have discovered that if energy is to be conserved, the energy carried in an electric wave per unit area and per unit time (or what we have called the *intensity*) must be given by $\epsilon_0 c\overline{E^2}$. If we call the intensity $\overline{S}$, we have

$$\overline{S} = \left\{\begin{matrix}\text{intensity}\\ \text{or}\\ \text{energy/area/time}\end{matrix}\right\} = \epsilon_0 c\overline{E^2}, \tag{31.28}$$

where the *bar* means the *time average*. We have a nice bonus result from our theory of the refractive index!

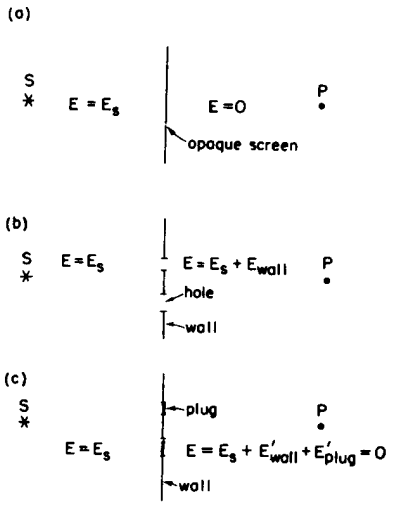


Fig. 31–6. Diffraction by a screen.

31–6 Diffraction of light by a screen

It is now a good time to take up a somewhat different matter which we can handle with the machinery of this chapter. In the last chapter we said that when you have an opaque screen and the light can come through some holes, the distribution of intensity—the diffraction pattern—could be obtained by imagining instead that the holes are replaced by sources (oscillators) uniformly distributed over the hole. In other words, the diffracted wave is the same as though the hole were a new source. We have to explain the reason for that, because the hole is, of course, just where there are *no* sources, where there are *no* accelerating charges.

Let us first ask: "What *is* an opaque screen?" Suppose we have a completely opaque screen between a source S and an observer at P, as in Fig. 31–6(a). If the screen is "opaque" there is no field at P. Why is there no field there? According to the basic principles we should obtain the field at P as the field E_s of the source delayed, plus the field from all the other charges around. But, as we have seen above, the charges in the screen will be set in motion by the field E_s, and these motions generate a new field which, if the screen is opaque, must *exactly cancel* the field E_s on the back side of the screen. You say: "What a miracle that it balances *exactly*! Suppose it was not exactly right!" If it were not exactly right (remember that this opaque screen has some thickness), the field toward the rear part of the screen would not be exactly zero. So, not being zero, it would set into motion some other charges in the material of the screen, and thus make a little more field, trying to get the total balanced out. So if we make the screen thick enough, there is no residual field, because there is enough opportunity to finally get the thing quieted down. In terms of our formulas above we would say that the

screen has a large and imaginary index, so the wave is absorbed exponentially as it goes through. You know, of course, that a thin enough sheet of the most opaque material, even gold, is transparent.

Now let us see what happens with an opaque screen which has holes in it, as in Fig. 31–6(b). What do we expect for the field at P? The field at P can be represented as a sum of two parts—the field due to the source S plus the field due to the wall, i.e., due to the motions of the charges in the walls. We might expect the motions of the charges in the walls to be complicated, but we can find out *what fields they produce* in a rather simple way.

Suppose that we were to take the same screen, but plug up the holes, as indicated in part (c) of the figure. We imagine that the plugs are of exactly the same material as the wall. Mind you, the plugs go where the holes were in case (b). Now let us calculate the field at P. The field at P is certainly zero in case (c), but it is *also* equal to the field from the source plus the field due to all the motions of the atoms in the walls and in the plugs. We can write the following equations:

$$\textit{Case (b)}: \quad E_{\text{at } P} = E_s + E_{\text{wall}},$$
$$\textit{Case (c)}: \quad E'_{\text{at } P} = 0 = E_s + E'_{\text{wall}} + E'_{\text{plug}},$$

where the primes refer to the case where the plugs are in place, but E_s is, of course, the same in both cases. Now if we subtract the two equations, we get

$$E_{\text{at } P} = (E_{\text{wall}} - E'_{\text{wall}}) - E'_{\text{plug}}.$$

Now if the holes are not too small (say many wavelengths across), we would not expect the presence of the plugs to change the fields which arrive at the walls except possibly for a little bit around the edges of the holes. Neglecting this small effect, we can set $E_{\text{wall}} = E'_{\text{wall}}$ and obtain that

$$E_{\text{at } P} = -E'_{\text{plug}}.$$

We have the result that the field at P *when there are holes* in a screen (case *b*) is the same (except for sign) as the field that is produced by *that part* of a complete opaque wall which is *located where the holes are!* (The sign is not too interesting, since we are usually interested in intensity which is proportional to the square of the field.) It seems like an amazing backwards-forwards argument. It is, however, not only true (approximately for not too small holes), but useful, and is the justification for the usual theory of diffraction.

The field E'_{plug} is computed in any particular case by remembering that the motion of the charges *everywhere* in the screen is just that which will cancel out the field E_s on the back of the screen. Once we know these motions, we add the radiation fields at P due just to the charges in the plugs.

We remark again that this theory of diffraction is only approximate, and will be good only if the holes are not too small. For holes which are too small the E'_{plug} term will be small and then the difference between E'_{wall} and E_{wall} (which difference we have taken to be zero) may be comparable to or larger than the small E'_{plug} term, and our approximation will no longer be valid.

32

Radiation Damping. Light Scattering

32–1 Radiation resistance

32–2 The rate of radiation of energy

32–3 Radiation damping

32–4 Independent sources

32–5 Scattering of light

32–1 Radiation resistance

In the last chapter we learned that when a system is oscillating, energy is carried away, and we deduced a formula for the energy which is radiated by an oscillating system. If we know the electric field, then the average of the square of the field times $\epsilon_0 c$ is the amount of energy that passes per square meter per second through a surface normal to the direction in which the radiation is going:

$$S = \epsilon_0 c \langle E^2 \rangle. \tag{32.1}$$

Any oscillating charge radiates energy; for instance, a driven antenna radiates energy. If the system radiates energy, then in order to account for the conservation of energy we must find that power is being delivered along the wires which lead into the antenna. That is, to the driving circuit the antenna acts like a *resistance*, or a place where energy can be "lost" (the energy is not really lost, it is really radiated out, but so far as the circuit is concerned, the energy is lost). In an ordinary resistance, the energy which is "lost" passes into heat; in this case the energy which is "lost" goes out into space. But from the standpoint of circuit theory, without considering *where* the energy goes, the net effect on the circuit is the same—energy is "lost" from that circuit. Therefore the antenna appears to the generator as having a resistance, even though it may be made with perfectly good copper. In fact, if it is well built it will appear as almost a pure resistance, with very little inductance or capacitance, because we would like to radiate as much energy as possible out of the antenna. This resistance that an antenna shows is called the *radiation resistance*.

If a current I is going to the antenna, then the average rate at which power is delivered to the antenna is the average of the square of the current times the resistance. The rate at which power is *radiated* by the antenna is proportional to the square of the current in the antenna, of course, because all the fields are proportional to the currents, and the energy liberated is proportional to the square of the field. The coefficient of proportionality between radiated power and $\langle I^2 \rangle$ is the radiation resistance.

An interesting question is, what is this radiation resistance due to? Let us take a simple example: let us say that currents are driven up and down in an antenna. We find that we have to put work in, if the antenna is to radiate energy. If we take a charged body and accelerate it up and down it radiates energy; if it were not charged it would not radiate energy. It is one thing to calculate from the conservation of energy that energy is lost, but another thing to answer the question, *against what force* are we doing the work? That is an interesting and very difficult question which has never been completely and satisfactorily answered for electrons, although it has been for antennas. What happens is this: in an antenna, the fields produced by the moving charges in one part of the antenna react on the moving charges in another part of the antenna. We can calculate these forces and find out how much work they do, and so find the right rule for the radiation resistance. When we say "We can calculate—" that is not quite right—*we* cannot, because we have not yet studied the laws of electricity at short distances; only at large distances do we know what the electric field is. We saw the formula (28.3), but at present it is too complicated for *us* to calculate the fields inside the wave zone. Of course, since conservation of energy is valid, we can calculate the result all right without knowing the fields at short distances. (As a matter of fact, by using this argument backwards it turns out that one can find the formula for the forces at short

distances only by knowing the field at very large distances, by using the laws of conservation of energy, but we shall not go into that here.)

The problem in the case of a single electron is this: if there is only one charge, what can the force act on? It has been proposed, in the old classical theory, that the charge was a little ball, and that one part of the charge acted on the other part. Because of the delay in the action across the tiny electron, the force is not exactly in phase with the motion. That is, if we have the electron standing still, we know that "action equals reaction." So the various internal forces are equal, and there is no net force. But if the electron is accelerating, then because of the time delay across it, the force which is acting on the front from the back is not exactly the same as the force on the back from the front, because of the delay in the effect. This delay in the timing makes for a lack of balance, so, as a net effect, the thing holds itself back by its bootstraps! This model of the origin of the resistance to acceleration, the radiation resistance of a moving charge, has run into many difficulties, because our present view of the electron is that it is *not* a "little ball"; this problem has never been solved. Nevertheless we can calculate exactly, of course, what the net radiation resistance force must be, i.e., how much loss there must be when we accelerate a charge, in spite of not knowing directly the mechanism of how that force works.

32–2 The rate of radiation of energy

Now we shall calculate the total energy radiated by an accelerating charge. To keep the discussion general, we shall take the case of a charge accelerating any which way, but nonrelativistically. At a moment when the acceleration is, say, vertical, we know that the electric field that is generated is the charge multiplied by the projection of the retarded acceleration, divided by the distance. So we know the electric field at any point, and we therefore know the square of the electric field and thus the energy $\epsilon_0 c E^2$ leaving through a unit area per second.

The quantity $\epsilon_0 c$ appears quite often in expressions involving radiowave propagation. Its reciprocal is called the *impedance of a vacuum*, and it is an easy number to remember: it has the value $1/\epsilon_0 c = 377$ ohms. So the power in watts per square meter is equal to the average of the field squared, divided by 377.

Using our expression (29.1) for the electric field, we find that

$$S = \frac{q^2 a'^2 \sin^2 \theta}{16\pi^2 \epsilon_0 r^2 c^3} \tag{32.2}$$

is the power per square meter radiated in the direction θ. We notice that it goes inversely as the square of the distance, as we said before. Now suppose we wanted the total energy radiated in all directions: then we must integrate (32.2) over all directions. First we multiply by the area, to find the amount that flows within a little angle $d\theta$ (Fig. 32–1). We need the area of a spherical section. The way to think of it is this: if r is the radius, then the width of the annular segment is $r\,d\theta$, and the circumference is $2\pi r \sin\theta$, because $r \sin\theta$ is the radius of the circle. So the area of the little piece of the sphere is $2\pi r \sin\theta$ times $r\,d\theta$:

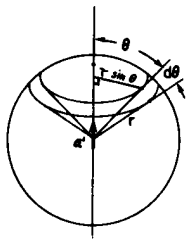


Fig. 32–1. The area of a spherical segment is $2\pi r \sin\theta \cdot r\,d\theta$.

$$dA = 2\pi r^2 \sin\theta\, d\theta. \tag{32.3}$$

By multiplying the flux [(32.2), the power per square meter] by the area in square meters included in the small angle $d\theta$, we find the amount of energy that is liberated in this direction between θ and $\theta + d\theta$; then we integrate that over all the angles θ from 0 to 180°:

$$P = \int S\,dA = \frac{q^2 a'^2}{8\pi\epsilon_0 c^3} \int_0^\pi \sin^3\theta\, d\theta. \tag{32.4}$$

By writing $\sin^3\theta = (1 - \cos^2\theta)\sin\theta$ it is not hard to show that $\int_0^\pi \sin^3\theta\, d\theta = 4/3$. Using that fact, we finally get

$$P = \frac{q^2 a'^2}{6\pi\epsilon_0 c^3}. \tag{32.5}$$

This expression deserves some remarks. First of all, since the vector $\mathbf{a}'$ had a certain direction, the a'^2 in (32.5) would be the square of the vector $\mathbf{a}'$, that is, $\mathbf{a}' \cdot \mathbf{a}'$, the length of the vector, squared. Secondly, the flux (32.2) was calculated using the retarded acceleration; that is, the acceleration at the time at which the energy now passing through the sphere was radiated. We might like to say that this energy was in fact liberated at this earlier time. This is not exactly true; it is only an approximate idea. The exact time when the energy is liberated cannot be defined precisely. All we can really calculate precisely is what happens in a complete motion, like an oscillation or something, where the acceleration finally ceases. Then what we find is that the total energy flux per cycle is the average of acceleration squared, for a complete cycle. This is what should really appear in (32.5). Or, if it is a motion with an acceleration that is initially and finally zero, then the total energy that has flown out is the time integral of (32.5).

To illustrate the consequences of formula (32.5) when we have an oscillating system, let us see what happens if the displacement x of the charge is oscillating so that the acceleration a is $-\omega^2 x_0 e^{i\omega t}$. The average of the acceleration squared over a cycle (remember that we have to be very careful when we square things that are written in complex notation—it really is the cosine, and the average of $\cos^2 \omega t$ is one-half) thus is

$$\langle a'^2 \rangle = \tfrac{1}{2}\omega^4 x_0^2.$$

Therefore

$$P = \frac{q^2\omega^4 x_0^2}{12\pi\epsilon_0 c^3}. \tag{32.6}$$

The formulas we are now discussing are relatively advanced and more or less modern; they date from the beginning of the twentieth century, and they are very famous. Because of their historical value, it is important for us to be able to read about them in older books. In fact, the older books also used a system of units different from our present mks system. However, all these complications can be straightened out in the final formulas dealing with electrons by the following rule: The quantity $q_e^2/4\pi\epsilon_0$, where q_e is the electronic charge (in coulombs), has, historically, been written as e^2. It is very easy to calculate that e in the mks system is numerically equal to 1.5188×10^{-14}, because we know that, numerically, $q_e = 1.60206 \times 10^{-19}$ and $1/4\pi\epsilon_0 = 8.98748 \times 10^9$. Therefore we shall often use the convenient abbreviation

$$e^2 = \frac{q_e^2}{4\pi\epsilon_0}. \tag{32.7}$$

If we use the above numerical value of e in the older formulas and treat them as though they were written in mks units, we will get the right numerical results. For example, the older form of (32.5) is $P = \frac{2}{3}e^2a^2/c^3$. Again, the potential energy of a proton and an electron at distance r is $q_e^2/4\pi\epsilon_0 r$ or e^2/r, with $e = 1.5188 \times 10^{-14}$ mks.

32–3 Radiation damping

Now the fact that an oscillator loses a certain energy would mean that if we had a charge on the end of a spring (or an electron in an atom) which has a natural frequency ω_0, and we start it oscillating and let it go, it will not oscillate forever, even if it is in empty space millions of miles from anything. There is no oil, no resistance, in an ordinary sense; no "viscosity." But nevertheless it will not oscillate, as we might once have said, "forever," because if it is charged it is radiating energy, and therefore the oscillation will slowly die out. How slowly? What is the Q of such an oscillator, caused by the electromagnetic effects, the so-called radiation resistance or radiation damping of the oscillator? The Q of any oscillating system is the total energy content of the oscillator at any time divided by the energy loss per radian:

$$Q = \frac{W}{dW/d\phi}.$$

Or (another way to write it), since $dW/d\phi = (dW/dt)/(d\phi/dt) = (dW/dt)/\omega$,

$$Q = \frac{\omega W}{dW/dt}. \tag{32.8}$$

If for a given Q this tells us how the energy of the oscillation dies out, $dW/dt = -(\omega/Q)W$, which has the solution $W = W_0 e^{-\omega t/Q}$ if W_0 is the initial energy (at $t = 0$).

To find the Q for a radiator, we go back to (32.8) and use (32.6) for dW/dt.

Now what do we use for the energy W of the oscillator? The kinetic energy of the oscillator is $\frac{1}{2}mv^2$, and the *mean* kinetic energy is $m\omega^2 x_0^2/4$. But we remember that for the total energy of an oscillator, on the average half is kinetic and half is potential energy, and so we double our result, and find for the total energy of the oscillator

$$W = \tfrac{1}{2}m\omega^2 x_0^2. \tag{32.9}$$

What do we use for the frequency in our formulas? We use the natural frequency ω_0 because, for all practical purposes, that is the frequency at which our atom is radiating, and for m we use the electron mass m_e. Then, making the necessary divisions and cancellations, the formula comes down to

$$\frac{1}{Q} = \frac{4\pi e^2}{3\lambda m_e c^2}. \tag{32.10}$$

(In order to see it better and in a more historical form we write it using our abbreviation $q_e^2/4\pi\epsilon_0 = e^2$, and the factor ω_0/c which was left over has been written as $2\pi/\lambda$.) Since Q is dimensionless, the combination $e^2/m_e c^2$ must be a property only of the electron charge and mass, an intrinsic property of the electron, and it must be a *length*. It has been given a name, the *classical electron radius*, because the early atomic models, which were invented to explain the radiation resistance on the basis of the force of one part of the electron acting on the other parts, all needed to have an electron whose dimensions were of this general order of magnitude. However, this quantity no longer has the significance that we believe that the electron really *has* such a radius. Numerically, the magnitude of the radius is

$$r_0 = \frac{e^2}{m_e c^2} = 2.82 \times 10^{-15}\ \text{m}. \tag{32.11}$$

Now let us actually calculate the Q of an atom that is emitting light—let us say a sodium atom. For a sodium atom, the wavelength is roughly 6000 angstroms, in the yellow part of the visible spectrum, and this is a typical wavelength. Thus

$$Q = \frac{3\lambda}{4\pi r_0} \approx 5 \times 10^7, \tag{32.12}$$

so the Q of an atom is of the order 10^8. This means that an atomic oscillator will oscillate for 10^8 radians or about 10^7 oscillations, before its energy falls by a factor $1/e$. The frequency of oscillation of light corresponding to 6000 angstroms, $\nu = c/\lambda$, is on the order of 10^{15} cycles/sec, and therefore the lifetime, the time it takes for the energy of a radiating atom to die out by a factor $1/e$, is on the order of 10^{-8} sec. In ordinary circumstances, freely emitting atoms usually take about this long to radiate. This is valid only for atoms which are in empty space, not being disturbed in any way. If the electron is in a solid and it has to hit other atoms or other electrons, then there are additional resistances and different damping.

The effective resistance term γ in the resistance law for the oscillator can be found from the relation $1/Q = \gamma/\omega_0$, and we remember that the size of γ determines how wide the resonance curve is (Fig. 23–2). Thus we have just computed the *widths of spectral lines* for freely radiating atoms! Since $\lambda = 2\pi c/\omega$, we find that

$$\begin{aligned}\Delta\lambda &= 2\pi c\,\Delta\omega/\omega^2 = 2\pi c\gamma/\omega_0^2 = 2\pi c/Q\omega_0 \\ &= \lambda/Q = 4\pi r_0/3 = 1.18 \times 10^{-14}\ \text{m}.\end{aligned} \tag{32.13}$$

32–4 Independent sources

In preparation for our second topic, the scattering of light, we must now discuss a certain feature of the phenomenon of interference that we neglected to discuss previously. This is the question of when interference does *not* occur. If we have two sources S_1 and S_2, with amplitudes A_1 and A_2, and we make an observation in a certain direction in which the phases of arrival of the two signals are ϕ_1 and ϕ_2 (a combination of the actual time of oscillation and the delayed time, depending on the position of observation), then the energy that we receive can be found by compounding the two complex number vectors A_1 and A_2, one at angle ϕ_1 and the other at angle ϕ_2 (as we did in Chapter 30) and we find that the resultant energy is proportional to

$$A_R^2 = A_1^2 + A_2^2 + 2A_1A_2\cos(\phi_1 - \phi_2). \tag{32.14}$$

Now if the cross term $2A_1A_2\cos(\phi_1 - \phi_2)$ were not there, then the total energy that would be received in a given direction would simply be the sum of the energies, $A_1^2 + A_2^2$, that would be liberated by each source separately, which is what we usually expect. That is, the combined intensity of light shining on something from two sources is the sum of the intensities of the two lights. On the other hand, if we have things set just right and we have a cross term, it is not such a sum, because there is also some interference. If there are circumstances in which this term is of no importance, then we would say the interference is apparently lost. Of course, in nature it is always there, but we may not be able to detect it.

Let us consider some examples. Suppose, first, that the two sources are 7,000,000,000 wavelengths apart, not an impossible arrangement. Then in a given direction it is true that there is a very definite value of these phase differences. But, on the other hand, if we move just a hair in one direction, a few wavelengths, which is no distance at all (our eye already has a hole in it that is so large that we are averaging the effects over a range very wide compared with one wavelength) then we change the relative phase, and the cosine changes very rapidly. If we take the *average* of the intensity over a little region, then the cosine, which goes plus, minus, plus, minus, as we move around, averages to zero.

So if we average over regions where the phase varies very rapidly with position, we get no interference.

Another example. Suppose that the two sources are two independent radio oscillators—not a single oscillator being fed by two wires, which guarantees that the phases are kept together, but two independent sources—and that they are not *precisely* tuned at the same frequency (it is very hard to make them at exactly the same frequency without actually wiring them together). In this case we have what we call two *independent* sources. Of course, since the frequencies are not exactly equal, although they started in phase, one of them begins to get a little ahead of the other, and pretty soon they are out of phase, and then it gets still further ahead, and pretty soon they are in phase again. So the phase difference between the two is gradually drifting with time, but if our observation is so crude that we cannot see that little time, if we average over a much longer time, then although the intensity swells and falls like what we call "beats" in sound, if these swellings and fallings are too rapid for our equipment to follow, then again this term averages out.

In other words, in any circumstance in which the phase shift averages out, we get no interference!

One finds many books which say that two distinct light sources never interfere. This is not a statement of physics, but is merely a statement of the degree of sensitivity of the technique of the experiments at the time the book was written. What happens in a light source is that first one atom radiates, then another atom radiates, and so forth, and we have just seen that atoms radiate a train of waves only for about 10^{-8} sec; after 10^{-8} sec, some atom has probably taken over, then another atom takes over, and so on. So the phases can really only stay the same for about 10^{-8} sec. Therefore, if we average for very much more than 10^{-8} sec, we do not see an interference from two different sources, because they cannot hold their phases steady for longer than 10^{-8} sec. With photocells, very high-speed detection

is possible, and one can show that there is an interference which varies with time, up and down, in about 10^{-8} sec. But most detection equipment, of course, does not look at such fine time intervals, and thus sees no interference. Certainly with the eye, which has a tenth-of-a-second averaging time, there is no chance whatever of seeing an interference between two different ordinary sources.

Recently it has become possible to make light sources which get around this effect by making all the atoms emit *together* in time. The device which does this is a very complicated thing, and has to be understood in a quantum-mechanical way. It is called a *laser*, and it is possible to produce from a laser a source in which the interference frequency, the time at which the phase is kept constant, is very much longer than 10^{-8} sec. It can be of the order of a hundredth, a tenth, or even one second, and so, with ordinary photocells, one can pick up the frequency between two different lasers. One can easily detect the pulsing of the beats between two laser sources. Soon, no doubt, someone will be able to demonstrate two sources shining on a wall, in which the beats are so slow that one can *see* the wall get bright and dark!

Another case in which the interference averages out is that in which, instead of having only *two* sources, we have *many*. In this case, we would write the expression for A_R^2 as the sum of a whole lot of amplitudes, complex numbers, squared, and we would get the square of each one, all added together, plus cross terms between every pair, and if the circumstances are such that the latter average out, then there will be no effects of interference. It may be that the various sources are located in such random positions that, although the phase difference between A_2 and A_3 is also definite, it is very different from that between A_1 and A_2, etc. So we would get a whole lot of cosines, many plus, many minus, all averaging out.

So it is that in many circumstances we do not see the effects of interference, but see only a collective, total intensity equal to the sum of all the intensities.

32–5 Scattering of light

The above leads us to an effect which occurs in air as a consequence of the irregular positions of the atoms. When we were discussing the index of refraction, we saw that an incoming beam of light will make the atoms radiate again. The electric field of the incoming beam drives the electrons up and down, and they radiate because of their acceleration. This scattered radiation combines to give a beam in the same direction as the incoming beam, but of somewhat different phase, and this is the origin of the index of refraction.

But what can we say about the amount of re-radiated light in some other direction? Ordinarily, if the atoms are very beautifully located in a nice pattern, it is easy to show that we get nothing in other directions, because we are adding a lot of vectors with their phases always changing, and the result comes to zero. But if the objects are *randomly* located, then the total intensity in any direction is the *sum* of the intensities that are scattered by each atom, as we have just discussed. Furthermore, the atoms in a gas are in actual motion, so that although the relative phase of two atoms is a definite amount now, later the phase would be quite different, and therefore *each* cosine term will average out. Therefore, to find out how much light is scattered in a given direction by a gas, we merely study the effects of *one atom* and multiply the intensity it radiates by the number of atoms.

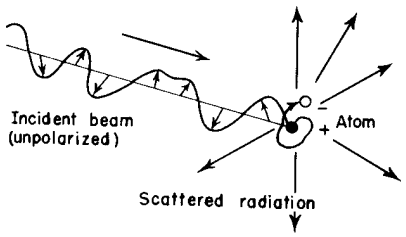


Fig. 32–2. A beam of radiation falls on an atom and causes the charges (electrons) in the atom to move. The moving electrons in turn radiate in various directions.

Earlier, we remarked that the phenomenon of scattering of light of this nature is the origin of the blue of the sky. The sunlight goes through the air, and when we look to one side of the sun—say at 90° to the beam—we see blue light; what we now have to calculate is *how much* light we see and *why it is blue*.

If the incident beam has the electric field $\mathbf{E} = \mathbf{E}_0 e^{i\omega t}$ at the point where the atom is located, we know that an electron in the atom will vibrate up and down in response to this $\mathbf{E}$ (Fig. 32–2). From Eq. (23.8), the amplitude will be

$$\hat{\mathbf{x}} = \frac{q_e \mathbf{E}_0}{m(\omega_0^2 - \omega^2 + i\omega\gamma)}. \tag{32.15}$$

We could include the damping and the possibility that the atom acts like several oscillators of different frequency and sum over the various frequencies, but for simplicity let us just take one oscillator and neglect the damping. Then the response to the external electric field, which we have already used in the calculation of the index of refraction, is simply

$$\hat{\mathbf{x}} = \frac{q_e\hat{\mathbf{E}}_0}{m(\omega_0^2 - \omega^2)}. \tag{32.16}$$

We could now easily calculate the intensity of light that is emitted in various directions, using formula (32.2) and the acceleration corresponding to the above $\hat{\mathbf{x}}$.

Rather than do this, however, we shall simply calculate the *total amount* of light scattered in *all* directions, just to save time. The total amount of light energy per second, scattered in all directions by the single atom, is of course given by Eq. (32.7) So, putting together the various pieces and regrouping them, we get

$$\begin{aligned} P &= [(q_e^2\omega^4/12\pi\epsilon_0c^3)q_e^2E_0^2/m_e^2(\omega^2 - \omega_0^2)^2] \\ &= (\tfrac{1}{2}\epsilon_0cE_0^2)(8\pi/3)(q_e^4/16\pi^2\epsilon_0^2m_e^2c^4)[\omega^4/(\omega^2 - \omega_0^2)^2] \\ &= (\tfrac{1}{2}\epsilon_0cE_0^2)(8\pi r_0^2/3)[\omega^4/(\omega^2 - \omega_0^2)^2] \end{aligned} \tag{32.17}$$

for the total scattered power, radiated in all directions.

We have written the result in the above form because it is then easy to remember: First, the total energy that is scattered is proportional to the square of the incident field. What does that mean? Obviously, the square of the incident field is proportional to the energy which is coming in per second. In fact, the energy incident per square meter per second is ϵ_0c times the average $\langle E^2\rangle$ of the square of the electric field, and if E_0 is the maximum value of E, then $\langle E^2\rangle = \frac{1}{2}E_0^2$. In other words, the total energy scattered is proportional to the energy per square meter that comes in; the brighter the sunlight that is shining in the sky, the brighter the sky is going to look.

Next, what *fraction* of the incoming light is scattered? Let us imagine a "target" with a certain area, let us say σ, in the beam (not a real, material target, because this would diffract light, and so on; we mean an imaginary area drawn in space). The total amount of energy that would pass through this surface σ in a given circumstance is proportional both to the incoming intensity and to σ, and would be

$$P = (\tfrac{1}{2}\epsilon_0cE_0^2)\sigma. \tag{32.18}$$

Now we invent an idea: we say that the atom scatters a total amount of intensity which is the amount which would fall on a certain geometrical area, and we give the answer by giving that area. That answer, then, is independent of the incident intensity; it gives the ratio of the energy scattered to the energy incident per square meter. In other words, the ratio

$$\frac{\text{total energy scattered per second}}{\text{energy incident per square meter per second}} \text{ is an } \textit{area}.$$

The significance of this area is that, if all the energy that impinged on that area were to be spewed in all directions, then that is the amount of energy that would be scattered by the atom.

This area is called a *cross section for scattering;* the idea of cross section is used constantly, whenever some phenomenon occurs in proportion to the intensity of a beam. In such cases one always describes the amount of the phenomenon by saying what the effective area would have to be to pick up that much of the beam. It does not mean in any way that this oscillator actually *has* such an area. If there were nothing present but a free electron shaking up and down there would be no area directly associated with it, physically. It is merely a way of expressing the answer to a certain kind of problem; it tells us what area the incident beam would

have to hit in order to account for that much energy coming off. Thus, for our case,

$$\sigma_s = \frac{8\pi r_0^2}{3} \cdot \frac{\omega^4}{(\omega^2 - \omega_0^2)^2} \tag{32.19}$$

(the subscript s is for "scattering").

Let us look at some examples. First, if we go to a very low natural frequency ω_0, or to completely unbound electrons, for which $\omega_0 = 0$, then the frequency ω cancels out and the cross section is a constant. This low-frequency limit, or the free electron cross section, is known as the *Thompson scattering cross section.* It is an area whose dimensions are approximately 10^{-15} meter, more or less, on a side, i.e., 10^{-30} square meter, which is rather small!

On the other hand, if we take the case of light in the air, we remember that for air the natural frequencies of the oscillators are higher than the frequency of the light that we use. This means that, to a first approximation, we can disregard ω^2 in the denominator, and we find that the scattering is proportional to the *fourth power* of the frequency. That is to say, light which is of higher frequency by, say, a factor of two, is *sixteen times* more intensely scattered, which is a quite sizable difference. This means that blue light, which has about twice the frequency of the reddish end of the spectrum, is scattered to a far greater extent than red light. Thus when we look at the sky it looks that glorious blue that we see all the time!

There are several points to be made about the above results. One interesting question is, why do we ever see the *clouds*? Where do the clouds come from? Everybody knows it is the condensation of water vapor. But, of course, the water vapor is already in the atmosphere *before* it condenses, so why don't we see it then? After it condenses it is perfectly obvious. It wasn't there, now it *is* there. So the mystery of where the clouds come from is not really such a childish mystery as "Where does the water come from, Daddy?," but has to be explained.

We have just explained that every atom scatters light, and of course the water vapor will scatter light, too. The mystery is why, when the water is condensed into clouds, does it scatter such a *tremendously greater amount* of light?

Consider what would happen if, instead of a single atom, we had an agglomerate of atoms, say two, very close together compared with the wavelength of the light. Remember, atoms are only an angstrom or so across, while the wavelength of light is some 5000 angstroms, so when they form a clump, a few atoms together, they can *be* very close together compared with the wavelength of light. Then when the electric field acts, *both of the atoms will move together.* The electric field that is scattered will then be the sum of the two electric fields in phase, i.e., double the amplitude that there was with a single atom, and the *energy* which is scattered is therefore *four times* what it is with a single atom, not twice! So lumps of atoms radiate or scatter more energy than they do as single atoms. Our argument that the phases are independent is based on the assumption that there is a real and large difference in phase between any two atoms, which is true only if they are several wavelengths apart and randomly spaced, or moving. But if they are right next to each other, they necessarily scatter in phase, and they have a coherent interference which produces an increase in the scattering.

If we have N atoms in a lump, which is a tiny droplet of water, then each one will be driven by the electric field in about the same way as before (the effect of one atom on the other is not important; it is just to get the idea anyway) and the amplitude of scattering from each one is the same, so the total field which is scattered is N-fold increased. The *intensity* of the light which is scattered is then the square, or N^2-fold, increased. We would have expected, if the atoms were spread out in space, only N times as much as 1, whereas we get N^2 times as much as 1! That is to say, the scattering of water in lumps of N molecules each is N times more intense than the scattering of the single atoms. So as the water agglomerates the scattering increases. Does it increase *ad infinitum*? No! When does this analysis begin to fail? How many atoms can we put together before we cannot drive this argument any further? *Answer*: If the water drop gets so big that from one end to the other is a wavelength or so, then the atoms are no longer all in

phase because they are too far apart. So as we keep increasing the size of the droplets we get more and more scattering, until such a time that a drop gets about the size of a wavelength, and then the scattering does not increase anywhere nearly as rapidly as the drop gets bigger. Furthermore, the blue disappears, because for long wavelengths the drops can be bigger, before this limit is reached, than they can be for short wavelengths. Although the short waves scatter more per atom than the long waves, there is a bigger enhancement for the red end of the spectrum than for the blue end when all the drops are bigger than the wavelength, so the color is shifted from the blue toward the red.

Now we can make an experiment that demonstrates this. We can make particles that are very small at first, and then gradually grow in size. We use a solution of sodium thiosulfate (hypo) with sulphuric acid, which precipitates very fine grains of sulphur. As the sulphur precipitates, the grains first start very small, and the scattering is a little bluish. As it precipitates more it gets more intense, and then it will get whitish as the particles get bigger. In addition, the light which goes straight through will have the blue taken out. That is why the sunset is red, of course, because the light that comes through a lot of air to the eye has had a lot of blue light scattered out, so it is yellow-red.

Finally, there is one other important feature which really belongs in the next chapter, on polarization, but it is so interesting that we point it out now. This is that the electric field of the scattered light tends to vibrate in a particular direction. The electric field in the incoming light is oscillating in some way, and the driven oscillator goes in this same direction, and if we are situated about at right angles to the beam, we will see *polarized* light, that is to say, light in which the electric field is going only one way. In general, the atoms can vibrate in any direction at right angles to the beam, but if they are driven directly toward or away from us, we do not see it. So if the incoming light has an electric field which changes and oscillates in any direction, which we call unpolarized light, then the light which is coming out at 90° to the beam vibrates in only one direction! (See Fig. 32–3.)

There is a substance called polaroid which has the property that when light goes through it, only the piece of the electric field which is along one particular axis can get through. We can use this to test for polarization, and indeed we find the light scattered by the hypo solution to be strongly polarized.

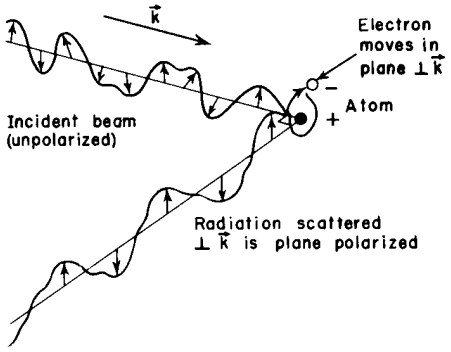


Fig. 32–3. Illustration of the origin of the polarization of radiation scattered at right angles to the incident beam.

33

Polarization

33–1 The electric vector of light

33–2 Polarization of scattered light

33–3 Birefringence

33–4 Polarizers

33–5 Optical activity

33–6 The intensity of reflected light

33–7 Anomalous refraction

33–1 The electric vector of light

In this chapter we shall consider those phenomena which depend on the fact that the electric field that describes the light is a vector. In previous chapters we have not been concerned with the direction of oscillation of the electric field, except to note that the electric vector lies in a plane perpendicular to the direction of propagation. The particular direction in this plane has not concerned us. We now consider those phenomena whose central feature is the particular direction of oscillation of the electric field.

In ideally monochromatic light, the electric field must oscillate at a definite frequency, but since the x-component and the y-component can oscillate independently at a definite frequency, we must first consider the resultant effect produced by superposing two independent oscillations at right angles to each other. What kind of electric field is made up of an x-component and a y-component which oscillate at the same frequency? If one adds to an x-vibration a certain amount of y-vibration at the same phase, the result is a vibration in a new direction in the xy-plane. Figure 33–1 illustrates the superposition of different amplitudes for the x-vibration and the y-vibration. But the resultants shown in Fig. 33–1 are not the only possibilities; in all of these cases we have assumed that the x-vibration and the y-vibration are *in phase*, but it does not have to be that way. It could be that the x-vibration and the y-vibration are out of phase.

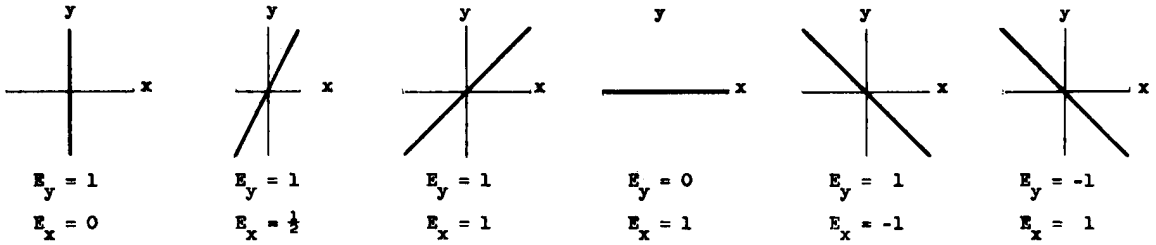


Fig. 33–1. Superposition of x-vibrations and y-vibrations in phase.

When the x-vibration and the y-vibration are not in phase, the electric field vector moves around in an ellipse, and we can illustrate this in a familiar way. If we hang a ball from a support by a long string, so that it can swing freely in a horizontal plane, it will execute sinusoidal oscillations. If we imagine horizontal x- and y-coordinates with their origin at the rest position of the ball, the ball can swing in either the x- or y-direction with the same pendulum frequency. By selecting the proper initial displacement and initial velocity, we can set the ball in oscillation along either the x-axis or the y-axis, or along any straight line in the xy-plane. These motions of the ball are analogous to the oscillations of the electric field vector illustrated in Fig. 33–1. In each instance, since the x-vibrations and the y-vibrations reach their maxima and minima at the same time, the x- and y-oscillations are in phase. But we know that the most general motion of the ball is motion in an ellipse, which corresponds to oscillations in which the x- and y-directions are *not* in the same phase. The superposition of x- and y-vibrations which are not in phase is illustrated in Fig. 33–2 for a variety of angles between the phase of the x-vibration and that of the y-vibration. The general result is that the electric vector moves around an ellipse. The motion in a straight line is a particular

case corresponding to a phase difference of zero (or an integral multiple of π); motion in a circle corresponds to equal amplitudes with a phase difference of 90° (or any odd integral multiple of $\pi/2$).

In Fig. 33–2 we have labeled the electric field vectors in the x- and y-directions with complex numbers, which are a convenient representation in which to express the phase difference. Do not confuse the real and imaginary components of the complex electric vector in this notation with the x- and y-coordinates of the field. The x- and y-coordinates plotted in Fig. 33–1 and Fig. 33–2 are actual electric fields that we can measure. The real and imaginary components of a complex electric field vector are only a mathematical convenience and have no physical significance.

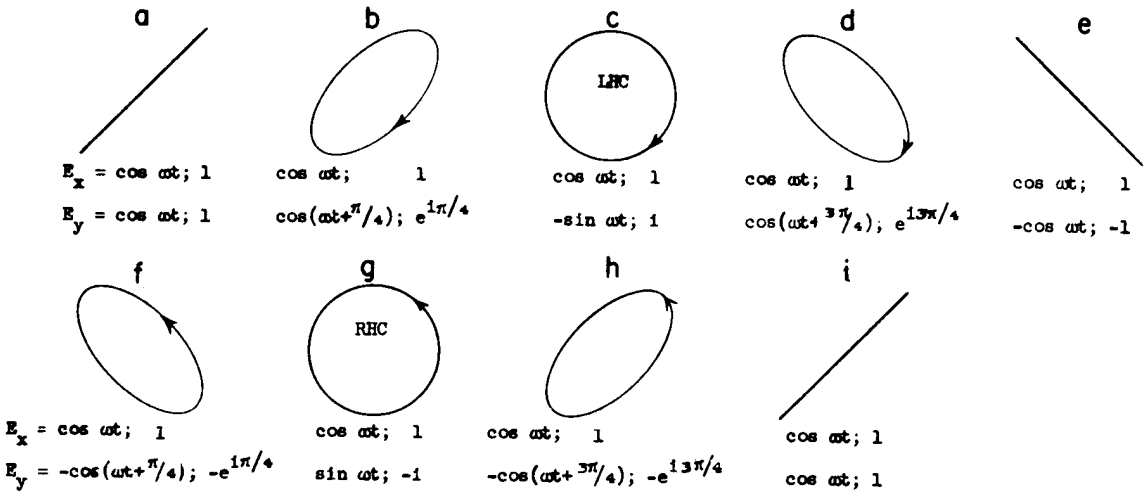


Fig. 33–2. Superposition of x-vibrations and y-vibrations with equal amplitudes but various relative phases. The components E_x and E_y are expressed in both real and complex notations.

Now for some terminology. Light is *linearly polarized* (sometimes called plane polarized) when the electric field oscillates on a straight line; Fig. 33–1 illustrates linear polarization. When the end of the electric field vector travels in an ellipse, the light is *elliptically polarized.* When the end of the electric field vector travels around a circle, we have *circular polarization.* If the end of the electric vector, when we look at it as the light comes straight toward us, goes around in a counterclockwise direction, we call it right-hand circular polarization. Figure 33–2(g) illustrates right-hand circular polarization, and Fig. 33–2(c) shows left-hand circular polarization. In both cases the light is coming out of the paper. Our convention for labeling left-hand and right-hand circular polarization is consistent with that which is used today for all the other particles in physics which exhibit polarization (e.g., electrons). However, in some books on optics the opposite conventions are used, so one must be careful.

We have considered linearly, circularly, and elliptically polarized light, which covers everything except for the case of *unpolarized* light. Now how can the light be unpolarized when we know that it must vibrate in one or another of these ellipses? If the light is not absolutely monochromatic, or if the x- and y-phases are not kept perfectly together, so that the electric vector first vibrates in one direction, then in another, the polarization is constantly changing. Remember that one atom emits during 10^{-8} sec, and if one atom emits a certain polarization, and then another atom emits light with a different polarization, the polarizations will change every 10^{-8} sec. If the polarization changes more rapidly than we can detect it, then we call the light unpolarized, because all the effects of the polarization average out. None of the interference effects of polarization would show up with unpolarized light. But as we see from the definition, light is unpolarized only if we are unable to find out whether the light is polarized or not.

33–2 Polarization of scattered light

The first example of the polarization effect that we have already discussed is the scattering of light. Consider a beam of light, for example from the sun, shining on the air. The electric field will produce oscillations of charges in the air, and motion of these charges will radiate light with its maximum intensity in a plane normal to the direction of vibration of the charges. The beam from the sun is unpolarized, so the direction of polarization changes constantly, and the direction of vibration of the charges in the air changes constantly. If we consider light scattered at 90°, the vibration of the charged particles radiates to the observer only when the vibration is perpendicular to the observer's line of sight, and then light will be polarized along the direction of vibration. So scattering is an example of one means of producing polarization.

33–3 Birefringence

Another interesting effect of polarization is the fact that there are substances for which the index of refraction is different for light linearly polarized in one direction and linearly polarized in another. Suppose that we had some material which consisted of long, nonspherical molecules, longer than they are wide, and suppose that these molecules were arranged in the substance with their long axes parallel. Then what happens when the oscillating electric field passes through this substance? Suppose that because of the structure of the molecule, the electrons in the substance respond more easily to oscillations in the direction parallel to the axes of the molecules than they would respond if the electric field tries to push them at right angles to the molecular axis. In this way we expect a different response for polarization in one direction than for polarization at right angles to that direction. Let us call the direction of the axes of the molecules the *optic axis*. When the polarization is in the direction of the optic axis the index of refraction is different than it would be if the direction of polarization were at right angles to it. Such a substance is called *birefringent*. It has two refrangibilities, i.e., two indexes of refraction, depending on the direction of the polarization inside the substance. What kind of a substance can be birefringent? In a birefringent substance there must be a certain amount of lining up, for one reason or another, of unsymmetrical molecules. Certainly a cubic crystal, which has the symmetry of a cube, cannot be birefringent. But long needlelike crystals undoubtedly contain molecules that are asymmetric, and one observes this effect very easily.

Let us see what effects we would expect if we were to shine polarized light through a plate of a birefringent substance. If the polarization is parallel to the optic axis, the light will go through with one velocity; if the polarization is perpendicular to the axis, the light is transmitted with a different velocity. An interesting situation arises when, say, light is linearly polarized at 45° to the optic axis. Now the 45° polarization, we have already noticed, can be represented as a superposition of the x- and the y-polarizations of equal amplitude and in phase, as shown in Fig. 33–2(a). Since the x- and y-polarizations travel with different velocities, their phases change at a different rate as the light passes through the substance. So, although at the start the x- and y-vibrations are in phase, inside the material the phase difference between x- and y-vibrations is proportional to the depth in the substance. As the light proceeds through the material the polarization changes as shown in the series of diagrams in Fig. 33–2. If the thickness of the plate is just right to introduce a 90° phase shift between the x- and y-polarizations, as in Fig. 33–2(c), the light will come out circularly polarized. Such a thickness is called a quarter-wave plate, because it introduces a quarter-cycle phase difference between the x- and the y-polarizations. If linearly polarized light is sent through two quarter-wave plates, it will come out plane-polarized again, but at right angles to the original direction, as we can see from Fig. 33–2(e).

One can easily illustrate this phenomenon with a piece of cellophane. Cellophane is made of long, fibrous molecules, and is not isotropic, since the fibers lie preferentially in a certain direction. To demonstrate birefringence we need a

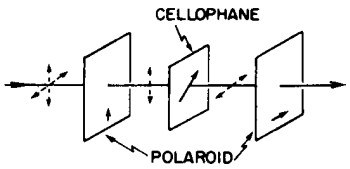


Fig. 33–3. An experimental demonstration of the birefringence of cellophane. The electric vectors in the light are indicated by the dotted lines. The pass axes of the polaroid sheets and optic axes of the cellophane are indicated by arrows. The incident beam is unpolarized.

beam of linearly polarized light, and we can obtain this conveniently by passing unpolarized light through a sheet of polaroid. Polaroid, which we will discuss later in more detail, has the useful property that it transmits light that is linearly polarized parallel to the axis of the polaroid with very little absorption, but light polarized in a direction perpendicular to the axis of the polaroid is strongly absorbed. When we pass unpolarized light through a sheet of polaroid, only that part of the unpolarized beam which is vibrating parallel to the axis of the polaroid gets through, so that the transmitted beam is linearly polarized. This same property of polaroid is also useful in detecting the direction of polarization of a linearly polarized beam, or in determining whether a beam is linearly polarized or not. One simply passes the beam of light through the polaroid sheet and rotates the polaroid in the plane normal to the beam. If the beam is linearly polarized, it will not be transmitted through the sheet when the axis of the polaroid is normal to the direction of polarization. The transmitted beam is only slightly attenuated when the axis of the polaroid sheet is rotated through 90°. If the transmitted intensity is independent of the orientation of the polaroid, the beam is not linearly polarized.

To demonstrate the birefringence of cellophane, we use two sheets of polaroid, as shown in Fig. 33–3. The first gives us a linearly polarized beam which we pass through the cellophane and then through the second polaroid sheet, which serves to detect any effect the cellophane may have had on the polarized light passing through it. If we first set the axes of the two polaroid sheets perpendicular to each other and remove the cellophane, no light will be transmitted through the second polaroid. If we now introduce the cellophane between the two polaroid sheets, and rotate the sheet about the beam axis, we observe that in general the cellophane makes it possible for some light to pass through the second polaroid. However, there are two orientations of the cellophane sheet, at right angles to each other, which permit no light to pass through the second polaroid. These orientations in which linearly polarized light is transmitted through the cellophane with no effect on the direction of polarization must be the directions parallel and perpendicular to the optic axis of the cellophane sheet.

We suppose that the light passes through the cellophane with two different velocities in these two different orientations, but it is transmitted without changing the direction of polarization. When the cellophane is turned halfway between these two orientations, as shown in Fig. 33–3, we see that the light transmitted through the second polaroid is bright.

It just happens that ordinary cellophane used in commercial packaging is very close to a half-wave thickness for most of the colors in white light. Such a sheet will turn the axis of linearly polarized light through 90° if the incident linearly polarized beam makes an angle of 45° with the optic axis, so that the beam emerging from the cellophane is then vibrating in the right direction to pass through the second polaroid sheet.

If we use white light in our demonstration, the cellophane sheet will be of the proper half-wave thickness only for a particular component of the white light, and the transmitted beam will have the color of this component. The color transmitted depends on the thickness of the cellophane sheet, and we can vary the effective thickness of the cellophane by tilting it so that the light passes through the cellophane at an angle, consequently through a longer path in the cellophane. As the sheet is tilted the transmitted color changes. With cellophane of different thicknesses one can construct filters that will transmit different colors. These filters have the interesting property that they transmit one color when the two polaroid sheets have their axes perpendicular, and the complementary color when the axes of the two polaroid sheets are parallel.

Another interesting application of aligned molecules is quite practical. Certain plastics are composed of very long and complicated molecules all twisted together. When the plastic is solidified very carefully, the molecules are all twisted in a mass, so that there are as many aligned in one direction as another, and so the plastic is not particularly birefringent. Usually there are strains and stresses introduced when the material is solidified, so the material is not perfectly homo-

geneous. However, if we apply tension to a piece of this plastic material, it is as if we were pulling a whole tangle of strings, and there will be more strings preferentially aligned parallel to the tension than in any other direction. So when a stress is applied to certain plastics, they become birefringent, and one can see the effects of the birefringence by passing polarized light through the plastic. If we examine the transmitted light through a polaroid sheet, patterns of light and dark fringes will be observed (in color, if white light is used). The patterns move as stress is applied to the sample, and by counting the fringes and seeing where most of them are, one can determine what the stress is. Engineers use this phenomenon as a means of finding the stresses in odd-shaped pieces that are difficult to calculate.

Another interesting example of a way of obtaining birefringence is by means of a liquid substance. Consider a liquid composed of long asymmetric molecules which carry a plus or minus average charge near the ends of the molecule, so that the molecule is an electric dipole. In the collisions in the liquid the molecules will ordinarily be randomly oriented, with as many molecules pointed in one direction as in another. If we apply an electric field the molecules will tend to line up, and the moment they line up the liquid becomes birefringent. With two polaroid sheets and a transparent cell containing such a polar liquid, we can devise an arrangement with the property that light is transmitted only when the electric field is applied. So we have an electrical switch for light, which is called a *Kerr cell.* This effect, that an electric field can produce birefringence in certain liquids, is called the Kerr effect.

33–4 Polarizers

So far we have considered substances in which the refractive index is different for light polarized in different directions. Of very practical value are those crystals and other substances in which not only the index, but also the coefficient of absorption, is different for light polarized in different directions. By the same arguments which supported the idea of birefringence, it is understandable that absorption can vary with the direction in which the charges are forced to vibrate in an anisotropic substance. Tourmaline is an old, famous example and polaroid is another. Polaroid consists of a thin layer of small crystals of herapathite (a salt of iodine and quinine), all aligned with their axes parallel. These crystals absorb light when the oscillations are in one direction, and they do not absorb appreciably when the oscillations are in the other direction.

Suppose that we send light into a polaroid sheet polarized linearly at an angle θ to the passing direction. What intensity will come through? This incident light can be resolved into a component perpendicular to the pass direction which is proportional to sin θ, and a component along the pass direction which is proportional to cos θ. The amplitude which comes out of the polaroid is only the cosine θ part; the sin θ component is absorbed. The amplitude which passes through the polaroid is smaller than the amplitude which entered, by a factor cos θ. The energy which passes through the polaroid, i.e., the intensity of the light, is proportional to the square of cos θ. $\text{Cos}^2\,\theta$, then, is the intensity transmitted when the light enters polarized at an angle θ to the pass direction. The absorbed intensity, of course, is $\sin^2\theta$.

An interesting paradox is presented by the following situation. We know that is is not possible to send a beam of light through two polaroid sheets with their axes crossed at right angles. But if we place a third polaroid sheet *between* the first two, with its pass axis at 45° to the crossed axes, some light is transmitted. We know that polaroid absorbs light, it does not create anything. Nevertheless, the addition of a third polaroid at 45° allows more light to get through. The analysis of this phenomenon is left as an exercise for the student.

One of the most interesting examples of polarization is not in complicated crystals or difficult substances, but in one of the simplest and most familiar of situations—the reflection of light from a surface. Believe it or not, when light is reflected from a glass surface it may be polarized, and the physical explanation of this is very simple. It was discovered empirically by Brewster that light reflected

from a surface is completely polarized if the reflected beam and the beam refracted into the material form a right angle. The situation is illustrated in Fig. 33–4. If the incident beam is polarized in the plane of incidence, there will be no reflection at all. Only if the incident beam is polarized normal to the plane of incidence will it be reflected. The reason is very easy to understand. In the reflecting material the light is polarized transversely, and we know that it is the motion of the charges in the material which generates the emergent beam, which we call the reflected beam. The source of this so-called reflected light is not simply that the incident beam is reflected; our deeper understanding of this phenomenon tells us that the incident beam drives an oscillation of the charges in the material, which in turn generates the reflected beam. From Fig. 33–4 it is clear that only oscillations normal to the paper can radiate in the direction of reflection, and consequently the reflected beam will be polarized normal to the plane of incidence. If the incident beam is polarized in the plane of incidence, there will be no reflected light.

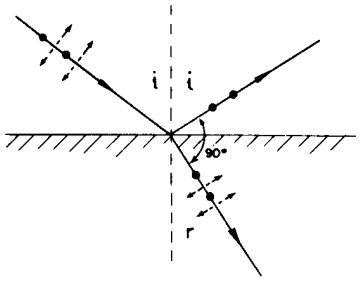


Fig. 33–4. Reflection of linearly polarized light at Brewster's angle. The polarization direction is indicated by dashed arrows; round dots indicate polarization normal to the paper.

This phenomenon is readily demonstrated by reflecting a linearly polarized beam from a flat piece of glass. If the glass is turned to present different angles of incidence to the polarized beam, sharp attenuation of the reflected intensity is observed when the angle of incidence passes through Brewster's angle. This attenuation is observed only if the plane of polarization lies in the plane of incidence. If the plane of polarization is normal to the plane of incidence, the usual reflected intensity is observed at all angles.

33–5 Optical activity

Another most remarkable effect of polarization is observed in materials composed of molecules which do not have reflection symmetry: molecules shaped something like a corkscrew, or like a gloved hand, or any shape which, if viewed through a mirror, would be reversed in the same way that a left-hand glove reflects as a right-hand glove. Suppose all of the molecules in the substance are the same, i.e., none is a mirror image of any other. Such a substance may show an interesting effect called optical activity, whereby as linearly polarized light passes through the substance, the direction of polarization rotates about the beam axis.

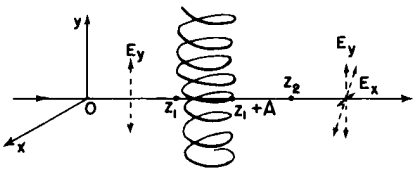


Fig. 33–5. A molecule with a shape that is not symmetric when reflected in a mirror. A beam of light, linearly polarized in the y-direction, falls on the molecule.

To understand the phenomenon of optical activity requires some calculation, but we can see qualitatively how the effect might come about, without actually carrying out the calculations. Consider an asymmetric molecule in the shape of a spiral, as shown in Fig. 33–5. Molecules need not actually be shaped like a corkscrew in order to exhibit optical activity, but this is a simple shape which we shall take as a typical example of those that do not have reflection symmetry. When a light beam linearly polarized along the y-direction falls on this molecule, the electric field will drive charges up and down the helix, thereby generating a current in the y-direction and radiating an electric field E_y polarized in the y-direction. However, if the electrons are constrained to move along the spiral, they must also move in the x-direction as they are driven up and down. When a current is flowing up the spiral, it is also flowing into the paper at $z = z_1$ and out of the paper at $z = z_1 + A$, if A is the diameter of our molecular spiral. One might suppose that the current in the x-direction would produce no net radiation, since the currents are in opposite directions on opposite sides of the spiral. However, if we consider the x-components of the electric field arriving at $z = z_2$, we see that the field radiated by the current at $z = z_1 + A$ and the field radiated from $z = z_1$ arrive at z_2 separated in time by the amount A/c, and thus separated in phase by $\pi + \omega A/c$. Since the phase difference is not exactly π, the two fields do not cancel exactly, and we are left with a small x-component in the electric field generated by the motion of the electrons in the molecule, whereas the driving electric field had only a y-component. This small x-component, added to the large y-component, produces a resultant field that is tilted slightly with respect to the y-axis, the original direction of polarization. As the light moves through the material, the direction of polarization rotates about the beam axis. By drawing a few examples and considering the currents that will be set in motion by an incident

electric field, one can convince himself that the existence of optical activity and the sign of the rotation are independent of the orientation of the molecules.

Corn syrup is a common substance which possesses optical activity. The phenomenon is easily demonstrated with a polaroid sheet to produce a linearly polarized beam, a transmission cell containing corn syrup, and a second polaroid sheet to detect the rotation of the direction of polarization as the light passes through the corn syrup.

33–6 The intensity of reflected light

Let us now consider quantitatively the reflection coefficient as a function of angle. Figure 33–6(a) shows a beam of light striking a glass surface, where it is partly reflected and partly refracted into the glass. Let us suppose that the incident beam, of unit amplitude, is linearly polarized normal to the plane of the paper. We will call the amplitude of the reflected wave b, and the amplitude of the refracted wave a. The refracted and reflected waves will, of course, be linearly polarized, and the electric field vectors of the incident, reflected, and refracted waves are all parallel to each other. Figure 33–6(b) shows the same situation, but now we suppose that the incident wave, of unit amplitude, is polarized in the plane of the paper. Now let us call the amplitude of the reflected and refracted wave B and A, respectively.

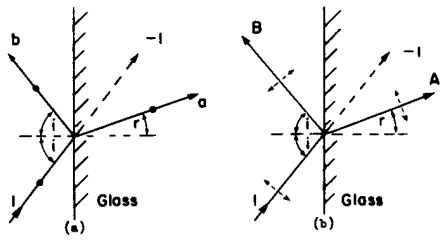


Fig. 33–6. An incident wave of unit amplitude is reflected and refracted at a glass surface. In (a) the incident wave is linearly polarized normal to the plane of the paper. In (b) the incident wave is linearly polarized in the direction shown by the dotted electric vector.

We wish to calculate how strong the reflection is in the two situations illustrated in Fig. 33–6(a) and 33–6(b). We already know that when the angle between the reflected beam and refracted beam is a right angle, there will be no reflected wave in Fig. 33–6(b), but let us see if we cannot get a quantitative answer—an exact formula for B and b as a function of the angle of incidence, i.

The principle that we must understand is as follows. The currents that are generated in the glass produce two waves. First, they produce the reflected wave. Moreover, we know that if there were no currents generated in the glass, the incident wave would continue straight into the glass. Remember that all the sources in the world make the net field. The source of the incident light beam produces a field of unit amplitude which would move into the glass along the dotted line in the figure. This field is not observed, and therefore the currents generated in the glass must produce a field of amplitude -1 which moves along the dotted line. Using this fact, we will calculate the amplitude of the refracted waves, a and A.

In Fig. 33–6(a) we see that the field of amplitude b is radiated by the motion of charges inside the glass which are responding to a field a inside the glass, and that therefore b is proportional to a. We might suppose that since our two figures are exactly the same, except for the direction of polarization, the ratio B/A would be the same as the ratio b/a. This is not quite true, however, because in Fig. 33–6(b) the polarization directions are not all parallel to each other, as they are in Fig. 33–6(a). It is only the component of A which is perpendicular to B, $A\cos(i + r)$, which is effective in producing B. The correct expression for the proportionality is then

$$\frac{b}{a} = \frac{B}{A\cos(i + r)}. \tag{33.1}$$

Now we use a trick. We know that in both (a) and (b) of Fig. 33–6 the electric field in the glass must produce oscillations of the charges which generate a field of amplitude -1, polarized parallel to the incident beam, and moving in the direction of the dotted line. But we see from part (b) of the figure that only the component of A that is normal to the dashed line has the right polarization to produce this field, whereas in Fig. 33–6(a) the full amplitude a is effective, since the polarization of wave a is parallel to the polarization of the wave of amplitude -1. Therefore we can write

$$\frac{A\cos(i - r)}{a} = \frac{-1}{-1}, \tag{33.2}$$

since the two amplitudes on the left side of Eq. (33.2) each produce the wave of amplitude -1.

Dividing Eq. (33.1) by Eq. (33.2), we obtain

$$\frac{B}{b} = \frac{\cos(i + r)}{\cos(i - r)}, \tag{33.3}$$

a result which we can check against what we already know. If we set $(i + r) = 90°$, Eq. (33.3) gives $B = 0$, as Brewster says it should be, so our results so far are at least not obviously wrong.

We have assumed unit amplitudes for the incident waves, so that $|B|^2/1^2$ is the reflection coefficient for waves polarized in the plane of incidence, and $|b|^2/1^2$ is the reflection coefficient for waves polarized normal to the plane of incidence. The ratio of these two reflection coefficients is determined by Eq. (33.3).

Now we perform a miracle, and compute not just the ratio, but each coefficient $|B|^2$ and $|b|^2$ individually! We know from the conservation of energy that the energy in the refracted wave must be equal to the incident energy minus the energy in the reflected wave, $1 - |B|^2$ in one case, $1 - |b|^2$ in the other. Furthermore, the energy which passes into the glass in Fig. 33–6(b) is to the energy which passes into the glass in Fig. 33–6(a) as the ratio of the squares of the refracted amplitudes, $|A|^2/|a|^2$. One might ask whether we really know how to compute the energy inside the glass, because, after all, there are energies of motion of the atoms in addition to the energy in the electric field. But it is obvious that all of the various contributions to the total energy will be proportional to the square of the amplitude of the electric field. Therefore we can write

$$\frac{1 - |B|^2}{1 - |b|^2} = \frac{|A|^2}{|a|^2}. \tag{33.4}$$

We now substitute Eq. (33.2) to eliminate A/a from the expression above, and express B in terms of b by means of Eq. (33.3):

$$\frac{1 - |b|^2 \dfrac{\cos^2(i + r)}{\cos^2(i - r)}}{1 - |b|^2} = \frac{1}{\cos^2(i - r)}. \tag{33.5}$$

This equation contains only one unknown amplitude, b. Solving for $|b|^2$, we obtain

$$|b|^2 = \frac{\sin^2(i - r)}{\sin^2(i + r)} \tag{33.6}$$

and, with the aid of (33.3),

$$|B|^2 = \frac{\tan^2(i - r)}{\tan^2(i + r)}. \tag{33.7}$$

So we have found the reflection coefficient $|b|^2$ for an incident wave polarized perpendicular to the plane of incidence, and also the reflection coefficient $|B|^2$ for an incident wave polarized in the plane of incidence!

It is possible to go on with arguments of this nature and deduce that b is real. To prove this, one must consider a case where light is coming from both sides of the glass surface at the same time, a situation not easy to arrange experimentally, but fun to analyze theoretically. If we analyze this general case, we can prove that b must be real, and therefore, in fact, that $b = \pm \sin(i - r)/\sin(i + r)$. It is even possible to determine the sign by considering the case of a very, very thin layer in which there is reflection from the front and from the back surfaces, and calculating how much light is reflected. We know how much light should be reflected by a thin layer, because we know how much current is generated, and we have even worked out the fields produced by such currents.

One can show by these arguments that

$$b = -\frac{\sin(i - r)}{\sin(i + r)}, \qquad B = -\frac{\tan(i - r)}{\tan(i + r)}. \tag{33.8}$$

These expressions for the reflection coefficients as a function of the angles of incidence and refraction are called Fresnel's reflection formulas.

If we consider the limit as the angles i and r go to zero, we find, for the case of normal incidence, that $B^2 \approx b^2 \approx (i - r)^2/(i + r)^2$ for both polarizations,

since the sines are practically equal to the angles, as are also the tangents. But we know that $\sin i/\sin r = n$, and when the angles are small, $i/r \approx n$. It is thus easy to show that the coefficient of reflection for normal incidence is

$$B^2 = b^2 = \frac{(n-1)^2}{(n+1)^2}.$$

It is interesting to find out how much light is reflected at normal incidence from the surface of water, for example. For water, n is 4/3, so that the reflection coefficient is $(1/7)^2 \approx 2\%$. At normal incidence, only two percent of the light is reflected from the surface of water.

33–7 Anomalous refraction

The last polarization effect we shall consider was actually one of the first to be discovered: anomalous refraction. Sailors visiting Iceland brought back to Europe crystals of Iceland spar ($CaCO_3$) which had the amusing property of making anything seen through the crystal appear doubled, i.e., as two images. This came to the attention of Huygens, and played an important role in the discovery of polarization. As is often the case, the phenomena which are discovered first are the hardest, ultimately, to explain. It is only after we understand a physical concept thoroughly that we can carefully select those phenomena which most clearly and simply demonstrate the concept.

Anomalous refraction is a particular case of the same birefringence that we considered earlier. Anomalous refraction comes about when the optic axis, the long axis of our asymmetric molecules, is *not* parallel to the surface of the crystal. In Fig. 33–7 are drawn two pieces of birefringent material, with the optic axis as shown. In the upper figure, the incident beam falling on the material is linearly polarized in a direction perpendicular to the optic axis of the material. When this beam strikes the surface of the material, each point on the surface acts as a source of a wave which travels into the crystal with velocity $v_\perp$, the velocity of light in the crystal when the plane of polarization is normal to the optic axis. The wavefront is just the envelope or locus of all these little spherical waves, and this wavefront moves straight through the crystal and out the other side. This is just the ordinary behavior we would expect, and this ray is called the *ordinary ray*.

In the lower figure the linearly polarized light falling on the crystal has its direction of polarization turned through 90°, so that the optic axis lies in the plane of polarization. When we now consider the little waves originating at any point on the surface of the crystal, we see that they do not spread out as spherical waves. Light travelling along the optic axis travels with velocity $v_\perp$ because the polarization is perpendicular to the optic axis, whereas the light travelling perpendicular to the optic axis travels with velocity $v_\parallel$ because the polarization is parallel to the optic axis. In a birefringent material $v_\parallel \neq v_\perp$, and in the figure $v_\parallel < v_\perp$. A more complete analysis will show that the waves spread out on the surface of an ellipsoid, with the optic axis as major axis of the ellipsoid. The envelope of all these elliptical waves is the wavefront which proceeds through the crystal in the direction shown. Again, at the back surface the beam will be deflected just as it was at the front surface, so that the light emerges parallel to the incident beam, but displaced from it. Clearly, this beam does not follow Snell's law, but goes in an extraordinary direction. It is therefore called the *extraordinary ray*.

When an unpolarized beam strikes an anomalously refracting crystal, it is separated into an ordinary ray, which travels straight through in the normal manner, and an extraordinary ray which is displaced as it passes through the crystal. These two emergent rays are linearly polarized at right angles to each other. That this is true can be readily demonstrated with a sheet of polaroid to analyze the polarization of the emergent rays. We can also demonstrate that our interpretation of this phenomenon is correct by sending linearly polarized light into the crystal. By properly orienting the direction of polarization of the incident beam, we can make this light go straight through without splitting, or we can make it go through without splitting but with a displacement.

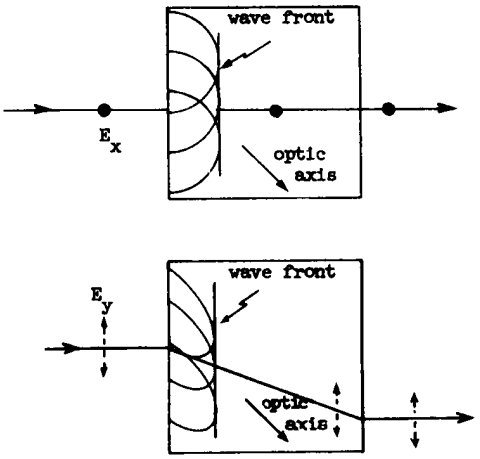


Fig. 33–7. The upper diagram shows the path of the ordinary ray through a doubly refracting crystal. The extraordinary ray is shown in the lower diagram. The optic axis lies in the plane of the paper.

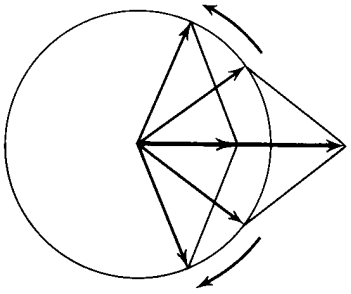


Fig. 33–8. Two oppositely rotating vectors of equal amplitude add to produce a vector in a fixed direction, but with an oscillating amplitude.

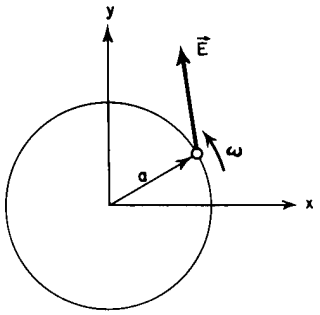


Fig. 33–9. A charge moving in a circle in response to circularly polarized light.

We have represented all the various polarization cases in Figs. 33–1 and 33–2 as superpositions of two special polarization cases, namely x and y in various amounts and phases. Other pairs could equally well have been used. Polarization along any two perpendicular axes x', y' inclined to x and y would serve as well [for example, any polarization can be made up of superpositions of cases (a) and (e) of Fig. 33–2]. It is interesting, however, that this idea can be extended to other cases also. For example, any *linear* polarization can be made up by superposing suitable amounts at suitable phases of right and left *circular* polarizations [cases (c) and (g) of Fig. 33–2], since two equal vectors rotating in opposite directions add to give a single vector oscillating in a straight line (Fig. 33–8). If the phase of one is shifted relative to the other, the line is inclined. Thus all the pictures of Fig. 33–1 could be labeled "the superposition of equal amounts of right and left circularly polarized light at various relative phases." As the left slips behind the right in phase, the direction of the linear polarization changes. Therefore optically active materials are, in a sense, birefringent. Their properties can be described by saying that they have different indexes for right- and left-hand circularly polarized light. Superposition of right and left circularly polarized light of different intensities produces elliptically polarized light.

Circularly polarized light has another interesting property—it carries *angular momentum* (about the direction of propagation). To illustrate this, suppose that such light falls on an atom represented by a harmonic oscillator that can be displaced equally well in any direction in the plane xy. Then the x-displacement of the electron will respond to the E_x component of the field, while the y-component responds, equally, to the equal E_y component of the field but 90° behind in phase. That is, the responding electron goes around in a circle, with angular velocity ω, in response to the rotating electric field of the light (Fig. 33–9). Depending on the damping characteristics of the response of the oscillator, the direction of the displacement **a** of the electron, and the direction of the force $q_e\mathbf{E}$ on it need not be the same but they rotate around together. The **E** may have a component at right angles to **a**, so work is done on the system and a torque τ is exerted. The work done per second is $\tau\omega$. Over a period of time T the energy absorbed is $\tau\omega T$, while τT is the angular momentum delivered to the matter absorbing the energy. We see therefore that *a beam of right circularly polarized light containing a total energy* $\mathcal{E}$ *carries an angular momentum* (*with vector directed along the direction of propagation*) $\mathcal{E}/\omega$. For when this beam is absorbed that angular momentum is delivered to the absorber. Left-hand circular light carries angular momentum of the opposite sign, $-\mathcal{E}/\omega$.

34

Relativistic Effects in Radiation

34–1 Moving sources

34–2 Finding the "apparent" motion

34–3 Synchrotron radiation

34–4 Cosmic synchrotron radiation

34–5 Bremsstrahlung

34–6 The Doppler effect

34–7 The ω, k four-vector

34–8 Aberration

34–9 The momentum of light

34–1 Moving sources

In the present chapter we shall describe a number of miscellaneous effects in connection with radiation, and then we shall be finished with the classical theory of light propagation. In our analysis of light, we have gone rather far and into considerable detail. The only phenomena of any consequence associated with electromagnetic radiation that we have not discussed is what happens if radiowaves are contained in a box with reflecting walls, the size of the box being comparable to a wavelength, or are transmitted down a long tube. The phenomena of so-called *cavity resonators* and *waveguides* we shall discuss later; we shall first use another physical example—sound—and then we shall return to this subject. Except for this, the present chapter is our last consideration of the classical theory of light.

We can summarize all the effects that we shall now discuss by remarking that they have to do with the effects of *moving sources*. We no longer assume that the source is localized, with all its motion being at a relatively low speed near a fixed point.

We recall that the fundamental laws of electrodynamics say that, at large distances from a moving charge, the electric field is given by the formula

$$\mathbf{E} = -\frac{q}{4\pi\epsilon_0 c^2}\frac{d^2\mathbf{e}_{R'}}{dt^2}. \tag{34.1}$$

The second derivative of the unit vector $\mathbf{e}_{R'}$ which points in the apparent direction of the charge, is the determining feature of the electric field. This unit vector does not point toward the *present* position of the charge, of course, but rather in the direction that the charge would seem to be, if the information travels only at the finite speed c from the charge to the observer.

Associated with the electric field is a magnetic field, always at right angles to the electric field and at right angles to the apparent direction of the source, given by the formula

$$\mathbf{B} = -\mathbf{e}_{R'} \times \mathbf{E}/c. \tag{34.2}$$

Until now we have considered only the case in which motions are nonrelativistic in speed, so that there is no appreciable motion in the direction of the source to be considered. Now we shall be more general and study the case where the motion is at an arbitrary velocity, and see what different effects may be expected in those circumstances. We shall let the motion be at an arbitrary speed, but of course we shall still assume that the detector is very far from the source.

We already know from our discussion in Chapter 28 that the only things that count in $d^2\mathbf{e}_{R'}/dt^2$ are the changes in the *direction* of $\mathbf{e}_{R'}$. Let the coordinates of the charge be (x, y, z), with z measured along the direction of observation (Fig. 34–1). At a given moment in time, say the moment τ, the three components of the position are $x(\tau)$, $y(\tau)$, and $z(\tau)$. The distance R is very nearly equal to $R(\tau) = R_0 + z(\tau)$. Now the direction of the vector $\mathbf{e}_{R'}$ depends mainly on x and y, but hardly at all upon z: the transverse components of the unit vector are x/R and y/R, and when we differentiate these components we get things like R^2 in the denominator:

$$\frac{d(x/R)}{dt} = \frac{dx/dt}{R} - \frac{dz}{dt}\frac{x}{R^2}.$$

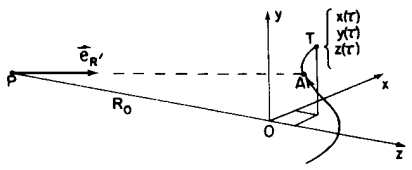


Fig. 34–1. The path of a moving charge. The true position at the time τ is at T, but the retarded position is at A.

So, when we are far enough away the only terms we have to worry about are the variations of x and y. Thus we take out the factor R_0 and get

$$E_x = -\frac{q}{4\pi\epsilon_0 c^2 R_0}\frac{d^2x'}{dt^2},$$

$$E_y = -\frac{q}{4\pi\epsilon_0 c^2 R_0}\frac{d^2y'}{dt^2}, \tag{34.3}$$

where R_0 is the distance, more or less, to q; let us take it as the distance OP to the origin of the coordinates (x, y, z). Thus the electric field is a constant multiplied by a very simple thing, the second derivatives of the x- and y-coordinates. (We could put it more mathematically by calling x and y the *transverse* components of the position vector $\mathbf{r}$ of the charge, but this would not add to the clarity.)

Of course, we realize that the coordinates must be measured at the retarded time. Here we find that $z(\tau)$ *does* affect the retardation. What time is the retarded time? If the time of observation is called t (the time at P) then the time τ to which this corresponds at A is not the time t, but is delayed by the total distance that the light has to go, divided by the speed of light. In the first approximation, this delay is R_0/c, a constant (an uninteresting feature), but in the next approximation we must include the effects of the position in the z-direction at the time τ, because if q is a little farther back, there is a little more retardation. This is an effect that we have neglected before, and it is the only change needed in order to make our results valid for all speeds.

What we must now do is to choose a certain value of t and calculate the value of τ from it, and thus find out where x and y are at that τ. These are then the retarded x and y, which we call x' and y', whose second derivatives determine the field. Thus τ is determined by

$$t = \tau + \frac{R_0}{c} + \frac{z(\tau)}{c}$$

and

$$x'(t) = x(\tau), \qquad y'(t) = y(\tau). \tag{34.4}$$

Now these are complicated equations, but it is easy enough to make a geometrical picture to describe their solution. This picture will give us a good qualitative feeling for how things work, but it still takes a lot of detailed mathematics to deduce the precise results of a complicated problem.

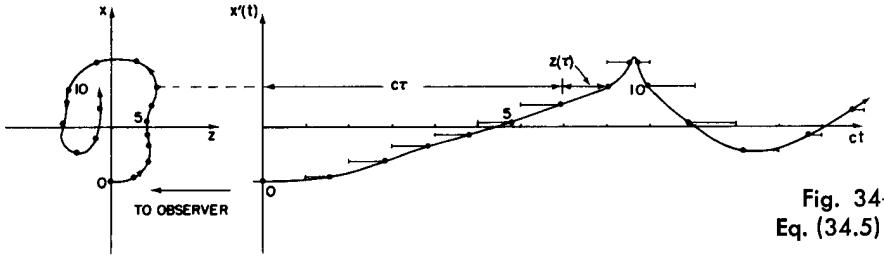


Fig. 34–2. A geometrical solution of Eq. (34.5) to find $x'(t)$.

34–2 Finding the "apparent" motion

The above equation has an interesting simplification. If we disregard the uninteresting constant delay R_0/c, which just means that we must change the origin of t by a constant, then it says that

$$ct = c\tau + z(\tau), \qquad x' = x(\tau), \qquad y' = y(\tau). \tag{34.5}$$

Now we need to find x' and y' as functions of t, not τ, and we can do this in the following way: Eq. (34.5) says that we should take the actual motion and add a constant (the speed of light) times τ. What that turns out to mean is shown in Fig. 34–2. We take the actual motion of the charge (shown at left) and imagine that as it is going around it is being swept away from the point P at the speed c (there are no contractions from relativity or anything like that; this is just a mathematical addition of the $c\tau$). In this way we get a new motion, in which the line-

of-sight coordinate is ct, as shown at the right. (The figure shows the result for a rather complicated motion in a plane, but of course the motion may not be in one plane—it may be even more complicated than motion in a plane.) The point is that the horizontal (i.e., line-of-sight) distance now is no longer the old z, but is $z + c\tau$, and therefore is ct. Thus we have found a picture of the curve, x' (and y') against t! All we have to do to find the field is to look at the acceleration of this curve, i.e., to differentiate it twice. So the final answer is: in order to find the electric field for a moving charge, take the motion of the charge and translate it back at the speed c to "open it out"; then the curve, so drawn, is a curve of the x' and y' positions of the function of t. The acceleration of this curve gives the electric field as a function of t. Or, if we wish, we can now imagine that this whole "rigid" curve moves forward at the speed c through the plane of sight, so that the point of intersection with the plane of sight has the coordinates x' and y'. The acceleration of this point makes the electric field. This solution is just as exact as the formula we started with—it is simply a geometrical representation.

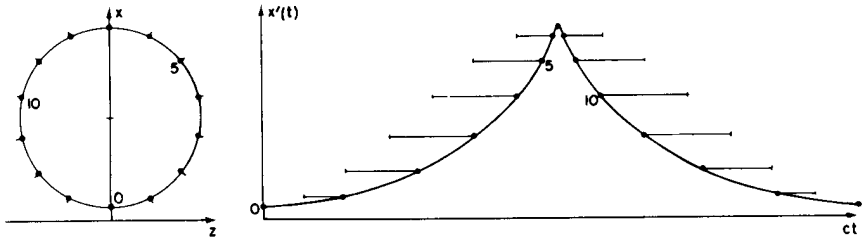


Fig. 34–3. The $x'(t)$ curve for a particle moving at constant speed $v = 0.94c$, a circle.

If the motion is relatively slow, for instance if we have an oscillator just going up and down slowly, then when we shoot that motion away at the speed of light, we would get, of course, a simple cosine curve, and that gives a formula we have been looking at for a long time: it gives the field produced by an oscillating charge. A more interesting example is an electron moving rapidly, very nearly at the speed of light, in a circle. If we look in the plane of the circle, the retarded $x'(t)$ appears as shown in Fig. 34–3. What is this curve? If we imagine a radius vector from the center of the circle to the charge, and if we extend this radial line a little bit past the charge, just a shade if it is going fast, then we come to a point on the line that goes at the speed of light. Therefore, when we translate the motion back at the speed of light, that corresponds to having a wheel with a charge on it rolling backward (without slipping) at the speed c; thus we find a curve which is very close to a cycloid—it is called a *hypocycloid.* If the charge is going very nearly at the speed of light, the "cusps" are very sharp indeed; if it went at exactly the speed of light, they would be actual cusps, infinitely sharp. "Infinitely sharp" is interesting; it means that near a cusp the second derivative is enormous. Once in each cycle we get a sharp pulse of electric field. This is not at all what we would get from a nonrelativistic motion, where each time the charge goes around there is an oscillation which is of about the same "strength" all the time. Instead, there are very sharp pulses of electric field spaced at time intervals $1/T_0$ apart, where T_0 is the period of revolution. These strong electric fields are emitted in a narrow cone in the direction of motion of the charge. When the charge is moving away from P, there is very little curvature and there is very little radiated field in the direction of P.

34–3 Synchrotron radiation

We have very fast electrons moving in circular paths in the synchrotron; they are travelling at very nearly the speed c, and it is possible to see the above radiation as actual *light*! Let us discuss this in more detail.

In the synchrotron we have electrons which go around in circles in a uniform magnetic field. First, let us see why they go in circles. From Eq. (12.10), we know that the force on a particle in a magnetic field is given by

$$\mathbf{F} = q\mathbf{v} \times \mathbf{B}, \tag{34.6}$$

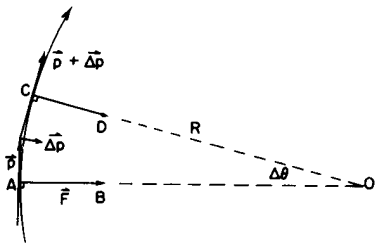


Fig. 34–4. A charged particle moves in a circular (or helical) path in a uniform magnetic field.

and it is at right angles both to the field and to the velocity. As usual, the force is equal to the rate of change of momentum with time. If the field is directed upward out of the paper, the momentum of the particle and the force on it are as shown in Fig. 34–4. Since the force is at right angles to the velocity, the kinetic energy, and therefore the speed, remains *constant*. All the magnetic field does is to change the *direction of motion*. In a short time Δt, the momentum vector changes at right angles to itself by an amount $\Delta\mathbf{p} = \mathbf{F}\Delta t$, and therefore $\mathbf{p}$ turns through an angle $\Delta\theta = \Delta p/p = qvB\,\Delta t/p$, since $|F| = qvB$. But in this same time the particle has gone a distance $\Delta s = v\,\Delta t$. Evidently, the two lines AB and CD will intersect at a point O such that $OA = OC = R$, where $\Delta s = R\,\Delta\theta$. Combining this with the previous expressions, we find $R\,\Delta\theta/\Delta t = R\omega = v = qvBR/p$, from which we find

$$p = qBR \tag{34.7}$$

and

$$\omega = qvB/p. \tag{34.8}$$

Since this same argument can be applied during the next instant, the next, and so on, we conclude that the particle must be moving in a *circle* of radius R, with angular velocity ω.

The result that the momentum of the particle is equal to a charge times the radius times the magnetic field is a very important law that is used a great deal. It is important for practical purposes because if we have elementary particles which all have the same charge and we observe them in a magnetic field, we can measure the radii of curvature of their orbits and, knowing the magnetic field, thus determine the momenta of the particles. If we multiply both sides of Eq. (34.7) by c, and express q in terms of the electronic charge, we can measure the momentum in units of the *electron volt*. In those units our formula is

$$pc(\text{ev}) = 3 \times 10^8(q/q_e)BR, \tag{34.9}$$

where B, R, and the speed of light are all expressed in the mks system, the latter being 3×10^8, numerically.

The mks unit of magnetic field is called a *weber per square meter*. There is an older unit which is still in common use, called a *gauss*. One weber/m^{-2} is equal to 10^4 gauss. To give an idea of how big magnetic fields are, the strongest magnetic field that one can usually make in iron is about 1.5×10^4 gauss; beyond that, the advantage of using iron disappears. Today, electromagnets wound with superconducting wire are able to produce steady fields of over 10^5 gauss strength—that is, 10 mks units. The field of the earth is a few tenths of a gauss at the equator.

Returning to Eq. (34.9), we could imagine the synchrotron running at a billion electron volts, so pc would be 10^9 for a billion electron volts. (We shall come back to the energy in just a moment.) Then, if we had a B corresponding to, say, 10,000 gauss, which is a good substantial field, one mks unit, then we see that R would have to be 3.3 meters. The actual radius of the Caltech synchrotron is 3.7 meters, the field is a little bigger, and the energy is 1.5 billion, but it is the same idea. So now we have a feeling for why the synchrotron has the size it has.

We have calculated the momentum, but we know that the total energy, including the rest energy, is given by $W = \sqrt{p^2c^2 + m^2c^4}$, and for an electron the rest energy corresponding to mc^2 is 0.511×10^6 ev, so when pc is 10^9 ev we can neglect mc^2, and so for all practical purposes $W = pc$ when the speeds are relativistic. It is practically the same to say the energy of an electron is a billion electron volts as to say the momentum times c is a billion electron volts. If $W = 10^9$ ev, it is easy to show that the speed differs from the speed of light by but one part in eight million!

We turn now to the radiation emitted by such a particle. A particle moving on a circle of radius 3.3 meters, or 20 meters circumference, goes around once in roughly the time it takes light to go 20 meters. So the wavelength that should be emitted by such a particle would be 20 meters—in the shortwave radio region. But because of the piling up effect that we have been discussing (Fig. 34–3), and because the distance by which we must extend the radius to reach the speed c is

only one part in eight million of the radius, the cusps of the hypocycloid are enormously sharp compared with the distance between them. The acceleration, which involves a second derivative with respect to time, gets twice the "compression factor" of 8×10^6 because the time scale is reduced by eight million twice in the neighborhood of the cusp. Thus we might expect the effective wavelength to be much shorter, to the extent of 64 times 10^{12} smaller than 20 meters, and that corresponds to the x-ray region. (Actually, the cusp itself is not the entire determining factor; one must also include a certain region about the cusp. This changes the factor to the 3/2 power instead of the square, but still leaves us above the optical region.) Thus, even though a slowly moving electron would have radiated 20-meter radiowaves, the relativistic effect cuts down the wavelength so much that we can *see* it! Clearly, the light should be *polarized*, with the electric field perpendicular to the uniform magnetic field.

To further appreciate what we would observe, suppose that we were to take such light (to simplify things, because these pulses are so far apart in time, we shall just take one pulse) and direct it onto a diffraction grating, which is a lot of scattering wires. After this pulse comes away from the grating, what do we see? (We should see red light, blue light, and so on, if we see any light at all.) What *do* we see? The pulse strikes the grating head-on, and all the oscillators in the grating, together, are violently moved up and then back down again, just once. They then produce effects in various directions, as shown in Fig. 34–5. But the point P is closer to one end of the grating than to the other, so at this point the electric field arrives first from wire A, next from B, and so on; finally, the pulse from the last wire arrives. In short, the sum of the reflections from all the successive wires is as shown in Fig. 34–6(a); it is an electric field which is a series of pulses, and it is very like a sine wave whose wavelength is the distance between the pulses, just as it would be for monochromatic light striking the grating! So, we get colored light all right. But, by the same argument, will we not get light from any kind of a "pulse"? No. Suppose that the curve were much smoother; then we would add all the scattered waves together, separated by a small time between them (Fig. 34–6b). Then we see that the field would not shake at all, it would be a very smooth curve, because each pulse does not vary much in the time interval between pulses.

The electromagnetic radiation emitted by relativistic charged particles circulating in a magnetic field is called *synchrotron radiation*. It is so named for obvious reasons, but it is not limited specifically to synchrotrons, or even to earthbound laboratories. It is exciting and interesting that it also occurs in nature!

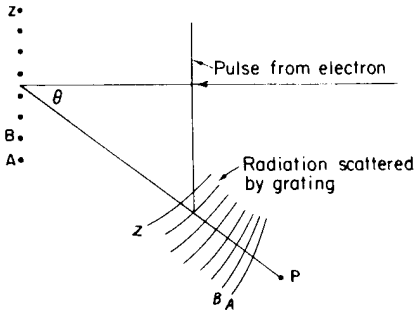


Fig. 34–5. The light which strikes a grating as a single, sharp pulse is scattered in various directions as different colors.

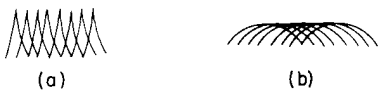


Fig. 34–6. The total electric field due to a series of (a) sharp pulses and (b) smooth pulses.



Fig. 34–7. The crab nebula as seen in all colors (no filter).

34–4 Cosmic synchrotron radiation

In the year 1054 the Chinese and Japanese civilizations were among the most advanced in the world; they were conscious of the external universe, and they recorded, most remarkably, an explosive bright star in that year. (It is amazing that none of the European monks, writing all the books of the middle ages, even bothered to write that a star exploded in the sky, but they did not.) Today we may take a picture of that star, and what we see is shown in Fig. 34–7. On the outside is a big mass of red filaments, which is produced by the atoms of the thin gas "ringing" at their natural frequencies; this makes a bright line spectrum with different frequencies in it. The red happens in this case to be due to nitrogen. On the other hand, in the central region is a mysterious, fuzzy patch of light in a *continuous* distribution of frequency, i.e., there are no special frequencies associated with particular atoms. Yet this is not dust "lit up" by nearby stars, which is one way by which one can get a continuous spectrum. We can see stars through it, so it is transparent, but it is *emitting* light.

In Fig. 34–8 we look at the same object, using light in a region of the spectrum which has no bright spectral line, so that we see only the central region. But in this case, also, polarizers have been put on the telescope, and the two views correspond to two orientations 90° apart. We see that the pictures are different! That is to say, the light is polarized. The reason, presumably, is that there is a local magnetic field, and many very energetic electrons are going around in that magnetic field.

We have just illustrated how the electrons could go around the field in a circle. We can add to this, of course, any uniform motion in the direction of the field, since the force, $q\mathbf{v} \times \mathbf{B}$, has no component in this direction and, as we have already remarked, the synchrotron radiation is evidently polarized in a direction at right angles to the projection of the magnetic field onto the plane of sight.

Putting these two facts together, we see that in a region where one picture is bright and the other one is black, the light must have its electric field completely polarized in one direction. This means that there is a magnetic field at right angles to this direction, while in other regions, where there is a strong emission in the other picture, the magnetic field must be the other way. If we look carefully at Fig. 34–8, we may notice that there is, roughly speaking, a general set of "lines" that go one way in one picture and at right angles to this in the other. The pictures show a kind of fibrous structure. Presumably, the magnetic field lines will tend to extend relatively long distances in their own direction, and so, presumably, there are long regions of magnetic field with all the electrons spiralling one way, while in another region the field is the other way and the electrons are also spiralling that way.

What keeps the electron energy so high for so long a time? After all, it is 900 years since the explosion—how can they keep going so fast? How they maintain their energy and how this whole thing keeps going is still not thoroughly understood.

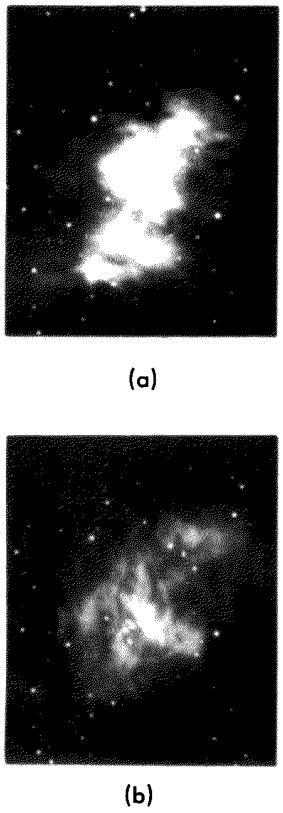


Fig. 34–8. The crab nebula as seen through a blue filter and a polaroid. (a) Electric vector vertical. (b) Electric vector horizontal.

34–5 Bremsstrahlung

We shall next remark briefly on one other interesting effect of a very fast-moving particle that radiates energy. The idea is very similar to the one we have just discussed. Suppose that there are charged particles in a piece of matter and a very fast electron, say, comes by (Fig. 34–9). Then, because of the electric field around the atomic nucleus the electron is pulled, accelerated, so that the curve of its motion has a slight kink or bend in it. If the electron is travelling at very nearly the speed of light, what is the electric field produced in the direction c? Remember our rule: we take the actual motion, translate it backwards at speed c, and that gives us a curve whose curvature measures the electric field. It was coming toward us at the speed v, so we get a backward motion, with the whole picture compressed into a smaller distance in proportion as $c - v$ is smaller than c. So, if $1 - v/c \ll 1$, there is a very sharp and rapid curvature at B', and when we take the second derivative of that we get a very high field in the direction of the motion. So when very energetic electrons move through matter they spit radiation in a forward direction. This is called *bremsstrahlung*. As a matter of fact, the synchrotron

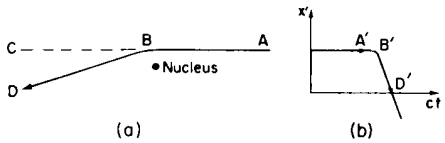


Fig. 34–9. A fast electron passing near a nucleus radiates energy in the direction of its motion.

is used, not so much to make high-energy electrons (actually if we could get them out of the machine more conveniently we would not say this) as to make very energetic photons—gamma rays—by passing the energetic electrons through a solid tungsten "target," and letting them radiate photons from this bremsstrahlung effect.

34–6 The Doppler effect

Now we go on to consider some other examples of the effects of moving sources. Let us suppose that the source is a stationary atom which is oscillating at one of its natural frequencies, ω_0. Then we know that the frequency of the light we would observe is ω_0. But now let us take another example, in which we have a similar oscillator oscillating with a frequency ω_1, and at the same time the whole atom, the whole oscillator, is moving along in a direction toward the observer at velocity v. Then the actual motion in space, of course, is as shown in Fig. 34–10(a). Now we play our usual game, we add $c\tau$; that is to say, we translate the whole curve backward and we find then that it oscillates as in Fig. 34–10(b). In a given amount of time τ, when the oscillator would have gone a distance $v\tau$, on the x' vs. ct diagram it goes a distance $(c - v)\tau$. So all the oscillations of frequency ω_1 in the time $\Delta\tau$ are now found in the interval $\Delta\tau = (1 - v/c)\,\Delta\tau$; they are squashed together, and as this curve comes by us at speed c, we will see light of a *higher frequency*, higher by just the compression factor $(1 - v/c)$. Thus we observe

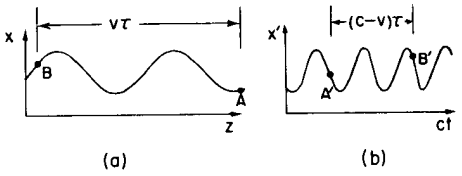


Fig. 34–10. The $x - z$ and $x' - t$ curves of a moving oscillator.

$$\omega = \frac{\omega_1}{1 - v/c}. \tag{34.10}$$

We can, of course, analyze this situation in various other ways. Suppose that the atom were emitting, instead of sine waves, a series of pulses, pip, pip, pip, pip, at a certain frequency ω_1. At what frequency would they be received by us? The first one that arrives has a certain delay, but the next one is delayed less because in the meantime the atom moves closer to the receiver. Therefore, the time between the "pips" is decreased by the motion. If we analyze the geometry of the situation, we find that the frequency of the pips is increased by the factor $1/(1 - v/c)$.

Is $\omega = \omega_0/(1 - v/c)$, then, the frequency that would be observed if we took an ordinary atom, which had a natural frequency ω_0, and moved it toward the receiver at speed v? No; as we well know, the natural frequency ω_1 of a moving atom is not the same as that measured when it is standing still, because of the relativistic dilation in the rate of passage of time. Thus if ω_0 were the true natural frequency, then the modified natural frequency ω_1 would be

$$\omega_1 = \omega_0 \sqrt{1 - v^2/c^2}. \tag{34.11}$$

Therefore the observed frequency ω is

$$\omega = \frac{\omega_0 \sqrt{1 - v^2/c^2}}{1 - v/c}. \tag{34.12}$$

The shift in frequency observed in the above situation is called the *Doppler effect*: if something moves toward us the light it emits appears more violet, and if it moves away it appears more red.

We shall now give two more derivations of this same interesting and important result. Suppose, now, that the *source* is standing still and is emitting waves at frequency ω_0, while the *observer* is moving with speed v toward the source. After a certain period of time t the observer will have moved to a new position, a distance vt from where he was at $t = 0$. How many radians of phase will he have seen go by? A certain number, $\omega_0 t$, went past any fixed point, and in addition the observer has swept past some more by his own motion, namely a number vtk_0 (the number of radians per meter times the distance). So the total number of radians in the time t, or the observed frequency, would be $\omega_1 = \omega_0 + k_0 v$. We have made this analysis from the point of view of a man at rest; we would like to know how it would look to the man who is moving. Here we have to worry again about the difference in clock rate for the two observers, and this time that means that we have to *divide* by $\sqrt{1 - v^2/c^2}$. So if k_0 is the wave number, the number of radians

per meter in the direction of motion, and ω_0 is the frequency, then the observed frequency for a moving man is

$$\omega = \frac{\omega_0 + k_0 v}{\sqrt{1 - v^2/c^2}}. \tag{34.13}$$

For the case of light, we know that $k_0 = \omega_0/c$. So, in this particular problem, the equation would read

$$\omega = \frac{\omega_0(1 + v/c)}{\sqrt{1 - v^2/c^2}}, \tag{34.14}$$

which looks completely unlike formula (34.12)! Is the frequency that we would observe if we move toward a source different than the frequency that we would see if the source moved toward us? Of course not! The theory of relativity says that these two must be *exactly equal.* If we were expert enough mathematicians we would probably recognize that these two mathematical expressions *are* exactly equal! In fact, the *necessary* equality of the two expressions is one of the ways by which some people like to demonstrate that relativity requires a time dilation, because if we did not put those square-root factors in, they would no longer be equal.

Since we know about relativity, let us analyze it in still a third way, which may appear a little more general. (It is really the same thing, since it makes no difference *how* we do it!) According to the relativity theory there is a relationship between position and time as observed by one man and position and time as seen by another who is moving relative to him. We wrote down those relationships long ago (Chapter 16). This is the *Lorentz transformation* and its inverse:

$$\begin{aligned} x' &= \frac{x + vt}{\sqrt{1 - v^2/c^2}}, \qquad x = \frac{x' - vt'}{\sqrt{1 - v^2/c^2}}, \\ t' &= \frac{t + vx/c^2}{\sqrt{1 - v^2/c^2}}, \qquad t = \frac{t' - vx'/c^2}{\sqrt{1 - v^2/c^2}}. \end{aligned} \tag{34.15}$$

If we were standing still on the ground, the form of a wave would be $\cos(\omega t - kx)$; all the nodes and maxima and minima would follow this form. But what would a man in motion, observing the same physical wave, see? Where the field is zero, the positions of all the nodes are the same (when the field is *zero, everyone* measures the field as zero); that is a relativistic invariant. So the form is the same for the other man too, except that we must transform it into his frame of reference:

$$\cos(\omega t - kx) = \cos\left[\omega \frac{t' - vx'/c^2}{\sqrt{1 - v^2/c^2}} - k \frac{x' - vt'}{\sqrt{1 - v^2/c^2}}\right].$$

If we regroup the terms inside the brackets, we get

$$\begin{aligned} \cos(\omega t - kx) &= \cos\left[\underbrace{\frac{\omega + kv}{\sqrt{1 - v^2/c^2}}}\, t' - \underbrace{\frac{k + v\omega/c^2}{\sqrt{1 - v^2/c^2}}}\, x'\right] \\ &= \cos[\quad \omega' \quad t' - \quad k' \quad x']. \end{aligned} \tag{34.16}$$

This is again a wave, a cosine wave, in which there is a certain frequency ω', a constant multiplying t', and some other constant, k', multiplying x'. We call k' the wave number, or the number of waves per meter, for the other man. Therefore the other man will see a new frequency and a new wave number given by

$$\omega' = \frac{\omega + kv}{\sqrt{1 - v^2/c^2}}, \tag{34.17}$$

$$k' = \frac{k + \omega v/c^2}{\sqrt{1 - v^2/c^2}}. \tag{34.18}$$

If we look at (34.17), we see that it is the same formula (34.13), that we obtained by a more physical argument.

34–7 The ω, k four-vector

The relationships indicated in Eqs. (34.17) and (34.18) are very interesting, because these say that the new frequency ω' is a combination of the old frequency ω and the old wave number k, and that the new wave number is a combination of the old wave number and frequency. Now the wave number is the rate of change of phase with distance, and the frequency is the rate of change of phase with time, and in these expressions we see a close analogy with the Lorentz transformation of the position and time: if ω is thought of as being like t, and k is thought of as being like x divided by c^2, then the new ω' will be like t', and the new k' will be like x'/c^2. That is to say, *under the Lorentz transformation ω and k transform the same way as do t and x.* They constitute what we call a *four-vector*; when a quantity has four components transforming like time and space, it is a four-vector. Everything seems all right, then, except for one little thing: we said that a four-vector has to have *four components*; where are the other two components? We have seen that ω and k are like time and space in one space direction, but not in all directions, and so we must next study the problem of the propagation of light in three space dimensions, not just in one direction, as we have been doing up until now.

Suppose that we have a coordinate system, x, y, z, and a wave which is travelling along and whose wavefronts are as shown in Fig. 34–11. The wavelength of the wave is λ, but the direction of motion of the wave does not happen to be in the direction of one of the axes. What is the formula for such a wave? The answer is clearly $\cos(\omega t - ks)$, where $k = 2\pi/\lambda$ and s is the distance along the direction of motion of the wave—the component of the spatial position in the direction of motion. Let us put it this way: if $\mathbf{r}$ is the vector position of a point in space, then s is $\mathbf{r} \cdot \mathbf{e}_k$, where $\mathbf{e}_k$ is a unit vector in the direction of motion. That is, s is just $r \cos(\mathbf{r}, \mathbf{e}_k)$, the component of distance in the direction of motion. Therefore our wave is $\cos(\omega t - k\mathbf{e}_k \cdot \mathbf{r})$.

Now it turns out to be very convenient to define a vector $\mathbf{k}$, which is called the *wave vector*, which has a magnitude equal to the wave number, $2\pi/\lambda$, and is pointed in the direction of propagation of the waves:

$$\mathbf{k} = 2\pi\mathbf{e}_k/\lambda = k\mathbf{e}_k. \tag{34.19}$$

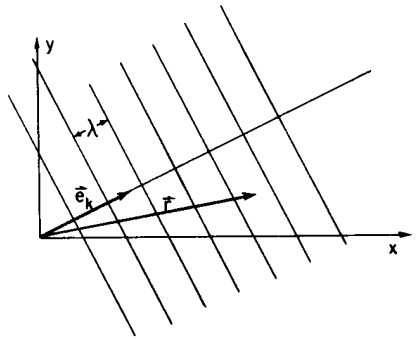


Fig. 34–11. A plane wave travelling in an oblique direction.

Using this vector, our wave can be written as $\cos(\omega t - \mathbf{k} \cdot \mathbf{r})$, or as $\cos(\omega t - k_x x - k_y y - k_z z)$. What is the significance of a component of k, say k_x? Clearly, k_x is the rate of change of phase with respect to x. Referring to Fig. 34–11, we see that the phase changes as we change x, just as if there were a wave along x, *but of a longer wavelength.* The "wavelength in the x-direction" is longer than a natural, true wavelength by the secant of the angle α between the actual direction of propagation and the x-axis:

$$\lambda_x = \lambda/\cos\alpha. \tag{34.20}$$

Therefore the rate of change of phase, which is proportional to the *reciprocal* of λ_x, is *smaller* by the factor $\cos\alpha$; that is just how k_x would vary—it would be the magnitude of $\mathbf{k}$, times the cosine of the angle between $\mathbf{k}$ and the x-axis!

That, then, is the nature of the wave vector that we use to represent a wave in three dimensions. The four quantities ω, k_x, k_y, k_z transform in relativity as a four-vector, where ω corresponds to the time, and k_x, k_y, k_z correspond to the x-, y-, and z-components of the four-vector.

In our previous discussion of special relativity (Chapter 17), we learned that there are ways of making relativistic dot products with four-vectors. If we use the position vector x_μ, where μ stands for the four components (time and three space ones), and if we call the wave vector k_μ, where the index μ again has four values, time and three space ones, then the dot product of x_μ and k_μ is written $\Sigma' k_\mu x_\mu$ (see Chapter 17). This dot product is an invariant, independent of the coordinate system; what is it equal to? By the definition of this dot product in four dimensions, it is

$$\Sigma' k_\mu x_\mu = \omega t - k_x x - k_y y - k_z z. \tag{34.21}$$

We know from our study of vectors that $\Sigma' k_\mu x_\mu$ is invariant under the Lorentz transformation, since k_μ is a four-vector. But this quantity is precisely what appears inside the cosine for a plane wave, and it *ought* to be invariant under a Lorentz transformation. We cannot have a formula with something that changes inside the cosine, since we know that the phase of the wave cannot change when we change the coordinate system.

34–8 Aberration

In deriving Eqs. (34.17) and (34.18), we have taken a simple example where **k** happened to be in a direction of motion, but of course we can generalize it to other cases also. For example, suppose there is a source sending out light in a certain direction from the point of view of a man at rest, but we are moving along on the earth, say (Fig. 34–12). From which direction does the light appear to come? To find out, we will have to write down the four components of k_μ and apply the Lorentz transformation. The answer, however, can be found by the following argument: we have to point our telescope at an angle to see the light. Why? Because light is coming down at the speed c, and we are moving sidewise at the speed v, so the telescope has to be tilted forward so that as the light comes down it goes "straight" down the tube. It is very easy to see that the horizontal distance is vt when the vertical distance is ct, and therefore, if θ' is the angle of tilt, $\tan\theta' = v/c$. How nice! How nice, indeed—except for one little thing: θ' is *not* the angle at which we would have to set the telescope *relative to the earth*, because we made our analysis from the point of view of a "fixed" observer. When we said the horizontal distance is vt, the man on the earth would have found a different distance, since he measured with a "squashed" ruler. It turns out that, because of that contraction effect,

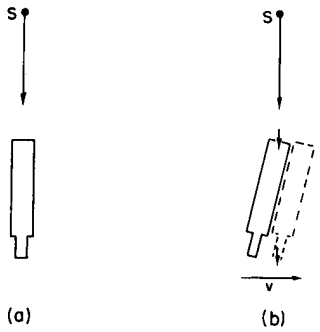


Fig. 34–12. A distant source S is viewed by (a) a stationary telescope, and (b) a laterally moving telescope.

$$\tan\theta = \frac{v/c}{\sqrt{1 - v^2/c^2}}, \tag{34.22}$$

which is equivalent to

$$\sin\theta = v/c. \tag{34.23}$$

It will be instructive for the student to derive this result, using the Lorentz transformation.

This effect, that a telescope has to be tilted, is called *aberration*, and it has been observed. *How* can we observe it? Who can say where a given star *should* be? Suppose we *do* have to look in the wrong direction to see a star; how do we know it is the wrong direction? Because the earth goes around the sun. Today we have to point the telescope one way; six months later we have to tilt the telescope the other way. That is how we can tell that there is such an effect.

34–9 The momentum of light

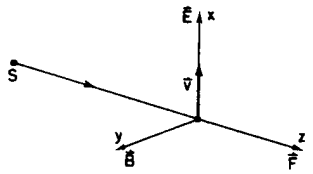


Fig. 34–13. The magnetic force on a charge which is driven by the electric field is in the direction of the light beam.

Now we turn to a different topic. We have never, in all our discussion of the past few chapters, said anything about the effects of the *magnetic* field that is associated with light. Ordinarily, the effects of the magnetic field are very small, but there is one interesting and important effect which is a consequence of the magnetic field. Suppose that light is coming from a source and is acting on a charge and driving that charge up and down. We will suppose that the electric field is in the x-direction, so the motion of the charge is also in the x-direction: it has a position x and a velocity v, as shown in Fig. 34–13. The magnetic field is at right angles to the electric field. Now as the electric field acts on the charge and moves it up and down, what does the magnetic field do? The magnetic field acts on the charge (say an electron) only when it is moving; but the electron *is* moving, it is driven by the electric field, so the two of them work together: While the thing is going up and down it has a velocity and there is a force on it, B times v times q; but in which *direction* is this force? *It is in the direction of the propagation of light.* Therefore, when light is shining on a charge and it is oscillating in response to that

charge, there is a driving force in the direction of the light beam. This is called *radiation pressure* or light pressure.

Let us determine how strong the radiation pressure is. Evidently it is $F = qvB$ or, since everything is oscillating, it is the *time average* of this, $\langle F \rangle$. From (34.2) the strength of the magnetic field is the same as the strength of the electric field divided by c, so we need to find the average of the electric field, times the velocity, times the charge, times $1/c$: $\langle F \rangle = q\langle vE \rangle/c$. But the charge q times the field E is the electric force on a charge, and the force on the charge times the velocity is the work dW/dt being done on the charge! Therefore the force, the "pushing momentum," that is delivered per second by the light, is equal to $1/c$ times the *energy absorbed* from the light per second! That is a general rule, since we did not say how strong the oscillator was, or whether some of the charges cancel out. *In any circumstance where light is being absorbed, there is a pressure.* The momentum that the light delivers is always equal to the energy that is absorbed, divided by c:

$$\langle F \rangle = \frac{dW/dt}{c}. \tag{34.24}$$

That light carries energy we already know. We now understand that it also carries *momentum*, and further, that the momentum carried is always $1/c$ times the energy.

When light is emitted from a source there is a recoil effect: the same thing in reverse. If an atom is emitting an energy W in some direction, then there is a recoil momentum $p = W/c$. If light is *reflected* normally from a mirror, we get twice the force.

That is as far as we shall go using the classical theory of light. Of course we know that there is a quantum theory, and that in many respects light acts like a particle. The energy of a light-particle is a constant times the frequency:

$$W = h\nu = \hbar\omega \tag{34.25}$$

We now appreciate that light also carries a momentum equal to the energy divided by c, so it is also true that these effective particles, these *photons*, carry a momentum

$$p = W/c = \hbar\omega/c = \hbar k. \tag{34.26}$$

The *direction* of the momentum is, of course, the direction of propagation of the light. So, to put it in vector form,

$$W = \hbar\omega, \qquad \mathbf{p} = \hbar\mathbf{k}. \tag{34.27}$$

We also know, of course, that the energy and momentum of a particle should form a four-vector. We have just discovered that ω and $\mathbf{k}$ form a four-vector. Therefore it is a good thing that (34.27) has the same constant in both cases; it means that the quantum theory and the theory of relativity are mutually consistent.

Equation (34.27) can be written more elegantly as $p_\mu = \hbar k_\mu$, a relativistic equation, for a particle associated with a wave. Although we have discussed this only for photons, for which k (the magnitude of $\mathbf{k}$) equals ω/c and $p = W/c$, the relation is much more general. In quantum mechanics *all* particles, not only photons, exhibit wavelike properties, but the frequency and wave number of the waves is related to the energy and momentum of particles by (34.27) (called the deBroglie relations) even when p is not equal to W/c.

In the last chapter we saw that a beam of right or left circularly polarized light also carries *angular momentum* in an amount proportional to the energy $\mathcal{E}$ of the wave. In the quantum picture, a beam of circularly polarized light is regarded as a stream of photons, each carrying an angular momentum $\pm\hbar$ along the direction of propagation. That is what becomes of polarization in the corpuscular point of view—the photons carry angular momentum like spinning rifle bullets. But this "bullet" picture is really as incomplete as the "wave" picture, and we shall have to discuss these ideas more fully in a later chapter on Quantum Behavior.

35

Color Vision

35–1 The human eye

35–2 Color depends on intensity

35–3 Measuring the color sensation

35–4 The chromaticity diagram

35–5 The mechanism of color vision

35–6 Physiochemistry of color vision

35–1 The human eye

The phenomenon of colors depends partly on the physical world. We discuss the colors of soap films and so on as being produced by interference. But also, of course, it depends on the eye, or what happens behind the eye, in the brain. Physics characterizes the light that enters the eye, but after that, our sensations are the result of photochemical-neural processes and psychological responses.

There are many interesting phenomena associated with vision which involve a mixture of physical phenomena and physiological processes, and the full appreciation of natural phenomena, as we *see* them, must go beyond physics in the usual sense. We make no apologies for making these excursions into other fields, because the separation of fields, as we have emphasized, is merely a human convenience, and an unnatural thing. Nature is not interested in our separations, and many of the interesting phenomena bridge the gaps between fields.

In Chapter 3 we have already discussed the relation of physics to the other sciences in general terms, but now we are going to look in some detail at a specific field in which physics and other sciences are very, very closely interrelated. That area is *vision*. In particular, we shall discuss *color vision*. In the present chapter we shall discuss mainly the observable phenomena of human vision, and in the next chapter we shall consider the physiological aspects of vision, both in man and in other animals.

It all begins with the eye; so, in order to understand what phenomena we see, some knowledge of the eye is required. In the next chapter we shall discuss in some detail how the various parts of the eye work, and how they are interconnected with the nervous system. For the present, we shall describe only briefly how the eye functions (Fig. 35–1).

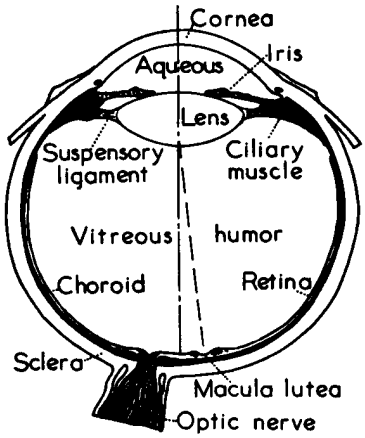


Fig. 35–1. The eye.

Light enters the eye through the *cornea*; we have already discussed how it is bent and is imaged on a layer called the *retina* in the back of the eye, so that different parts of the retina receive light from different parts of the visual field outside. The retina is not absolutely uniform: there is a place, a spot, in the center of our field of view which we use when we are trying to see things very carefully, and at which we have the greatest acuity of vision; it is called the *fovea* or *macula*. The side parts of the eye, as we can immediately appreciate from our experience in looking at things, are not as effective for seeing detail as is the center of the eye. There is also a spot in the retina where the nerves carrying all the information run out; that is a blind spot. There is no sensitive part of the retina here, and it is possible to demonstrate that if we close, say, the left eye and look straight at something, and then move a finger or another small object slowly out of the field of view it suddenly disappears somewhere. The only practical use of this fact that we know of is that some physiologist became quite a favorite in the court of a king of France by pointing this out to him; in the boring sessions that he had with his courtiers, the king could amuse himself by "cutting off their heads" by looking at one and watching another's head disappear.

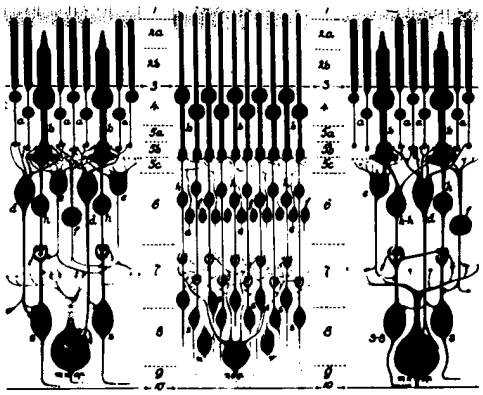


Fig. 35–2. The structure of the retina. (Light enters from below.)

Figure 35–2 shows a magnified view of the inside of the retina in somewhat schematic form. In different parts of the retina there are different kinds of structures. The objects that occur more densely near the periphery of the retina are called *rods*. Closer to the fovea, we find, besides these rod cells, also *cone* cells. We shall describe the structure of these cells later. As we get close to the fovea, the number of cones increases, and in the fovea itself there are in fact nothing but cone cells, packed very tightly, so tightly that the cone cells are much finer, or narrower

here than anywhere else. So we must appreciate that we see with the cones right in the middle of the field of view, but as we go to the periphery we have the other cells, the rods. Now the interesting thing is that in the retina each of the cells which is sensitive to light is not connected by a fiber directly to the optic nerve, but is connected to many other cells, which are themselves connected to each other. There are several kinds of cells: there are cells that carry the information toward the optic nerve, but there are others that are mainly interconnected "horizontally." There are essentially four kinds of cells, but we shall not go into these details now. The main thing we emphasize is that the light signal is already being "thought about." That is to say, the information from the various cells does not immediately go to the brain, spot for spot, but in the retina a certain amount of the information has already been digested, by a combining of the information from several visual receptors. It is important to understand that some brain-function phenomena occur in the eye itself.

35–2 Color depends on intensity

One of the most striking phenomena of vision is the dark adaptation of the eye. If we go into the dark from a brightly lighted room, we cannot see very well for a while, but gradually things become more and more apparent, and eventually we can see something where we could see nothing before. If the intensity of the light is very low, the things that we see have *no color*. It is known that this dark-adapted vision is amost entirely due to the rods, while the vision in bright light is due to the cones. As a result, there are a number of phenomena that we can easily appreciate because of this transfer of function from the cones and rods together, to just the rods.

There are many situations in which, if the light intensity were stronger, we could see color, and we would find these things quite beautiful. One example is that through a telescope we nearly always see "black and white" images of faint nebulae, but W. C. Miller of the Mt. Wilson and Palomar Observatories had the patience to make *color* pictures of some of these objects. Nobody has ever really seen these colors with the eye, but they are not artificial colors, it is merely that the light intensity is not strong enough for the cones in our eye to see them. Among the more spectacular such objects are the ring nebula and the Crab nebula. The former shows a beautiful blue inner part, with a bright red outer halo, and the latter shows a general bluish haze permeated by bright red-orange filaments.

In the bright light, apparently, the rods are at very low sensitivity but, in the dark, as time goes on they pick up their ability to see light. The variations in light intensity for which one can adapt is over a million to one. Nature does not do all this with just one kind of cell, but she passes her job from bright-light-seeing cells, the color-seeing cells, the cones, to low-intensity, dark-adapted cells, the rods. Among the interesting consequences of this shift is, first, that there is no color, and second, that there is a difference in the relative brightness of differently colored objects. It turns out that the rods see better toward the blue than the cones do, and the cones can see, for example, deep red light, while the rods find that absolutely impossible to see. So red light is black so far as the rods are concerned. Thus two pieces of colored paper, say blue and red, in which the red might be even brighter than the blue in good light, will, in the dark, appear completely reversed. It is a very striking effect. If we are in the dark and can find a magazine or something that has colors and, before we know for sure what the colors are, we judge the lighter and darker areas, and if we then carry the magazine into the light, we may see this very remarkable shift between which was the brightest color and which was not. The phenomenon is called the *Purkinje effect*.

In Fig. 35–3, the dashed curve represents the sensitivity of the eye in the dark, i.e., using the rods, while the solid curve represents it in the light. We see that the peak sensitivity of the rods is in the green region and that of the cones is more in the yellow region. If there is a red-colored page (red is about 650 mμ) we can see it if it is brightly lighted, but in the dark it is almost invisible.

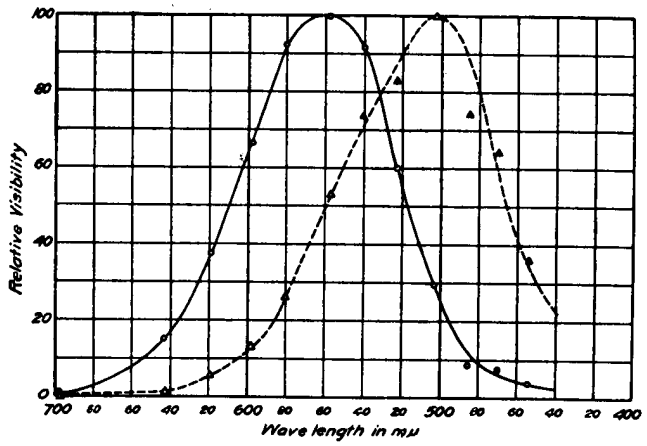


Fig. 35–3. The spectral sensitivity of the eye. Dashed curve, rods; solid curve, cones.

Another effect of the fact that rods take over in the dark, and that there are no rods in the fovea, is that when we look straight at something in the dark, our vision is not quite as acute as when we look to one side. A faint star or nebula can sometimes be seen better by looking a little to one side than directly at it, because we do not have sensitive rods in the middle of the fovea.

Another interesting effect of the fact that the number of cones decreases as we go farther to the side of the field of view is that even in a bright light color disappears as the object goes far to one side. The way to test that is to look in some particular fixed direction, let a friend walk in from one side with colored cards, and try to decide what color they are before they are right in front of you. One finds that he can see that the cards are there long before he can determine the color. When doing this, it is advisable to come in from the side opposite the blind spot, because it is otherwise rather confusing to almost see the color, then not see anything, then to see the color again.

Another interesting phenomenon is that the periphery of the retina is very sensitive to motion. Although we cannot see very well from the corner of our eye, if a little bug moves and we do not expect anything to be moving over there, we are immediately sensitive to it. We are all "wired up" to look for something jiggling to the side of the field.

35–3 Measuring the color sensation

Now we go to the cone vision, to the brighter vision, and we come to the question which is most characteristic of cone vision, and that is color. As we know, white light can be split by a prism into a whole spectrum of wavelengths which appear to us to have different colors; that is what colors are, of course: appearances. Any source of light can be analyzed by a grating or a prism, and one can determine the spectral distribution, i.e., the "amount" of each wavelength. A certain light may have a lot of blue, considerable red, very little yellow, and so on. That is all very precise in the sense of physics, but the question is, what *color* will it appear to be? It is evident that the different colors depend somehow upon the spectral distribution of the light, but the problem is to find what characteristics of the spectral distribution produce the various sensations. For example, what do we have to do to get a green color? We all know that we can simply take a piece of the spectrum which is green. But is that the *only* way to get green, or orange, or any other color?

Is there more than one spectral distribution which produces the same apparent visual effect? The answer is, definitely *yes*. There is a very limited number of visual effects, in fact just a three-dimensional manifold of them, as we shall shortly see, but there is an infinite number of different curves that we can draw for the light that comes from different sources. Now the question we have to discuss is, under what conditions do different distributions of light appear as exactly the same color to the eye?

The most powerful psycho-physical technique in color judgment is to use the eye as a *null instrument*. That is, we do not try to define what consititutes a green sensation, or to measure in what circumstances we get a green sensation, because it turns out that this is extremely complicated. Instead, we study the conditions under which two stimuli are *indistinguishable*. Then we do not have to decide whether two people see the same sensation in different circumstances, but only whether, if for one person two sensations are the same, they are also the same for another. We do not have to decide whether, when one sees something green, what it feels like inside is the same as what it feels like inside someone else when he sees something green; we do not know anything about that.

To illustrate the possibilities, we may use a series of four projector lamps which have filters on them, and whose brightnesses are continuously adjustable over a wide range: one has a red filter and makes a spot of red light on the screen, the next one has a green filter and makes a green spot, the third one has a blue filter, and the fourth one is a white circle with a black spot in the middle of it. Now if we turn on some red light, and next to it put some green, we see that in the area of overlap it produces a sensation which is not what we call reddish green, but a new color, yellow in this particular case. By changing the proportions of the red and the green, we can go through various shades of orange and so forth. If we have set it for a certain yellow, we can also obtain that same yellow, not by mixing these two colors but by mixing some other ones, perhaps a yellow filter with white light, or something like that, to get the same sensation. In other words, it is possible to make various colors in more than one way by mixing the lights from various filters.

What we have just discovered may be expressed analytically as follows. A particular yellow, for example, can be represented by a certain symbol Y, which is the "sum" of certain amounts of red-filtered light (R) and green-filtered light (G). By using two numbers, say r and g, to describe how bright the (R) and (G) are, we can write a formula for this yellow:

$$Y = rR + gG. \tag{35.1}$$

The question is, can we make *all* the different colors by adding together two or three lights of different, fixed colors? Let us see what can be done in that connection. We certainly cannot get all the different colors by mixing only red and green, because, for instance, blue never appears in such a mixture. However, by putting in some blue the central region, where all three spots overlap, may be made to appear to be a fairly nice white. By mixing the various colors and looking at this central region, we find that we can get a considerable range of colors in that region by changing the proportions, and so it is not impossible that *all* the colors can be made by mixing these three colored lights. We shall discuss to what extent this is true; it is in fact essentially correct, and we shall shortly see how to define the proposition better.

In order to illustrate our point, we move the spots on the screen so that they all fall on top of each other, and then we try to match a particular color which appears in the annular ring made by the fourth lamp. What we once thought was "white" coming from the fourth lamp now appears yellowish. We may try to match that by adjusting the red and green and blue as best we can by a kind of trial and error, and we find that we can approach rather closely this particular shade of "cream" color. So it is not hard to believe that we can make all colors. We shall try to make yellow in a moment, but before we do that, there is one color that might be very hard to make. People who give lectures on color make all the "bright" colors, but they never make *brown*, and it is hard to recall ever having seen brown light. As a matter of fact, this color is never used for any stage effect, one never sees a spotlight with brown light; so we think it might be impossible to make brown. In order to find out whether it is possible to make brown, we point out that brown light is merely something that we are not used to seeing without its background. As a matter of fact, we can make it by mixing some red and yellow. To prove that we are looking at brown light, we merely increase the brightness of the annular

background against which we see the very same light, and we see that that is, in fact, what we call brown! Brown is always a dark color next to a lighter background. We can easily change the character of the brown. For example, if we take some green out we get a reddish brown, apparently a chocolatey reddish brown, and if we put more green into it, in proportion, we get that horrible color which all the uniforms of the Army are made of, but the light from that color is not so horrible by itself; it is of yellowish green, but seen against a light background.

Now we put a yellow filter in front of the fourth light and try to match that. (The intensity must of course be within the range of the various lamps; we cannot match something which is too bright, because we do not have enough power in the lamp.) But we *can* match the yellow; we use a green and red mixture, and put in a touch of blue to make it even more perfect. Perhaps we are ready to believe that, under good conditions, we can make a perfect match of any given color.

Now let us discuss the laws of color mixture. In the first place, we found that different spectral distributions can produce the same color; next, we saw that "any" color can be made by adding together three special colors, red, blue, and green. The most interesting feature of color mixing is this: if we have a certain light, which we may call X, and if it appears indistinguishable from Y, to the eye (it may be a different spectral distribution, but it *appears* indistinguishable), we call these colors "equal," in the sense that the eye sees them as equal, and we write

$$X = Y. \tag{35.2}$$

Here is one of the great laws of color: if two spectral distributions are indistinguishable, and we *add to each one* a certain light, say Z (if we write $X + Z$, this means that we shine both lights on the same patch), and then we take Y and add the same amount of the same other light, Z, *the new mixtures are also indistinguishable:*

$$X + Z = Y + Z. \tag{35.3}$$

We have just matched our yellow; if we now shine pink light on the whole thing, it will still match. So adding any other light to the matched lights leaves a match. In other words, we can summarize all these color phenomena by saying that once we have a match between two colored lights, seen next to each other in the same circumstances, then this match will remain, and one light can be substituted for the other light in any other color mixing situation. In fact, it turns out, and it is very important and interesting, that this matching of the color of lights is not dependent upon the characteristics of the eye at the moment of observation: we know that if we look for a long time at a bright red surface, or a bright red light, and then look at a white paper, it looks greenish, and other colors are also distorted by our having looked so long at the bright red. If we now have a match between, say, two yellows, and we look at them and make them match, then we look at a bright red surface for a long time, and then turn back to the yellow, it may not look yellow any more; I do not know what color it will look, but it will not look yellow. Nevertheless *the yellows will still look matched*, and so, as the eye adapts to various levels of intensity, the color match still works, with the obvious exception of when we go into the region where the intensity of the light gets so low that we have shifted from cones to rods; then the color match is no longer a color match, because we are using a different system.

The second principle of color mixing of lights is this: *any color at all can be made from three different colors*, in our case, red, green, and blue lights. By suitably mixing the three together we can make anything at all, as we demonstrated with our two examples. Further, these laws are very interesting mathematically. For those who are interested in the mathematics of the thing, it turns out as follows. Suppose that we take our three colors, which were red, green, and blue, but label them A, B, and C, and call them our *primary* colors. Then any color could be made by certain amounts of these three: say an amount a of color A, an amount b of color B, and an amount c of color C makes X:

$$X = aA + bB + cC. \tag{35.4}$$

Now suppose another color Y is made from the same three colors:

$$Y = a'A + b'B + c'C. \tag{35.5}$$

Then it turns out that the mixture of the two lights (it is one of the consequences of the laws that we have already mentioned) is obtained by taking the sum of the components of X and Y:

$$Z = X + Y = (a + a')A + (b + b')B + (c + c')C. \tag{35.6}$$

It is just like the mathematics of the addition of vectors, where (a, b, c) are the components of one vector, and (a', b', c') are those of another vector, and the new light Z is then the "sum" of the vectors. This subject has always appealed to physicists and mathematicians. In fact, Schrödinger wrote a wonderful paper on color vision in which he developed this theory of vector analysis as applied to the mixing of colors.

Now a question is, what are the correct primary colors to use? There is no such thing as "the" correct primary colors for the mixing of lights. There may be, for practical purposes, three paints that are more useful than others for getting a greater variety of mixed pigments, but we are not discussing that matter now. *Any three differently colored lights whatsoever** can always be mixed in the correct proportion to produce *any color whatsoever*. Can we demonstrate this fantastic fact? Instead of using red, green, and blue, let us use red, blue, and yellow in our projector. Can we use red, blue, and yellow to make, say, green?

By mixing these three colors in various proportions, we get quite an array of different colors, ranging over quite a spectrum. But as a matter of fact, after a lot of trial and error, we find that nothing ever looks like green. The question is, *can* we make green? The answer is yes. How? *By projecting some red onto the green*, then we can make a match with a certain mixture of yellow and blue! So we have matched them, except that we had to cheat by putting the red on the other side. But since we have some mathematical sophistication, we can appreciate that what we really showed was not that X could always be made, say, of red, blue, and yellow, but by putting the red on the other side we found that red plus X could be made out of blue and yellow. Putting it on the other side of the equation, we can interpret that as a *negative amount*, so if we will allow that the coefficients in equations like (35.4) can be both positive and negative, and if we interpret negative amounts to mean that we have to *add* those to the *other side*, then any color can be matched by any three, and there is no such thing as "the" fundamental primaries.

We may ask whether there are three colors that come only with positive amounts for all mixings. The answer is no. Every set of three primaries requires negative amounts for some colors, and therefore there is no unique way to define a primary. In elementary books they are said to be red, green, and blue, but that is merely because with these a *wider range* of colors is available without minus signs for some of the combinations.

35–4 The chromaticity diagram

Now let us discuss the combination of colors on a mathematical level as a geometrical proposition. If any one color is represented by Eq. (35.4), we can plot it as a vector in space by plotting along three axes the amounts a, b, and c, and then a certain color is a point. If another color is a', b', c', that color is located somewhere else. The sum of the two, as we know, is the color which comes from adding these as vectors. We can simplify this diagram and represent everything on a plane by the following observation: if we had a certain color light, and merely doubled a and b and c, that is, if we make them all stronger in the same ratio, it is the same color, but brighter. So if we agree to reduce everything to the *same light intensity*, then we can project everything onto a plane, and this has been done in Fig. 35–4. It follows that any color obtained by mixing a given two in some pro-

* Except, of course, if one of the three can be matched by mixing the other two.

portion will lie somewhere on a line drawn between the two points. For instance, a fifty-fifty mixture would appear halfway between them, and 1/4 of one and 3/4 of the other would appear 1/4 of the way from one point to the other, and so on. If we use a blue and a green and a red, as primaries, we see that all the colors that we can make with positive coefficients are inside the dotted triangle, which contains almost all of the colors that we can ever see, because all the colors that we can ever see are enclosed in the oddly shaped area bounded by the curve. Where did this area come from? Once somebody made a very careful match of all the colors that we can see against three special ones. But we do not have to check *all* colors that we can see, we only have to check the pure spectral colors, the lines of the spectrum. Any light can be considered as a sum of various positive amounts of various pure spectral colors—pure from the physical standpoint. A given light will have a certain amount of red, yellow, blue, and so on—spectral colors. So if we know how much of each of our three chosen primaries is needed to make each of these pure components, we can calculate how much of each is needed to make our given color. So, if we find out what the *color coefficients* of all the spectral colors are for any given three primary colors, then we can work out the whole color mixing table.

An example of such experimental results for mixing three lights together is given in Fig. 35–5. This figure shows the amount of each of three different particular primaries, red, green and blue, which is required to make each of the spectral colors. Red is at the left end of the spectrum, yellow is next, and so on, all the way to blue. Notice that at some points minus signs are necessary. It is from such data that it is possible to locate the position of all of the colors on a chart, where the x- and the y-coordinates are related to the amounts of the different primaries that are used. That is the way that the curved boundary line has been found. It is the locus of the pure spectral colors. Now any other color can be made by adding spectral lines, of course, and so we find that anything that can be produced by connecting one part of this curve to another is a color that is available in nature. The straight line connects the extreme violet end of the spectrum with the extreme red end. It is the locus of the purples. Inside the boundary are colors that can be made with lights, and outside it are colors that cannot be made with lights, and nobody has ever seen them (except, possibly, in after-images!).

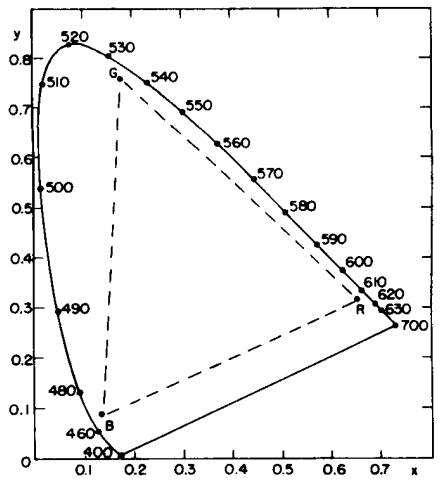


Fig. 35–4. The standard chromaticity diagram.

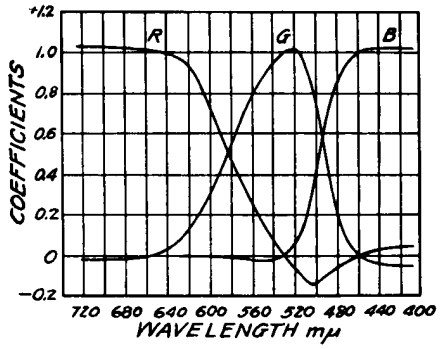


Fig. 35–5. The color coefficients of pure spectral colors in terms of a certain set of standard primary colors.

35–5 The mechanism of color vision

Now the next aspect of the matter is the question, *why* do colors behave in this way? The simplest theory, proposed by Young and Helmholtz, supposes that in the eye there are three different pigments which receive the light and that these have different absorption spectra, so that one pigment absorbs strongly, say, in the red, another absorbs strongly in the blue, another absorbs in the green. Then when we shine a light on them we will get different amounts of absorptions in the three regions, and these three pieces of information are somehow maneuvered in the brain or in the eye, or somewhere, to decide what the color is. It is easy to demonstrate that all of the rules of color mixing would be a consequence of this proposition. There has been considerable debate about the thing because the next problem, of course, is to find the absorption characteristics of each of the three pigments. It turns out, unfortunately, that because we can transform the color coordinates in any manner we want to, we can only find all kinds of linear combinations of absorption curves by the color-mixing experiments, but not the curves for the individual pigments. People have tried in various ways to obtain a specific curve which does describe some particular physical property of the eye. One such curve is called a *brightness curve*, demonstrated in Fig. 35–3. In this figure are two curves, one for eyes in the dark, the other for eyes in the light; the latter is the cone brightness curve. This is measured by finding what is the smallest amount of colored light we need in order to be able to just see it. This measures how sensitive the eye is in different spectral regions. There is another very interesting way to measure this. If we take two colors and make them appear in an area, by flickering back and forth from one to the other, we see a flicker if the frequency is too low. However, as the frequency increases,

the flicker will ultimately disappear at a certain frequency that depends on the brightness of the light, let us say at 16 repetitions per second. Now if we adjust the brightness or the intensity of one color against the other, there comes an intensity where the flicker at 16 cycles disappears. To get flicker with the brightness so adjusted, we have to go to a much lower frequency in order to see a flicker of the color. So, we get what we call a flicker of the brightness at a higher frequency and, at a lower frequency, a flicker of the color. It is possible to match two colors for "equal brightness" by this flicker technique. The results are almost, but not exactly, the same as those obtained by measuring the threshold sensitivity of the eye for seeing weak light by the cones. Most workers use the flicker system as a definition of the brightness curve.

Now, if there are three color-sensitive pigments in the eye, the problem is to determine the shape of the absorption spectrum of each one. How? We know there are people who are color blind—eight percent of the male population, and one-half of one percent of the female population. Most of the people who are color blind or abnormal in color vision have a different degree of sensitivity than others to a variation of color, but they still need three colors to match. However, there are some who are called *dichromats*, for whom any color can be matched using only *two* primary colors. The obvious suggestion, then, is to say that they are missing one of the three pigments. If we can find three kinds of color-blind dichromats who have different color-mixing rules, one kind should be missing the *red*, another the *green*, and another the *blue* pigmentation. By measuring all these types we can determine the three curves! It turns out that there *are* three types of dichromatic color blindness; there are two common types and a third very rare type, and from these three it has been possible to deduce the pigment absorption spectra.

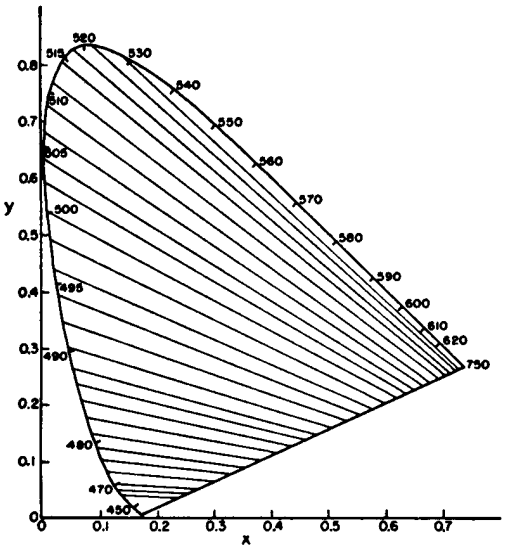


Fig. 35–6. Loci of colors confused by deuteranopes.

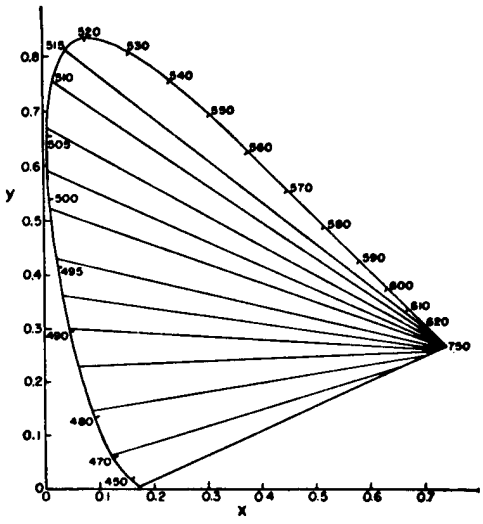


Fig. 35–7. Loci of colors confused by protanopes.

Figure 35–6 shows the color mixing of a particular type of color-blind person called a deuteranope. For him, the loci of constant colors are not points, but certain lines, along each of which the color appears to him to be the same. If the theory that he is missing one of the three pieces of information is right, all these lines should intersect at a point. If we carefully measure on this graph, they *do* intersect perfectly. Obviously, therefore, this has been made by a mathematician and does not represent real data! As a matter of fact, if we look at the latest paper with real data, it turns out that in the graph of Fig. 35–6, the point of focus of all the lines is not exactly at the right place. Using the lines in the above figure, we cannot find reasonable spectra; we need negative and positive absorptions in different regions. But using the new data of Yustova, it turns out that each of the absorption curves is everywhere positive.

Figure 35–7 shows a different kind of color blindness, that of the protanope, which has a focus near the red end of the boundary curve. Yustova gets approximately the same position in this case. Using the three different kinds of color blindness, the three pigment response curves have finally been determined, and are shown in Fig. 35–8. Finally? Perhaps. There *is* a question as to whether the three-pigment idea is right, whether color blindness results from lack of one pigment, and even whether the color-mix data on color blindness are right. Different workers get different results. This field is still very much under development.

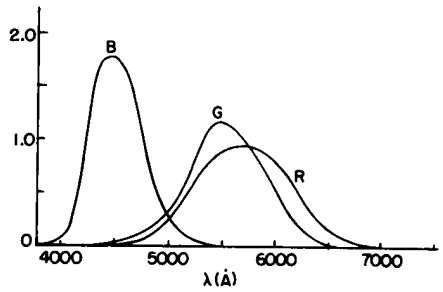


Fig. 35–8. The spectral sensitivity curves of a normal trichromat's receptors.

35–6 Physiochemistry of color vision

Now, what about checking these curves against actual pigments in the eye? The pigments that can be obtained from a retina consist mainly of a pigment called *visual purple*. The most remarkable features of this are, first, that it is in the eye of almost every vertebrate animal, and second, that its response curve fits beautifully with the sensitivity of the eye, as seen in Fig. 35–9, in which are plotted on the same scale the absorption of visual purple and the sensitivity of the dark-adapted eye. This pigment is evidently the pigment that we see with in the dark: visual purple is the pigment for the rods, and it has nothing to do with color vision. This fact was discovered in 1877. Even today it can be said that the color pigments of the cones have never been obtained in a test tube. In 1958 it could be said that the color pigments had never been seen at all. But since that time, two of them have been detected by Rushton by a very simple and beautiful technique.

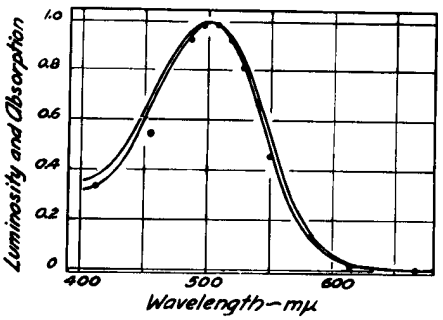


Fig. 35–9. The sensitivity curve of the dark-adapted eye, compared with the absorption curve of visual purple.

The trouble is, presumably, that since the eye is so weakly sensitive to bright light compared with light of low intensity, it needs a lot of visual purple to see with, but not much of the color pigments for seeing colors. Rushton's idea is to *leave the pigment in the eye*, and measure it anyway. What he does is this. There is an instrument called an opthalmoscope for sending light into the eye through the lens and then focusing the light that comes back out. With it one can measure how much is reflected. So one measures the reflection coefficient of light which has gone *twice* through the pigment (reflected by a back layer in the eyeball, and coming out through the pigment of the cone again). Nature is not always so beautifully designed. The cones are interestingly designed so that the light that comes into the cone bounces around and works its way down into the little sensitive points at the apex. The light goes right down into the sensitive point, bounces at the bottom and comes back out again, having traversed a considerable amount of the color-vision pigment; also, by looking at the fovea, where there are no rods, one is not confused by visual purple. But the color of the retina has been seen a long time ago: it is a sort of orangey pink; then there are all the blood vessels, and the color of the material at the back, and so on. How do we know when we are looking at the pigment? *Answer:* First we take a color-blind person, who has fewer pigments and for whom it is therefore easier to make the analysis. Second, the various pigments, like visual purple, have an intensity change when they are bleached by light; when we shine light on them they change their concentration. So, while looking at the absorption spectrum of the eye, Rushton put *another* beam in the whole eye, which changes the concentration of the pigment, and he measured the *change* in the spectrum, and the difference, of course, has nothing to do with the amount of blood or the color of the reflecting layers, and so on, but only the pigment, and in this manner Rushton obtained a curve for the pigment of the protanope eye, which is given in Fig. 35–10.

The second curve in Fig. 35–10 is a curve obtained with a normal eye. This was obtained by taking a normal eye and, having already determined what one pigment was, bleaching the other one in the red where the first one is insensitive. Red light has no effect on the protanope eye, but does in the normal eye, and thus one can obtain the curve for the missing pigment. The shape of one curve fits beautifully with Yustova's green curve, but the red curve is a little bit displaced. So perhaps we are getting on the right track. Or perhaps not—the latest work with deuteranopes does not show any definite pigment missing.

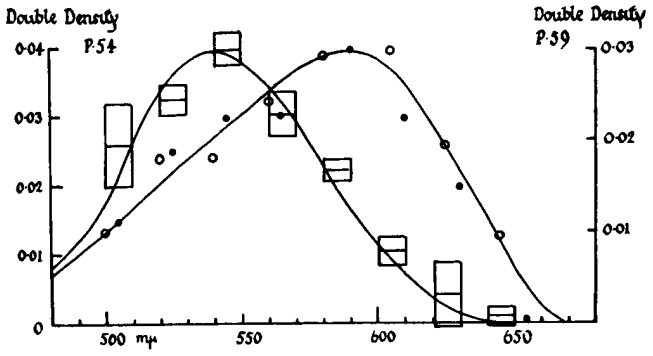


Fig. 35–10. Absorption spectrum of the color pigment of a protanope color-blind eye (squares) and a normal eye (dots).

Color is not a question of the physics of the light itself. Color is a *sensation*, and the sensation for different colors is different in different circumstances. For instance, if we have a pink light, made by superimposing crossing beams of white light and red light (all we can make with white and red is pink, obviously), we may show that white light may appear blue. If we place an object in the beams, it casts two shadows—one illuminated by the white light alone and the other by the red. For most people the "white" shadow of an object looks blue, but if we keep expanding this shadow until it covers the entire screen, we see that it suddenly appears white, not blue! We can get other effects of the same nature by mixing red, yellow, and white light. Red, yellow, and white light can produce only orangey yellows, and so on. So if we mix such lights roughly equally, we get only orange light. Nevertheless, by casting different kinds of shadows in the light, with various overlaps of colors, one gets quite a series of beautiful colors which are not in the light themselves (that is only orange), but in our *sensations*. We clearly *see* many different colors that are quite unlike the "physical" ones in the beam. It is very important to appreciate that a retina is already "thinking" about the light; it is comparing what it sees in one region with what it sees in another, although not consciously. What we know of how it does that is the subject of the next chapter.

BIBLIOGRAPHY

Committee on Colorimetry, Optical Society of America, *The Science of Color*, Thomas Y. Crowell Company, New York, 1953.

HECHT, S., S. SHLAER, and M. H. PIRENNE, "Energy, Quanta, and Vision," *Journal of General Physiology*, 1942, **25**, 819–840.

MORGAN, CLIFFORD and ELIOT STELLAR, *Physiological Psychology*, 2nd ed., McGraw-Hill Book Company, Inc., 1950.

NUBERG, N. D. and E. N. YUSTOVA, "Researches on Dichromatic Vision and the Spectral Sensitivity of the Receptors of Trichromats," presented at Symposium No. 8, *Visual Problems of Colour*, Vol. II, National Physical Laboratory, Teddington, England, September 1957. Published by Her Majesty's Stationery Office, London, 1958.

RUSHTON, W. A., "The Cone Pigments of the Human Fovea in Colour Blind and Normal," presented at Symposium No. 8, *Visual Problems of Colour*, Vol. I, National Physical Laboratory, Teddington, England, September 1957. Published by Her Majesty's Stationery Office, London, 1958.

WOODWORTH, ROBERT S., *Experimental Psychology*, Henry Holt and Company, New York, 1938. Revised edition, 1954, by Robert S. Woodworth and H. Schlosberg.

36

Mechanisms of Seeing

36–1 The sensation of color

36–2 The physiology of the eye

36–3 The rod cells

36–4 The compound (insect) eye

36–5 Other eyes

36–6 Neurology of vision

36–1 The sensation of color

In discussing the sense of sight, we have to realize that (outside of a gallery of modern art!) one does not see random spots of color or spots of light. When we look at an object we see a *man* or a *thing*; in other words, the brain interprets what we see. How it does that, no one knows, and it does it, of course, at a very high level. Although we evidently do learn to recognize what a man looks like after much experience, there are a number of features of vision which are more elementary but which also involve combining information from different parts of what we see. To help us understand how we make an interpretation of an entire image, it is worth while to study the earliest stages of the putting together of information from the different retinal cells. In the present chapter we shall concentrate mainly on that aspect of vision, although we shall also mention a number of side issues as we go along.

An example of the fact that we have an accumulation, at a very elementary level, of information from several parts of the eye at the same time, beyond our voluntary control or ability to learn, was that blue shadow which was produced by white light when both white and red were shining on the same screen. This effect at least involves the knowledge that the background of the screen is pink, even though, when we are looking at the blue shadow, it is only "white" light coming into a particular spot in the eye; somewhere, pieces of information have been put together. The more complete and familiar the context is, the more the eye will make corrections for peculiarities. In fact, Land has shown that if we mix that apparent blue and the red in various proportions, by using two photographic transparencies with absorption in front of the red and the white in different proportions, it can be made to represent a real scene, with real objects, rather faithfully. In this case we get a lot of intermediate apparent colors too, analogous to what we would get by mixing red and blue-green; it seems to be an almost complete set of colors, but if we look very hard at them, they are not so very good. Even so, it is surprising how much can be obtained from just red and white. The more the scene looks like a real situation, the more one is able to compensate for the fact that all the light is actually nothing but pink!

Another example is the appearance of "colors" in a black-and-white rotating disc, whose black and white areas are as shown in Fig. 36–1. When the disc is rotated, the variations of light and dark at any one radius are exactly the same; it is only the background that is different for the two kinds of "stripes." Yet one of the "rings" appears colored with one color and the other with another.* No one yet understands the reason for those colors, but it is clear that information is being put together at a very elementary level, in the eye itself, most likely.

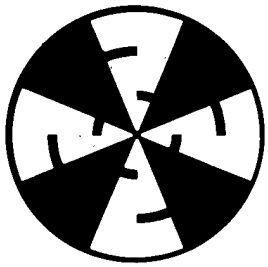


Fig. 36–1. When a disc like the above is spun, colors appear in only one of the two darker "rings." If the spin direction is reversed, the colors appear in the other ring.

Almost all present-day theories of color vision agree that the color-mixing data indicate that there are only three pigments in the cones of the eye, and that it is the spectral absorption in those three pigments that fundamentally produces the color sense. But the total sensation that is associated with the absorption characteristics of the three pigments acting together is not necessarily the sum of the individual sensations. We all agree that yellow simply *does not* seem to be reddish green; in fact it might be a tremendous surprise to most people to discover that light is, in fact, a mixture of colors, because presumably the sensation of light

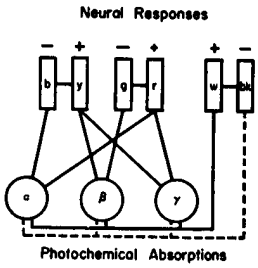
* The colors depend on speed of rotation, on the brightness of illumination, and to some extent on who looks at them and how intently he stares at them.

is due to some other process than a simple mixture like a chord in music, where the three notes are there at the same time and if we listen hard we can hear them individually. We cannot look hard and see the red and the green.

The earliest theories of vision said that there are three pigments and three kinds of cones, each kind containing one pigment; that a nerve runs from each cone to the brain, so that the three pieces of information are carried to the brain; and then in the brain, anything can happen. This, of course, is an incomplete idea: it does no good to discover that the information is carried along the optic nerve to the brain, because we have not even started to solve the problem. We must ask more basic questions: Does it make any difference *where* the information is put together? Is it important that it be carried right up into the brain in the optic nerve, or could the retina do some analysis first? We have seen a picture of the retina as an extremely complicated thing with lots of interconnections (Fig. 35–2) and it might make some analyses.

As a matter of fact, people who study anatomy and the development of the eye have shown that the retina is, in fact, the brain: in the development of the embryo, a piece of the brain comes out in front, and long fibers grow back, connecting the eyes to the brain. The retina is organized in just the way the brain is organized and, as someone has beautifully put it, "The brain has developed a way to look out upon the world." The eye is a piece of brain that is touching light, so to speak, on the outside. So it is not at all unlikely that some analysis of the color has already been made in the retina.

This gives us a very interesting opportunity. None of the other senses involves such a large amount of calculation, so to speak, before the signal gets into a nerve that one can make measurements on. The calculations for all the rest of the senses usually happen in the brain itself, where it is very difficult to get at specific places to make measurements, because there are so many interconnections. Here, with the visual sense, we have the light, three layers of cells making calculations, and the results of the calculations being transmitted through the optic nerve. So we have the first chance to observe physiologically how, perhaps, the first layers of the brain work in their first steps. It is thus of double interest, not simply interesting for vision, but interesting to the whole problem of physiology.



$y\text{-}b = k_1(\beta+\gamma-2\alpha)$
$r\text{-}g = k_2(\alpha+\gamma-2\beta)$
$w\text{-}bk = k_3(\alpha+\gamma+\beta)-k_4(\alpha+\beta+\gamma)$

Fig. 36–2. Neural connections according to an "opponent" theory of color vision.

The fact that there are three pigments does not mean that there must be three kinds of sensations. One of the other theories of color vision has it that there are really opposing color schemes (Fig. 36–2). That is, one of the nerve fibers carries a lot of impulses if there is yellow being seen, and less than usual for blue. Another nerve fiber carries green and red information in the same way, and another, white and black. In other words, in this theory someone has already started to make a guess as to the system of wiring, the method of calculation.

The problems we are trying to solve by guessing at these first calculations are questions about the apparent colors that are seen on a pink background, what happens when the eye is adapted to different colors, and also the so-called psychological phenomena. The psychological phenomena are of the nature, for instance, that white does not "feel" like red and yellow and blue, and this theory was advanced because the psychologists say that there are *four* apparent pure colors: "There are four stimuli which have a remarkable capacity to evoke psychologically simple blue, yellow, green, and red hues respectively. Unlike sienna, magenta, purple, or most of the discriminable colors, these simple hues are unmixed in the sense that none partakes of the nature of the other; specifically, blue is not yellowish, reddish, or greenish, and so on; these are psychologically primary hues." That is a psychological fact, so-called. To find out from what evidence this psychological fact was deduced, we must search very hard indeed through all the literature. In the modern literature all we find on the subject are repeats of the same statement, or of one by a German psychologist, who uses as one of his authorities Leonardo da Vinci, who, of course, we all know was a great artist. He says, "Leonardo thought there were five colors." Then, looking still further, we find, in a still older book, the evidence for the subject. The book says something like this: "Purple is reddish-blue, orange is reddish-yellow, but can red be seen as purplish-orange? Are not red and yellow more unitary than purple or orange? The average person, asked to state which

colors are unitary, names red, yellow, and blue, these three, and some observers add a fourth, green. Psychologists are accustomed to accept the four as salient hues." So that is the situation in the psychological analysis of this matter: if everybody says there are three, and somebody says there are four, and they want it to be four, it will be four. That shows the difficulty with psychological researches. It is clear that we have such feelings, but it is very difficult to obtain much information about them.

So the other direction to go is the physiological direction, to find out experimentally what actually happens in the brain, the eye, the retina, or wherever, and perhaps to discover that some combinations of impulses from various cells move along certain nerve fibers. Incidentally, primary pigments do not have to be in separate cells; one could have cells in which are mixtures of the various pigments, cells with the red and the green pigments, cells with all three (the information of all three is then white information), and so on. There are many ways of hooking the system up, and we have to find out which way nature has used. It would be hoped, ultimately, that when we understand the physiological connections we will have a little bit of understanding of some of those aspects of the psychology, so we look in that direction.

36–2 The physiology of the eye

We begin by talking not only about color vision, but about vision in general, just to remind ourselves about the interconnections in the retina, shown in Fig. 35–2. The retina is really like the surface of the brain. Although the actual picture through a microscope is a little more complicated looking than this somewhat schematized drawing, by careful analysis one can see all these interconnections. There is no question that one part of the surface of the retina is connected to other parts, and that the information that comes out on the long axons, which produce the optic nerve, are combinations of information from many cells. There are three layers of cells in the succession of function: there are retinal cells, which are the ones that the light affects, an intermediate cell which takes information from a single or a few retinal cells and gives it out again to several cells in a third layer of cells and carries it to the brain. There are all kinds of cross connections between cells in the layers.

We now turn to some aspects of the structure and performance of the eye (see Fig. 35–1). The focusing of the light is accomplished mainly by the cornea, by the fact that it has a curved surface which "bends" the light. This is why we cannot see clearly under water, because we then do not have enough difference between the index of the cornea, which is 1.37, and that of the water, which is 1.33. Behind the cornea is water, practically, with an index of 1.33, and behind that is a lens which has a very interesting structure: it is a series of layers, like an onion, except that it is all transparent, and it has an index of 1.40 in the middle and 1.38 at the outside. (It would be nice if we could make optical glass in which we could adjust the index throughout, for then we would not have to curve it as much as we do when we have a uniform index.) Furthermore, the shape of the cornea is not that of a sphere. A spherical lens has a certain amount of spherical aberration. The cornea is "flatter" at the outside than is a sphere, in just such a manner that the spherical aberration is less for the cornea than it would be if we put a spherical lens in there! The light is focused by the cornea-lens system onto the retina. As we look at things that are closer and farther away, the lens tightens and loosens and changes the focus to adjust for the different distances. To adjust for the total amount of light there is the iris, which is what we call the color of the eye, brown or blue, depending on who it is; as the amount of light increases and decreases, the iris moves in and out.

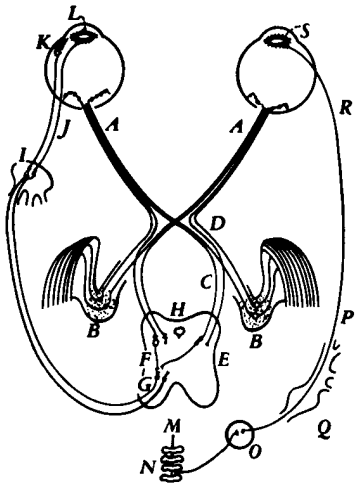


Fig. 36–3. The neural interconnections for the mechanical operation of the eyes.

Let us now look at the neural machinery for controlling the accommodation of the lens, the motion of the eye, the muscles which turn the eye in the socket, and the iris, shown schematically in Fig. 36–3. Of all the information that comes out of the optic nerve A, the great majority is divided into one of two bundles (which we will talk about later) and thence to the brain. But there are a few fibers, of

interest to us now, which do not run directly to the visual cortex of the brain where we "see" the images, but instead go into the mid-brain *H*. These are the fibers which measure the average light and make adjustment for the iris; or, if the image looks foggy, they try to correct the lens; or, if there is a double image, they try to adjust the eye for binocular vision. At any rate, they go through the mid-brain and feed back into the eye. At *K* are the muscles which run the accommodation of the lens, and at *L* another one that runs into the iris. The iris has two muscle systems. One is a circular muscle *L* which, when it is excited, pulls in and closes down the iris; it acts very rapidly and the nerves are directly connected from the brain through short axons into the iris. The opposite muscles are radial muscles, so that, when the things get dark and the circular muscle relaxes, these radial muscles pull out. Here we have, as in many places in the body, a pair of muscles which work in opposite directions, and in almost every such case the nerve systems which control the two are very delicately adjusted, so that when signals are sent in to tighten one, signals are automatically sent in to loosen the other. The iris is a peculiar exception: the nerves which make the iris contract are the ones we have already described, but the nerves which make the iris *expand* come out from no one knows exactly where, go down into the spinal cord back of the chest, into the thoracic sections, out of the spinal cord, up through the neck ganglia, and all the way around and back up into the head in order to run the other end of the iris. In fact, the signal goes through a completely different nervous system, not the central nervous system at all, but the sympathetic nervous system, so it is a very strange way of making things go.

We have already emphasized another strange thing about the eye, that the light-sensitive cells are on the wrong side, so that the light has to go through several layers of other cells before it gets to the receptors—it is built inside out! So some of the features are wonderful and some are apparently stupid.

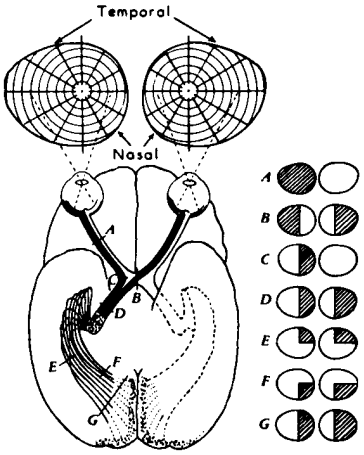


Fig. 36–4. The neural connections from the eyes to the visual cortex.

Figure 36–4 shows the connections of the eye to the part of the brain which is most directly concerned with the visual process. The optic nerve fibers run into a certain area just beyond *D*, called the lateral geniculate, whereupon they run out to a section of the brain called the visual cortex. Notice that some of the fibers from each eye are sent over to the other side of the brain, so the picture formed is incomplete. The optic nerves from the left side of the right eye run across the optic chiasma *B*, while the ones on the left side of the left eye come around and go this same way. So the left side of the brain receives all the information which comes from the left side of the eyeball of each eye, i.e., on the right side of the visual field, while the right side of the brain sees the left side of the visual field. This is the manner in which the information from each of the two eyes is put together in order to tell how far away things are. This is the system of binocular vision.

The connections between the retina and the visual cortex are interesting. If a spot in the retina is excised or destroyed in any way, then the whole fiber will die, and we can thereby find out where it is connected. It turns out that, essentially, the connections are one to one—for each spot in the retina there is one spot in the visual cortex—and spots that are very close together in the retina are very close together in the visual cortex. So the visual cortex still represents the spatial arrangement of the rods and cones, but of course much distorted. Things which are in the center of the field, which occupy a very small part of the retina, are expanded over many, many cells in the visual cortex. It is clear that it is useful to have things which are originally close together, still close together. The most remarkable aspect of the matter, however, is the following. The place where one would think it would be most important to have things close together would be right in the middle of the visual field. Believe it or not, the up-and-down line in our visual field as we look at something is of such a nature that the information from all the points on the right side of that line is going into the left side of the brain, and information from the points on the left side is going into the right side of the brain, and the way this area is made, there is a cut right down through the middle, so that the things that are very close together right in the middle are very far apart in the brain! Somehow, the information has to go from one side of the brain to the other through some other channels, which is quite surprising.

The question of how this network ever gets "wired" together is very interesting. The problem of how much is already wired and how much is learned is an old one. It used to be thought long ago that perhaps it does not have to be wired carefully at all, it is only just roughly interconnected, and then, by experience, the young child learns that when a thing is "up there" it produces some sensation in the brain. (Doctors always tell us what the young child "feels," but how do *they* know what a child feels at the age of one?) The child, at the age of one, supposedly sees that an object is "up there," gets a certain sensation, and learns to reach "there," because when he reaches "here," it does not work. That approach probably is not correct, because we already see that in many cases there are these special detailed interconnections. More illuminating are some most remarkable experiments done with a salamander. (Incidentally, with the salamander there is a direct crossover connection, without the optic chiasma, because the eyes are on each side of the head and have no common area. Salamanders do not have binocular vision.) The experiment is this. We can cut the optic nerve in a salamander and the nerve will grow out from the eyes again. Thousands and thousands of cell fibers will thus re-establish themselves. Now, in the optic nerve the fibers do not stay adjacent to each other—it is like a great, sloppily made telephone cable, all the fibers twisting and turning, but when it gets to the brain they are all sorted out again. When we cut the optic nerve of the salamander, the interesting question is, will it ever get straightened out? The answer is remarkable: yes. If we cut the optic nerve of the salamander and it grows back, the salamander has good visual acuity again. However, if we cut the optic nerve and *turn the eye upside down* and let it grow back again, it has good visual acuity all right, but it has a terrible error: when the salamander sees a fly "up here," it jumps at it "down there," and it never learns. Therefore there is some mysterious way by which the thousands and thousands of fibers find their right places in the brain.

This problem of how much is wired in, and how much is not, is an important problem in the theory of the development of creatures. The answer is not known, but is being studied intensively.

The same experiment in the case of a goldfish shows that there is a terrible knot, like a great scar or complication, in the optic nerve where we cut it, but in spite of all this the fibers grow back to their right places in the brain.

In order to do this, as they grow into the old channels of the optic nerve they must make several decisions about the direction in which they should grow. How do they do this? There seem to be chemical clues that different fibers respond to differently. Think of the enormous number of growing fibers, each of which is an individual differing in some way from its neighbors; in responding to whatever the chemical clues are, it responds in a unique enough way to find its proper place for ultimate connection in the brain! This is an interesting—a fantastic—thing. It is one of the great recently discovered phenomena of biology and is undoubtedly connected to many older unsolved problems of growth, organization, and development of organisms, and particularly of embryos.

One other interesting phenomenon has to do with the motion of the eye. The eyes must be moved in order to make the two images coincide in different circumstances. These motions are of different kinds: one is to follow something, which requires that both eyes must go in the same direction, right or left, and the other is to point them toward the same place at various distances away, which requires that they must move oppositely. The nerves going into the muscles of the eye are already wired up for just such purposes. There is one set of nerves which will pull the muscles on the inside of one eye and the outside of the other, and relax the opposite muscles, so that the two eyes move together. There is another center where an excitation will cause the eyes to move in toward each other from parallel. Either eye can be turned out to the corner if the other eye moves toward the nose, but it is impossible consciously or unconsciously to turn both eyes *out* at the same time, not because there are no *muscles*, but because there is no way to send a signal to turn both eyes out, unless we have had an accident or there is something the matter, for instance if a nerve has been cut. Although the muscles of one eye can certainly steer that eye about, not even a Yogi is able to move *both* eyes out freely

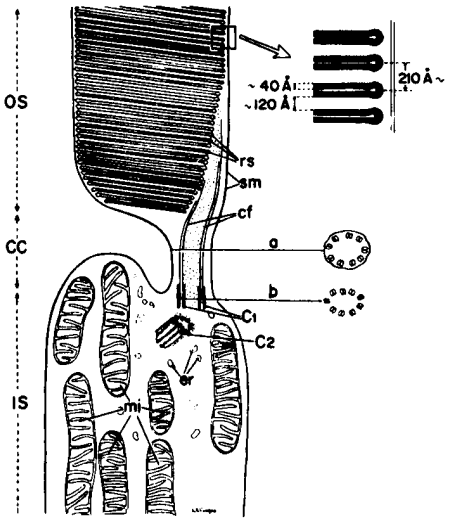


Fig. 36–5. Electron micrograph of a rod cell.

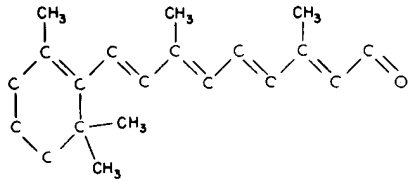


Fig. 36–6. The structure of retinene.

under voluntary control, because there does not seem to be any way to do it. We are already wired to a certain extent. This is an important point, because most of the earlier books on anatomy and psychology, and so on, do not appreciate or do not emphasize the fact that we are so completely wired already—they say that everything is just learned.

36–3 The rod cells

Let us now examine in more detail what happens in the rod cells. Figure 36–5 shows an electron micrograph of the middle of a rod cell (the rod cell keeps going up out of the field). There are layer after layer of plane structures, shown magnified at the right, which contain the substance rhodopsin (visual purple), the dye, or pigment, which produces the effects of vision in the rods. The rhodopsin, which is the pigment, is a big protein which contains a special group called retinene, which can be taken off the protein, and which is, undoubtedly, the main cause of the absorption of light. We do not understand the reason for the planes, but it is very likely that there is some reason for holding all the rhodopsin molecules parallel. The chemistry of the thing has been worked out to a large extent, but there might be some physics to it. It may be that all of the molecules are arranged in some kind of a row so that when one is excited an electron which is generated, say, may run all the way down to some place at the end to get the signal out, or something. This subject is very important, and has not been worked out. It is a field in which both biochemistry and solid state physics, or something like it, will ultimately be used.

This kind of a structure, with layers, appears in other circumstances where light is important, for example in the chloroplast in plants, where the light causes photosynthesis. If we magnify those, we find the same thing with almost the same kind of layers, but there we have chlorophyll, of course, instead of retinene. The chemical form of retinene is shown in Fig. 36–6. It has a series of alternate double bonds along the side chain, which is characteristic of nearly all strongly absorbing organic substances, like chlorophyll, blood, and so on. This substance is impossible for human beings to manufacture in their own cells—we have to eat it. So we eat it in the form of a special substance, which is exactly the same as retinene except that there is a hydrogen tied on the right end; it is called vitamin A, and if we do not eat enough of it, we do not get a supply of retinene, and the eye becomes what we call *night blind*, because there is then not enough pigment in the rhodopsin to see with the rods at night.

The reason why such a series of double bonds absorbs light very strongly is also known. We may just give a hint: The alternating series of double bonds is called a *conjugated* double bond; a double bond means that there is an extra electron there, and this extra electron is easily shifted to the right or left. When light strikes this molecule, the electron of each double bond is shifted over by one step. All the electrons in the whole chain shift, like a string of dominoes falling over, and though each one moves only a little distance (we would expect that, in a single atom, we could move the electron only a little distance), the net effect is the same as though the one at the end was moved over to the other end! It is the same as though one electron went the whole distance back and forth, and so, in this manner, we get a much stronger absorption under the influence of the electric field, than if we could only move the electron a distance which is associated with one atom. So, since it is easy to move the electrons back and forth, retinene absorbs light very strongly; that is the machinery of the physical-chemical end of it.

36–4 The compound (insect) eye

Let us now return to biology. The human eye is not the only kind of eye. In the vertebrates, almost all eyes are essentially like the human eye. However, in the lower animals there are many other kinds of eyes: eye spots, various eye cups, and other less sensitive things, which we have no time to discuss. But there is one other highly developed eye among the invertebrates, the *compound* eye of the insect. (Most insects having large compound eyes also have various additional simpler eyes as well.) A bee is an insect whose vision has been studied very carefully. It is

easy to study the properties of the vision of bees because they are attracted to honey, and we can make experiments in which we identify the honey by putting it on blue paper or red paper, and see which one they come to. By this method some very interesting things have been discovered about the vision of the bee.

In the first place, in trying to measure how acutely bees could see the color difference between two pieces of "white" paper, some researchers found they were not very good, and others found they were fantastically good. Even if the two pieces of white paper were almost exactly the same, the bees could still tell the difference. The experimenters used zinc white for one piece of paper and lead white for the other, and although these look exactly the same to us, the bee could easily distinguish them, because they reflect a different amount in the ultraviolet. In this way it was discovered that the bee's eye is sensitive over a wider range of the spectrum than is our own. Our eye works from 7000 angstroms to 4000 angstroms, from red to violet, but the bee's can see down to 3000 angstroms into the ultraviolet! This makes for a number of different interesting effects. In the first place, bees can distinguish between many flowers which to us look alike. Of course, we must realize that the colors of flowers are not designed for *our* eyes, but for the bee; they are signals to attract the bees to a specific flower. We all know that there are many "white" flowers. Apparently white is not very interesting to the bees, because it turns out that all of the white flowers have different proportions of reflection in the *ultraviolet*; they do not reflect one hundred percent of the ultraviolet as would a true white. All the light is not coming back, the ultraviolet is missing, and that is a color, just as, for us, if the blue is missing, it comes out yellow. So, all the flowers are colored for the bees. However, we also know that red cannot be seen by bees. Thus we might expect that all red flowers should look black to the bee. Not so! A careful study of red flowers shows, first, that even with our own eye we can see that a great majority of red flowers have a bluish tinge because they are mainly reflecting an additional amount in the blue, which is the part that the bee sees. Furthermore, experiments also show that flowers vary in their reflection of the ultraviolet over different parts of the petals, and so on. So if we could see the flowers as bees see them they would be even more beautiful and varied!

It has been shown, however, that there are a few red flowers which do *not* reflect in the blue or in the ultraviolet, and *would*, therefore, appear black to the bee! This was of quite some concern to the people who worry about this matter, because black does not seem like an interesting color, since it is hard to tell from a dirty old shadow. It actually turned out that these flowers were *not* visited by bees, these are the flowers that are visited by *hummingbirds*, and hummingbirds *can* see the red!

Another interesting aspect of the vision of the bee is that bees can apparently tell the direction of the sun by looking at a patch of blue sky, without seeing the sun itself. We cannot easily do this. If we look out the window at the sky and see that it is blue, in which direction is the sun? The bee can tell, because the bee is quite sensitive to the *polarization* of light, and the scattered light of the sky is polarized.* There is still some debate about how this sensitivity operates. Whether it is because the reflections of the light are different in different circumstances, or the bee's eye is directly sensitive, is not yet known.†

It is also said that the bee can notice flicker up to 200 oscillations per second, while we see it only up to 20. The motions of bees in the hives are very quick; the feet move and the wings vibrate, but it is very hard for us to see these motions with our eye. However, if we could see more rapidly we would be able to see the motion. It is probably very important to the bee that its eye has such a rapid response.

---

* The human eye also has a slight sensitivity to the polarization of light, and one *can* learn to tell the direction of the sun! The phenomenon that is involved here is called *Haidinger's brush;* it is a faint, yellowish hourglass-like pattern seen at the center of the visual field when one looks at a broad, featureless expanse using polarizing glasses. It can also be seen in the blue sky without polarizing glasses if one rotates his head back and forth about the axis of vision.

† Evidence obtained since this lecture was given indicates that the eye is directly sensitive.



Fig. 36–7. The structure of an ommatidium (a single cell of a compound eye).

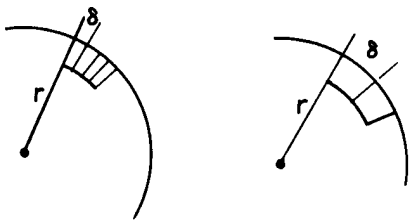


Fig. 36–8. Schematic view of packing of ommatidia in the eye of a bee.

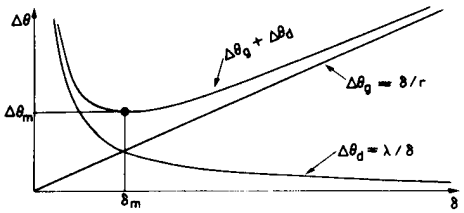


Fig. 36–9. The optimum size for an ommatidium is δ_m.

Now let us discuss the visual acuity we could expect from the bee. The eye of a bee is a compound eye, and it is made of a large number of special cells called *ommatidia*, which are arranged conically on the surface of a sphere (roughly) on the outside of the bee's head. Figure 36–7 shows a picture of one such ommatidium. At the top there is a transparent area, a kind of "lens," but actually it is more like a filter or light pipe to make the light come down along the narrow fiber, which is where the absorption presumably occurs. Out of the other end of it comes the nerve fiber. The central fiber is surrounded on its sides by six cells which, in fact, have secreted the fiber. That is enough description for our purposes; the point is that it is a conical thing and many can fit next to each other all over the surface of the eye of the bee.

Now let us discuss the resolution of the eye of the bee. If we draw lines (Fig. 36–8) to represent the ommatidia on the surface, which we suppose is a sphere of radius r, we may actually *calculate* how wide each ommatidium is by using our brains, and assuming that evolution is as clever as we are! If we have a very large ommatidium we do not have much resolution. That is, one cell gets a piece of information from one direction, and the adjacent cell gets a piece of information from another direction, and so on, and the bee cannot see things in between very well. So the uncertainty of visual acuity in the eye will surely correspond to an angle, the angle of the end of the ommatidium relative to the center of curvature of the eye. (The eye cells, of course, exist only at the surface of the sphere; inside that is the head of the bee.) This angle, from one ommatidium to the next, is, of course, the diameter of the ommatidia divided by the radius of the eye surface:

$$\Delta\theta_g = \delta/r. \tag{36.1}$$

So, we may say, "The finer we make the δ, the more the visual acuity. So why doesn't the bee just use very, very fine ommatidia?" *Answer*: We know enough physics to realize that if we are trying to get light down into a narrow slot, we cannot see accurately in a given direction because of the diffraction effect. The light that comes from several directions can enter and, due to diffraction, we will get light coming in at angle $\Delta\theta_d$ such that

$$\Delta\theta_d = \lambda/\delta. \tag{36.2}$$

Now we see that if we make the δ too small, then each ommatidium does not look in only one direction, because of diffraction! If we make them too big, each one sees in a definite direction, but there are not enough of them to get a good view of the scene. So we adjust the distance d in order to make minimal the total effect of these two. If we add the two together, and find the place where the sum has a minimum (Fig. 36–9), we find that

$$\frac{d(\Delta\theta_g + \Delta\theta_d)}{d\delta} = 0 = \frac{1}{r} - \frac{\lambda}{\delta^2}, \tag{36.3}$$

which gives us a distance

$$\delta = \sqrt{\lambda r}. \tag{36.4}$$

If we guess that r is about 3 millimeters, take the light that the bee sees as 4000 angstroms, and put the two together and take the square root, we find

$$\begin{aligned}\delta &= (3 \times 10^{-3} \times 4 \times 10^{-7})^{1/2}\ \text{m} \\ &= 3.5 \times 10^{-5}\ \text{m} = 35\,\mu.\end{aligned} \tag{36.5}$$

The book says the diameter is $30\,\mu$, so that is rather good agreement! So, apparently, it really works, and we can understand what determines the size of the bee's eye! It is also easy to put the above number back in and find out how good the bee's eye actually is in angular resolution; it is very poor relative to our own. We can see things that are thirty times smaller in apparent size than the bee; the bee has a rather fuzzy out-of-focus image relative to what we can see. Nevertheless it is all right, and it is the best they can do. We might ask why the bees do not develop

a good eye like our own, with a lens and so on. There are several interesting reasons. In the first place, the bee is too small; if it had an eye like ours, but on his scale, the opening would be about 30 μ in size and diffraction would be so important that it would not be able to see very well anyway. The eye is not good if it is too small. Secondly, if it were as big as the bee's head, then the eye would occupy the whole head of the bee. The beauty of the compound eye is that it takes up no space, it is just a very thin layer on the surface of the bee. So when we argue that they should have done it our way, we must remember that they had their own problems!

36–5 Other eyes

Besides the bees, many other animals can see color. Fish, butterflies, birds, and reptiles can see color, but it is believed that most mammals cannot. The primates can see color. The birds certainly see color, and that accounts for the colors of birds. There would be no point in having such brilliantly colored males if the females could not notice it! That is, the evolution of the sexual "whatever it is" that the birds have is a result of the female being able to see color. So next time we look at a peacock and think of what a brilliant display of gorgeous color it is, and how delicate all the colors are, and what a wonderful aesthetic sense it takes to appreciate all that, we should not compliment the peacock, but should compliment the visual acuity and aesthetic sense of the pea*hen*, because that is what has generated the beautiful scene!

All invertebrates have poorly developed eyes or compound eyes, but all the vertebrates have eyes very similar to our own, with one exception. If we consider the highest form of animal, we usually say, "Here we are!," but if we take a less prejudiced point of view and restrict ourselves to the invertebrates, so that we cannot include ourselves, and ask what is the highest invertebrate animal, most zoologists agree that the *octopus* is the highest animal! It is very interesting that, besides the development of its brain and its reactions and so on, which are rather good for an invertebrate, it has also developed, independently, a different eye. It is not a compound eye or an eye spot—it has a cornea, it has lids, it has an iris, it has a lens, it has two regions of water, it has a retina behind. It is essentially the same as the eye of the vertebrates! It is a remarkable example of a coincidence in evolution where nature has twice discovered the same solution to a problem, with one slight improvement. In the octopus it also turns out, amazingly, that the retina is a piece of the brain that has come out in the same way in its embryonic development as is true for vertebrates, but the interesting thing which is different is that the cells which are sensitive to light are on the *inside*, and the cells which do the calculation are in back of them, rather than "inside out," as in our eye. So we see, at least, that there is no good reason for its being inside out. The other time nature tried it, she got it straightened out! (See Fig. 36–10.) The biggest eyes in the world are those of the giant squid; they have been found up to 15 inches in diameter!

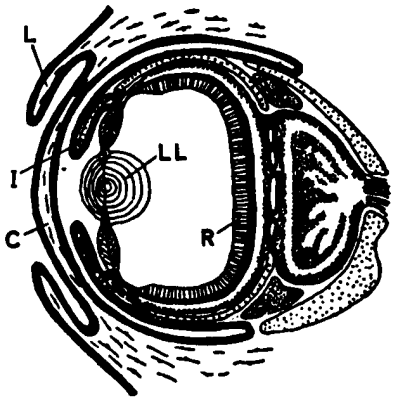


Fig. 36–10. The eye of an octopus.

36–6 Neurology of vision

One of the main points of our subject is the interconnection of information from one part of the eye to the other. Let us consider the compound eye of the horseshoe crab, on which considerable experimentation has been done. First of all, we must appreciate what kind of information can come along nerves. A nerve carries a kind of disturbance which has an electrical effect that is easy to detect, a kind of wavelike disturbance which runs down the nerve and produces an effect at the other end: a long piece of the nerve cell, called the axon, carries the information along, and a certain kind of impulse, called a "spike," goes along if it is excited at one end. When one spike goes down the nerve, another cannot immediately follow. All the spikes are of the same size, so it is not that we get *higher* spikes when the thing is more strongly excited, but that we get *more spikes per second.* The *size* of the spike is determined by the fiber. It is important to appreciate this in order to see what happens next.

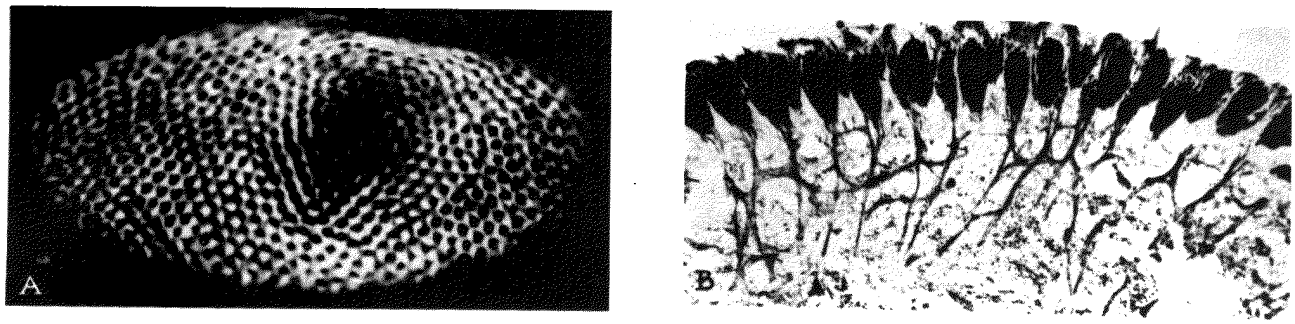


Fig. 36–11. The compound eye of the horseshoe crab. (a) Normal view. (b) Cross section.

Figures 36–7, 11, 12, 13 reprinted with permission from Goldsmith, *Sensory Communications*, W. A. Rosenblith, ed. Copyright 1961, Massachusetts Institute of Technology.

Figure 36–11(a) shows the compound eye of the horseshoe crab; it is not very much of an eye, it has only about a thousand ommatidia. Figure 36–11(b) is a cross section through the system; one can see the ommatidia, with the nerve fibers that run out of them and go into the brain. But note that even in a horseshoe crab there are little interconnections. They are much less elaborate than in the human eye, and it gives us a chance to study a simpler example.

Let us now look at the experiments which have been done by putting fine electrodes into the optic nerve of the horseshoe crab, and shining light on only one of the ommatidia, which is easy to do with lenses. If we turn a light on at some instant t_0, and measure the electrical pulses that come out, we find that there is a slight delay and then a rapid series of discharges which gradually slow down to a uniform rate, as shown in Fig. 36–12(a). When the light goes out, the discharge stops. Now it is very interesting that if, while our amplifier is connected to this same nerve fiber, we shine light on a *different* ommatidium nothing happens; no signal.

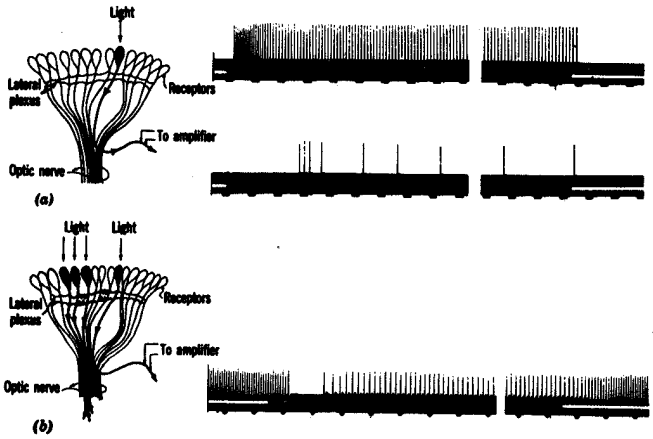


Fig. 36–12. The response to light of the nerve fibers of the eye of the horseshoe crab.

Now we do another experiment: we shine the light on the original ommatidium and get the same response, but if we now turn light on another one nearby as well, the pulses are interrupted briefly and then run at a much lower rate (Fig. 36–12b). The rate of one is inhibited by the impulses which are coming out of the other! In other words, each nerve fiber carries the information from one ommatidium, but the amount that it carries is inhibited by the signals from the others. So, for example, if the whole eye is more or less uniformly illuminated, the information coming from any one ommatidium will be relatively weak, because it is inhibited by so many. In fact the inhibition is additive—if we shine light on several nearby ommatidia the inhibition is very great. The inhibition is greater when the ommatidia are closer, and if the ommatidia are far enough away from one another, inhibition is practically zero. So it is additive and depends on the distance; here is a first example of information from different parts of the eye being combined in the eye itself. We can see, perhaps, if we think about it awhile, that this is a device to *enhance contrast* at the edges of objects, because if a part of the scene is light and a part is black, then the ommatidia in the lighted area give impulses that are

inhibited by all the other light in the neighborhood, so it is relatively weak. On the other hand, an ommatidium at the boundary which is given a "white" impulse is also inhibited by others in the neighborhood, but there are not as many of them, since some are black; the net signal is therefore stronger. The result would be a curve, something like that of Fig. 36–13. The crab will see an enhancement of the contour.

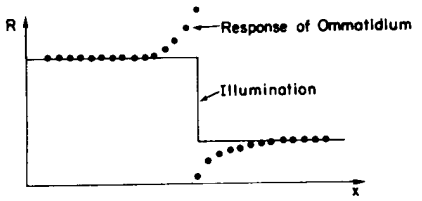


Fig. 36–13. The net response of horseshoe crab ommatidia near a sharp change in illumination.

The fact that there is an enhancement of contours has long been known; in fact it is a remarkable thing that has been commented on by psychologists many times. In order to draw an object, we have only to draw its outline. How used we are to looking at pictures that have only the outline! What is the outline? The outline is only the edge difference between light and dark or one color and another. It is not something definite. It is not, believe it or not, that every object has a line around it! There is no such line. It is only in our own psychological makeup that there is a line; we are beginning to understand the reasons why the "line" is enough of a clue to get the whole thing. Presumably our own eye works in some similar manner—much more complicated, but similar.

Finally, we shall briefly describe the more elaborate work, the beautiful, advanced work that has been done on the frog. Doing a corresponding experiment on a frog, by putting very fine, beautifully built needlelike probes into the optic nerve of a frog, one can obtain the signals that are going along one particular axon and, just as in the case of the horseshoe crab, we find that the information does not depend on just one spot in the eye, but is a sum of information over several spots.

The most recent picture of the operation of the frog's eye is the following. One can find four different kinds of optic nerve fibers, in the sense that there are four different kinds of responses. These experiments were not done by shining on-and-off impulses of light, because that is not what a frog sees. A frog just sits there and his eyes never move, unless the lily pad is flopping back and forth, and in that case his eyes wobble just right so that the image stays put. He does not turn his eyes. If anything moves in his field of vision, like a little bug (he has to be able to see something small moving in the fixed background), it turns out that there are four different kinds of fibers which discharge, whose properties are summarized in Table 36–1. Sustained edge detection, nonerasable, means that if we bring an object with an edge into the field of view of the frog, then there are a lot of impulses in this particular fiber while the object is moving, but they die down to a sustained impulse that continues as long as the edge is there, even if it is standing still. If we turn out the light, the impulses stop. If we turn it on again while the edge is still in view, they start again. They are not erasable. Another kind of fiber is very similar, except that if the edge is straight it does not work. It must be a convex edge with dark behind it! How complicated must be the system of interconnections in the retina of the eye of the frog in order for it to understand that a convex surface has moved in! Furthermore, although this fiber does sustain somewhat, it does not sustain as long as the other, and if we turn out the light and turn it on again it does *not* build up again. It depends on the moving in of the convex

Table 36–1

Types of response in optic nerve fibers of a frog

Type	Speed	Angular field
1. Sustained edge detection (nonerasable)	0.2–0.5 m/sec	1°
2. Convex edge detection (erasable)	0.5 m/sec	2°–3°
3. Changing contrast detection	1–2 m/sec	7°–10°
4. Dimming detection	Up to $\frac{1}{2}$ m/sec	Up to 15°
5. Darkness detection	?	Very large

surface. The eye sees it move in and remembers that it is there, but if we merely turn out the light for a moment, it simply forgets it and no longer sees it.

Another example is change-in-contrast detection. If there is an edge moving in or out there are pulses, but if the thing stands still there are no pulses at all.

Then there is a dimming detector. If the light intensity is going down it creates pulses, but if it stays down or stays up, the impulse stops; it only works while the light is dimming.

Then, finally, there are a few fibers which are dark detectors—a most amazing thing—they fire all the time! If we increase the light, they fire less rapidly, but all the time. If we decrease the light, they fire more rapidly, all the time. In the dark they fire like mad, perpetually saying, "It is dark! It is dark! It is dark!"

Now these responses seem to be rather complicated to classify, and we might wonder whether perhaps the experiments are being misinterpreted. But it is very interesting that these same classes are very clearly separated in the anatomy of the frog! By other measurements, after these responses had been classified (*afterwards*, that is what is important about this), it was discovered that the *speed* of the signals on the different fibers was not the same, so here was another, independent way to check which kind of a fiber we have found!

Another interesting question is from how big an area is one particular fiber making its calculations? The answer is different for the different classes.

Figure 36–14 shows the surface of the so-called tectum of a frog, where the nerves come into the brain from the optic nerve. All the nerve fibers coming in from the optic nerve make connections in various layers of the tectum. This layered structure is analogous to the retina; that is partly why we know that the brain and retina are very similar. Now, by taking an electrode and moving it down in succession through the layers, we can find out which kinds of optic nerves end where, and the beautiful and wonderful result is that the different kinds of fibers end in different layers! The first ones end in number 1 type, the second in number 2, the threes and fives end in the same place, and deepest of all is number four. (What a coincidence, they got the numbers almost in the right order! No, that is why they numbered them that way, the first paper had the numbers in a different order!)

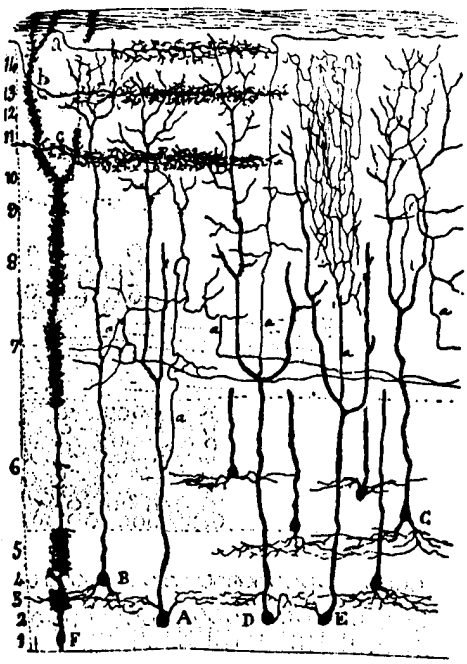


Fig. 36–14. The tectum of a frog.

We may briefly summarize what we have just learned this way: There are three pigments, presumably. There may be many different kinds of receptor cells containing the three pigments in different proportions, but there are many cross connections which may permit additions and subtractions through addition and reinforcement in the nervous system. So before we really understand color vision, we will have to understand the final sensation. This subject is still an open one, but these researches with microelectrodes and so on will perhaps ultimately give us more information on how we see color.

BIBLIOGRAPHY

Committee on Colorimetry, Optical Society of America, *The Science of Color*, Thomas Y. Crowell Company, New York, 1953.

"Mechanisms of Vision," 2nd Supplement to *Journal of General Physiology*, Vol. 43, No. 6, Part 2, July 1960, Rockefeller Institute Press.

SPECIFIC ARTICLES:

DeRobertis, E., "Some Observations on the Ultrastructure and Morphogenesis of Photoreceptors," pp. 1–15.

Hurvich, L. M. and D. Jameson, "Perceived Color, Induction Effects, and Opponent-Response Mechanisms," pp. 63–80.

Rosenblith, W. A., ed., *Sensory Communication*, Massachusetts Institute of Technology Press, Cambridge, Mass., 1961.

"Sight, Sense of," *Encyclopaedia Britannica*, Vol. 20, 1957, pp. 628–635.

37

Quantum Behavior

37–1 Atomic mechanics

37–2 An experiment with bullets

37–3 An experiment with waves

37–4 An experiment with electrons

37–5 The interference of electron waves

37–6 Watching the electrons

37–7 First principles of quantum mechanics

37–8 The uncertainty principle

37–1 Atomic mechanics

In the last few chapters we have treated the essential ideas necessary for an understanding of most of the important phenomena of light—or electromagnetic radiation in general. (We have left a few special topics for next year. Specifically, the theory of the index of dense materials and total internal reflection.) What we have dealt with is called the "classical theory" of electric waves, which turns out to be a completely adequate description of nature for a large number of effects. We have not had to worry yet about the fact that light energy comes in lumps or "photons."

We would like to take up as our next subject the problem of the behavior of relatively large pieces of matter—their mechanical and thermal properties, for instance. In discussing these, we will find that the "classical" (or older) theory fails almost immediately, because matter is really made up of atomic-sized particles. Still, we will deal only with the classical part, because that is the only part that we can understand using the classical mechanics we have been learning. But we shall not be very successful. We shall find that in the case of matter, unlike the case of light, we shall be in difficulty relatively soon. We could, of course, continuously skirt away from the atomic effects, but we shall instead interpose here a short excursion in which we will describe the basic ideas of the quantum properties of matter, i.e., the quantum ideas of atomic physics, so that you will have some feeling for what it is we are leaving out. For we will have to leave out some important subjects that we cannot avoid coming close to.

So we will give now the *introduction* to the subject of quantum mechanics, but will not be able actually to get into the subject until much later.

"Quantum mechanics" is the description of the behavior of matter in all its details and, in particular, of the happenings on an atomic scale. Things on a very small scale behave like nothing that you have any direct experience about. They do not behave like waves, they do not behave like particles, they do not behave like clouds, or billiard balls, or weights on springs, or like anything that you have ever seen.

Newton thought that light was made up of particles, but then it was discovered, as we have seen here, that it behaves like a wave. Later, however (in the beginning of the twentieth century) it was found that light did indeed sometimes behave like a particle. Historically, the electron, for example, was thought to behave like a particle, and then it was found that in many respects it behaved like a wave. So it really behaves like neither. Now we have given up. We say: "It is like *neither*."

There is one lucky break, however—electrons behave just like light. The quantum behavior of atomic objects (electrons, protons, neutrons, photons, and so on) is the same for all, they are all "particle waves," or whatever you want to call them. So what we learn about the properties of electrons (which we shall use for our examples) will apply also to all "particles," including photons of light.

The gradual accumulation of information about atomic and small-scale behavior during the first quarter of this century, which gave some indications about how small things do behave, produced an increasing confusion which was finally resolved in 1926 and 1927 by Schrödinger, Heisenberg, and Born. They finally obtained a consistent description of the behavior of matter on a small scale. We take up the main features of that description in this chapter.

Because atomic behavior is so unlike ordinary experience, it is very difficult to get used to and it appears peculiar and mysterious to everyone, both to the

novice and to the experienced physicist. Even the experts do not understand it the way they would like to, and it is perfectly reasonable that they should not, because all of direct, human experience and of human intuition applies to large objects. We know how large objects will act, but things on a small scale just do not act that way. So we have to learn about them in a sort of abstract or imaginative fashion and not by connection with our direct experience.

In this chapter we shall tackle immediately the basic element of the mysterious behavior in its most strange form. We choose to examine a phenomenon which is impossible, *absolutely* impossible, to explain in any classical way, and which has in it the heart of quantum mechanics. In reality, it contains the *only* mystery. We cannot explain the mystery in the sense of "explaining" how it works. We will *tell* you how it works. In telling you how it works we will have told you about the basic peculiarities of all quantum mechanics.

37–2 An experiment with bullets

To try to understand the quantum behavior of electrons, we shall compare and contrast their behavior, in a particular experimental setup, with the more familiar behavior of particles like bullets, and with the behavior of waves like water waves. We consider first the behavior of bullets in the experimental setup shown diagrammatically in Fig. 37–1. We have a machine gun that shoots a stream of bullets. It is not a very good gun, in that it sprays the bullets (randomly) over a fairly large angular spread, as indicated in the figure. In front of the gun we have a wall (made of armor plate) that has in it two holes just about big enough to let a bullet through. Beyond the wall is a backstop (say a thick wall of wood) which will "absorb" the bullets when they hit it. In front of the wall we have an object which we shall call a "detector" of bullets. It might be a box containing sand. Any bullet that enters the detector will be stopped and accumulated. When we wish, we can empty the box and count the number of bullets that have been caught. The detector can be moved back and forth (in what we will call the x-direction). With this apparatus, we can find out experimentally the answer to the question: "What is the probability that a bullet which passes through the holes in the wall will arrive at the backstop at the distance x from the center?" First, you should realize that we should talk about probability, because we cannot say definitely where any particular bullet will go. A bullet which happens to hit one of the holes may bounce off the edges of the hole, and may end up anywhere at all. By "probability" we mean the chance that the bullet will arrive at the detector, which we can measure by counting the number which arrive at the detector in a certain time and then taking the ratio of this number to the *total* number that hit the backstop during that time. Or, if we assume that the gun always shoots at the same rate during the measurements, the probability we want is just proportional to the number that reach the detector in some standard time interval.

For our present purposes we would like to imagine a somewhat idealized experiment in which the bullets are not real bullets, but are *indestructible* bullets—they cannot break in half. In our experiment we find that bullets always arrive in lumps, and when we find something in the detector, it is always one whole bullet. If the rate at which the machine gun fires is made very low, we find that at any given moment either nothing arrives, or one and only one—exactly one—bullet arrives at the backstop. Also, the size of the lump certainly does not depend on the rate of firing of the gun. We shall say: "Bullets *always* arrive in identical lumps." What we measure with our detector is the probability of arrival of a lump. And we measure the probability as a function of x. The result of such measurements with this apparatus (we have not yet done the experiment, so we are really imagining the result) are plotted in the graph drawn in part (c) of Fig. 37–1. In the graph we plot the probability to the right and x vertically, so that the x-scale fits the diagram of the apparatus. We call the probability P_{12} because the bullets may have come either through hole 1 or through hole 2. You will not be surprised that P_{12} is large near the middle of the graph but gets small if x is very large. You may wonder, however, why P_{12} has its maximum value at $x = 0$. We can understand

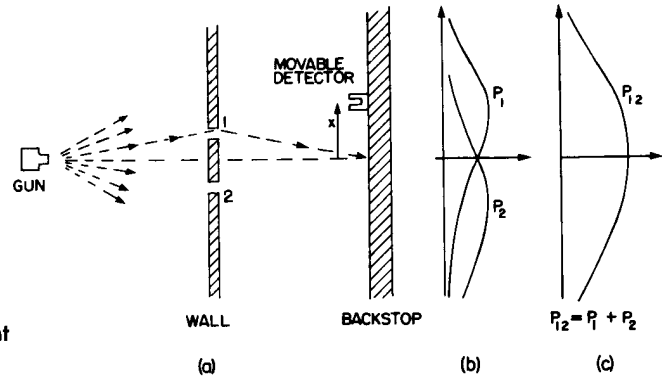


Fig. 37–1. Interference experiment with bullets.

this fact if we do our experiment again after covering up hole 2, and once more while covering up hole 1. When hole 2 is covered, bullets can pass only through hole 1, and we get the curve marked P_1 in part (b) of the figure. As you would expect, the maximum of P_1 occurs at the value of x which is on a straight line with the gun and hole 1. When hole 1 is closed, we get the symmetric curve P_2 drawn in the figure. P_2 is the probability distribution for bullets that pass through hole 2. Comparing parts (b) and (c) of Fig. 37–1, we find the important result that

$$P_{12} = P_1 + P_2. \tag{37.1}$$

The probabilities just add together. The effect with both holes open is the sum of the effects with each hole open alone. We shall call this result an observation of "*no interference*," for a reason that you will see later. So much for bullets. They come in lumps, and their probability of arrival shows no interference.

37-3

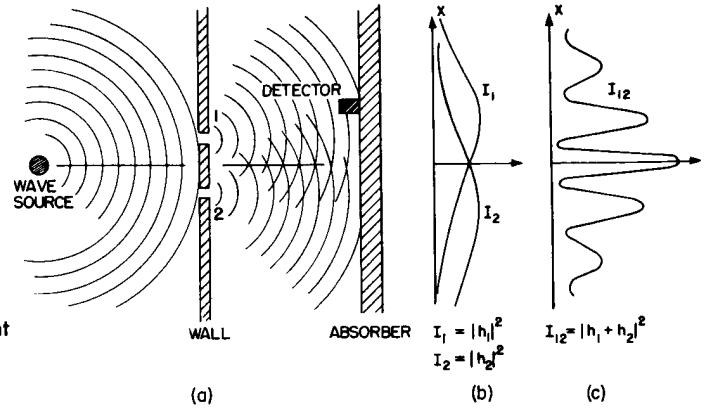


Fig. 37–2. Interference experiment with water waves.

37–3 An experiment with waves

Now we wish to consider an experiment with water waves. The apparatus is shown diagrammatically in Fig. 37–2. We have a shallow trough of water. A small object labeled the "wave source" is jiggled up and down by a motor and makes circular waves. To the right of the source we have again a wall with two holes, and beyond that is a second wall, which, to keep things simple, is an "absorber," so that there is no reflection of the waves that arrive there. This can be done by building a gradual sand "beach." In front of the beach we place a detector which can be moved back and forth in the x-direction, as before. The detector is now a device which measures the "intensity" of the wave motion. You can imagine a gadget which measures the height of the wave motion, but whose scale is calibrated in proportion to the *square* of the actual height, so that the reading is proportional to the intensity of the wave. Our detector reads, then, in proportion to the *energy* being carried by the wave—or rather, the rate at which energy is carried to the detector.

With our wave apparatus, the first thing to notice is that the intensity can have *any* size. If the source just moves a very small amount, then there is just a little bit of wave motion at the detector. When there is more motion at the source,

there is more intensity at the detector. The intensity of the wave can have any value at all. We would *not* say that there was any "lumpiness" in the wave intensity.

Now let us measure the wave intensity for various values of x (keeping the wave source operating always in the same way). We get the interesting-looking curve marked I_{12} in part (c) of the figure.

We have already worked out how such patterns can come about when we studied the interference of electric waves. In this case we would observe that the original wave is diffracted at the holes, and new circular waves spread out from each hole. If we cover one hole at a time and measure the intensity distribution at the absorber we find the rather simple intensity curves shown in part (b) of the figure. I_1 is the intensity of the wave from hole 1 (which we find by measuring when hole 2 is blocked off) and I_2 is the intensity of the wave from hole 2 (seen when hole 1 is blocked).

The intensity I_{12} observed when both holes are open is certainly *not* the sum of I_1 and I_2. We say that there is "interference" of the two waves. At some places (where the curve I_{12} has its maxima) the waves are "in phase" and the wave peaks add together to give a large amplitude and, therefore, a large intensity. We say that the two waves are "interfering constructively" at such places. There will be such constructive interference wherever the distance from the detector to one hole is a whole number of wavelengths larger (or shorter) than the distance from the detector to the other hole.

At those places where the two waves arrive at the detector with a phase difference of π (where they are "out of phase") the resulting wave motion at the detector will be the difference of the two amplitudes. The waves "interfere destructively," and we get a low value for the wave intensity. We expect such low values wherever the distance between hole 1 and the detector is different from the distance between hole 2 and the detector by an odd number of half-wavelengths. The low values of I_{12} in Fig. 37–2 correspond to the places where the two waves interfere destructively.

You will remember that the quantitative relationship between I_1, I_2, and I_{12} can be expressed in the following way: The instantaneous height of the water wave at the detector for the wave from hole 1 can be written as (the real part of) $\hat{h}_1 e^{i\omega t}$, where the "amplitude" $\hat{h}_1$ is, in general, a complex number. The intensity is proportional to the mean squared height or, when we use the complex numbers, to $|\hat{h}_1|^2$. Similarly, for hole 2 the height is $\hat{h}_2 e^{i\omega t}$ and the intensity is proportional to $|\hat{h}_2|^2$. When both holes are open, the wave heights add to give the height $(\hat{h}_1 + \hat{h}_2)e^{i\omega t}$ and the intensity $|\hat{h}_1 + \hat{h}_2|^2$. Omitting the constant of proportionality for our present purposes, the proper relations for *interfering waves* are

$$I_1 = |\hat{h}_1|^2, \qquad I_2 = |\hat{h}_2|^2, \qquad I_{12} = |\hat{h}_1 + \hat{h}_2|^2. \tag{37.2}$$

You will notice that the result is quite different from that obtained with bullets (Eq. 37.1). If we expand $|\hat{h}_1 + \hat{h}_2|^2$ we see that

$$|\hat{h}_1 + \hat{h}_2|^2 = |\hat{h}_1|^2 + |\hat{h}_2|^2 + 2|\hat{h}_1||\hat{h}_2| \cos \delta, \tag{37.3}$$

where δ is the phase difference between $\hat{h}_1$ and $\hat{h}_2$. In terms of the intensities, we could write

$$I_{12} = I_1 + I_2 + 2\sqrt{I_1 I_2} \cos \delta. \tag{37.4}$$

The last term in (37.4) is the "interference term." So much for water waves. The intensity can have any value, and it shows inteference.

37–4 An experiment with electrons

Now we imagine a similar experiment with electrons. It is shown diagrammatically in Fig. 37–3. We make an electron gun which consists of a tungsten wire heated by an electric current and surrounded by a metal box with a hole in it. If the wire is at a negative voltage with respect to the box, electrons emitted by the wire will be accelerated toward the walls and some will pass through the hole. All the electrons which come out of the gun will have (nearly) the same energy. In front of the gun is again a wall (just a thin metal plate) with two holes in it.

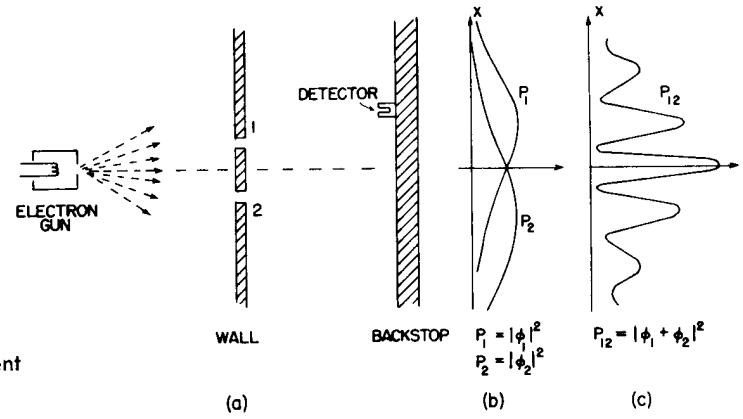


Fig. 37–3. Interference experiment with electrons.

Beyond the wall is another plate which will serve as a "backstop." In front of the backstop we place a movable detector. The detector might be a geiger counter or, perhaps better, an electron multiplier, which is connected to a loudspeaker.

We should say right away that you should not try to set up this experiment (as you could have done with the two we have already described). This experiment has never been done in just this way. The trouble is that the apparatus would have to be made on an impossibly small scale to show the effects we are interested in. We are doing a "thought experiment," which we have chosen because it is easy to think about. We know the results that *would* be obtained because there *are* many experiments that have been done, in which the scale and the proportions have been chosen to show the effects we shall describe.

The first thing we notice with our electron experiment is that we hear sharp "clicks" from the detector (that is, from the loudspeaker). And all "clicks" are the same. There are *no* "half-clicks."

We would also notice that the "clicks" come very erratically. Something like: click . . . . . click-click . . . click . . . . . click . . . . click-click . . . . . click . . . , etc., just as you have, no doubt, heard a geiger counter operating. If we count the clicks which arrive in a sufficiently long time—say for many minutes—and then count again for another equal period, we find that the two numbers are very nearly the same. So we can speak of the *average rate* at which the clicks are heard (so-and-so-many clicks per minute on the average).

As we move the detector around, the *rate* at which the clicks appear is faster or slower, but the size (loudness) of each click is always the same. If we lower the temperature of the wire in the gun the rate of clicking slows down, but still each click sounds the same. We would notice also that if we put two separate detectors at the backstop, one *or* the other would click, but never both at once. (Except that once in a while, if there were two clicks very close together in time, our ear might not sense the separation.) We conclude, therefore, that whatever arrives at the backstop arrives in "lumps." All the "lumps" are the same size: only whole "lumps" arrive, and they arrive one at a time at the backstop. We shall say: "Electrons always arrive in identical lumps."

Just as for our experiment with bullets, we can now proceed to find experimentally the answer to the question: "What is the relative probability that an electron 'lump' will arrive at the backstop at various distances x from the center?" As before, we obtain the relative probability by observing the rate of clicks, holding the operation of the gun constant. The probability that lumps will arrive at a particular x is proportional to the average rate of clicks at that x.

The result of our experiment is the interesting curve marked P_{12} in part (c) of Fig. 37–3. Yes! That is the way electrons go.

37–5 The interference of electron waves

Now let us try to analyze the curve of Fig. 37–3 to see whether we can understand the behavior of the electrons. The first thing we would say is that since they come in lumps, each lump, which we may as well call an electron, has come either through hole 1 or through hole 2. Let us write this in the form of a "Proposition":

Proposition A: Each electron *either* goes through hole 1 *or* it goes through hole 2.

Assuming Proposition A, all electrons that arrive at the backstop can be divided into two classes: (1) those that come through hole 1, and (2) those that come through hole 2. So our observed curve must be the sum of the effects of the electrons which come through hole 1 and the electrons which come through hole 2. Let us check this idea by experiment. First, we will make a measurement for those electrons that come through hole 1. We block off hole 2 and make our counts of the clicks from the detector. From the clicking rate, we get P_1. The result of the measurement is shown by the curve marked P_1 in part (b) of Fig. 37–3. The result seems quite reasonable. In a similar way, we measure P_2, the probability distribution for the electrons that come through hole 2. The result of this measurement is also drawn in the figure.

The result P_{12} obtained with *both* holes open is clearly not the sum of P_1 and P_2, the probabilities for each hole alone. In analogy with our water-wave experiment, we say: "There is interference."

$$\textit{For electrons:} \qquad P_{12} \neq P_1 + P_2. \tag{37.5}$$

How can such an interference come about? Perhaps we should say: "Well, that means, presumably, that it is *not true* that the lumps go either through hole 1 or hole 2, because if they did, the probabilities should add. Perhaps they go in a more complicated way. They split in half and . . ." But no! They cannot, they always arrive in lumps . . . "Well, perhaps some of them go through 1, and then they go around through 2, and then around a few more times, or by some other complicated path . . . then by closing hole 2, we changed the chance that an electron that *started out* through hole 1 would finally get to the backstop . . ." But notice! There are some points at which very few electrons arrive when *both* holes are open, but which receive many electrons if we close one hole, so *closing* one hole *increased* the number from the other. Notice, however, that at the center of the pattern, P_{12} is more than twice as large as $P_1 + P_2$. It is as though closing one hole *decreased* the number of electrons which come through the other hole. It seems hard to explain *both* effects by proposing that the electrons travel in complicated paths.

It is all quite mysterious. And the more you look at it the more mysterious it seems. Many ideas have been concocted to try to explain the curve for P_{12} in terms of individual electrons going around in complicated ways through the holes. None of them has succeeded. None of them can get the right curve for P_{12} in terms of P_1 and P_2.

Yet, surprisingly enough, the *mathematics* for relating P_1 and P_2 to P_{12} is extremely simple. For P_{12} is just like the curve I_{12} of Fig. 37–2, and *that* was simple. What is going on at the backstop can be described by two complex numbers that we can call $\hat{\phi}_1$ and $\hat{\phi}_2$ (they are functions of x, of course). The absolute square of $\hat{\phi}_1$ gives the effect with only hole 1 open. That is, $P_1 = |\hat{\phi}_1|^2$. The effect with only hole 2 open is given by $\hat{\phi}_2$ in the same way. That is, $P_2 = |\hat{\phi}_2|^2$. And the combined effect of the two holes is just $P_{12} = |\hat{\phi}_1 + \hat{\phi}_2|^2$. The *mathematics* is the same as that we had for the water waves! (It is hard to see how one could get such a simple result from a complicated game of electrons going back and forth through the plate on some strange trajectory.)

We conclude the following: The electrons arrive in lumps, like particles, and the probability of arrival of these lumps is distributed like the distribution of intensity of a wave. It is in this sense that an electron behaves "sometimes like a particle and sometimes like a wave."

Incidentally, when we were dealing with classical waves we defined the intensity as the mean over time of the square of the wave amplitude, and we used complex numbers as a mathematical trick to simplify the analysis. But in quantum mechanics it turns out that the amplitudes *must* be represented by complex numbers. The real parts alone will not do. That is a technical point, for the moment, because the formulas look just the same.

Since the probability of arrival through both holes is given so simply, although it is not equal to $(P_1 + P_2)$, that is really all there is to say. But there are a large number of subtleties involved in the fact that nature does work this way. We would like to illustrate some of these subtleties for you now. First, since the number that arrives at a particular point is *not* equal to the number that arrives through 1 plus the number that arrives through 2, as we would have concluded from Proposition A, undoubtedly we should conclude that *Proposition A is false*. It is *not* true that the electrons go *either* through hole 1 or hole 2. But that conclusion can be tested by another experiment.

37–6 Watching the electrons

We shall now try the following experiment. To our electron apparatus we add a very strong light source, placed behind the wall and between the two holes, as shown in Fig. 37–4. We know that electric charges scatter light. So when an electron passes, however it does pass, on its way to the detector, it will scatter some light to our eye, and we can *see* where the electron goes. If, for instance, an electron were to take the path via hole 2 that is sketched in Fig. 37–4, we should see a flash of light coming from the vicinity of the place marked A in the figure. If an electron passes through hole 1 we would expect to see a flash from the vicinity of the upper hole. If it should happen that we get light from both places at the same time, because the electron divides in half . . . Let us just do the experiment!

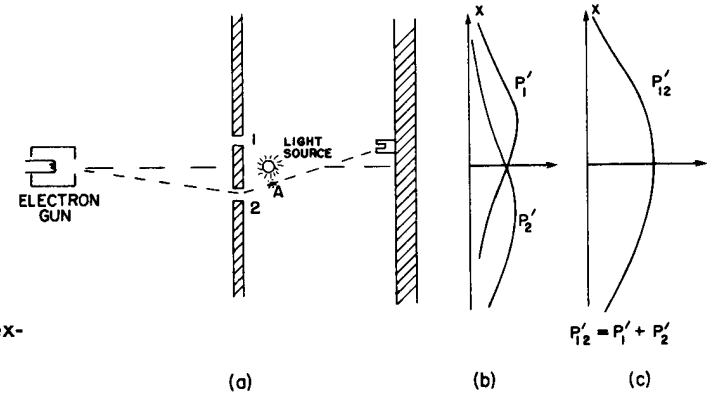


Fig. 37–4. A different electron experiment.

Here is what we see: *every* time that we hear a "click" from our electron detector (at the backstop), we *also see* a flash of light *either* near hole 1 *or* near hole 2, but *never* both at once! And we observe the same result no matter where we put the detector. From this observation we conclude that when we look at the electrons we find that the electrons go either through one hole or the other. Experimentally, Proposition A is necessarily true.

What, then, is wrong with our argument *against* Proposition A? Why *isn't* P_{12} just equal to $P_1 + P_2$? Back to experiment! Let us keep track of the electrons and find out what they are doing. For each position (x-location) of the detector we will count the electrons that arrive and *also* keep track of which hole they went through, by watching for the flashes. We can keep track of things this way: whenever we hear a "click" we will put a count in Column 1 if we see the flash near hole 1, and if we see the flash near hole 2, we will record a count in Column 2. Every electron which arrives is recorded in one of two classes: those which come through 1 and those which come through 2. From the number recorded in Column 1 we get the probability P_1' that an electron will arrive at the detector via hole 1; and from the number recorded in Column 2 we get P_2', the probability that an electron will arrive at the detector via hole 2. If we now repeat such a measurement for many values of x, we get the curves for P_1' and P_2' shown in part (b) of Fig. 37–4.

Well, that is not too surprising! We get for P_1' something quite similar to what we got before for P_1 by blocking off hole 2; and P_2' is similar to what we got by blocking hole 1. So there is *not* any complicated business like going through both holes. When we watch them, the electrons come through just as we would

expect them to come through. Whether the holes are closed or open, those which we see come through hole 1 are distributed in the same way whether hole 2 is open or closed.

But wait! What do we have *now* for the *total* probability, the probability that an electron will arrive at the detector by any route? We already have that information. We just pretend that we never looked at the light flashes, and we lump together the detector clicks which we have separated into the two columns. We *must* just *add* the numbers. For the probability that an electron will arrive at the backstop by passing through *either* hole, we do find $P'_{12} = P_1 + P_2$. That is, although we succeeded in watching which hole our electrons come through, we no longer get the old interference curve P_{12}, but a new one, P'_{12}, showing no interference! If we turn out the light P_{12} is restored.

We must conclude that *when we look at the electrons* the distribution of them on the screen is different than when we do not look. Perhaps it is turning on our light source that disturbs things? It must be that the electrons are very delicate, and the light, when it scatters off the electrons, gives them a jolt that changes their motion. We know that the electric field of the light acting on a charge will exert a force on it. So perhaps we *should* expect the motion to be changed. Anyway, the light exerts a big influence on the electrons. By trying to "watch" the electrons we have changed their motions. That is, the jolt given to the electron when the photon is scattered by it is such as to change the electron's motion enough so that if it *might* have gone to where P_{12} was at a maximum it will instead land where P_{12} was a minimum; that is why we no longer see the wavy interference effects.

You may be thinking: "Don't use such a bright source! Turn the brightness down! The light waves will then be weaker and will not disturb the electrons so much. Surely, by making the light dimmer and dimmer, eventually the wave will be weak enough that it will have a negligible effect." O.K. Let's try it. The first thing we observe is that the flashes of light scattered from the electrons as they pass by does *not* get weaker. *It is always the same-sized flash.* The only thing that happens as the light is made dimmer is that sometimes we hear a "click" from the detector but see *no flash at all.* The electron has gone by without being "seen." What we are observing is that light *also* acts like electrons, we *knew* that it was "wavy," but now we find that it is also "lumpy." It always arrives—or is scattered—in lumps that we call "photons." As we turn down the *intensity* of the light source we do not change the *size* of the photons, only the *rate* at which they are emitted. *That* explains why, when our source is dim, some electrons get by without being seen. There did not happen to be a photon around at the time the electron went through.

This is all a little discouraging. If it is true that whenever we "see" the electron we see the same-sized flash, then those electrons we see are *always* the disturbed ones. Let us try the experiment with a dim light anyway. Now whenever we hear a click in the detector we will keep a count in three columns: in Column (1) those electrons seen by hole 1, in Column (2) those electrons seen by hole 2, and in Column (3) those electrons not seen at all. When we work up our data (computing the probabilities) we find these results: Those "seen by hole 1" have a distribution like P'_1; those "seen by hole 2" have a distribution like P'_2 (so that those "seen by either hole 1 or 2" have a distribution like P'_{12}); and those "not seen at all" have a "wavy" distribution just like P_{12} of Fig. 37–3! *If the electrons are not seen, we have interference!*

That is understandable. When we do not see the electron, no photon disturbs it, and when we do see it, a photon has disturbed it. There is always the same amount of disturbance because the light photons all produce the same-sized effects and the effect of the photons being scattered is enough to smear out any interference effect.

Is there not *some* way we can see the electrons without disturbing them? We learned in an earlier chapter that the momentum carried by a "photon" is inversely proportional to its wavelength ($p = h/\lambda$). Certainly the jolt given to the electron when the photon is scattered toward our eye depends on the momentum that photon carries. Aha! If we want to disturb the electrons only

slightly we should not have lowered the *intensity* of the light, we should have lowered its *frequency* (the same as increasing its wavelength). Let us use light of a redder color. We could even use infrared light, or radiowaves (like radar), and "see" where the electron went with the help of some equipment that can "see" light of these longer wavelengths. If we use "gentler" light perhaps we can avoid disturbing the electrons so much.

Let us try the experiment with longer waves. We shall keep repeating our experiment, each time with light of a longer wavelength. At first, nothing seems to change. The results are the same. Then a terrible thing happens. You remember that when we discussed the microscope we pointed out that, due to the *wave nature* of the light, there is a limitation on how close two spots can be and still be seen as two separate spots. This distance is of the order of the wavelength of light. So now, when we make the wavelength longer than the distance between our holes, we see a *big* fuzzy flash when the light is scattered by the electrons. We can no longer tell which hole the electron went through! We just know it went somewhere! And it is just with light of this color that we find that the jolts given to the electron are small enough so that P'_{12} begins to look like P_{12}—that we begin to get some interference effect. And it is only for wavelengths much longer than the separation of the two holes (when we have no chance at all of telling where the electron went) that the disturbance due to the light gets sufficiently small that we again get the curve P_{12} shown in Fig. 37–3.

In our experiment we find that it is impossible to arrange the light in such a way that one can tell which hole the electron went through, and at the same time not disturb the pattern. It was suggested by Heisenberg that the then new laws of nature could only be consistent if there were some basic limitation on our experimental capabilities not previously recognized. He proposed, as a general principle, his *uncertainty principle*, which we can state in terms of our experiment as follows: "It is impossible to design an apparatus to determine which hole the electron passes through, that will not at the same time disturb the electrons enough to destroy the interference pattern." If an apparatus is capable of determining which hole the electron goes through, it *cannot* be so delicate that it does not disturb the pattern in an essential way. No one has ever found (or even thought of) a way around the uncertainty principle. So we must assume that it describes a basic characteristic of nature.

The complete theory of quantum mechanics which we now use to describe atoms and, in fact, all matter depends on the correctness of the uncertainty principle. Since quantum mechanics is such a successful theory, our belief in the uncertainty principle is reinforced. But if a way to "beat" the uncertainty principle were ever discovered, quantum mechanics would give inconsistent results and would have to be discarded as a valid theory of nature.

"Well," you say, "what about Proposition A? It is true, or is it *not* true, that the electron either goes through hole 1 or it goes through hole 2?" The only answer that can be given is that we have found from experiment that there is a certain special way that we have to think in order that we do not get into inconsistencies. What we must say (to avoid making wrong predictions) is the following. If one looks at the holes or, more accurately, if one has a piece of apparatus which is capable of determining whether the electrons go through hole 1 or hole 2, then one *can* say that it goes either through hole 1 or hole 2. *But*, when one does *not* try to tell which way the electron goes, when there is nothing in the experiment to disturb the electrons, then one may *not* say that an electron goes either through hole 1 or hole 2. If one does say that, and starts to make any deductions from the statement, he will make errors in the analysis. This is the logical tightrope on which we must walk if we wish to describe nature successfully.

If the motion of all matter—as well as electrons—must be described in terms of waves, what about the bullets in our first experiment? Why didn't we see an interference pattern there? It turns out that for the bullets the wavelengths were so tiny that the interference patterns became very fine. So fine, in fact, that with any

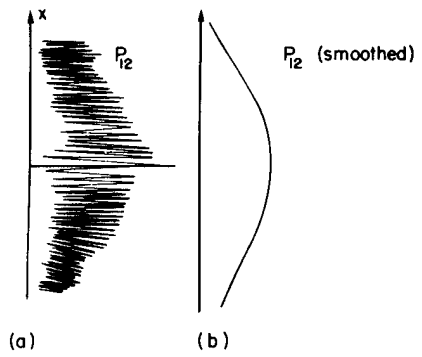


Fig. 37–5. Interference pattern with bullets: (a) actual (schematic), (b) observed.

detector of finite size one could not distinguish the separate maxima and minima. What we saw was only a kind of average, which is the classical curve. In Fig. 37–5 we have tried to indicate schematically what happens with large-scale objects. Part (a) of the figure shows the probability distribution one might predict for bullets, using quantum mechanics. The rapid wiggles are supposed to represent the interference pattern one gets for waves of very short wavelength. Any physical detector, however, straddles several wiggles of the probability curve, so that the measurements show the smooth curve drawn in part (b) of the figure.

37–7 First principles of quantum mechanics

We will now write a summary of the main conclusions of our experiments. We will, however, put the results in a form which makes them true for a general class of such experiments. We can write our summary more simply if we first define an "ideal experiment" as one in which there are no uncertain external influences, i.e., no jiggling or other things going on that we cannot take into account. We would be quite precise if we said: "An ideal experiment is one in which all of the initial and final conditions of the experiment are completely specified." What we will call "an event" is, in general, just a specific set of initial and final conditions. (For example: "an electron leaves the gun, arrives at the detector, and nothing else happens.") Now for our summary.

SUMMARY

(1) The probability of an event in an ideal experiment is given by the square of the absolute value of a complex number ϕ which is called the probability amplitude.

$$\begin{aligned} P &= \text{probability},\\ \phi &= \text{probability amplitude},\\ P &= |\phi|^2. \end{aligned} \tag{37.6}$$

(2) When an event can occur in several alternative ways, the probability amplitude for the event is the sum of the probability amplitudes for each way considered separately. There is interference.

$$\begin{aligned} \phi &= \phi_1 + \phi_2,\\ P &= |\phi_1 + \phi_2|^2. \end{aligned} \tag{37.7}$$

(3) If an experiment is performed which is capable of determining whether one or another alternative is actually taken, the probability of the event is the sum of the probabilities for each alternative. The interference is lost.

$$P = P_1 + P_2. \tag{37.8}$$

One might still like to ask: "How does it work? What is the machinery behind the law?" No one has found any machinery behind the law. No one can "explain" any more than we have just "explained." No one will give you any deeper representation of the situation. We have no ideas about a more basic mechanism from which these results can be deduced.

We would like to emphasize a very important difference between classical and quantum mechanics. We have been talking about the probability that an electron will arrive in a given circumstance. We have implied that in our experimental arrangement (or even in the best possible one) it would be impossible to predict exactly what would happen. We can only predict the odds! This would mean, if it were true, that physics has given up on the problem of trying to predict exactly what will happen in a definite circumstance. Yes! physics *has* given up. *We do not know how to predict what would happen in a given circumstance*, and we believe now that it is impossible, that the only thing that can be predicted is the probability of different events. It must be recognized that this is a retrenchment in our earlier ideal of understanding nature. It may be a backward step, but no one has seen a way to avoid it.

We make now a few remarks on a suggestion that has sometimes been made to try to avoid the description we have given: "Perhaps the electron has some kind of internal works—some inner variables—that we do not yet know about. Perhaps that is why we cannot predict what will happen. If we could look more closely at the electron we could be able to tell where it would end up." So far as we know, that is impossible. We would still be in difficulty. Suppose we were to assume that inside the electron there is some kind of machinery that determines where it is going to end up. That machine must *also* determine which hole it is going to go through on its way. But we must not forget that what is inside the electron should not be dependent on what *we* do, and in particular upon whether we open or close one of the holes. So if an electron, before it starts, has already made up its mind (a) which hole it is going to use, and (b) where it is going to land, we should find P_1 for those electrons that have chosen hole 1, P_2 for those that have chosen hole 2, *and necessarily* the sum $P_1 + P_2$ for those that arrive through the two holes. There seems to be no way around this. But we have verified experimentally that that is not the case. And no one has figured a way out of this puzzle. So at the present time we must limit ourselves to computing probabilities. We say "at the present time," but we suspect very strongly that it is something that will be with us forever—that it is impossible to beat that puzzle—that this is the way nature really *is*.

37–8 The uncertainty principle

This is the way Heisenberg stated the uncertainty principle originally: If you make the measurement on any object, and you can determine the x-component of its momentum with an uncertainty Δp, you cannot, at the same time, know its x-position more accurately than $\Delta x = h/\Delta p$. The uncertainties in the position and momentum at any instant must have their product greater than Planck's constant. This is a special case of the uncertainty principle that was stated above more generally. The more general statement was that one cannot design equipment in any way to determine which of two alternatives is taken, without, at the same time, destroying the pattern of interference.

Let us show for one particular case that the kind of relation given by Heisenberg must be true in order to keep from getting into trouble. We imagine a modification of the experiment of Fig. 37–3, in which the wall with the holes consists of a plate mounted on rollers so that it can move freely up and down (in the x-direction), as shown in Fig. 37–6. By watching the motion of the plate carefully we can try to tell which hole an electron goes through. Imagine what happens when the detector is placed at $x = 0$. We would expect that an electron which passes through hole 1 must be deflected downward by the plate to reach the detector. Since the vertical component of the electron momentum is changed, the plate must recoil with an equal momentum in the opposite direction. The plate will get an upward kick. If the electron goes through the lower hole, the plate should feel a downward kick. It is clear that for every position of the detector, the momentum received by the plate will have a different value for a traversal via hole 1 than for a traversal via hole 2. So! Without disturbing the electrons *at all*, but just by watching the *plate*, we can tell which path the electron used.

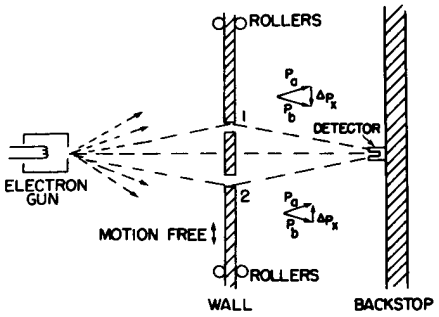


Fig. 37–6. An experiment in which the recoil of the wall is measured.

Now in order to do this it is necessary to know what the momentum of the screen is, before the electron goes through. So when we measure the momentum after the electron goes by, we can figure out how much the plate's momentum has changed. But remember, according to the uncertainty principle we cannot at the same time know the position of the plate with an arbitrary accuracy. But if we do not know exactly *where* the plate is we cannot say precisely where the two holes are. They will be in a different place for every electron that goes through. This means that the center of our interference pattern will have a different location for each electron. The wiggles of the interference pattern will be smeared out. We shall show quantitatively in the next chapter that if we determine the momentum of the plate sufficiently accurately to determine from the recoil measurement which hole was used, then the uncertainty in the x-position of the plate will, according to the un-

certainty principle, be enough to shift the pattern observed at the detector up and down in the x-direction about the distance from a maximum to its nearest minimum. Such a random shift is just enough to smear out the pattern so that no interference is observed.

The uncertainty principle "protects" quantum mechanics. Heisenberg recognized that if it were possible to measure the momentum and the position simultaneously with a greater accuracy, the quantum mechanics would collapse. So he proposed that it must be impossible. Then people sat down and tried to figure out ways of doing it, and nobody could figure out a way to measure the position and the momentum of anything—a screen, an electron, a billiard ball, anything—with any greater accuracy. Quantum mechanics maintains its perilous but accurate existence.

38

The Relation of Wave and Particle Viewpoints

38–1 Probability wave amplitudes

38–2 Measurement of position and momentum

38–3 Crystal diffraction

38–4 The size of an atom

38–5 Energy levels

38–6 Philosophical implications

38–1 Probability wave amplitudes

In this chapter we shall discuss the relationship of the wave and particle viewpoints. We already know, from the last chapter, that neither the wave viewpoint nor the particle viewpoint is correct. Usually we have tried to present things accurately, or at least precisely enough that they will not have to be changed when we learn more—it may be extended, but it will not be changed! But when we try to talk about the wave picture or the particle picture, both are approximate, and both will change. Therefore what we learn in this chapter will not be accurate in a certain sense; it is a kind of half-intuitive argument that will be made more precise later, but certain things will be changed a little bit when we interpret them correctly in quantum mechanics. The reason for doing such a thing, of course, is that we are not going to go directly into quantum mechanics, but we want to have at least some idea of the kinds of effects that we will find. Furthermore, all our experiences are with waves and with particles, and so it is rather handy to use the wave and particle ideas to get some understanding of what happens in given circumstances before we know the complete mathematics of the quantum-mechanical amplitudes. We shall try to illustrate the weakest places as we go along, but most of it is very nearly correct—it is just a matter of interpretation.

First of all, we know that the new way of representing the world in quantum mechanics—the new framework—is to give an amplitude for every event that can occur, and if the event involves the reception of one particle then we can give the amplitude to find that one particle at different places and at different times. The probability of finding the particle is then proportional to the absolute square of the amplitude. In general, the amplitude to find a particle in different places at different times varies with position and time.

In a special case the amplitude varies sinusoidally in space and time like $e^{i(\omega t-\mathbf{k}\cdot\mathbf{r})}$ (do not forget that these amplitudes are complex numbers, not real numbers) and involves a definite frequency ω and wave number $\mathbf{k}$. Then it turns out that this corresponds to a classical limiting situation where we would have believed that we have a particle whose energy E was known and is related to the frequency by

$$E = \hbar\omega, \tag{38.1}$$

and whose momentum $\mathbf{p}$ is also known and is related to the wave number by

$$\mathbf{p} = \hbar\mathbf{k}. \tag{38.2}$$

This means that the idea of a particle is limited. The idea of a particle—its location, its momentum, etc.—which we use so much, is in certain ways unsatisfactory. For instance, if an amplitude to find a particle at different places is given by $e^{i(\omega t-\mathbf{k}\cdot\mathbf{r})}$, whose absolute square is a constant, that would mean that the probability of finding a particle is the same at all points. That means we do not know *where* it is—it can be anywhere—there is a great uncertainty in its location.

On the other hand, if the position of a particle is more or less well known and we can predict it fairly accurately, then the probability of finding it in different places must be confined to a certain region, whose length we call Δx. Outside this region, the probability is zero. Now this probability is the absolute square of an amplitude, and if the absolute square is zero, the amplitude is also zero, so that

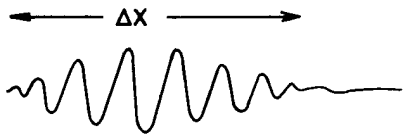


Fig. 38–1. A wave packet of length Δx.

we have a wave train whose length is Δx (Fig. 38–1), and the wavelength (the distance between nodes of the waves in the train) of that wave train is what corresponds to the particle momentum.

Here we encounter a strange thing about waves; a very simple thing which has nothing to do with quantum mechanics strictly. It is something that anybody who works with waves, even if he knows no quantum mechanics, knows: namely, *we cannot define a unique wavelength for a short wave train.* Such a wave train does not *have* a definite wavelength; there is an indefiniteness in the wave number that is related to the finite length of the train, and thus there is an indefiniteness in the momentum.

38–2 Measurement of position and momentum

Let us consider two examples of this idea—to see the reason why there is an uncertainty in the position and/or the momentum, if quantum mechanics is right. We have also seen before that if there were not such a thing—if it were possible to measure the position and the momentum of anything simultaneously—we would have a paradox; it is fortunate that we do not have such a paradox, and the fact that such an uncertainty comes naturally from the wave picture shows that everything is mutually consistent.

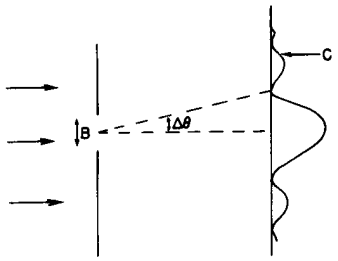


Fig. 38–2. Diffraction of particles passing through a slit.

Here is one example which shows the relationship between the position and the momentum in a circumstance that is easy to understand. Suppose we have a single slit, and particles are coming from very far away with a certain energy—so that they are all coming essentially horizontally (Fig. 38–2). We are going to concentrate on the vertical components of momentum. All of these particles have a certain horizontal momentum p_0, say, in a classical sense. So, in the classical sense, the vertical momentum p_y, before the particle goes through the hole, is definitely known. The particle is moving neither up nor down, because it came from a source that is far away—and so the vertical momentum is of course zero. But now let us suppose that it goes through a hole whose width is B. Then after it has come out through the hole, we know the position vertically—the y position—with considerable accuracy—namely $\pm B$. That is, the uncertainty in position, Δy, is of order B. Now we might also want to say, since we know the momentum is absolutely horizontal, that Δp_y is zero; but that is wrong. We *once* knew the momentum was horizontal, but we do not know it any more. Before the particles passed through the hole, we did not know their vertical positions. Now that we have found the vertical position by having the particle come through the hole, we have lost our information on the vertical momentum! Why? According to the wave theory, there is a spreading out, or diffraction, of the waves after they go through the slit, just as for light. Therefore there is a certain probability that particles coming out of the slit are not coming exactly straight. The pattern is spread out by the diffraction effect, and the angle of spread, which we can define as the angle of the first minimum, is a measure of the uncertainty in the final angle.

How does the pattern become spread? To say it is spread means that there is some chance for the particle to be moving up or down, that is, to have a component of momentum up or down. We say *chance* and *particle* because we can detect this diffraction pattern with a particle counter, and when the counter receives the particle, say at C in Fig. 38–2, it receives the *entire* particle, so that, in a classical sense, the particle has a vertical momentum, in order to get from the slit up to C.

To get a rough idea of the spread of the momentum, the vertical momentum p_y has a spread which is equal to $p_0 \Delta\theta$, where p_0 is the horizontal momentum. And how big is $\Delta\theta$ in the spread-out pattern? We know that the first minimum occurs at an angle $\Delta\theta$ such that the waves from one edge of the slit have to travel one wavelength farther than the waves from the other side—we worked that out before (Chapter 30). Therefore $\Delta\theta$ is λ/B, and so Δp_y in this experiment is $p_0\lambda/B$. Note that if we make B smaller and make a more accurate measurement of the position of the particle, the diffraction pattern gets wider. Remember, when we closed the slits on the experiment with the microwaves, we had more intensity farther out. So the narrower we make the slit, the wider the pattern gets, and the

more is the likelihood that we would find that the particle has sidewise momentum. Thus the uncertainty in the vertical momentum is inversely proportional to the uncertainty of y. In fact, we see that the product of the two is equal to $p_0\lambda$. But λ is the wavelength and p_0 is the momentum, and in accordance with quantum mechanics, the wavelength times the momentum is Planck's constant h. So we obtain the rule that the uncertainties in the vertical momentum and in the vertical position have a product of the order h:

$$\Delta y\, \Delta p_y \approx h. \tag{38.3}$$

We cannot prepare a system in which we know the vertical position of a particle and can predict how it will move vertically with greater certainty than given by (38.3). That is, the uncertainty in the vertical momentum must exceed $h/\Delta y$, where Δy is the uncertainty in our knowledge of the position.

Sometimes people say quantum mechanics is all wrong. When the particle arrived from the left, its vertical momentum was zero. And now that it has gone through the slit, its position is known. Both position and momentum seem to be known with arbitrary accuracy. It is quite true that we can receive a particle, and on reception determine what its position is and what its momentum would have had to have been to have gotten there. That is true, but that is not what the uncertainty relation (38.3) refers to. Equation (38.3) refers to the *predictability* of a situation, not remarks about the *past*. It does no good to say "I knew what the momentum was before it went through the slit, and now I know the position," because now the momentum knowledge is lost. The fact that it went through the slit no longer permits us to predict the vertical momentum. We are talking about a predictive theory, not just measurements after the fact. So we must talk about what we can predict.

Now let us take the thing the other way around. Let us take another example of the same phenomenon, a little more quantitatively. In the previous example we measured the momentum by a classical method. Namely, we considered the direction and the velocity and the angles, etc., so we got the momentum by classical analysis. But since momentum is related to wave number, there exists in nature still another way to measure the momentum of a particle—photon or otherwise—which has no classical analog, because it uses Eq. (38.2). We measure the *wavelengths of the waves*. Let us try to measure momentum in this way.

Suppose we have a grating with a large number of lines (Fig. 38–3), and send a beam of particles at the grating. We have often discussed this problem: if the particles have a definite momentum, then we get a very sharp pattern in a certain direction, because of the interference. And we have also talked about how accurately we can determine that momentum, that is to say, what the resolving power of such a grating is. Rather than derive it again, we refer to Chapter 30, where we found that the relative uncertainty in the wavelength that can be measured with a given grating is $1/Nm$, where N is the number of lines on the grating and m is the order of the diffraction pattern. That is,

$$\Delta\lambda/\lambda = 1/Nm. \tag{38.4}$$

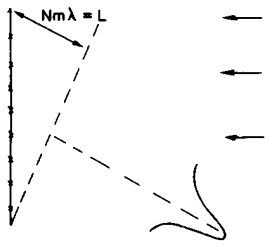


Fig. 38–3. Determination of momentum by using a diffraction grating.

Now formula (38.4) can be rewritten as

$$\Delta\lambda/\lambda^2 = 1/Nm\lambda = 1/L, \tag{38.5}$$

where L is the distance shown in Fig. 38–3. This distance is the difference between the total distance that the particle or wave or whatever it is has to travel if it is reflected from the bottom of the grating, and the distance that it has to travel if it is reflected from the top of the grating. That is, the waves which form the diffraction pattern are waves which come from different parts of the grating. The first ones that arrive come from the bottom end of the grating, from the beginning of the wave train, and the rest of them come from later parts of the wave train, coming from different parts of the grating, until the last one finally arrives, and that involves a point in the wave train a distance L behind the first point. So in order that we shall have a sharp line in our spectrum corresponding to a definite momentum,

with an uncertainty given by (38.4), we have to have a wave train of at least length L. If the wave train is too short we are not using the entire grating. The waves which form the spectrum are being reflected from only a very short sector of the grating if the wave train is too short, and the grating will not work right—we will find a big angular spread. In order to get a narrower one, we need to use the whole grating, so that at least at some moment the whole wave train is scattering simultaneously from all parts of the grating. Thus the wave train must be of length L in order to have an uncertainty in the wavelength less than that given by (38.5). Incidentally,

$$\Delta\lambda/\lambda^2 = \Delta(1/\lambda) = \Delta k/2\pi. \tag{38.6}$$

Therefore

$$\Delta k = 2\pi/L, \tag{38.7}$$

where L is the length of the wave train.

This means that if we have a wave train whose length is less than L, the uncertainty in the wave number must exceed $2\pi/L$. Or the uncertainty in a wave number times the length of the wave train—we will call that for a moment Δx—exceeds 2π. We call it Δx because that is the uncertainty in the location of the particle. If the wave train exists only in a finite length, then that is where we could find the particle, within an uncertainty Δx. Now this property of waves, that the length of the wave train times the uncertainty of the wave number associated with it is at least 2π, is a property that is known to everyone who studies them. It has nothing to do with quantum mechanics. It is simply that if we have a finite train, we cannot count the waves in it very precisely. Let us try another way to see the reason for that.

Suppose that we have a finite train of length L; then because of the way it has to decrease at the ends, as in Fig. 38–1, the number of waves in the length L is uncertain by something like ± 1. But the number of waves in L is $kL/2\pi$. Thus k is uncertain, and we again get the result (38.7), a property merely of waves. The same thing works whether the waves are in space and k is the number of radians per centimeter and L is the length of the train, or the waves are in time and ω is the number of oscillations per second and T is the "length" in time that the wave train comes in. That is, if we have a wave train lasting only for a certain finite time T, then the uncertainty in the frequency is given by

$$\Delta\omega = 2\pi/T. \tag{38.8}$$

We have tried to emphasize that these are properties of waves alone, and they are well known, for example, in the theory of sound.

The point is that in quantum mechanics we interpret the wave number as being a measure of the momentum of a particle, with the rule that $p = \hbar k$, so that relation (38.7) tells us that $\Delta p \approx h/\Delta x$. This, then, is a limitation of the classical idea of momentum. (Naturally, it has to be limited in some ways if we are going to represent particles by waves!) It is nice that we have found a rule that gives us some idea of when there is a failure of classical ideas.

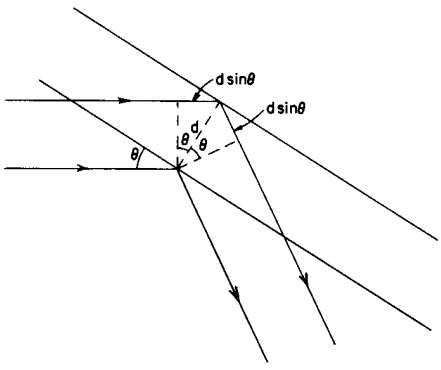


Fig. 38–4. Scattering of waves by crystal planes.

38–3 Crystal diffraction

Next let us consider the reflection of particle waves from a crystal. A crystal is a thick thing which has a whole lot of similar atoms—we will include some complications later—in a nice array. The question is how to set the array so that we get a strong reflected maximum in a given direction for a given beam of, say, light (x-rays), electrons, neutrons, or anything else. In order to obtain a strong reflection, the scattering from all of the atoms must be in phase. There cannot be equal numbers in phase and out of phase, or the waves will cancel out. The way to arrange things is to find the regions of constant phase, as we have already explained; they are planes which make equal angles with the initial and final directions (Fig. 38–4).

If we consider two parallel planes, as in Fig. 38–4, the waves scattered from the two planes will be in phase provided the difference in distance travelled by a wave-

front is an integral number of wavelengths. This difference can be seen to be $2d \sin \theta$, where d is the perpendicular distance between the planes. Thus the condition for coherent reflection is

$$2d \sin \theta = n\lambda \qquad (n = 1, 2, \ldots). \tag{38.9}$$

If, for example, the crystal is such that the atoms happen to lie on planes obeying condition (38.9) with $n = 1$, then there will be a strong reflection. If, on the other hand, there are other atoms of the same nature (equal in density) halfway between, then the intermediate planes will also scatter equally strongly and will interfere with the others and produce no effect. So d in (38.9) must refer to *adjacent* planes; we cannot take a plane five layers farther back and use this formula!

As a matter of interest, actual crystals are not usually as simple as a single kind of atom repeated in a certain way. Instead, if we make a two-dimensional analog, they are much like wallpaper, in which there is some kind of figure which repeats all over the wallpaper. By "figure" we mean, in the case of atoms, some arrangement—calcium and a carbon and three oxygens, etc., for calcium carbonate, and so on—which may involve a relatively large number of atoms. But whatever it is, the figure is repeated in a pattern. This basic figure is called a *unit cell.*

The basic pattern of repetition defines what we call the *lattice type*; the lattice type can be immediately determined by looking at the reflections and seeing what their symmetry is. In other words, where we find any reflections *at all* determines the lattice type, but in order to determine what is in each of the elements of the lattice one must take into account the *intensity* of the scattering at the various directions. *Which* directions scatter depends on the type of lattice, but *how strongly* each scatters is determined by what is inside each unit cell, and in that way the structure of crystals is worked out.

Two photographs of x-ray diffraction patterns are shown in Figs. 38–5 and 38–6; they illustrate scattering from rock salt and myoglobin, respectively.

Incidentally, an interesting thing happens if the spacings of the nearest planes are less than $\lambda/2$. In this case (38.9) has no solution for n. Thus if λ is bigger than twice the distance between adjacent planes then there is no side diffraction pattern, and the light—or whatever it is—will go right through the material without bouncing off or getting lost. So in the case of light, where λ is much bigger than the spacing, of course it does go through and there is no pattern of reflection from the planes of the crystal.

This fact also has an interesting consequence in the case of piles which make neutrons (these are obviously particles, for anybody's money!). If we take these neutrons and let them into a long block of graphite, the neutrons diffuse and work their way along (Fig. 38–7). They diffuse because they are bounced by the atoms, but strictly, in the wave theory, they are bounced by the atoms because of diffraction from the crystal planes. It turns out that if we take a very long piece of graphite, the neutrons that come out the far end are all of long wavelength! In fact, if one plots the intensity as a function of wavelength, we get nothing except for wavelengths longer than a certain minimum (Fig. 38–8). In other words, we can get very slow neutrons that way. Only the slowest neutrons come through; they are not diffracted or scattered by the crystal planes of the graphite, but keep going right through like light through glass, and are not scattered out the sides. There are many other demonstrations of the reality of neutron waves and waves of other particles.

38–4 The size of an atom

We now consider another application of the uncertainty relation, Eq. (38.3). It must not be taken too seriously; the idea is right but the analysis is not very accurate. The idea has to do with the determination of the size of atoms, and the fact that, classically, the electrons would radiate light and spiral in until they settle down right on top of the nucleus. But that cannot be right quantum-mechanically because then we would know where each electron was and how fast it was moving.

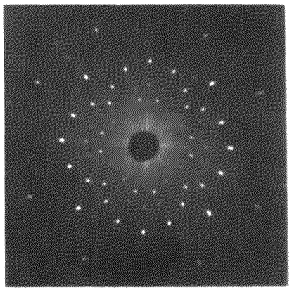


Figure 38–5

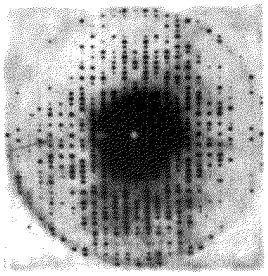


Figure 38–6

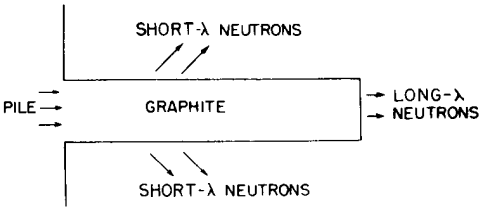


Fig. 38–7. Diffusion of pile neutrons through graphite block.

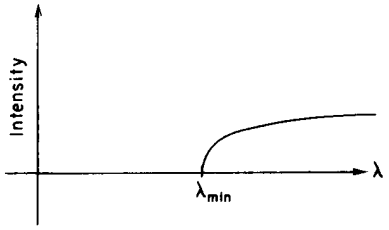


Fig. 38–8. Intensity of neutrons out of graphite rod as function of wavelength.

Suppose we have a hydrogen atom, and measure the position of the electron; we must not be able to predict exactly where the electron will be, or the momentum spread will then turn out to be infinite. Every time we look at the electron, it is somewhere, but it has an amplitude to be in different places so there is a probability of it being found in different places. These places cannot all be at the nucleus; we shall suppose there is a spread in position of order a. That is, the distance of the electron from the nucleus is usually about a. We shall determine a by minimizing the total energy of the atom.

The spread in momentum is roughly h/a because of the uncertainty relation, so that if we try to measure the momentum of the electron in some manner, such as by scattering x-rays off it and looking for the Doppler effect from a moving scatterer, we would expect not to get zero every time—the electron is not standing still—but the momenta must be of the order $p \approx h/a$. Then the kinetic energy is roughly $\frac{1}{2}mv^2 = p^2/2m = h^2/2ma^2$. (In a sense, this is a kind of dimensional analysis to find out in what way the kinetic energy depends upon Planck's constant, upon m, and upon the size of the atom. We need not trust our answer to within factors like 2, π, etc. We have not even defined a very precisely.) Now the potential energy is minus e^2 over the distance from the center, say $-e^2/a$, where, we remember, e^2 is the charge of an electron squared, divided by $4\pi\epsilon_0$. Now the point is that the potential energy is reduced if a gets smaller, but the smaller a is, the higher the momentum required, because of the uncertainty principle, and therefore the higher the kinetic energy. The total energy is

$$E = h^2/2ma^2 - e^2/a. \tag{38.10}$$

We do not know what a is, but we know that the atom is going to arrange itself to make some kind of compromise so that the energy is as little as possible. In order to minimize E, we differentiate with respect to a, set the derivative equal to zero, and solve for a. The derivative of E is

$$dE/da = -h^2/ma^3 + e^2/a^2, \tag{38.11}$$

and setting $dE/da = 0$ gives for a the value

$$\begin{aligned} a_0 = h^2/me^2 &= 0.528 \text{ angstrom} \\ &= 0.528 \times 10^{-10} \text{ meter}. \end{aligned} \tag{38.12}$$

This particular distance is called the *Bohr radius*, and we have thus learned that atomic dimensions are of the order of angstroms, which is right: This is pretty good—in fact, it is amazing, since until now we have had no basis for understanding the size of atoms! Atoms are completely impossible from the classical point of view, since the electrons would spiral into the nucleus.

Now if we put the value (38.12) for a_0 into (38.10) to find the energy, it comes out

$$E_0 = -e^2/2a_0 = -me^4/2h^2 = -13.6 \text{ ev}. \tag{38.13}$$

What does a negative energy mean? It means that the electron has less energy when it is in the atom than when it is free. It means it is bound. It means it takes energy to kick the electron out; it takes energy of the order of 13.6 ev to ionize a hydrogen atom. We have no reason to think that it is not two or three times this—or half of this—or $(1/\pi)$ times this, because we have used such a sloppy argument. However, we have cheated, we have used all the constants in such a way that it happens to come out the right number! This number, 13.6 electron volts, is called a Rydberg of energy; it is the ionization energy of hydrogen.

So we now understand why we do not fall through the floor. As we walk, our shoes with their masses of atoms push against the floor with *its* mass of atoms. In order to squash the atoms closer together, the electrons would be confined to a smaller space and, by the uncertainty principle, their momenta would have to be higher on the average, and that means high energy; the resistance to atomic compression is a quantum-mechanical effect and not a classical effect. Classically,

we would expect that if we were to draw all the electrons and protons closer together, the energy would be reduced still further, and the best arrangement of positive and negative charges in classical physics is all on top of each other. This was well known in classical physics and was a puzzle because of the existence of the atom. Of course, the early scientists invented some ways out of the trouble—but never mind, we have the *right* way out, now! (Maybe.)

Incidentally, although we have no reason to understand it at the moment, in a situation where there are many electrons it turns out that they try to keep away from each other. If one electron is occupying a certain space, then another does not occupy the same space. More precisely, there are two spin cases, so that two can sit on top of each other, one spinning one way and one the other way. But after that we cannot put any more there. We have to put others in another place, and that is the real reason that matter has strength. If we could put all the electrons in the same place it would condense even more than it does. It is the fact that the electrons cannot all get on top of each other that makes tables and everything else solid.

Obviously, in order to understand the properties of matter, we will have to use quantum mechanics and not be satisfied with classical mechanics.

38–5 Energy levels

We have talked about the atom in its lowest possible energy condition, but it turns out that the electron can do other things. It can jiggle and wiggle in a more energetic manner, and so there are many different possible motions for the atom. According to quantum mechanics, in a stationary condition there can only be definite energies for an atom. We make a diagram (Fig. 38–9) in which we plot the energy vertically, and we make a horizontal line for each allowed value of the energy. When the electron is free, i.e., when its energy is positive, it can have any energy; it can be moving at any speed. But bound energies are not arbitrary. The atom must have one or another out of a set of allowed values, such as those in Fig. 38–9.

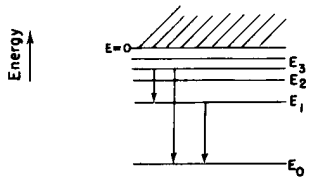


Fig. 38–9. Energy diagram for an atom, showing several possible transitions.

Now let us call the allowed values of the energy E_0, E_1, E_2, E_3. If an atom is initially in one of these "excited states," E_1, E_2, *etc.*, it does not remain in that state forever. Sooner or later it drops to a lower state and radiates energy in the form of light. The frequency of the light that is emitted is determined by conservation of energy plus the quantum-mechanical understanding that the frequency of the light is related to the energy of the light by (38.1). Therefore the frequency of the light which is liberated in a transition from energy E_3 to energy E_1 (for example) is

$$\omega_{31} = (E_3 - E_1)/\hbar. \tag{38.14}$$

This, then, is a characteristic frequency of the atom and defines a spectral emission line. Another possible transition would be from E_3 to E_0. That would have a different frequency

$$\omega_{30} = (E_3 - E_0)/\hbar. \tag{38.15}$$

Another possibility is that if the atom were excited to the state E_1 it could drop to the ground state E_0, emitting a photon of frequency

$$\omega_{10} = (E_1 - E_0)/\hbar. \tag{38.16}$$

The reason we bring up three transitions is to point out an interesting relationship. It is easy to see from (38.14), (38.15), and (38.16) that

$$\omega_{30} = \omega_{31} + \omega_{10}. \tag{38.17}$$

In general, if we find two spectral lines, we shall expect to find another line at the sum of the frequencies (or the difference in the frequencies), and that all the lines can be understood by finding a series of levels such that every line corresponds to the difference in energy of some pair of levels. This remarkable coincidence in

spectral frequencies was noted before quantum mechanics was discovered, and it is called the *Ritz combination principle*. This is again a mystery from the point of view of classical mechanics. Let us not belabor the point that classical mechanics is a failure in the atomic domain; we seem to have demonstrated that pretty well.

We have already talked about quantum mechanics as being represented by amplitudes which behave like waves, with certain frequencies and wave numbers. Let us observe how it comes about from the point of view of amplitudes that the atom has definite energy states. This is something we cannot understand from what has been said so far, but we are all familiar with the fact that confined waves have definite frequencies. For instance, if sound is confined to an organ pipe, or anything like that, then there is more than one way that the sound can vibrate, but for each such way there is a definite frequency. Thus an object in which the waves are confined has certain resonance frequencies. It is therefore a property of waves in a confined space—a subject which we will discuss in detail with formulas later on—that they exist only at definite frequencies. And since the general relation exists between frequencies of the amplitude and energy, we are not surprised to find definite energies associated with electrons bound in atoms.

38–6 Philosophical implications

Let us consider briefly some philosophical implications of quantum mechanics. As always, there are two aspects of the problem: one is the philosophical implication for physics, and the other is the extrapolation of philosophical matters to other fields. When philosophical ideas associated with science are dragged into another field, they are usually completely distorted. Therefore we shall confine our remarks as much as possible to physics itself.

First of all, the most interesting aspect is the idea of the uncertainty principle; making an observation affects the phenomenon. It has always been known that making observations affects a phenomenon, but the point is that the effect cannot be disregarded or minimized or decreased arbitrarily by rearranging the apparatus. When we look for a certain phenomenon we cannot help but disturb it in a certain minimum way, and *the disturbance is necessary for the consistency of the viewpoint.* The observer was sometimes important in prequantum physics, but only in a rather trivial sense. The problem has been raised: if a tree falls in a forest and there is nobody there to hear it, does it make a noise? A *real* tree falling in a *real* forest makes a sound, of course, even if nobody is there. Even if no one is present to hear it, there are other traces left. The sound will shake some leaves, and if we were careful enough we might find somewhere that some thorn had rubbed against a leaf and made a tiny scratch that could not be explained unless we assumed the leaf were vibrating. So in a certain sense we would have to admit that there is a sound made. We might ask: was there a *sensation* of sound? No, sensations have to do, presumably, with consciousness. And whether ants are conscious and whether there were ants in the forest, or whether the tree was conscious, we do not know. Let us leave the problem in that form.

Another thing that people have emphasized since quantum mechanics was developed is the idea that we should not speak about those things which we cannot measure. (Actually relativity theory also said this.) Unless a thing can be defined by measurement, it has no place in a theory. And since an accurate value of the momentum of a localized particle cannot be defined by measurement it therefore has no place in the theory. The idea that this is what was the matter with classical theory *is a false position.* It is a careless analysis of the situation. Just because we cannot *measure* position and momentum precisely does not *a priori* mean that we *cannot* talk about them. It only means that we *need* not talk about them. The situation in the sciences is this: A concept or an idea which cannot be measured or cannot be referred directly to experiment may or may not be useful. It need not exist in a theory. In other words, suppose we compare the classical theory of the world with the quantum theory of the world, and suppose that it is true experimentally that we can measure position and momentum only imprecisely. The question is whether the *ideas* of the exact position of a particle and the exact

momentum of a particle are valid or not. The classical theory admits the ideas; the quantum theory does not. This does not in itself mean that classical physics is wrong. When the new quantum mechanics was discovered, the classical people—which included everybody except Heisenberg, Schrödinger, and Born—said: "Look, your theory is not any good because you cannot answer certain questions like: what is the exact position of a particle?, which hole does it go through?, and some others." Heisenberg's answer was: "I do not need to answer such questions because you cannot ask such a question experimentally." It is that we do not *have* to. Consider two theories (a) and (b); (a) contains an idea that cannot be checked directly but which is used in the analysis, and the other, (b), does not contain the idea. If they disagree in their predictions, one could not claim that (b) is false because it cannot explain this idea that is in (a), because that idea is one of the things that cannot be checked directly. It is always good to know which ideas cannot be checked directly, but it is not necessary to remove them all. It is not true that we can pursue science completely by using only those concepts which are directly subject to experiment.

In quantum mechanics itself there is a wave function amplitude, there is a potential, and there are many constructs that we cannot measure directly. The basis of a science is its ability to *predict*. To predict means to tell what will happen in an experiment that has never been done. How can we do that? By assuming that we know what is there, independent of the experiment. We must extrapolate the experiments to a region where they have not been done. We must take our concepts and extend them to places where they have not yet been checked. If we do not do that, we have no prediction. So it was perfectly sensible for the classical physicists to go happily along and suppose that the position—which obviously means something for a baseball—meant something also for an electron. It was not stupidity. It was a sensible procedure. Today we say that the law of relativity is supposed to be true at all energies, but someday somebody may come along and say how stupid we were. We do not know where we are "stupid" until we "stick our neck out," and so the whole idea is to put our neck out. And the only way to find out that we are wrong is to find out *what* our predictions are. It is absolutely necessary to make constructs.

We have already made a few remarks about the indeterminacy of quantum mechanics. That is, that we are unable now to predict what will happen in physics in a given physical circumstance which is arranged as carefully as possible. If we have an atom that is in an excited state and so is going to emit a photon, we cannot say *when* it will emit the photon. It has a certain amplitude to emit the photon at any time, and we can predict only a probability for emission; we cannot predict the future exactly. This has given rise to all kinds of nonsense and questions on the meaning of freedom of will, and of the idea that the world is uncertain.

Of course we must emphasize that classical physics is also indeterminate, in a sense. It is usually thought that this indeterminacy, that we cannot predict the future, is an important quantum-mechanical thing, and this is said to explain the behavior of the mind, feelings of free will, etc. But if the world *were* classical—if the laws of mechanics were classical—it is not quite obvious that the mind would not feel more or less the same. It is true classically that if we knew the position and the velocity of every particle in the world, or in a box of gas, we could predict exactly what would happen. And therefore the classical world is deterministic. Suppose, however, that we have a finite accuracy and do not know *exactly* where just one atom is, say to one part in a billion. Then as it goes along it hits another atom, and because we did not know the position better than to one part in a billion, we find an even larger error in the position after the collision. And that is amplified, of course, in the next collision, so that if we start with only a tiny error it rapidly magnifies to a very great uncertainty. To give an example: if water falls over a dam, it splashes. If we stand nearby, every now and then a drop will land on our nose. This appears to be completely random, yet such a behavior would be predicted by purely classical laws. The exact position of all the drops depends upon the precise wigglings of the water before it goes over the dam. How? The tiniest irregularities are magnified in falling, so that we get complete randomness. Ob-

viously, we cannot really predict the position of the drops unless we know the motion of the water *absolutely exactly.*

Speaking more precisely, given an arbitrary accuracy, no matter how precise, one can find a time long enough that we cannot make predictions valid for that long a time. Now the point is that this length of time is not very large. It is not that the time is millions of years if the accuracy is one part in a billion. The time goes, in fact, only logarithmically with the error, and it turns out that in only a very, very tiny time we lose all our information. If the accuracy is taken to be one part in billions and billions and billions—no matter how many billions we wish, provided we do stop somewhere—then we can find a time less than the time it took to state the accuracy—after which we can no longer predict what is going to happen! It is therefore not fair to say that from the apparent freedom and indeterminacy of the human mind, we should have realized that classical "deterministic" physics could not ever hope to understand it, and to welcome quantum mechanics as a release from a "completely mechanistic" universe. For already in classical mechanics there was indeterminability from a practical point of view.

39

The Kinetic Theory of Gases

39–1 Properties of matter
39–2 The pressure of a gas
39–3 Compressibility of radiation
39–4 Temperature and kinetic energy
39–5 The ideal gas law

39–1 Properties of matter

With this chapter we begin a new subject which will occupy us for some time. It is the first part of the analysis of the properties of matter from the physical point of view, in which, recognizing that matter is made out of a great many atoms, or elementary parts, which interact electrically and obey the laws of mechanics, we try to understand why various aggregates of atoms behave the way they do.

It is obvious that this is a difficult subject, and we emphasize at the beginning that it is in fact an *extremely* difficult subject, and that we have to deal with it differently than we have dealt with the other subjects so far. In the case of mechanics and in the case of light, we were able to begin with a precise statement of some laws, like Newton's laws, or the formula for the field produced by an accelerating charge, from which a whole host of phenomena could be essentially understood, and which would produce a basis for our understanding of mechanics and of light from that time on. That is, we may learn more later, but we do not learn different physics, we only learn better methods of mathematical analysis to deal with the situation.

We cannot use this approach effectively in studying the properties of matter. We can discuss matter only in a most elementary way; it is much too complicated a subject to analyze directly from its specific basic laws, which are none other than the laws of mechanics and electricity. But these are a bit too far away from the properties we wish to study; it takes too many steps to get from Newton's laws to the properties of matter, and these steps are, in themselves, fairly complicated. We will now start to take some of these steps, but while many of our analyses will be quite accurate, they will eventually get less and less accurate. We will have only a rough understanding of the properties of matter.

One of the reasons that we have to perform the analysis so imperfectly is that the mathematics of it requires a deep understanding of the theory of probability; we are not going to want to know where every atom is actually moving, but rather, how many move here and there on the average, and what the odds are for different effects. So this subject involves a knowledge of the theory of probability, and our mathematics is not yet quite ready and we do not want to strain it too hard.

Secondly, and more important from a physical standpoint, the actual behavior of the atoms is not according to classical mechanics, but according to quantum mechanics, and a correct understanding of the subject cannot be attained until we understand quantum mechanics. Here, unlike the case of billiard balls and automobiles, the difference between the classical mechanical laws and the quantum-mechanical laws is very important and very significant, so that many things that we will deduce by classical physics will be fundamentally incorrect. Therefore there will be certain things to be partially unlearned; however, we shall indicate in every case when a result is incorrect, so that we will know just where the "edges" are. One of the reasons for discussing quantum mechanics in the preceding chapters was to give an idea as to why, more or less, classical mechanics is incorrect in the various directions.

Why do we deal with the subject now at all? Why not wait a half a year, or a year, until we know the mathematics of probability better, and we learn a little quantum mechanics, and then we can do it in a more fundamental way? The answer is that it is a difficult subject, and the best way to learn is to do it slowly! The first thing to do is to get some idea, more or less, of what ought to happen in

different circumstances, and then, later, when we know the laws better, we will formulate them better.

Anyone who wants to analyze the properties of matter in a real problem might want to start by writing down the fundamental equations and then try to solve them mathematically. Although there are people who try to use such an approach, these people are the failures in this field; the real successes come to those who start from a *physical* point of view, people who have a rough idea where they are going and then begin by making the right kind of approximations, knowing what is big and what is small in a given complicated situation. These problems are so complicated that even an elementary understanding, although inaccurate and incomplete, is worth while having, and so the subject will be one that we shall go over again and again, each time with more and more accuracy, as we go through our course in physics.

Another reason for beginning the subject right now is that we have already used many of these ideas in, for example, chemistry, and we have even heard of some of them in high school. It is interesting to know the physical basis for these things.

As an interesting example, we all know that equal volumes of gases, at the same pressure and temperature, contain the same number of molecules. The law of multiple proportions, that when two gases combine in a chemical reaction the volumes needed always stand in simple integral proportions, was understood ultimately by Avogadro to mean that equal volumes have equal numbers of atoms. Now *why* do they have equal numbers of atoms? Can we deduce from Newton's laws that the number of atoms should be equal? We shall address ourselves to that specific matter in this chapter. In succeeding chapters, we shall discuss various other phenomena involving pressures, volumes, temperature, and heat.

We shall also find that the subject can be attacked from a nonatomic point of view, and that there are many interrelationships of the properties of substances. For instance, when we compress something, it heats; if we heat it, it expands. There is a relationship between these two facts which can be deduced independently of the machinery underneath. This subject is called *thermodynamics*. The deepest understanding of thermodynamics comes, of course, from understanding the actual machinery underneath, and that is what we shall do: we shall take the atomic viewpoint from the beginning and use it to understand the various properties of matter and the laws of thermodynamics.

Let us, then, discuss the properties of gases from the standpoint of Newton's laws of mechanics.

39–2 The pressure of a gas

First, we know that a gas exerts a pressure, and we must clearly understand what this is due to. If our ears were a few times more sensitive, we would hear a perpetual rushing noise. Evolution has not developed the ear to that point, because it would be useless if it were so much more sensitive—we would hear a perpetual racket. The reason is that the eardrum is in contact with the air, and air is a lot of molecules in perpetual motion and these bang against the eardrums. In banging against the eardrums they make an irregular tattoo—boom, boom, boom—which we do not hear because the atoms are so small, and the sensitivity of the ear is not quite enough to notice it. The result of this perpetual bombardment is to push the drum away, but of course there is an equal perpetual bombardment of atoms on the other side of the eardrum, so the net force on it is zero. If we were to take the air away from one side, or change the relative amounts of air on the two sides, the eardrum would then be pushed one way or the other, because the amount of bombardment on one side would be greater than on the other. We sometimes feel this uncomfortable effect when we go up too fast in an elevator or an airplane, especially if we also have a bad cold (when we have a cold, inflammation closes the tube which connects the air on the inside of the eardrum with the outside air through the throat, so that the two pressures cannot readily equalize).

In considering how to analyze the situation quantitatively, we imagine that we have a volume of gas in a box, at one end of which is a piston which can be moved (Fig. 39–1). We would like to find out what force on the piston results from the fact that there are atoms in this box. The volume of the box is V, and as the atoms move around inside the box with various velocities they bang against the piston. Suppose there is nothing, a vacuum, on the outside of the piston. What of it? If the piston were left alone, and nobody held onto it, each time it got banged it would pick up a little momentum and it would gradually get pushed out of the box. So in order to keep it from being pushed out of the box, we have to hold it with a force F. The problem is, how much force? One way of expressing the force is to talk about the force per unit area: if A is the area of the piston, then the force on the piston will be written as a number times the area. We define the pressure, then, as equal to the force that we have to apply on a piston, divided by the area of the piston:

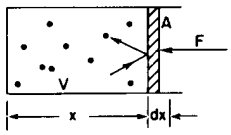


Fig. 39–1. Atoms of a gas in a box with a frictionless piston.

$$P = F/A. \tag{39.1}$$

To make sure we understand the idea (we have to derive it for another purpose anyway), the differential *work* dW done on the gas in compressing it by moving the piston in a differential amount $-dx$ would be the force times the distance that we compress it, which, according to (39.1), would be the pressure times the area, times the distance, which is equal to minus the pressure times the change in the volume:

$$dW = F(-dx) = -PA\,dx = -P\,dV. \tag{39.2}$$

(The area A times the distance dx is the volume change.) The minus sign is there because, as we compress it, we *decrease* the volume; if we think about it we can see that if a gas is compressed, work is done *on* it.

How much force do we have to apply to balance the banging of the molecules? The piston receives from each collision a certain amount of momentum. A certain amount of momentum per second will pour into the piston, and it will start to move. To keep it from moving, we must pour back into it the same amount of momentum per second from our force. Of course, the force *is* the amount of momentum per second that we must pour in. There is another way to put it: if we let go of the piston it will pick up speed because of the bombardments; with each collision we get a little more speed, and the speed thus accelerates. The rate at which the piston picks up speed, or accelerates, is proportional to the force on it. So we see that the force, which we already have said is the pressure times the area, is equal to the momentum per second delivered to the piston by the colliding molecules.

To calculate the momentum per second is easy—we can do it in two parts: first, we find the momentum delivered to the piston by one particular atom in a collision with the piston, then we have to multiply by the number of collisions per second that the atoms have with the wall. The force will be the product of these two factors. Now let us see what the two factors are: In the first place, we shall suppose that the piston is a perfect "reflector" for the atoms. If it is not, the whole theory is wrong, and the piston will start to heat up and things will change, but eventually, when equilibrium has set in, the net result is that the collisions are effectively perfectly elastic. On the average, every particle that comes in leaves with the same energy. So we shall imagine that the gas is in a steady condition, and we lose no energy to the piston because the piston is standing still. In those circumstances, if a particle comes in with a certain speed, it comes out with the same speed and, we will say, with the same mass.

If $\mathbf{v}$ is the velocity of an atom, and v_x is the x-component of $\mathbf{v}$, then mv_x is the x-component of momentum "in"; but we also have an equal component of momentum "out," and so the total momentum delivered to the piston by the particle, in one collision, is $2mv_x$, because it is "reflected."

Now, we need the number of collisions made by the atoms in a second, or in a certain amount of time dt; then we divide by dt. How many atoms are hitting? Let us suppose that there are N atoms in the volume V, or $n = N/V$ in each unit volume. To find how many atoms hit the piston, we note that, given a certain

amount of time t, if a particle has a certain velocity toward the piston it will hit during the time t, provided it is close enough. If it is too far away, it goes only part way toward the piston in the time t, but does not reach the piston. Therefore it is clear that only those molecules which are within a distance $v_x t$ from the piston are going to hit the piston in the time t. Thus the number of collisions in a time t is equal to the number of atoms which are in the region within a distance $v_x t$, and since the area of the piston is A, the *volume* occupied by the atoms which are going to hit the piston is $v_x tA$. But the *number* of atoms that are going to hit the piston is that volume times the number of atoms per unit volume, $nv_x tA$. Of course we do not want the number that hit in a time t, we want the number that hit per second, so we divide by the time t, to get $nv_x A$. (This time t could be made very short; if we feel we want to be more elegant, we call it dt, then differentiate, but it is the same thing.)

So we find that the force is

$$F = nv_x A \cdot 2mv_x. \tag{39.3}$$

See, the force *is* proportional to the area, if we keep the particle density fixed as we change the area! The pressure is then

$$P = 2nmv_x^2. \tag{39.4}$$

Now we notice a little trouble with this analysis: First, all the molecules do not have the same velocity, and they do not move in the same direction. So, all the v_x^2's are different! So what we must do, of course, is to take an *average* of the v_x^2's, since each one makes its own contribution. What we want is the square of v_x, averaged over all the molecules:

$$P = nm\langle v_x^2\rangle. \tag{39.5}$$

Did we forget to include the factor 2? No; of all the atoms, only half are headed toward the piston. The other half are headed the other way, and if we take $\langle v_x^2\rangle$, we are averaging the *negative* v_x's squared, as well as the positive v_x's. So when we just take $\langle v_x^2\rangle$, without looking, we are getting twice as much as we want. The average of v_x^2, for positive v_x, is equal to the average of v_x^2 for *all* v_x, times one half.

Now as the atoms bounce around, it is clear that there is nothing special about the "x-direction"; the atoms may also be moving up and down, back and forth, in and out. Therefore it is going to be true that $\langle v_x^2\rangle$, the average motion of the atoms in one direction, and the average in the other two directions, are all going to be equal:

$$\langle v_x^2\rangle = \langle v_y^2\rangle = \langle v_z^2\rangle. \tag{39.6}$$

It is only a matter of rather tricky mathematics to notice, therefore, that they are each equal to one-third of their sum, which is of course the square of the magnitude of the velocity:

$$\langle v_x^2\rangle = \tfrac{1}{3}\langle v_x^2 + v_y^2 + v_z^2\rangle = \langle v^2\rangle/3. \tag{39.7}$$

This has the advantage that we do not have to worry about any particular direction, and so we write our pressure formula again in this form:

$$P = (\tfrac{2}{3})n\langle mv^2/2\rangle. \tag{39.8}$$

The reason we wrote the last factor as $\langle mv^2/2\rangle$ is that this is the *kinetic energy* of the center-of-mass motion of the molecule. We find, therefore, that

$$PV = N(\tfrac{2}{3})\langle mv^2/2\rangle. \tag{39.9}$$

With this equation we can calculate how much the pressure is, if we know the speeds.

As a very simple example let us take helium gas, or any other gas, like mercury vapor, or potassium vapor of high enough temperature, or argon, in which all the molecules are single atoms, for which we may suppose that there is no internal

motion in the atom. If we had a complex molecule, there might be some internal motion, mutual vibrations, or something. We suppose that we may disregard that; this is actually a serious matter that we will have to come back to, but it turns out to be all right. We suppose that the internal motion of the atoms can be disregarded, and therefore, for this purpose, that the kinetic energy of the center-of-mass motion is all the energy there is. So for a monatomic gas, the kinetic energy is the total energy. In general, we are going to call U the total energy (it is sometimes called the total *internal* energy—we may wonder why, since there is no *external* energy to a gas), i.e., all the energy of all the molecules in the gas, or the object, whatever it is.

For a monatomic gas we will suppose that the total energy U is equal to a number of atoms times the average kinetic energy of each, because we are disregarding any possibility of excitation or motion inside the atoms themselves. Then, in these circumstances, we would have

$$PV = \tfrac{2}{3}U. \tag{39.10}$$

Incidentally, we can stop here and find the answer to the following question: Suppose that we take a can of gas and compress the gas slowly, how much pressure do we need to squeeze the volume down? It is easy to find out, since the pressure is $\frac{2}{3}$ the energy divided by V. As we squeeze it down, we do work on the gas and we thereby increase the energy U. So we are going to have some kind of a differential equation: If we start out in a given circumstance with a certain energy and a certain volume, we then know the pressure. Now we start to squeeze, but the moment we do, the energy U increases and the volume V decreases, so the pressure goes up.

So, we have to solve a differential equation, and we will solve it in a moment. We must first emphasize, however, that as we are compressing this gas, we are supposing that all the work goes into increasing the energy of the atoms inside. We may ask, "Isn't that necessary? Where else could it go?" It turns out that it *can* go another place. There are what we call "heat leaks" through the walls: the hot (i.e., fast-moving) atoms that bombard the walls, heat the walls, and energy goes away. We shall suppose for the present that this is not the case.

For somewhat wider generality, although we are still making some very special assumptions about our gas, we shall write, not $PV = \frac{2}{3}U$, but

$$PV = (\gamma - 1)U. \tag{39.11}$$

It is written $(\gamma - 1)$ times U for conventional reasons, because we will deal with a few other cases later where the number in front of U will not be $\frac{2}{3}$, but will be a different number. So, in order to do the thing in general, we call it $\gamma - 1$, because people have been calling it that for almost one hundred years. This γ, then, is $\frac{5}{3}$, because $\frac{5}{3} - 1$ is $\frac{2}{3}$ for a monatomic gas like helium.

We have already noticed that when we compress a gas the work done is $-P\,dV$. A compression in which there is no heat energy added or removed is called an *adiabatic* compression, from the Greek *a* (not) + *dia* (through) + *bainein* (to go). (The word adiabatic is used in physics in several ways, and it is sometimes hard to see what is common about them.) That is, for an adiabatic compression all the work done goes into changing the internal energy. That is the key—that there are no other losses of energy—for then we have $P\,dV = -dU$. But since $U = PV/(\gamma - 1)$, we may write

$$dU = (P\,dV + V\,dP)/(\gamma - 1). \tag{39.12}$$

So we have $P\,dV = -(P\,dV + V\,dP)/(\gamma - 1)$, or, rearranging the terms, $\gamma P\,dV = -V\,dP$, or

$$(\gamma\,dV/V) + (dP/P) = 0. \tag{39.13}$$

Fortunately, assuming that γ is constant, as it is for a monatomic gas, we can integrate this: it gives $\gamma \ln V + \ln P = \ln C$, where $\ln C$ is the constant of integration. If we take the exponential of both sides, we get the law

$$PV^\gamma = C \text{ (a constant)}. \tag{39.14}$$

In other words, under adiabatic conditions, where the temperature rises as we compress because no heat is being lost, the pressure times the volume to the $\frac{5}{3}$ power is a constant for a monatomic gas! Although we derived it theoretically, this *is*, in fact, the way monatomic gases behave experimentally.

39–3 Compressibility of radiation

We may give one other example of the kinetic theory of a gas, one which is not used in chemistry so much, but is used in astronomy. We have a large number of photons in a box in which the temperature is very high. (The box is, of course, the gas in a very hot star. The sun is not hot enough; there are still too many atoms, but at still higher temperatures in certain very hot stars, we may neglect the atoms and suppose that the only objects that we have in the box are photons.) Now then, a photon has a certain momentum $\mathbf{p}$. (We always find that we are in terrible trouble when we do kinetic theory: p is the pressure, but p is the momentum; v is the volume, but v is the velocity; T is the temperature, but T is the kinetic energy or the time or the torque; one must keep one's wits about one!) This $\mathbf{p}$ is momentum, it is a vector. Going through the same analysis as before, it is the x-component of the vector $\mathbf{p}$ which generates the "kick," and twice the x-component of the vector $\mathbf{p}$ is the momentum which is given in the kick. Thus $2p_x$ replaces $2mv_x$, and in evaluating the number of collisions, v_x is still v_x, so when we get all the way through, we find that the pressure in Eq. (39.4) is, instead,

$$P = 2np_xv_x. \tag{39.15}$$

Then, in the averaging, it becomes n times the average of p_xv_x (the same factor of 2) and, finally, putting in the other two directions, we find

$$PV = N\langle \mathbf{p} \cdot \mathbf{v} \rangle/3. \tag{39.16}$$

This checks with the formula (39.9), because the momentum is $m\mathbf{v}$; it is a little more general, that is all. The pressure times the volume is the total number of atoms times $\frac{1}{3}(\mathbf{p} \cdot \mathbf{v})$, averaged.

Now, for photons, what is $\mathbf{p} \cdot \mathbf{v}$? The momentum and the velocity are in the same direction, and the velocity is the speed of light, so this is the momentum of each of the objects, times the speed of light. The momentum times the speed of light of every photon is its energy: $E = pc$, so these terms are the *energies* of each of the photons, and we should, of course, take an average energy, times the number of photons. So we have $\frac{1}{3}$ of the energy inside the gas:

$$PV = U/3 \text{ (photon gas)}. \tag{39.17}$$

For photons, then, since we have $\frac{1}{3}$ in front, $(\gamma - 1)$ in (39.11) is $\frac{1}{3}$, or $\gamma = \frac{4}{3}$, and we have discovered that radiation in a box obeys the law

$$PV^{4/3} = C. \tag{39.18}$$

So we know the compressibility of radiation! That is what is used in an analysis of the contribution of radiation pressure in a star, that is how we calculate it, and how it changes when we compress it. What wonderful things are already within our power!

39–4 Temperature and kinetic energy

So far we have not dealt with *temperature*; we have purposely been avoiding the temperature. As we compress a gas, we know that the energy of the molecules increases, and we are used to saying that the gas gets hotter; we would like to understand what this has to do with the temperature. If we try to do the experiment, not adiabatically but at what we call *constant temperature*, what are we doing? We know that if we take two boxes of gas and let them sit next to each other long enough, even if at the start they were at what we call different temperatures, they

will in the end come to the same temperature. Now what does that mean? That means that they get to a condition that they would get to if we left them alone long enough! What we mean by equal temperature is just that—the final condition when things have been sitting around interacting with each other long enough.

Let us consider, now, what happens if we have two gases in containers separated by a movable piston as in Fig. 39–2 (just for simplicity we shall take two monatomic gases, say helium and neon). In container (1) the atoms have mass m_1, velocity v_1, and there are n_1 per unit volume, and in the other container the atoms have mass m_2, velocity v_2, there are n_2 atoms per unit volume. What are the conditions for equilibrium?

Obviously, the bombardment from the left side must be such that it moves the piston to the right and compresses the other gas until its pressure builds up, and the thing will thus slosh back and forth, and will gradually come to rest at a place where the pressures are equal on both sides. So we can arrange that the pressures are equal; that just means that the internal energies per unit volume are equal, or that the numbers n times the average kinetic energies on each side are equal. What we have to try to prove, eventually, is that the *numbers themselves* are equal. So far, all we know is that the numbers times the kinetic energies are equal,

$$n_1\langle m_1 v_1^2/2\rangle = n_2\langle m_2 v_2^2/2\rangle,$$

from (39.8), because the pressures are equal. We must realize that this is not the only condition over the long run, but something else must happen more slowly as the true complete equilibrium corresponding to equal temperatures sets in.

To see the idea, suppose that the pressure on the left side were developed by having a very high density but a low velocity. By having a large n and a small v, we can get the same pressure as by having a small n and a large v. The atoms may be moving slowly but be packed nearly solidly, or there may be fewer but they are hitting harder. Will it stay like that forever? At first we might think so, but then we think again and find we have forgotten one important point. That is, that the intermediate piston does not receive a steady pressure; it wiggles, just like the eardrum that we were first talking about, because the bangings are not absolutely uniform. There is not a perpetual, steady pressure, but a tattoo—the pressure varies, and so the thing jiggles. Suppose that the atoms on the right side are not jiggling much, but those on the left are few and far between and very energetic. The piston will, now and then, get a big impulse from the left, and will be driven against the slow atoms on the right, giving them more speed. (As each atom collides with the piston, it either gains or loses energy, depending upon whether the piston is moving one way or the other when the atom strikes it.) So, as a result of the collisions, the piston finds itself jiggling, jiggling, jiggling, and this shakes the other gas—it gives energy to the other atoms, and they build up faster motions, until they balance the jiggling that the piston is giving to them. The system comes to some equilibrium where the piston is moving at such a mean square speed that it picks up energy from the atoms at about the same rate as it puts energy back into them. So the piston picks up a certain mean irregularity in speed, and it is our problem to find it. When we do find it, we can solve our problem better, because the gases will adjust their velocities until the rate at which they are trying to pour energy into each other through the piston will become equal.

It is quite difficult to figure out the details of the piston in this particular circumstance; although it is ideally simple to understand, it turns out to be a little harder to analyze. Before we analyze that, let us analyze another problem in which we have a box of gas but now we have two different kinds of molecules in it, having masses m_1 and m_2, velocities v_1 and v_2, and so forth; there is now a much more intimate relationship. If all of the No. 2 molecules are standing still, that condition is not going to last, because they get kicked by the No. 1 molecules and so pick up speed. If they are all going much faster than the No. 1 molecules, then maybe that will not last either—they will pass the energy back to the No. 1 molecules. So when both gases are in the same box, the problem is to find the rule that determines the relative speeds of the two.

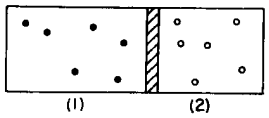


Fig. 39–2. Atoms of two different monatomic gases are separated by a movable piston.

This is still a very difficult problem, but we will solve it as follows. First we consider the following sub-problem (again this is one of those cases where—never mind the derivation—in the end the result is very simple to remember, but the derivation is just ingenious). Let us suppose that we have two molecules, of different mass, colliding, and that the collision is viewed on the center-of-mass (CM) system. In order to remove a complication, we look at the collision in the CM. As we know from the laws of collision, by the conservation of momentum and energy, after the molecules collide the only way they can move is such that each maintains its own original speed—and they just change their *direction*. So we have an average collision that looks like that in Fig. 39–3. Suppose, for a moment, that we watch all the collisions with the CM at rest. Suppose we imagine that they are all initially moving horizontally. Of course, after the first collision some of them are moving at an angle. In other words, if they were all going horizontally, then at least some would later be moving vertically. Now in some other collision, they would be coming in from another direction, and then they would be twisted at still another angle. So even if they were completely organized in the beginning, they would get sprayed around at all angles, and then the sprayed ones would get sprayed some more, and sprayed some more, and sprayed some more. Ultimately, what will be the distribution? *Answer: It will be equally likely to find any pair moving in any direction in space.* After that further collisions could not change the distribution.

They are equally likely to go in all directions, but how do we say that? There is of course *no* likelihood that they will go in any specific direction, because a specific direction is too exact, so we have to talk about per unit "something." The idea is that any area on a sphere centered at a collision point will have just as many molecules going through it as go through any other equal area on the sphere. So the result of the collisions will be to distribute the directions so that equal areas on a sphere will have equal probabilities.

Incidentally, if we just want to discuss the original direction and some other direction an angle θ from it, it is an interesting property that the differential area of a sphere of unit radius is $\sin\theta\, d\theta$ times 2π, and that is the same as the differential of $\cos\theta$. So what it means is that the cosine of the angle θ between any two directions is equally likely to be anything from -1 to $+1$.

Next, we have to worry about the actual case, where we do not have the collision in the CM system, but we have two atoms which are coming together with vector velocities $\mathbf{v}_1$ and $\mathbf{v}_2$. What happens now? We can analyze this collision with the vector velocities $\mathbf{v}_1$ and $\mathbf{v}_2$ in the following way: We first say that there is a certain CM; the velocity of the CM is given by the "average" velocity, with weights proportional to the masses, so the velocity of the CM is $\mathbf{v}_{CM} = (m_1\mathbf{v}_1 + m_2\mathbf{v}_2)/(m_1 + m_2)$. If we watch this collision in the CM system, then we see a collision just like that in Fig. 39–3, with a certain relative velocity $\mathbf{w}$ coming in. The relative velocity is just $\mathbf{v}_1 - \mathbf{v}_2$. Now the idea is that, first, the whole CM is moving, and in the CM there is a relative velocity $\mathbf{w}$, and the molecules collide and come off in some new direction. All this happens while the CM keeps right on moving, without any change.

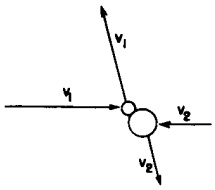


Fig. 39–3. A collision between unequal molecules, viewed in the CM system.

Now then, what is the distribution resulting from this? From our previous argument we conclude this: that at equilibrium, *all directions for* $\mathbf{w}$ *are equally likely, relative to the direction of the motion of the CM.** There will be no particular correlation, in the end, between the direction of the motion of the relative velocity and that of the motion of the CM. Of course, if there were, the collisions would spray it about, so it is all sprayed around. So the cosine of the angle between $\mathbf{w}$ and $\mathbf{v}_{CM}$ is zero on the average. That is,

$$\langle \mathbf{w} \cdot \mathbf{v}_{CM} \rangle = 0. \tag{39.19}$$

* This argument, which was the one used by Maxwell, involves some subtleties. Although the conclusion is correct, the result does *not* follow purely from the considerations of symmetry that we used before, since, by going to a reference frame moving through the gas, we may find a distorted velocity distribution. We have not found a simple proof of this result.

But $\mathbf{w} \cdot \mathbf{v}_{CM}$ can be expressed in terms of $\mathbf{v}_1$ and $\mathbf{v}_2$ as well:

$$\mathbf{w} \cdot \mathbf{v}_{CM} = \frac{(\mathbf{v}_1 - \mathbf{v}_2) \cdot (m_1\mathbf{v}_1 + m_2\mathbf{v}_2)}{m_1 + m_2}$$
$$= \frac{(m_1 v_1^2 - m_2 v_2^2) + (m_2 - m_1)(\mathbf{v}_1 \cdot \mathbf{v}_2)}{m_1 + m_2}. \qquad (39.20)$$

First, let us look at the $\mathbf{v}_1 \cdot \mathbf{v}_2$; what is the average of $\mathbf{v}_1 \cdot \mathbf{v}_2$? That is, what is the average of the component of velocity of one molecule in the direction of another? Surely there is just as much likelihood of finding any given molecule moving one way as another. *The average of the velocity v_2 in any direction is zero.* Certainly, then, in the direction of $\mathbf{v}_1$, $\mathbf{v}_2$ has zero average. So, the average of $\mathbf{v}_1 \cdot \mathbf{v}_2$ is zero! Therefore, we conclude that the average of $m_1 v_1^2$ must be equal to the average of $m_2 v_2^2$. That is, *the average kinetic energy of the two must be equal:*

$$\tfrac{1}{2}m_1 v_1^2 = \tfrac{1}{2}m_2 v_2^2. \qquad (39.21)$$

If we have two kinds of atoms in a gas, it can be shown, and we presume to have shown it, that the average of the kinetic energy of one is the same as the average of the kinetic energy of the other, when they are both in the same gas in the same box in equilibrium. That means that the heavy ones will move slower than the light ones; this is easily shown by experimentation with "atoms" of different masses in an air trough.

Now we would like to go one step further, and say that if we have two different gases *separated* in a box, they will also have equal average kinetic energy when they have finally come to equilibrium, even though they are not in the same box. We can make the argument in a number of ways. One way is to argue that if we have a fixed partition with a tiny hole in it (Fig. 39–4) so that one gas could leak out through the holes while the other could not, because the molecules are too big, and these had attained equilibrium, then we know that in one part, where they are mixed, they have the same average kinetic energy, but some come through the hole without loss of kinetic energy, so the average kinetic energy in the pure gas and in the mixture must be the same. That is not too satisfactory, because maybe there are no holes, for this kind of molecule, that separate one kind from the other.

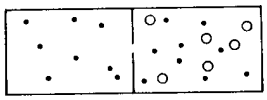


Fig. 39–4. Two gases in a box with a semipermeable membrane.

Let us now go back to the piston problem. We can give an argument which shows that the kinetic energy of this piston must also be $\frac{1}{2}m_2 v_2^2$. Actually, that would be the kinetic energy due to the purely horizontal motion of the piston, so, forgetting its up and down motion, it will have to be the same as $\frac{1}{2}m_2 v_{2_x}^2$. Likewise, from the equilibrium on the other side, we can prove that the kinetic energy of the piston is $\frac{1}{2}m_1 v_{1_x}^2$. Although this is not in the middle of the gas, but is on one side of the gas, we can still make the argument, although it is a little more difficult, that the average kinetic energy of the piston and of the gas molecules are equal as a result of all the collisions.

If this still does not satisfy us, we may make an artificial example by which the equilibrium is generated by an object which can be hit on all sides. Suppose that we have a short rod with a ball on each end sticking through the piston, on a frictionless sliding universal joint. Each ball is round, like one of the molecules, and can be hit on all sides. This whole object has a certain total mass, m. Now, we have the gas molecules with mass m_1 and mass m_2 as before. The result of the collisions, by the analysis that was made before, is that the kinetic energy of m because of collisions with the molecules on one side must be $\frac{1}{2}m_1 v_1^2$, on the average. Likewise, because of the collisions with molecules on the other side, it has to be $\frac{1}{2}m_2 v_2^2$ on the average. So, therefore, both sides have to have the *same* kinetic energy when they are in thermal equilibrium. So, although we only proved it for a mixture of gases, it is easily extended to the case where there are two different, separate gases at the same temperature.

Thus *when we have two gases at the same temperature, the mean kinetic energy of the CM motions are equal.*

The mean molecular kinetic energy is a property only of the "temperature." Being a property of the "temperature," *and not of the gas*, we can use it as a *definition* of the temperature. The mean kinetic energy of a molecule is thus some

function of the temperature. But who is to tell us what scale to use for the temperature? We may arbitrarily *define* the scale of temperature so that the mean energy is linearly proportional to the temperature. The best way to do it would be to call the mean energy itself "the temperature." That would be the simplest possible function. Unfortunately, the scale of temperature has been chosen differently, so instead of calling it temperature directly we use a constant conversion factor between the energy of a molecule and a degree of absolute temperature called a degree Kelvin. The constant of proportionality is $k = 1.38 \times 10^{-23}$ joule for every degree Kelvin.* So if T is absolute temperature, our definition says that the mean molecular kinetic energy is $\frac{3}{2}kT$. (The $\frac{3}{2}$ is put in as a matter of convenience, so as to get rid of it somewhere else.)

We point out that the kinetic energy associated with the component of motion in any particular direction is only $\frac{1}{2}kT$. The three independent directions that are involved make it $\frac{3}{2}kT$.

39–5 The ideal gas law

Now, of course, we can put our definition of temperature into Eq. (39.9) and so find the law for the pressure of gases as a function of the temperature: it is that the pressure times the volume is equal to the total number of atoms times the universal constant k, times the temperature:

$$PV = NkT. \tag{39.22}$$

Furthermore, at the same temperature and pressure and volume, the *number of atoms* is determined; it too is a universal constant! So equal volumes of different gases, at the same pressure and temperature, have the same number of molecules, because of Newton's laws. That is an amazing conclusion!

In practice, when dealing with molecules, because the numbers are so large, the chemists have artificially chosen a specific number, a very large number, and called it something else. They have a number which they call a *mole*. A mole is merely a handy number. Why they did not choose 10^{24} objects, so it would come out even, is a historical question. They happened to choose, for the convenient number of objects on which they standardize, $N_0 = 6.02 \times 10^{23}$ objects, and this is called a mole of objects. So instead of measuring the number of molecules in units, they measure in terms of numbers of moles.† In terms of N_0 we can write the number of moles, times the number of atoms in a mole, times kT, and if we want to, we can take the number of atoms in a mole times k, which is a mole's worth of k, and call it something else, and we do—we call it R. A mole's worth of k is 8.317 joules: $R = N_0k = 8.317\,\text{j} \cdot \text{mole}^{-1} \cdot {}^\circ\text{K}^{-1}$. Thus we also find the gas law written as the number of moles (also called N) times RT, or the number of atoms, times kT:

$$PV = NRT. \tag{39.23}$$

It is the same thing, just a different scale for measuring numbers. We use 1 as a unit, and chemists use 6×10^{23} as a unit!

We now make one more remark about our gas law, and that has to do with the law for objects other than monatomic molecules. We have dealt only with the CM motion of the atoms of a monatomic gas. What happens if there are forces present? First, consider the case that the piston is held by a horizontal spring, and there are forces on it. The exchange of jiggling motion between atoms and piston at any moment does not depend on where the piston is at that moment, of course. The equilibrium conditions are the same. No matter where the piston is, its speed of motion must be such that it passes energy to the molecules in just

---

* The centigrade scale is just this Kelvin scale with a zero chosen at 273.16 °K, so $T = 273.16 +$ centigrade temperature.

† What the chemists call molecular weights are the masses in grams of a mole of a molecule. The mole is defined so that the mass of a mole of carbon atoms of isotope 12 (i.e., having 6 protons and 6 neutrons in the nucleus) is exactly 12 grams.

the right way. So it makes no difference about the spring. The *speed* at which the piston has to move, on the average, is the same. So our theorem, that the mean value of the kinetic energy in one direction is $\frac{1}{2}kT$, *is true whether there are forces present or not.*

Consider, for example, a diatomic molecule composed of atoms m_A and m_B. What we have proved is that the motion of the CM of part A and that of part B are such that $\langle\frac{1}{2}m_Av_A^2\rangle = \langle\frac{1}{2}m_Bv_B^2\rangle = \frac{3}{2}kT$. How can this be, if they are held together? Although they are held together, when they are spinning and turning in there, when something hits them, exchanging energy with them, *the only thing that counts is how fast they are moving.* That alone determines how fast they exchange energy in collisions. At the particular instant, the force is not an essential point. Therefore the same principle is right, even when there are forces.

Let us prove, finally, that the gas law is consistent also with a disregard of the internal motion. We did not really include the internal motions before; we just treated a monatomic gas. But we shall now show that an entire object, considered as a single body of total mass M, has a velocity of the CM such that

$$\tfrac{1}{2}Mv_{\text{CM}}^2 = \tfrac{3}{2}kT. \tag{39.24}$$

In other words, we can consider either the separate pieces or the whole thing! Let us see the reason for that: The mass of the diatomic molecule is $M = m_A + m_B$, and the velocity of the center of mass is equal to $\mathbf{v}_{\text{CM}} = (m_A\mathbf{v}_A + m_B\mathbf{v}_B)/M$. Now we need $\langle v_{\text{CM}}^2\rangle$. If we square $\mathbf{v}_{\text{CM}}$, we get

$$v_{\text{CM}}^2 = \frac{m_A^2v_A^2 + 2m_Am_B\mathbf{v}_A\cdot\mathbf{v}_B + m_B^2v_B^2}{M^2}.$$

Now we multiply $\frac{1}{2}M$ and take the average, and thus we get

$$\begin{aligned}\tfrac{1}{2}Mv_{\text{CM}}^2 &= \frac{m_A\frac{3}{2}kT + 2m_Am_B\langle\mathbf{v}_A\cdot\mathbf{v}_B\rangle + m_B\frac{3}{2}kT}{M}\\ &= \tfrac{3}{2}kT + \frac{2m_Am_B\langle\mathbf{v}_A\cdot\mathbf{v}_B\rangle}{M}.\end{aligned}$$

(We have used the fact that $(m_A + m_B)/M = 1$.) Now what is $\langle\mathbf{v}_A\cdot\mathbf{v}_B\rangle$? (It had better be zero!) To find out, let us use our assumption that the relative velocity, $\mathbf{w} = \mathbf{v}_A - \mathbf{v}_B$ is not any more likely to point in one direction than in another—that is, that its average component in any direction is zero. Thus we assume that

$$\langle\mathbf{w}\cdot\mathbf{v}_{\text{CM}}\rangle = 0.$$

But what *is* $\mathbf{w}\cdot\mathbf{v}_{\text{CM}}$? It is

$$\begin{aligned}\mathbf{w}\cdot\mathbf{v}_{\text{CM}} &= \frac{(\mathbf{v}_A - \mathbf{v}_B)\cdot(m_A\mathbf{v}_A + m_B\mathbf{v}_B)}{M}\\ &= \frac{m_Av_A^2 + (m_B - m_A)(\mathbf{v}_A\cdot\mathbf{v}_B) - m_Bv_B^2}{M}.\end{aligned}$$

Therefore, since $\langle m_Av_A^2\rangle = \langle m_Bv_B^2\rangle$, the first and last terms cancel out on the average, and we are left with

$$(m_B - m_A)\langle\mathbf{v}_A\cdot\mathbf{v}_B\rangle = 0.$$

Thus if $m_A \neq m_B$, we find that $\langle\mathbf{v}_A\cdot\mathbf{v}_B\rangle = 0$, and therefore that the bodily motion of the entire molecule, regarded as a single particle of mass M, has a kinetic energy, on the average, equal to $\frac{3}{2}kT$.

Incidentally, we have also proved at the same time that the average kinetic energy of the *internal* motions of the diatomic molecule, disregarding the bodily motion of the CM, is $\frac{3}{2}kT$! For, the total kinetic energy of the parts of the molecule is $\frac{1}{2}m_Av_A^2 + \frac{1}{2}m_Bv_B^2$, whose average is $\frac{3}{2}kT + \frac{3}{2}kT$, or $3kT$. The kinetic energy of the center-of-mass motion is $\frac{3}{2}kT$, so the average kinetic energy of the rotational and vibratory motions of the two atoms inside the molecule is the difference, $\frac{3}{2}kT$.

The theorem concerning the average energy of the CM motion is general: for any object considered as a whole, with forces present or no, for every independent direction of motion that there is, the average kinetic energy in that motion is $\frac{1}{2}kT$. These "independent directions of motion" are sometimes called the *degrees of freedom* of the system. The number of degrees of freedom of a molecule composed of r atoms is $3r$, since each atom needs three coordinates to define its position. The entire kinetic energy of the molecule can be expressed either as the sum of the kinetic energies of the separate atoms, or as the sum of the kinetic energy of the CM motion plus the kinetic energy of the internal motions. The latter can sometimes be expressed as a sum of rotational kinetic energy of the molecule and vibrational energy, but this is an approximation. Our theorem, applied to the r-atom molecule, says that the molecule will have, on the average, $3rkT/2$ joules of kinetic energy, of which $\frac{3}{2}kT$ is kinetic energy of the center-of-mass motion of the entire molecule, and the rest, $\frac{3}{2}(r - 1)kT$, is internal vibrational and rotational kinetic energy.

40

The Principles of Statistical Mechanics

40–1 The exponential atmosphere

40–2 The Boltzmann law

40–3 Evaporation of a liquid

40–4 The distribution of molecular speeds

40–5 The specific heats of gases

40–6 The failure of classical physics

40–1 The exponential atmosphere

We have discussed some of the properties of large numbers of intercolliding atoms. The subject is called kinetic theory, a description of matter from the point of view of collisions between the atoms. Fundamentally, we assert that the gross properties of matter should be explainable in terms of the motion of its parts.

We limit ourselves for the present to conditions of thermal equilibrium, that is, to a subclass of all the phenomena of nature. The laws of mechanics which apply just to thermal equilibrium are called *statistical mechanics*, and in this section we want to become acquainted with some of the central theorems of this subject.

We already have one of the theorems of statistical mechanics, namely, the mean value of the kinetic energy for any motion at the absolute temperature T is $\frac{1}{2}kT$ for each independent motion, i.e., for each degree of freedom. That tells us something about the mean square velocities of the atoms. Our objective now is to learn more about the positions of the atoms, to discover how many of them are going to be in different places at thermal equilibrium, and also to go into a little more detail on the distribution of the velocities. Although we have the mean square velocity, we do not know how to answer a question such as how many of them are going three times faster than the root mean square, or how many of them are going one-quarter of the root mean square speed. Or have they all the same speed exactly?

So, these are the two questions that we shall try to answer: How are the molecules distributed in space when there are forces acting on them, and how are they distributed in velocity?

It turns out that the two questions are completely independent, and that the distribution of velocities is always the same. We already received a hint of the latter fact when we found that the average kinetic energy is the same, $\frac{1}{2}kT$ per degree of freedom, no matter what forces are acting on the molecules. The distribution of the velocities of the molecules is independent of the forces, because the collision rates do not depend upon the forces.

Let us begin with an example: the distribution of the molecules in an atmosphere like our own, but without the winds and other kinds of disturbance. Suppose that we have a column of gas extending to a great height, and at thermal equilibrium—unlike our atmosphere, which as we know gets colder as we go up. We could remark that if the temperature differed at different heights, we could demonstrate lack of equilibrium by connecting a rod to some balls at the bottom (Fig. 40–1), where they would pick up $\frac{1}{2}kT$ from the molecules there and would shake, via the rod, the balls at the top and those would shake the molecules at the top. So, ultimately, of course, the temperature becomes the same at all heights in a gravitational field.

If the temperature is the same at all heights, the problem is to discover by what law the atmosphere becomes tenuous as we go up. If N is the total number of molecules in a volume V of gas at pressure P, then we know $PV = NkT$, or $P = nkT$, where $n = N/V$ is the number of molecules per unit volume. In other words, if we know the number of molecules per unit volume, we know the pressure, and vice versa: they are proportional to each other, since the temperature is constant in this problem. But the pressure is not constant, it must increase as the altitude is reduced, because it has to hold, so to speak, the weight of all the gas above it. That is the clue by which we may determine how the pressure changes with height. If we take a unit area at height h, then the vertical force from below,

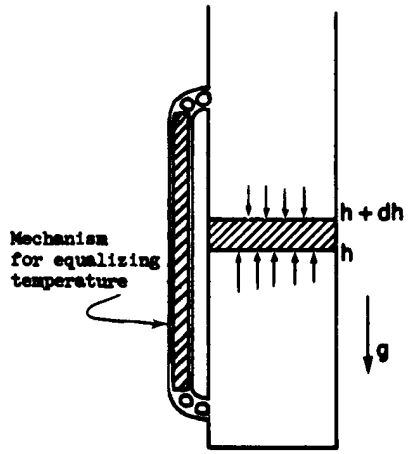


Fig. 40–1. The pressure at height h must exceed that at $h + dh$ by the weight of the intervening gas.

on this unit area, is the pressure P. The vertical force per unit area pushing down at a height $h + dh$ would be the same, in the absence of gravity, but here it is not, because the force from below must exceed the force from above by the weight of gas in the section between h and $h + dh$. Now mg is the force of gravity on each molecule, where g is the acceleration due to gravity, and $n\,dh$ is the total number of molecules in the unit section. So this gives us the differential equation $P_{h+dh} - P_h = dP = -mgn\,dh$. Since $P = nkT$, and T is constant, we can eliminate either P or n, say P, and get

$$\frac{dn}{dh} = -\frac{mg}{kT}n$$

for the differential equation, which tells us how the density goes down as we go up in energy.

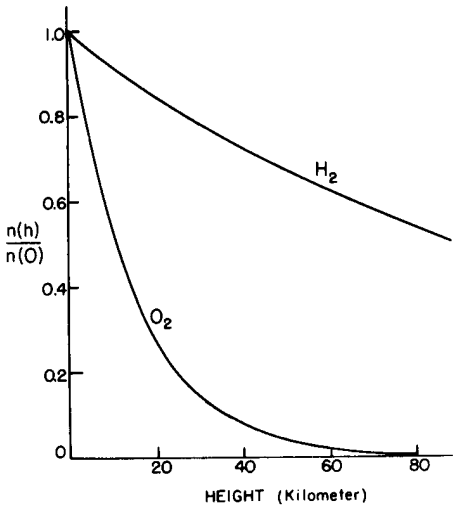


Fig. 40–2. The normalized density as a function of height in the earth's gravitational field for oxygen and for hydrogen, at constant temperature.

We thus have an equation for the particle density n, which varies with height, but which has a derivative which is proportional to itself. Now a function which has a derivative proportional to itself is an exponential, and the solution of this differential equation is

$$n = n_0 e^{-mgh/kT}. \tag{40.1}$$

Here the constant of integration, n_0, is obviously the density at $h = 0$ (which can be chosen anywhere), and the density goes down exponentially with height.

Note that if we have different kinds of molecules with different masses, they go down with different exponentials. The ones which were heavier would decrease with altitude faster than the light ones. Therefore we would expect that because oxygen is heavier than nitrogen, as we go higher and higher in an atmosphere with nitrogen and oxygen the proportion of nitrogen would increase. This does not really happen in our own atmosphere, at least at reasonable heights, because there is so much agitation which mixes the gases back together again. It is not an isothermal atmosphere. Nevertheless, there *is* a tendency for lighter materials, like hydrogen, to dominate at very great heights in the atmosphere, because the lowest masses continue to exist, while the other exponentials have all died out (Fig. 40–2).

40–2 The Boltzmann law

Here we note the interesting fact that the numerator in the exponent of Eq. (40.1) is the *potential energy* of an atom. Therefore we can also state this particular law as: the density at any point is proportional to

$$e^{-(\text{the potential energy of each atom}/kT)}.$$

That may be an accident, i.e., may be true only for this particular case of a uniform gravitational field. However, we can show that it is a more general proposition. Suppose that there were some kind of force other than gravity acting on the molecules in a gas. For example, the molecules may be charged electrically, and may be acted on by an electric field or another charge that attracts them. Or, because of the mutual attractions of the atoms for each other, or for the wall, or for a solid, or something, there is some force of attraction which varies with position and which acts on all the molecules. Now suppose, for simplicity, that the molecules are all the same, and that the force acts on each individual one, so that the total force on a piece of gas would be simply the number of molecules times the force on each one. To avoid unnecessary complication, let us choose a coordinate system with the x-axis in the direction of the force, **F**.

In the same manner as before, if we take two parallel planes in the gas, separated by a distance dx, then the force on each atom, times the n atoms per cm^3 (the generalization of the previous nmg), times dx, must be balanced by the pressure change: $Fn\,dx = dP = kT\,dn$. Or, to put this law in a form which will be useful to us later,

$$F = kT\frac{d}{dx}(\ln n). \tag{40.2}$$

For the present, observe that $-F\,dx$ is the work we would do in taking a molecule from x to $x + dx$, and if F comes from a potential, i.e., if the work done can be represented by a potential energy at all, then this would also be the difference in the potential energy (P.E.). The negative differential of potential energy is the work done, $F\,dx$, and we find that $d(\ln n) = -d(\text{P.E.})/kT$, or, after integrating,

$$n = (\text{constant})e^{-\text{P.E.}/kT}. \tag{40.3}$$

Therefore what we noticed in a special case turns out to be true in general. (What if F does not come from a potential? Then (40.2) has no solution at all. Energy can be generated, or lost by the atoms running around in cyclic paths for which the work done is not zero, and no equilbrium can be maintained at all. Thermal equilibrium cannot exist if the external forces on the atoms are not conservative.) Equation (40.3), known as *Boltzmann's law*, is another of the principles of statistical mechanics: that the probability of finding molecules in a given spatial arrangement varies exponentially with the negative of the potential energy of that arrangement, divided by kT.

This, then, could tell us the distribution of molecules: Suppose that we had a positive ion in a liquid, attracting negative ions around it, how many of them would be at different distances? If the potential energy is known as a function of distance, then the proportion of them at different distances is given by this law, and so on, through many applications.

40–3 Evaporation of a liquid

In more advanced statistical mechanics one tries to solve the following important problem. Consider an assembly of molecules which attract each other, and suppose that the force between any two, say i and j, depends only on their separation r_{ij}, and can be represented as the derivative of a potential function $V(r_{ij})$. Figure 40–3 shows a form such a function might have. For $r > r_0$, the energy decreases as the molecules come together, because they attract, and then the energy increases very sharply as they come still closer together, because they repel strongly, which is characteristic of the way molecules behave, roughly speaking.

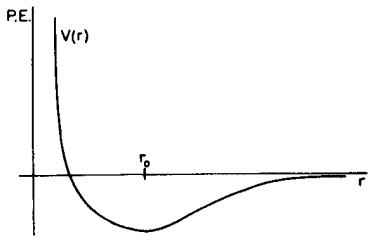


Fig. 40–3. A potential-energy function for two molecules, which depends only on their separation.

Now suppose we have a whole box full of such molecules, and we would like to know how they arrange themselves on the average. The answer is $e^{-\text{P.E.}/kT}$. The total potential energy in this case would be the sum over all the pairs, supposing that the forces are all in pairs (there may be three-body forces in more complicated things, but in electricity, for example, the potential energy is all in pairs). Then the probability for finding molecules in any particular combination of r_{ij}'s will be proportional to

$$\exp\Big[-\sum_{i,j} V(r_{ij})/kT\Big].$$

Now, if the temperature is very high, so that $kT \gg |V(r_0)|$, the exponent is relatively small almost everywhere, and the probability of finding a molecule is almost independent of position. Let us take the case of just two molecules: the $e^{-\text{P.E.}/kT}$ would be the probability of finding them at various mutual distances r. Clearly, where the potential goes most negative, the probability is largest, and where the potential goes toward infinity, the probability is almost zero, which occurs for very small distances. That means that for such atoms in a gas, there is no chance that they are on top of each other, since they repel so strongly. But there is a greater chance of finding them *per unit volume* at the point r_0 than at any other point. How much greater, depends on the temperature. If the temperature is very large compared with the difference in energy between $r = r_0$ and $r = \infty$, the exponential is always nearly unity. In this case, where the mean kinetic energy (about kT) greatly exceeds the potential energy, the forces do not make much difference. But as the temperature falls, the probability of finding the molecules at the preferred distance r_0 gradually increases relative to the probability of finding them elsewhere and, in fact, if kT is much less than $|V(r_0)|$, we have a relatively

large positive exponent in that neighborhood. In other words, in a given volume they are *much more* likely to be at the distance of minimum energy than far apart. As the temperature falls, the atoms fall together, clump in lumps, and reduce to liquids, and solids, and molecules, and as you heat them up they evaporate.

The requirements for the determination of exactly how things evaporate, exactly how things should happen in a given circumstance, involve the following. First, to discover the correct molecular-force law $V(r)$, which must come from something else, quantum mechanics, say, or experiment. But, given the law of force between the molecules, to discover what a billion molecules are going to do merely consists of studying the function $e^{-\Sigma V_{ij}/kT}$. Surprisingly enough, since it is such a simple function and such an easy idea, given the potential, the labor is *enormously complicated;* the difficulty is the tremendous number of variables.

In spite of such difficulties, the subject is quite exciting and interesting. It is often called an example of a "many-body problem," and it really has been a very interesting thing. In that single formula must be contained all the details, for example, about the solidification of gas, or the forms of the crystals that the solid can take, and people have been trying to squeeze it out, but the mathematical difficulties are very great, not in writing the law, but in dealing with so enormous a number of variables.

That then, is the distribution of particles in space. That is the end of classical statistical mechanics, practically speaking, because if we know the forces, we can, in principle, find the distribution in space, and the distribution of velocities is something that we can work out once and for all, and is not something that is different for the different cases. The great problems are in getting particular information out of our formal solution, and that is the main subject of classical statistical mechanics.

40–4 The distribution of molecular speeds

Now we go on to discuss the distribution of velocities, because sometimes it is interesting or useful to know how many of them are moving at different speeds. In order to do that, we may make use of the facts which we discovered with regard to the gas in the atmosphere. We take it to be a perfect gas, as we have already assumed in writing the potential energy, disregarding the energy of mutual attraction of the atoms. The only potential energy that we included in our first example was gravity. We would, of course, have something more complicated if there were forces between the atoms. Thus we assume that there are no forces between the atoms and, for a moment, disregard collisions also, returning later to the justification of this. Now we saw that there are fewer molecules at the height h than there are at the height 0; according to formula (40.1), they decrease exponentially with height. How can there be fewer at greater heights? After all, do not all the molecules which are moving up at height 0 arrive at h? No!, because some of those which are moving up at 0 are going too slowly, and cannot climb the potential hill to h. With that clue, we can calculate how many must be moving at various speeds, because from (40.1) we know *how many* are moving with less than enough speed to climb a given distance h. Those are just the ones that account for the fact that the density at h is lower than at 0.

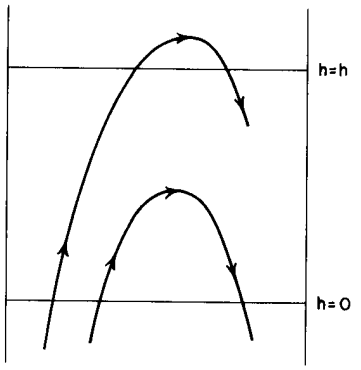


Fig. 40–4. Only those molecules moving up at $h = 0$ with sufficient velocity can arrive at height h.

Now let us put that idea a little more precisely: let us count how many molecules are passing from below to above the plane $h = 0$ (by calling it height $= 0$, we do not mean that there is a floor there; it is just a convenient label, and there is gas at negative h). These gas molecules are moving around in every direction, but some of them are moving through the plane, and at any moment a certain number per second of them are passing through the plane from below to above with different velocities. Now we note the following: if we call u the velocity which is just needed to get up to the height h (kinetic energy $mu^2/2 = mgh$), then the number of molecules per second which are passing upward through the lower plane in a vertical direction with velocity component greater than u is exactly the same as the number which pass through the upper plane with *any* upward velocity. Those molecules whose vertical velocity does not exceed u cannot get through the upper

plane. So therefore we see that

$$\text{Number passing } h = 0 \text{ with } v_z > u = \text{number passing } h = h \text{ with } v_z > 0.$$

But the number which pass through h with any velocity greater than 0 is less than the number which pass through the lower height with any velocity greater than 0, because the number of atoms is greater; that is all we need. We know already that the distribution of velocities is the same, after the argument we made earlier about the temperature being constant all the way through the atmosphere. So, since the velocity distributions are the same, and it is just that there are *more atoms* lower down, clearly the number $n_{>0}(h)$, passing with positive velocity at height h, and the number $n_{>0}(0)$, passing with positive velocity at height 0, are in the same ratio as the densities at the two heights, which is $e^{-mgh/kT}$. But $n_{>0}(h) = n_{>u}(0)$, and therefore we find that

$$\frac{n_{>u}(0)}{n_{>0}(0)} = e^{-mgh/kT} = e^{-mu^2/2kT},$$

since $\frac{1}{2}mu^2 = mgh$. Thus, in words, the number of molecules per unit area per second passing the height 0 with a z-component of velocity greater than u is $e^{-mu^2/2kT}$ times the total number that are passing through the plane with velocity greater than zero.

Now this is not only true at the arbitrarily chosen height 0, but of course it is true at any other height, and thus the distributions of velocities *are* all the same! (The final statement does not involve the height h, which appeared only in the intermediate argument.) The result is a general proposition that gives us the distribution of velocities. It tells us that if we drill a little hole in the side of a gas pipe, a very tiny hole, so that the collisions are few and far between, i.e., are farther apart than the diameter of the hole, then the particles which are coming out will have different velocities, but the fraction of particles which come out at a velocity greater than u is $e^{-mu^2/2kT}$.

Now we return to the question about the neglect of collisions: Why does it not make any difference? We could have pursued the same argument, not with a finite height h, but with an infinitesimal height h, which is so small that there would be no room for collisions between 0 and h. But that was not necessary: the argument is evidently based on an analysis of the energies involved, the conservation of energy, and in the collisions that occur there is an exchange of energies among the molecules. However, we do not really care whether we follow the same molecule if energy is merely exchanged with another molecule. So it turns out that even if the problem is analyzed more carefully (and it is more difficult, naturally, to do a rigorous job), it still makes no difference in the result.

It is interesting that the velocity distribution we have found is just

$$n_{>u} \propto e^{-\text{kinetic energy}/kT}. \tag{40.4}$$

This way of describing the distribution of velocities, by giving the number of molecules that pass a given area with a certain minimum z-component, is not the most convenient way of giving the velocity distribution. For instance, inside the gas, one more often wants to know how many molecules are moving with a z-component of velocity between two given values, and that, of course, is not directly given by Eq. (40.4). We would like to state our result in the more conventional form, even though what we already have written is quite general. *Note that it is not possible to say that any molecule has exactly some stated velocity;* none of them has a velocity *exactly equal* to 1.7962899173 meters per second. So in order to make a meaningful statement, we have to ask how many are to be found in some *range* of velocities. We have to say how many have velocities between 1.796 and 1.797, and so on. On mathematical terms, let $f(u)\,du$ be the fraction of all the molecules which have velocities between u and $u + du$ or, what is the same thing (if du is infinitesimal), all that have a velocity u with a range du. Figure 40–5 shows a possible form for the function $f(u)$, and the shaded part, of width du and mean height $f(u)$, represents this fraction $f(u)\,du$. That is, the ratio of the shaded

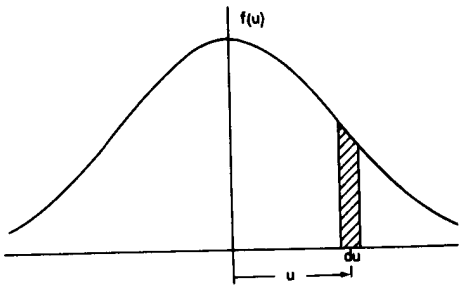


Fig. 40–5. A velocity distribution function. The shaded area is $f(u)\,du$, the fraction of particles having velocities within a range du about u.

area to the total area of the curve is the relative proportion of molecules with velocity u within du. If we define $f(u)$ so that the fraction having a velocity in this range is given directly by the shaded area, then the total area must be 100 percent of them, that is,

$$\int_{-\infty}^{\infty} f(u)\,du = 1. \tag{40.5}$$

Now we have only to get this distribution by comparing it with the theorem we derived before. First we ask, what is the number of molecules passing through an area per second with a velocity greater than u, expressed in terms of $f(u)$? At first we might think it is merely the integral of $\int_u^{\infty} f(u)\,du$, but it is not, because we want the number that are passing the area *per second*. The faster ones pass more often, so to speak, than the slower ones, and in order to express how many pass, you have to multiply by the velocity. (We discussed that in the previous chapter when we talked about the number of collisions.) In a given time t the total number which pass through the surface is all of those which have been able to arrive at the surface, and the number which arrive come from a distance ut. So the number of molecules which arrive is not simply the number which are there, but the number that are there per unit volume, multiplied by the distance that they sweep through in racing for the area through which they are supposed to go, and that distance is proportional to u. Thus we need the integral of u *times* $f(u)\,du$, an infinite integral with a lower limit u, and this must be the same as we found before, namely $e^{-mu^2/2kT}$, with a proportionality constant which we will get later:

$$\int_u^{\infty} uf(u)\,du = \text{const} \cdot e^{-mu^2/2kT}. \tag{40.6}$$

Now if we differentiate the integral with respect to u, we get the thing that is inside the integral, i.e., the integrand (with a minus sign, since u is the lower limit), and if we differentiate the other side, we get u times the same exponential (and some constants). The u's cancel and we find

$$f(u)\,du = Ce^{-mu^2/2kT}\,du. \tag{40.7}$$

We retain the du on both sides as a reminder that it is a *distribution*, and it tells what the proportion is for velocity between u and $u + du$.

The constant C must be so determined that the integral is unity, according to Eq. (40.5). Now we can prove* that

$$\int_{-\infty}^{\infty} e^{-x^2}\,dx = \sqrt{\pi}.$$

Using this fact, it is easy to find that $C = \sqrt{m/2\pi kT}$.

Since velocity and momentum are proportional, we may say that the distribution of momenta is also proportional to $e^{-\text{K.E.}/kT}$ per unit momentum range. It turns out that this theorem is true in relativity too, if it is in terms of momentum, while if it is in velocity it is not, so it is best to learn it in momentum instead of in velocity:

$$f(p)\,dp = Ce^{-\text{K.E.}/kT}\,dp. \tag{40.8}$$

So we find that the probabilities of different conditions of energy, kinetic and potential, are both given by $e^{-\text{energy}/kT}$, a very easy thing to remember and a rather beautiful proposition.

---

* To get the value of the integral, let

$$I = \int_{-\infty}^{\infty} e^{-x^2}\,dx.$$

Then

$$I^2 = \int_{-\infty}^{\infty} e^{-x^2}\,dx \cdot \int_{-\infty}^{\infty} e^{-y^2}\,dy = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-(x^2+y^2)}\,dy\,dx,$$

which is a double integral over the whole xy-plane. But this can also be written in polar coordinates as

$$I^2 = \int_0^{\infty} e^{-r^2} \cdot 2\pi r\,dr = \pi \int_0^{\infty} e^{-t}\,dt = \pi.$$

So far we have, of course, only the distribution of the velocities "vertically." We might want to ask, what is the probability that a molecule is moving in another direction? Of course these distributions are connected, and one can obtain the complete distribution from the one we have, because the complete distribution depends only on the square of the magnitude of the velocity, not upon the z-component. It must be something that is independent of direction, and there is only one function involved, the probability of different magnitudes. We have the distribution of the z-component, and therefore we can get the distribution of the other components from it. The result is that the probability is still proportional to $e^{-\text{K.E.}/kT}$, but now the kinetic energy involves three parts, $mv_x^2/2$, $mv_y^2/2$, and $mv_z^2/2$, summed in the exponent. Or we can write it as a product:

$$f(v_x, v_y, v_z)\,dv_x\,dv_y\,dv_z \propto e^{-mv_x^2/2kT}\cdot e^{-mv_y^2/2kT}\cdot e^{-mv_z^2/2kT}\,dv_x\,dv_y\,dv_z. \qquad (40.9)$$

You can see that this formula must be right because, first, it is a function only of v^2, as required, and second, the probabilities of various values of v_z obtained by integrating over all v_x and v_y is just (40.7). But this one function (40.9) can do both those things!

40–5 The specific heats of gases

Now we shall look at some ways to test the theory, and to see how successful is the classical theory of gases. We saw earlier that if U is the internal energy of N molecules, then $PV = NkT = (\gamma - 1)U$ holds, sometimes, for some gases, maybe. If it is a monatomic gas, we know this is also equal to $\frac{2}{3}$ of the kinetic energy of the center-of-mass motion of the atoms. If it is a monatomic gas, then the kinetic energy is equal to the internal energy, and therefore $\gamma - 1 = \frac{2}{3}$. But suppose it is, say, a more complicated molecule, that can spin and vibrate, and let us suppose (it turns out to be true according to classical mechanics) that the energies of the internal motions are also proportional to kT. Then at a given temperature, in addition to kinetic energy kT, it has internal vibrational or rotational energy. So the total U includes not just the internal kinetic energy, but also the rotational energy, and we get a different value of γ. Technically, the best way to measure γ is by measuring the specific heat, which is the change in energy with temperature. We will return to that approach later. For our present purposes, we may suppose γ is found experimentally from the PV^γ curve for adiabatic compression.

Let us make a calculation of γ for some cases. First, for a monatomic gas U is the total energy, the same as the kinetic energy, and we know already that γ should be $\frac{5}{3}$. For a diatomic gas, we may take, as an example, oxygen, hydrogen iodide, hydrogen, etc., and suppose that the diatomic gas can be represented as two atoms held together by some kind of force like the one of Fig. 40–3. We may also suppose, and it turns out to be quite true, that at the temperatures that are of interest for the diatomic gas, the pairs of atoms tend strongly to be separated by r_0, the distance of potential minimum. If this were not true, if the probability were not strongly varying enough to make the great majority sit near the bottom, we would have to remember that oxygen gas is a mixture of O_2 and single oxygen atoms in a nontrivial ratio. We know that there are, in fact, very few single oxygen atoms, which means that the potential energy minimum is very much greater in magnitude than kT, as we have seen. Since they are clustered strongly around r_0, the only part of the curve that is needed is the part near the minimum, which may be approximated by a parabola. A parabolic potential implies a harmonic oscillator, and in fact, to an excellent approximation, the oxygen molecule can be represented as two atoms connected by a spring.

Now what is the total energy of this molecule at temperature T? We know that for each of the two atoms, each of the kinetic energies should be $\frac{3}{2}kT$, so the kinetic energy of both of them is $\frac{3}{2}kT + \frac{3}{2}kT$. We can also put this in a different way: the same $\frac{3}{2}$ plus $\frac{3}{2}$ can also be looked at as kinetic energy of the center of mass

Table 40–1

Values of the specific heat ratio, γ, for various gases

Gas	T (°C)	γ
He	—180	1.660
Kr	19	1.68
Ar	15	1.668
H_2	100	1.404
O_2	100	1.399
HI	100	1.40
Br_2	300	1.32
I_2	185	1.30
NH_3	15	1.310
C_2H_6	15	1.22

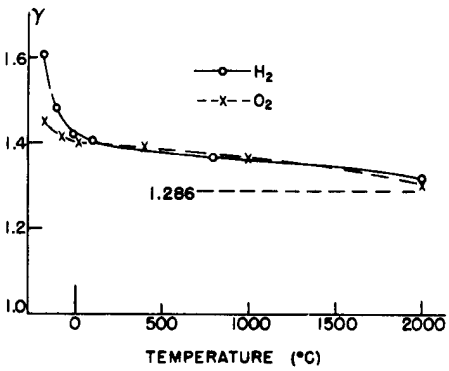


Fig. 40–6. Experimental values of γ as a function of temperature for hydrogen and oxygen. Classical theory predicts $\gamma = 1.286$, independent of temperature.

($\frac{3}{2}$), kinetic energy of rotation ($\frac{2}{2}$), and kinetic energy of vibration ($\frac{1}{2}$). We know that the kinetic energy of vibration is $\frac{1}{2}$, since there is just one dimension involved and each degree of freedom has $\frac{1}{2}kT$. Regarding the rotation, it can turn about either of two axes, so there are two independent motions. We assume that the atoms are some kind of points, and cannot spin about the line joining them; this is something to bear in mind, because if we get a disagreement, maybe that is where the trouble is. But we have one more thing, which is the *potential* energy of vibration; how much is that? In a harmonic oscillator the average kinetic energy and average potential energy are equal, and therefore the potential energy of vibration is $\frac{1}{2}kT$, also. The grand total of energy is $U = \frac{7}{2}kT$, or kT is $\frac{2}{7}U$ per atom. That means, then, that γ is $\frac{9}{7}$ instead of $\frac{5}{3}$, i.e., $\gamma = 1.286$.

We may compare these numbers with the relevant measured values shown in Table 40–1. Looking first at helium, which is a monatomic gas, we find very nearly $\frac{5}{3}$, and the error is probably experimental, although at such a low temperature there may be some forces between the atoms. Krypton and argon, both monatomic, agree also within the accuracy of the experiment.

We turn to the diatomic gases and find hydrogen with 1.404, which does not agree with the theory, 1.286. Oxygen, 1.399, is very similar, but again not in agreement. Hydrogen iodide again is similar at 1.40. It begins to look as though the right answer is 1.40, but it is not, because if we look further at bromine we see 1.32, and at iodine we see 1.30. Since 1.30 is reasonably close to 1.286, iodine may be said to agree rather well, but oxygen is far off. So here we have a dilemma. We have it right for one molecule, we do not have it right for another molecule, and we may need to be pretty ingenious in order to explain both.

Let us look further at a still more complicated molecule with large numbers of parts, for example, C_2H_6, which is ethane. It has eight different atoms, and they are all vibrating and rotating in various combinations, so the total amount of internal energy must be an enormous number of kT's, at least $12kT$ for kinetic energy alone, and $\gamma - 1$ must be very close to zero, or γ almost exactly 1. In fact, it *is* lower, but 1.22 is not so much lower, and is higher than the $1\frac{1}{12}$ calculated from the kinetic energy alone, and it is just not understandable!

Furthermore, the whole mystery is deep, because the diatomic molecule cannot be made rigid by a limit. Even if we made the couplings stiffer indefinitely, although it might not vibrate much, it would nevertheless keep vibrating. The vibrational energy inside is still kT, since it does not depend on the strength of the coupling. But if we could imagine *absolute* rigidity, stopping all vibration to eliminate a variable, then we would get $U = \frac{5}{2}kT$ and $\gamma = 1.40$ for the diatomic case. This looks good for H_2 or O_2. On the other hand, we would still have problems, because γ for either hydrogen or oxygen varies with temperature! From the measured values shown in Fig. 40–6, we see that for H_2, γ varies from about 1.6 at −185°C to 1.3 at 2000°C. The variation is more substantial in the case of hydrogen than for oxygen, but nevertheless, even in oxygen, γ tends definitely to go up as we go down in temperature.

40–6 The failure of classical physics

So, all in all, we might say that we have some difficulty. We might try some force law other than a spring, but it turns out that anything else will only make γ higher. If we include more forms of energy, γ approaches unity more closely, contradicting the facts. All the classical theoretical things that one can think of will only make it worse. The fact is that there are electrons in each atom, and we know from their spectra that there are internal motions; each of the electrons should have at least $\frac{1}{2}kT$ of kinetic energy, and something for the potential energy, so when these are added in, γ gets still smaller. It is ridiculous. It is wrong.

The first great paper on the dynamical theory of gases was by Maxwell in 1859. On the basis of ideas we have been discussing, he was able accurately to explain a great many known relations, such as Boyle's law, the diffusion theory, the viscosity of gases, and things we shall talk about in the next chapter. He listed all these great successes in a final summary, and at the end he said, "Finally, by

establishing a necessary relation between the motions of translation and rotation (he is talking about the $\frac{1}{2}kT$ theorem) of all particles not spherical, we proved that a system of such particles could not possibly satisfy the known relation between the two specific heats." He is referring to γ (which we shall see later is related to two ways of measuring specific heat), and he says we know we cannot get the right answer.

Ten years later, in a lecture, he said, "I have now put before you what I consider to be the greatest difficulty yet encountered by the molecular theory." These words represent the first discovery that the laws of classical physics were wrong. This was the first indication that there was something fundamentally impossible, because a rigorously proved theorem did not agree with experiment. About 1890, Jeans was to talk about this puzzle again. One often hears it said that physicists at the latter part of the nineteenth century thought they knew all the significant physical laws and that all they had to do was to calculate more decimal places. Someone may have said that once, and others copied it. But a thorough reading of the literature of the time shows they were all worrying about something. Jeans said about this puzzle that it is a very mysterious phenomenon, and it seems as though as the temperature falls, certain kinds of motions "freeze out."

If we could assume that the vibrational motion, say, did not exist at low temperature and did exist at high temperature, then we could imagine that a gas might exist at a temperature sufficiently low that vibrational motion does not occur, so $\gamma = 1.40$, or a higher temperature at which it begins to come in, so γ falls. The same might be argued for the rotation. If we can eliminate the rotation, say it "freezes out" at sufficiently low temperature, then we can understand the fact that the γ of hydrogen approaches 1.66 as we go down in temperature. How can we understand such a phenomenon? Of course that these motions "freeze out" cannot be understood by classical mechanics. It was only understood when quantum mechanics was discovered.

Without proof, we may state the results for statistical mechanics of the quantum-mechanical theory. We recall that according to quantum mechanics, a system which is bound by a potential, for the vibrations, for example, will have a discrete set of energy levels, i.e., states of different energy. Now the question is: how is statistical mechanics to be modified according to quantum-mechanical theory? It turns out, interestingly enough, that although most problems are more difficult in quantum mechanics than in classical mechanics, problems in statistical mechanics are much easier in quantum theory! The simple result we have in classical mechanics, that $n = n_0 e^{-\text{energy}/kT}$, becomes the following very important theorem: If the energies of the set of molecular states are called, say, E_0, E_1, E_2, $\ldots$, E_i, $\ldots$, then in thermal equilibrium the probability of finding a molecule in the particular state of having energy E_i is proportional to $e^{-E_i/kT}$. That gives the probability of being in various states. In other words, the relative chance, the probability, of being in state E_1 relative to the chance of being in state E_0, is

$$\frac{P_1}{P_0} = \frac{e^{-E_1/kT}}{e^{-E_0/kT}}, \tag{40.10}$$

which, of course, is the same as

$$n_1 = n_0 e^{-(E_1-E_0)/kT}, \tag{40.11}$$

since $P_1 = n_1/N$ and $P_0 = n_0/N$. So it is less likely to be in a higher energy state than in a lower one. The ratio of the number of atoms in the upper state to the number in the lower state is e raised to the power (minus the energy difference, over kT)—a very simple proposition.

Now it turns out that for a harmonic oscillator the energy levels are evenly spaced. Calling the lowest energy $E_0 = 0$ (it actually is not zero, it is a little different, but it does not matter if we shift all energies by a constant), the first one is then $E_1 = \hbar\omega$, and the second one is $2\hbar\omega$, and the third one is $3\hbar\omega$, and so on.

Now let us see what happens. We suppose we are studying the vibrations of a diatomic molecule, which we approximate as a harmonic oscillator. Let us ask

what is the relative chance of finding a molecule in state E_1 instead of in state E_0. The answer is that the chance of finding it in state E_1, relative to that of finding it in state E_0, goes down as $e^{-\hbar\omega/kT}$. Now suppose that kT is much less than $\hbar\omega$, and we have a low-temperature circumstance. Then the probability of its being in state E_1 is extremely small. Practically all the atoms are in state E_0. If we change the temperature but still keep it very small, then the chance of its being in state $E_1 = \hbar\omega$ remains infinitesimal—the energy of the oscillator remains nearly zero; it does not change with temperature so long as the temperature is much less than $\hbar\omega$. All oscillators are in the bottom state, and their motion is effectively "frozen"; *there is no contribution of it to the specific heat.* We can judge, then, from Table 40–1, that at 100°C, which is 373 degrees absolute, kT is much less than the vibrational energy in the oxygen or hydrogen molecules, but not so in the iodine molecule. The reason for the difference is that an iodine atom is very heavy, compared with hydrogen, and although the forces may be comparable in iodine and hydrogen, the iodine molecule is so heavy that the natural frequency of vibration is very low compared with the natural frequency of hydrogen. With $\hbar\omega$ higher than kT at room temperature for hydrogen, but lower for iodine, only the latter, iodine, exhibits the classical vibrational energy. As we increase the temperature of a gas, starting from a very low value of T, with the molecules almost all in their lowest state, they gradually begin to have an appreciable probability to be in the second state, and then in the next state, and so on. When the probability is appreciable for many states, the behavior of the gas approaches that given by classical physics, because the quantized states become nearly indistinguishable from a continuum of energies, and the system can have almost any energy. Thus, as the temperature rises, we should again get the results of classical physics, as indeed seems to be the case in Fig. 40–6. It is possible to show in the same way that the rotational states of atoms are also quantized, but the states are so much closer together that in ordinary circumstances kT is bigger than the spacing. Then many levels are excited, and the rotational kinetic energy in the system participates in the classical way. The one example where this is not quite true at room temperature is for hydrogen.

This is the first time that we have really deduced, by comparison with experiment, that there was something wrong with classical physics, and we have looked for a resolution of the difficulty in quantum mechanics in much the same way as it was done originally. It took 30 or 40 years before the next difficulty was discovered, and that had to do again with statistical mechanics, but this time the mechanics of a photon gas. That problem was solved by Planck, in the early years of this century.

41

The Brownian Movement

41–1 Equipartition of energy

41–2 Thermal equilibrium of radiation

41–3 Equipartition and the quantum oscillator

41–4 The random walk

41–1 Equipartition of energy

The Brownian movement was discovered in 1827 by Robert Brown, a botanist. While he was studying microscopic life, he noticed little particles of plant pollens jiggling around in the liquid he was looking at in the microscope, and he was wise enough to realize that these were not living, but were just little pieces of dirt moving around in the water. In fact he helped to demonstrate that this had nothing to do with life by getting from the ground an old piece of quartz in which there was some water trapped. It must have been trapped for millions and millions of years, but inside he could see the same motion. What one sees is that very tiny particles are jiggling all the time.

This was later proved to be one of the effects of *molecular motion*, and we can understand it qualitatively by thinking of a great push ball on a playing field, seen from a great distance, with a lot of people underneath, all pushing the ball in various directions. We cannot see the people because we imagine that we are too far away, but we can see the ball, and we notice that it moves around rather irregularly. We also know, from the theorems that we have discussed in previous chapters, that the mean kinetic energy of a small particle suspended in a liquid or a gas will be $\frac{3}{2}kT$ even though it is very heavy compared with a molecule. If it is very heavy, that means that the speeds are relatively slow, but it turns out, actually, that the speed is not really so slow. In fact, we cannot see the speed of such a particle very easily because although the mean kinetic energy is $\frac{3}{2}kT$, which represents a speed of a millimeter or so per second for an object a micron or two in diameter, this is very hard to see even in a microscope, because the particle continuously reverses its direction and does not get anywhere. How far it does get we will discuss at the end of the present chapter. This problem was first solved by Einstein at the beginning of the present century.

Incidentally, when we say that the mean kinetic energy of this particle is $\frac{3}{2}kT$, we claim to have derived this result from the kinetic theory, that is, from Newton's laws. We shall find that we can derive all kinds of things—marvelous things—from the kinetic theory, and it is most interesting that we can apparently get so much from so little. Of course we do not mean that Newton's laws are "little" —they are enough to do it, really—what we mean is that *we* did not do very much. How do we get so much out? The answer is that we have been perpetually making a certain important assumption, which is that if a given system is in thermal equilibrium at some temperature, it will also be in thermal equilibrium with *anything else* at the same temperature. For instance, if we wanted to see how a particle would move if it was really colliding with water, we could imagine that there was a gas present, composed of another kind of particle, little fine pellets that (we suppose) do not interact with water, but only hit the particle with "hard" collisions. Suppose the particle has a prong sticking out of it; all our pellets have to do is hit the prong. We know all about this imaginary gas of pellets at temperature T—it is an ideal gas. Water is complicated, but an ideal gas is simple. Now, *our particle has to be in equilibrium with the gas of pellets.* Therefore, the mean motion of the particle must be what we get for gaseous collisions, because if it were not moving at the right speed relative to the water but, say, was moving faster, that would mean that the pellets would pick up energy from it and get hotter than the water. But we had started them at the same temperature, and we assume that if a thing is once in equilibrium, it stays in equilibrium—parts of it do not get hotter and other parts colder, spontaneously.

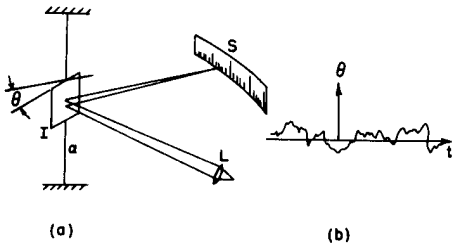


Fig. 41–1. (a) A sensitive light-beam galvanometer. Light from a source L is reflected from a small mirror onto a scale. (b) A schematic record of the reading of the scale as a function of the time.

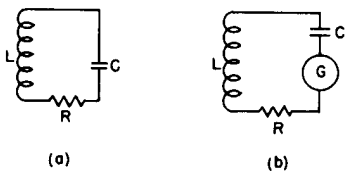


Fig. 41–2. A high-Q resonant circuit. (a) Actual circuit, at temperature T. (b) Artificial circuit, with an ideal (noiseless) resistance and a "noise generator" G.

This proposition is true and can be proved from the laws of mechanics, but the proof is very complicated and can be established only by using advanced mechanics. It is much easier to prove in quantum mechanics than it is in classical mechanics. It was proved first by Boltzmann, but for now we simply take it to be true, and then we can argue that our particle has to have $\frac{3}{2}kT$ of energy if it is hit with artificial pellets, so it also must have $\frac{3}{2}kT$ when it is being hit with water at the same temperature and we take away the pellets; so it is $\frac{3}{2}kT$. It is a strange line of argument, but perfectly valid.

In addition to the motion of colloidal particles for which the Brownian movement was first discovered, there are a number of other phenomena, both in the laboratory and in other situations, where one can see Brownian movement. If we are trying to build the most delicate possible equipment, say a very small mirror on a thin quartz fiber for a very sensitive ballistic galvanometer (Fig. 41–1), the mirror does not stay put, but jiggles all the time—*all the time*—so that when we shine a light on it and look at the position of the spot, we do not have a perfect instrument because the mirror is always jiggling. Why? Because the average kinetic energy of rotation of this mirror has to be, on the average, $\frac{1}{2}kT$.

What is the mean-square angle over which the mirror will wobble? Suppose we find the natural vibration period of the mirror by tapping on one side and seeing how long it takes to oscillate back and forth, and we also know the moment of inertia, I. We know the formula for the kinetic energy of rotation—it is given by Eq. (19.8): $T = \frac{1}{2}I\omega^2$. That is the kinetic energy, and the potential energy that goes with it will be proportional to the square of the angle—it is $V = \frac{1}{2}\alpha\theta^2$. But, if we know the period t_0 and calculate from that the natural frequency $\omega_0 = 2\pi/t_0$, then the potential energy is $V = \frac{1}{2}I\omega_0^2\theta^2$. Now we know that the average kinetic energy is $\frac{1}{2}kT$, but since it is a harmonic oscillator the average potential energy is also $\frac{1}{2}kT$. Thus

$$\tfrac{1}{2}I\omega_0^2\langle\theta^2\rangle = \tfrac{1}{2}kT,$$

or

$$\langle\theta^2\rangle = kT/I\omega_0^2. \tag{41.1}$$

In this way we can calculate the oscillations of a galvanometer mirror, and thereby find what the limitations of our instrument will be. If we want to have smaller oscillations, we have to cool the mirror. An interesting question is, *where* to cool it. This depends upon where it is getting its "kicks" from. If it is through the fiber, we cool it at the top—if the mirror is surrounded by a gas and is getting hit mostly by collisions in the gas, it is better to cool the gas. As a matter of fact, if we know where the *damping* of the oscillations comes from, it turns out that that is always the *source* of the fluctuations also, a point which we will come back to.

The same thing works, amazingly enough, in *electrical circuits*. Suppose that we are building a very sensitive, accurate amplifier for a definite frequency and have a resonant circuit (Fig. 41–2) in the input so as to make it very sensitive to this certain frequency, like a radio receiver, but a really good one. Suppose we wish to go down to the very lowest limit of things, so we take the voltage, say off the inductance, and send it into the rest of the amplifier. Of course, in any circuit like this, there is a certain amount of loss. It is not a perfect resonant circuit, but it is a very good one and there is a little resistance, say (we put the resistor in so we can see it, but it is supposed to be small). Now we would like to find out: How much does the voltage across the inductance fluctuate? *Answer:* We know that $\frac{1}{2}LI^2$ is the "kinetic energy"—the energy associated with a coil in a resonant circuit (Chapter 25). Therefore the mean value of $\frac{1}{2}LI^2$ is equal to $\frac{1}{2}kT$—that tells us what the rms current is and we can find out what the rms voltage is from the rms current. For if we want the voltage across the inductance the formula is $\hat{V}_L = i\omega L\hat{I}$, and the mean absolute square voltage on the inductance is $\langle V_L^2\rangle = L^2\omega_0^2\langle I^2\rangle$, and putting in $\frac{1}{2}L\langle I^2\rangle = \frac{1}{2}kT$, we obtain

$$\langle V_L^2\rangle = L\omega_0^2kT. \tag{41.2}$$

So now we can design circuits and tell when we are going to get what is called *Johnson noise*, the noise associated with thermal fluctuations!

Where do the fluctuations come from this time? They come again from the *resistor*—they come from the fact that the electrons in the resistor are jiggling around because they are in thermal equilibrium with the matter in the resistor, and they make fluctuations in the density of electrons. They thus make tiny electric fields which drive the resonant circuit.

Electrical engineers represent the answer in another way. Physically, the resistor is effectively the source of noise. However, we may replace the real circuit having an honest, true physical resistor which is making noise, by an artificial circuit which contains a little generator that is going to *represent* the noise, and now the resistor is otherwise ideal—no noise comes from it. All the noise is in the artificial generator. And so if we knew the characteristics of the noise generated by a resistor, if we had the formula for that, then we could calculate what the circuit is going to do in response to that noise. So, we need a formula for the noise fluctuations. Now the noise that is generated by the resistor is at all frequencies, since the resistor by itself is not resonant. Of course the resonant circuit only "listens" to the part that is near the right frequency, but the resistor has many different frequencies in it. We may describe how strong the generator is, as follows: The mean power that the resistor would absorb if it were connected directly across the noise generator would be $\langle E^2\rangle/R$, if E were the voltage from the generator. But we would like to know in more detail how much power there is at every frequency. There is very little power in any one frequency; it is a distribution. Let $P(\omega)\,d\omega$ be the power that the generator would deliver in the frequency range $d\omega$ into the very same resistor. Then we can prove (we shall prove it for another case, but the mathematics is exactly the same) that the power comes out

$$P(\omega)\,d\omega = (2/\pi)kT\,d\omega, \tag{41.3}$$

and is *independent of the resistance* when put this way.

41–2 Thermal equilibrium of radiation

Now we go on to consider a still more advanced and interesting proposition that is as follows. Suppose we have a charged oscillator like those we were talking about when we were discussing light, let us say an electron oscillating up and down in an atom. If it oscillates up and down, it radiates light. Now suppose that this oscillator is in a very thin gas of other atoms, and that from time to time the atoms collide with it. Then in equilibrium, after a long time, this oscillator will pick up energy such that its kinetic energy of oscillation is $\frac{1}{2}kT$, and since it is a harmonic oscillator, its entire energy of motion will become kT. That is, of course, a wrong description so far, because the oscillator carries *electric charge*, and if it has an energy kT it is shaking up and down and *radiating light*. Therefore it is impossible to have equilibrium of real matter alone without the charges in it emitting light, and as light is emitted, energy flows away, the oscillator loses its kT as time goes on, and thus the whole gas which is colliding with the oscillator gradually cools off. And that is, of course, the way a hot stove cools on a cold night by radiating the light into the sky, because the atoms are jiggling their charge and they continually radiate, and slowly, because of this radiation, the jiggling motion slows down.

On the other hand, if we enclose the whole thing in a box so that the light does not go away to infinity, then we *can* eventually get thermal equilibrium. We may either put the gas in a box where we can say that there are other radiators in the box walls sending light back or, to take a nicer example, we may suppose the box has mirror walls. It is easier to think about that case. Thus we assume that all the radiation that goes out from the oscillator keeps running around in the box. Then, of course, it is true that the oscillator starts to radiate, but pretty soon it can maintain its kT of kinetic energy in spite of the fact that it is radiating, because it is being illuminated, we may say, by its own light reflected from the walls of the box. That is, after a while there is a great deal of light rushing around in the box, and although the oscillator is radiating some, the light comes back and returns some of the energy that was radiated.

We shall now determine how much light there must be in such a box at temperature T in order that the shining of the light on this oscillator will generate just enough energy to account for the light it radiated.

Let the gas atoms be very few and far between, so that we have an ideal oscillator with no resistance except radiation resistance. Then we consider that at thermal equilibrium the oscillator is doing two things at the same time. First, it has a mean energy kT, and we calculate how much radiation it emits. Second, this radiation should be exactly the amount that would result because of the fact that the light shining on the oscillator is scattered. Since there is nowhere else the energy can go, this effective radiation is really just scattered light from the light that is in there.

Thus we first calculate the energy that is radiated by the oscillator per second, if the oscillator has a certain energy. (We borrow from Chapter 32 on radiation resistance a number of equations without going back over their derivation.) The energy radiated per radian divided by the energy of the oscillator is called $1/Q$ (Eq. 32.8): $1/Q = (dW/dt)/\omega_0 W$. Using the quantity γ, the damping constant, this can also be written as $1/Q = \gamma/\omega_0$, where ω_0 is the natural frequency of the oscillator—if gamma is very small, Q is very large. The energy radiated per second is then

$$\frac{dW}{dt} = \frac{\omega_0 W}{Q} = \frac{\omega_0 W\gamma}{\omega_0} = \gamma W. \tag{41.4}$$

The energy radiated per second is thus simply gamma times the energy of the oscillator. Now the oscillator should have an average energy kT, so we see that gamma kT is the average amount of energy radiated per second:

$$\langle dW/dt\rangle = \gamma kT. \tag{41.5}$$

Now we only have to know what gamma is. Gamma is easily found from Eq. (32.12). It is

$$\gamma = \frac{\omega_0}{Q} = \frac{2}{3}\frac{r_0\omega_0^2}{c}, \tag{41.6}$$

where $r_0 = e^2/mc^2$ is the classical electron radius, and we have set $\lambda = 2\pi c/\omega_0$.

Our final result for the average rate of radiation of light near the frequency ω_0 is therefore

$$\frac{dW}{dt} = \frac{2}{3}\frac{r_0\omega_0^2 kT}{c}. \tag{41.7}$$

Next we ask how much light must be shining on the oscillator. It must be enough that the energy absorbed from the light (and thereupon scattered) is just exactly this much. In other words, the emitted light is accounted for as *scattered* light from the light that is shining on the oscillator in the cavity. So we must now calculate how much light is scattered from the oscillator if there is a certain amount —unknown—of radiation incident on it. Let $I(\omega)\,d\omega$ be the amount of light energy there is at the frequency ω, within a certain range $d\omega$ (because there is no light at *exactly* a certain frequency; it is spread all over the spectrum). So $I(\omega)$ is a certain *spectral distribution* which we are now going to find—it is the color of a furnace at temperature T that we see when we open the door and look in the hole. Now how much light is absorbed? We worked out the amount of radiation absorbed from a given incident light beam, and we calculated it in terms of a *cross section.* It is just as though we said that all of the light that falls on a certain cross section is absorbed. So the total amount that is re-radiated (scattered) is the incident intensity $I(\omega)\,d\omega$ multiplied by the cross section σ.

The formula for the cross section which we derived (Eq. 31.19) did not have the damping included. It is not hard to go through the derivation again and put in the resistance term which we neglected. If we do that, and calculate the cross section the same way, we get

$$\sigma_s = \frac{8\pi r_0^2}{3}\left(\frac{\omega^4}{(\omega^2 - \omega_0^2)^2 + \gamma^2\omega^2}\right). \tag{41.8}$$

Now, as a function of frequency, σ_s is of significant size only for ω very near to the natural frequency ω_0. (Remember that the Q for a radiating oscillator is about 10^8.) The oscillator scatters very strongly when ω is equal to ω_0, and very weakly for other values of ω. Therefore we can replace ω by ω_0 and $\omega^2 - \omega_0^2$ by $2\omega_0(\omega - \omega_0)$, and we get

$$\sigma_s = \frac{2\pi r_0^2\omega_0^2}{3[(\omega - \omega_0)^2 + \gamma^2/4]}. \tag{41.9}$$

Now the whole curve is localized near $\omega = \omega_0$. (We do not really have to make any approximations, but it is much easier to do the integrals if we simplify the equation a bit.) Now we multiply the intensity in a given frequency range by the cross section of scattering, to get the amount of energy scattered in the range $d\omega$. The *total* energy scattered is then the integral of this for all ω. Thus

$$\begin{aligned}\frac{dW_s}{dt} &= \int_0^\infty I(\omega)\sigma_s(\omega)\,d\omega \\ &= \int_0^\infty \frac{2\pi r_0^2\omega_0^2 I(\omega)\,d\omega}{3[(\omega - \omega_0)^2 + \gamma^2/4]}.\end{aligned} \tag{41.10}$$

Now we set $dW_s/dt = 3\gamma kT$. Why *three*? Because when we made our analysis of the cross section in Chapter 32, we assumed that the polarization was such that the light could drive the oscillator. If we had used an oscillator which could move only in one direction, and the light, say, was polarized in the wrong way, it would not give any scattering. So we must either average the cross section of an oscillator which can go only in one direction, over all directions of incidence and polarization of the light or, more easily, we can imagine an oscillator which will follow the field no matter which way the field is pointing. Such an oscillator, which can oscillate equally in three directions, would have $3kT$ average energy because there are 3 degrees of freedom in that oscillator. So we should use $3\gamma kT$ because of the 3 degrees of freedom.

Now we have to do the integral. Let us suppose that the unknown spectral distribution $I(\omega)$ of the light is a smooth curve and does not vary very much across the very narrow frequency region where σ_s is peaked (Fig. 41–3). Then the only significant contribution comes when ω is very close to ω_0, within an amount gamma, which is very small. So therefore, although $I(\omega)$ may be an unknown and complicated function, the only place where it is important is near $\omega = \omega_0$, and there we may replace the smooth curve by a flat one—a "constant"—at the same height. In other words, we simply take $I(\omega)$ outside the integral sign and call it $I(\omega_0)$. We may also take the rest of the constants out in front of the integral, and what we have left is

$$\tfrac{2}{3}\pi r_0^2\omega_0^2 I(\omega_0)\int_0^\infty \frac{d\omega}{(\omega - \omega_0)^2 + \gamma^2/4} = 3\gamma kT. \tag{41.11}$$

Now, the integral should go from 0 to ∞, but 0 is so far from ω_0 that the curve is all finished by that time, so we go instead to minus ∞—it makes no difference and it is much easier to do the integral. The integral is an inverse tangent function of the form $\int dx/(x^2 + a^2)$. If we look it up in a book we see that it is equal to π/a. So what it comes to for our case is $2\pi/\gamma$. Therefore we get, with some rearranging,

$$I(\omega_0) = \frac{9\gamma^2 kT}{4\pi^2 r_0^2\omega_0^2}. \tag{41.12}$$

Then we substitute the formula (41.6) for gamma (do not worry about writing ω_0; since it is true of any ω_0, we may just call it ω) and the formula for $I(\omega)$ then comes out

$$I(\omega) = \frac{\omega^2 kT}{\pi^2 c^2}. \tag{41.13}$$

And that gives us the distribution of light in a hot furnace. It is called the *black-*

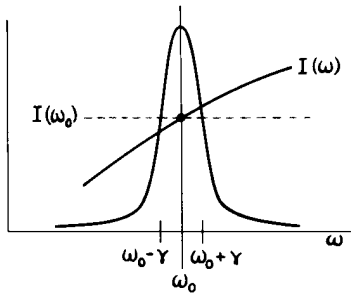


Fig. 41–3. The factors in the integrand (41.10). The peak is the resonance curve $1/(\omega - \omega_0)^2 + \gamma^2/4$. To a good approximation the factor $I(\omega)$ can be replaced by $I(\omega_0)$.

body radiation. Black, because the hole in the furnace that we look at is black when the temperature is zero.

Inside a closed box at temperature T, (41.13) is the distribution of energy of the radiation, according to classical theory. First, let us notice a remarkable feature of that expression. The charge of the oscillator, the mass of the oscillator, all properties specific to the oscillator, *cancel out*, because once we have reached equilibrium with one oscillator, we must be at equilibrium with any other oscillator of a different mass, or we will be in trouble. So this is an important kind of check on the proposition that equilibrium does not depend on what we are in equilibrium with, but *only on the temperature.* Now let us draw a picture of the $I(\omega)$ curve (Fig. 41–4). It tells us how much light we have at different frequencies.

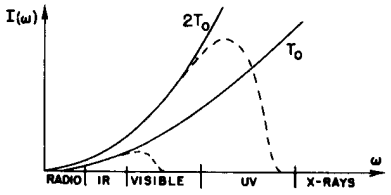


Fig. 41–4. The blackbody intensity distribution at two temperatures, according to classical physics (solid curves). The dashed curves show the actual distribution.

The amount of intensity that there is in our box, per unit frequency range, goes, as we see, as the square of the frequency, which means that if we have a box at any temperature at all, and if we look at the x-rays that are coming out, there will be a lot of them!

Of course we know this is false. When we open the furnace and take a look at it, we do not burn our eyes out from x-rays at all. It is completely false. Furthermore, the *total energy* in the box, the total of all this intensity summed over all frequencies, would be the area under this infinite curve. Therefore, something is fundamentally, powerfully, and absolutely wrong.

Thus was the classical theory *absolutely incapable* of correctly describing the distribution of light from a blackbody, just as it was incapable of correctly describing the specific heats of gases. Physicists went back and forth over this derivation from many different points of view, and there is no escape. This *is* the prediction of classical physics. Equation (41.13) is called *Rayleigh's law*, and it is the prediction of classical physics, and is obviously absurd.

41–3 Equipartition and the quantum oscillator

The difficulty above was another part of the continual problem of classical physics, which started with the difficulty of the specific heat of gases, and now has been focused on the distribution of light in a blackbody. Now, of course, at the time that theoreticians studied this thing, there were also many *measurements* of the actual curve. And it turned out that the correct curve looked like the dashed curves in Fig. 41–4. That is, the x-rays were not there. If we lower the temperature, the whole curve goes down in proportion to T, according to the classical theory, but the observed curve also cuts off sooner at a lower temperature. Thus the low-frequency end of the curve is right, but the high-frequency end is wrong. Why? When Sir James Jeans was worrying about the specific heats of gases, he noted that motions which have high frequency are "frozen out" as the temperature goes too low. That is, if the temperature is too low, if the frequency is too high, the oscillators *do not have* kT of energy on the average. Now recall how our derivation of (41.13) worked: It all depends on the energy of an oscillator at thermal equilibrium. What the kT of (41.5) was, and what the same kT in (41.13) is, is the mean energy of a harmonic oscillator of frequency ω at temperature T. Classically, this is kT, but experimentally, no!—not when the temperature is too low or the oscillator frequency is too high. And so the reason that the curve falls off is the same reason that the specific heats of gases fail. It is easier to study the blackbody curve than it is the specific heats of gases, which are so complicated, therefore our attention is focused on determining the true blackbody curve, because this curve is a curve which correctly tells us, at every frequency, what the average energy of harmonic oscillators actually is as a function of temperature.

Planck studied this curve. He first determined the answer empirically, by fitting the observed curve with a nice function that fitted very well. Thus he had an empirical formula for the average energy of a harmonic oscillator as a function of frequency. In other words, he had the *right* formula instead of kT, and then by fiddling around he found a simple derivation for it which involved a very peculiar assumption. That assumption was that *the harmonic oscillator can take up energies only* $\hbar\omega$ *at a time.* The idea that they can have *any energy at all* is false. Of course, that was the beginning of the end of classical mechanics.

The very first correctly determined quantum-mechanical formula will now be derived. Suppose that the permitted energy levels of a harmonic oscillator were equally spaced at $\hbar\omega_0$ apart, so that the oscillator could take on only these different energies (Fig. 41–5). Planck made a somewhat more complicated argument than the one that is being given here, because that was the very beginning of quantum mechanics and he had to prove some things. But we are going to take it as a fact (which he demonstrated in this case) that the probability of occupying a level of energy E is $P(E) = \alpha e^{-E/kT}$. If we go along with that, we will obtain the right result.

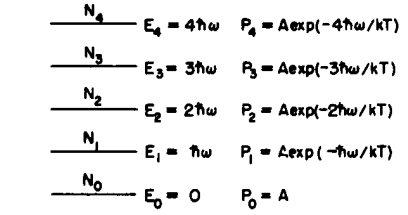


Fig. 41–5. The energy levels of a harmonic oscillator are equally spaced: $E_n = n\hbar\omega$.

Suppose now that we have a lot of oscillators, and each is a vibrator of frequency ω_0. Some of these vibrators will be in the bottom quantum state, some will be in the next one, and so forth. What we would like to know is the average energy of all these oscillators. To find out, let us calculate the total energy of all the oscillators and divide by the number of oscillators. That will be the average energy per oscillator in thermal equilibrium, and will also be the energy that is in equilibrium with the blackbody radiation and that should go in Eq. (41.13) in place of kT. Thus we let N_0 be the number of oscillators that are in the ground state (the lowest energy state); N_1 the number of oscillators in the state E_1; N_2 the number that are in state E_2; and so on. According to the hypothesis (which we have not proved) that in quantum mechanics the law that replaced the probability $e^{-\text{P.E.}/kT}$ or $e^{-\text{K.E.}/kT}$ in classical mechanics is that the probability goes down as $e^{-\Delta E/kT}$, where ΔE is the excess energy, we shall assume that the number N_1 that are in the first state will be the number N_0 that are in the ground state, times $e^{-\hbar\omega/kT}$. Similarly, N_2, the number of oscillators in the second state, is $N_2 = N_0e^{-2\hbar\omega/kT}$. To simplify the algebra, let us call $e^{-\hbar\omega/kT} = x$. Then we simply have $N_1 = N_0x$, $N_2 = N_0x^2, \ldots, N_n = N_0x^n$.

The total energy of all the oscillators must first be worked out. If an oscillator is in the ground state, there is no energy. If it is in the first state, the energy is $\hbar\omega_0$, and there are N_1 of them. So $N_1\hbar\omega$, or $\hbar\omega N_0x$ is how much energy we get from those. Those that are in the second state have $2\hbar\omega_0$, and there are N_2 of them, so $N_2 \cdot 2\hbar\omega = 2\hbar\omega N_0x^2$ is how much energy we get, and so on. Then we add it all together to get $E_{\text{tot}} = N_0\hbar\omega(0 + x + 2x^2 + 3x^3 + \ldots)$.

And now, how many oscillators are there? Of course, N_0 is the number that are in the ground state, N_1 in the first state, and so on, and we add them together: $N_{\text{tot}} = N_0(1 + x + x^2 + x^3 + \ldots)$. Thus the average energy is

$$\langle E\rangle = \frac{E_{\text{tot}}}{N_{\text{tot}}} = \frac{N_0\hbar\omega(0 + x + 2x^2 + 3x^3 + \cdots)}{N_0(1 + x + x^2 + \cdots)}. \tag{41.14}$$

Now the two sums which appear here we shall leave for the reader to play with and have some fun with. When we are all finished summing and substituting for x in the sum, we should get—if we make no mistakes in the sum—

$$\langle E\rangle = \frac{\hbar\omega}{e^{\hbar\omega/kT} - 1}. \tag{41.15}$$

This, then, was the first quantum-mechanical formula ever known, or ever discussed, and it was the beautiful culmination of decades of puzzlement. Maxwell knew that there was something wrong, and the problem was, what was *right*? Here is the quantitative answer of what is right instead of kT. This expression should, of course, approach kT as $\omega \to 0$ or as $T \to \infty$. See if you can prove that it does—learn how to do the mathematics.

This is the famous cutoff factor that Jeans was looking for, and if we use it instead of kT in (41.13), we obtain for the distribution of light in a black box

$$I(\omega)\,d\omega = \frac{\hbar\omega^3\,d\omega}{\pi^2c^2(e^{\hbar\omega/kT} - 1)}. \tag{41.16}$$

We see that for a large ω, even though we have ω^3 in the numerator, there is an e raised to a tremendous power in the denominator, so the curve comes down again and does not "blow up"—we do not get ultraviolet light and x-rays where we do not expect them!

One might complain that in our derivation of (41.16) we used the quantum theory for the energy levels of the harmonic oscillator, but the classical theory in determining the cross section σ_s. But the quantum theory of light interacting with a harmonic oscillator gives exactly the same result as that given by the classical theory. That, in fact, is why we were justified in spending so much time on our analysis of the index of refraction and the scattering of light, using a model of atoms like little oscillators—the quantum formulas are substantially the same.

Now let us return to the Johnson noise in a resistor. We have already remarked that the theory of this noise power is really the same theory as that of the classical blackbody distribution. In fact, rather amusingly, we have already said that if the resistance in a circuit were not a real resistance, but were an antenna (an antenna acts like a resistance because it radiates energy), a radiation resistance, it would be easy for us to calculate what the power would be. It would be just the power that runs into the antenna from the light that is all around, and we would get the same distribution, changed by only one or two factors. We can suppose that the resistor is a generator with an unknown power spectrum $P(\omega)$. The spectrum is determined by the fact that this same generator, connected to a resonant circuit of *any frequency*, as in Fig. 41–2(b), generates in the inductance a voltage of the magnitude given in Eq. (41.2). One is thus led to the same integral as in (41.10), and the same method works to give Eq. (41.3). For low temperatures the kT in (41.3) must of course be replaced by (41.15). The two theories (blackbody radiation and Johnson noise) are also closely related physically, for we may of course connect a resonant circuit to an *antenna*, so the resistance R is a pure *radiation resistance*. Since (41.2) does not depend on the physical origin of the resistance, we know the generator G for a real resistance and for radiation resistance is the same. What is the origin of the generated power $P(\omega)$ if the resistance R is only an ideal antenna in equilibrium with its environment at temperature T? It is the radiation $I(\omega)$ in the space at temperature T which impinges on the antenna and, as "received signals," makes an effective generator. Therefore one can deduce a direct relation of $P(\omega)$ and $I(\omega)$, leading then from (41.13) to (41.3).

All the things we have been talking about—the so-called Johnson noise and Planck's distribution, and the correct theory of the Brownian movement which we are about to describe—are developments of the first decade or so of this century. Now with those points and that history in mind, we return to the Brownian movement.

41–4 The random walk

Let us consider how the position of a jiggling particle should change with time, for very long times compared with the time between "kicks." Consider a little Brownian movement particle which is jiggling about because it is bombarded on all sides by irregularly jiggling water molecules. Query: After a given length of time, how far away is it likely to be from where it began? This problem was solved by Einstein and Smoluchowski. If we imagine that we divide the time into little intervals, let us say a hundredth of a second or so, then after the first hundredth of a second it moves here, and in the next hundredth it moves some more, in the next hundredth of a second it moves somewhere else, and so on. In terms of the rate of bombardment, a hundredth of a second is a very long time. The reader may easily verify that the number of collisions a single molecule of water receives in a second is about 10^{14}, so in a hundredth of a second it has 10^{12} collisions, which is a lot! Therefore, after a hundredth of a second it is not going to remember what happened before. In other words, the collisions are all *random*, so that one "step" is not related to the previous "step." It is like the famous drunken sailor problem: the sailor comes out of the bar and takes a sequence of steps, but each step is chosen at an arbitrary angle, at random (Fig. 41–6). The question is: After a long time, where is the sailor? Of course we do not know! It is impossible to say. What do we mean—he is just somewhere more or less random. Well then, on the average, where is he? *On the average, how far away from the bar has he gone?* We have already answered this question, because once we were discussing the

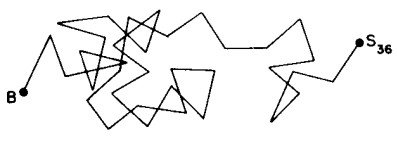


Fig. 41–6. A random walk of 36 steps of length l. How far is S_{36} from B? *Ans:* about $6l$ on the average.

superposition of light from a whole lot of different sources at different phases, and that meant adding a lot of arrows at different angles (Chapter 32). There we discovered that the mean square of the distance from one end to the other of the chain of random steps, which was the intensity of the light, is the sum of the intensities of the separate pieces. And so, by the same kind of mathematics, we can prove immediately that if $\mathbf{R}_N$ is the vector distance from the origin after N steps, the mean square of the distance from the origin is proportional to the number N of steps. That is, $\langle R_N^2 \rangle = NL^2$, where L is the length of each step. Since the number of steps is proportional to the time in our present problem, *the mean square distance is proportional to the time:*

$$\langle R^2 \rangle = \alpha t. \tag{41.17}$$

This does not mean that the *mean distance* is proportional to the time. If the mean distance were proportional to the time it would mean that the drifting is at a nice uniform velocity. The sailor *is* making some relatively sensible headway, but only such that his mean *square* distance is proportional to time. That is the characteristic of a random walk.

We may show very easily that in each successive step the square of the distance increases, on the average, by L^2. For if we write $\mathbf{R}_N = \mathbf{R}_{N-1} + \mathbf{L}$, we find that $\mathbf{R}_N^2$ is

$$\mathbf{R}_N \cdot \mathbf{R}_N = R_N^2 = R_{N-1}^2 + 2\mathbf{R}_{N-1} \cdot \mathbf{L} + L^2,$$

and averaging over many trials, we have $\langle R_N^2 \rangle = \langle R_{N-1}^2 \rangle + L^2$, since $\langle \mathbf{R}_{N-1} \cdot \mathbf{L} \rangle = 0$. Thus, by induction,

$$R_N^2 = NL^2. \tag{41.18}$$

Now we would like to calculate the coefficient α in Eq. (41.17), and to do so we must add a feature. We are going to suppose that if we were to put a force on this particle (having nothing to do with the Brownian movement—we are taking a side issue for the moment), then it would react in the following way against the force. First, there would be inertia. Let m be the coefficient of inertia, the effective mass of the object (not necessarily the same as the real mass of the real particle, because the water has to move around the particle if we pull on it). Thus if we talk about motion in one direction, there is a term like $m\,(d^2x/dt^2)$ on one side. And next, we want also to assume that if we kept a steady pull on the object, there would be a drag on it from the fluid, proportional to its velocity. Besides the inertia of the fluid, there is a resistance to flow due to the viscosity and the complexity of the fluid. It is absolutely essential that there *be* some irreversible losses, something like resistance, in order that there be fluctuations. There is no way to produce the kT unless there are also losses. The source of the fluctuations is very closely related to these losses. What the mechanism of this drag is, we will discuss soon—we shall talk about forces that are proportional to the velocity and where they come from. But let us suppose for now that there is such a resistance. Then the formula for the motion under an external force, when we are pulling on it in a normal manner, is

$$m\frac{d^2x}{dt^2} + \mu\frac{dx}{dt} = F_{\text{ext}}. \tag{41.19}$$

The quantity μ can be determined directly from experiment. For example, we can watch the drop fall under gravity. Then we know that the force is mg, and μ is mg divided by the speed of fall the drop ultimately acquires. Or we could put the drop in a centrifuge and see how fast it sediments. Or if it is charged, we can put an electric field on it. So μ is a measurable thing, not an artificial thing, and it is known for many types of colloidal particles, etc.

Now let us use the same formula in the case where the force is not external, but is equal to the irregular forces of the Brownian movement. We shall then try to determine the mean square distance that the object goes. Instead of taking the distances in three dimensions, let us take just one dimension, and find the mean of x^2, just to prepare ourselves. (Obviously the mean of x^2 is the same as the mean of y^2 is the same as the mean of z^2, and therefore the mean square of the distance

is just 3 times what we are going to calculate.) The x-component of the irregular forces is, of course, just as irregular as any other component. What is the rate of change of x^2? It is $d(x^2)/dt = 2x(dx/dt)$, so what we have to find is the average of the position times the velocity. We shall show that this is a constant, and that therefore the mean square radius will increase proportionally to the time, and at what rate. Now if we multiply Eq. (41.19) by x, $mx(d^2x/dt^2) + \mu x(dx/dt) = xF_x$. We want the time average of $x(dx/dt)$, so let us take the average of the whole equation, and study the three terms. Now what about x times the force? If the particle happens to have gone a certain distance x, then, since the irregular force is *completely* irregular and does not know where the particle started from, the next impulse can be in any direction relative to x. If x is positive, there is no reason why the average force should also be in that direction. It is just as likely to be one way as the other. The bombardment forces are not driving it in a definite direction. So the average value of x times F is zero. On the other hand, for the term $mx(d^2x/dt^2)$ we will have to be a little fancy, and write this as

$$mx\frac{d^2x}{dt^2} = m\frac{d[x(dx/dt)]}{dt} - m\left(\frac{dx}{dt}\right)^2.$$

Thus we put in these two terms and take the average of both. So let us see how much x times the velocity should be. Now x times the velocity has a mean that does not change with time, because when it gets to some position it has no remembrance of where it was before, so things are no longer changing with time. So this quantity, on the average, is zero. We have left the quantity mv^2, and that is the only thing we know: $mv^2/2$ has a mean value $\frac{1}{2}kT$. Therefore we find that

$$\left\langle mx\frac{d^2x}{dt^2}\right\rangle + \mu\left\langle x\frac{dx}{dt}\right\rangle = \langle xF_x\rangle$$

implies

$$-\langle mv^2\rangle + \frac{\mu}{2}\frac{d}{dt}\langle x^2\rangle = 0,$$

or

$$\frac{d\langle x^2\rangle}{dt} = 2\frac{kT}{\mu}. \tag{41.20}$$

Therefore the object has a mean square distance $\langle R^2\rangle$, at the end of a certain amount of t, equal to

$$\langle R^2\rangle = 6kT\frac{t}{\mu}. \tag{41.21}$$

And so we can actually determine *how far* the particles go! We first must determine how they react to a steady force, how fast they drift under a known force (to find μ), and then we can determine how far they go in their random motions. This equation was of considerable importance historically, because it was one of the first ways by which the constant k was determined. After all, we can measure μ, the time, how far the particles go, and we can take an average. The reason that the determination of k was important is that in the law $PV = RT$ for a mole, we know that R, which can also be measured, is equal to the number of atoms in a mole times k. A mole was originally defined as so and so many *grams* of oxygen—16 (now carbon is used), so the number of *atoms* in a mole was not known, originally. It is, of course, a very interesting and important problem. How big are atoms? How many are there? So one of the earliest determinations of the number of atoms was by the the determination of how far a dirty little particle would move if we watched it patiently under a microscope for a certain length of time. And thus Boltzmann's constant k and the Avogadro number N_0 were determined because R had already been measured.

42

Applications of Kinetic Theory

42–1 Evaporation

42–2 Thermionic emission

42–3 Thermal ionization

42–4 Chemical kinetics

42–5 Einstein's laws of radiation

42–1 Evaporation

In this chapter we shall discuss some further applications of kinetic theory. In the previous chapter we emphasized one particular aspect of kinetic theory, namely, that the average kinetic energy in any degree of freedom of a molecule or other object is $\frac{1}{2}kT$. The central feature of what we shall now discuss, on the other hand, is the fact that the probability of finding a particle in different places, per unit volume, varies as $e^{-\text{potential energy}/kT}$; we shall make a number of applications of this.

The phenomena which we want to study are relatively complicated: a liquid evaporating, or electrons in a metal coming out of the surface, or a chemical reaction in which there are a large number of atoms involved. In such cases it is no longer possible to make from the kinetic theory any simple and correct statements, because the situation is too complicated. Therefore, this chapter, except where otherwise emphasized, is quite inexact. The idea to be emphasized is only that we can understand, from the kinetic theory, *more or less* how things ought to behave. By using thermodynamic arguments, or some empirical measurements of certain critical quantities, we can get a more accurate representation of the phenomena.

However, it is very useful to know even only more or less why something behaves as it does, so that when the situation is a new one, or one that we have not yet started to analyze, we can say, more or less, what ought to happen. So this discussion is highly inaccurate but essentially right—right in idea, but a little bit simplified, let us say, in the specific details.

The first example that we shall consider is the evaporation of a liquid. Suppose we have a box with a large volume, partially filled with liquid in equilibrium and with the vapor at a certain temperature. We shall suppose that the molecules of the vapor are relatively far apart, and that inside the liquid, the molecules are packed close together. The problem is to find out how many molecules there are in the vapor phase, compared with the number there are in the liquid. How dense is the vapor at a given temperature, and how does it depend on the temperature?

Let us say that n equals the number of molecules per unit volume in the vapor. That number, of course, varies with the temperature. If we add heat, we get more evaporation. Now let another quantity, $1/V_a$, equal the number of atoms per unit volume in the liquid: We suppose that each molecule in the liquid occupies a certain volume, so that if there are more molecules of liquid, then all together they occupy a bigger volume. Thus if V_a is the volume occupied by one molecule, the number of molecules in a unit volume is a unit volume divided by the volume of each molecule. Furthermore, we suppose that there is a force of attraction between the molecules to hold them together in the liquid. Otherwise we cannot understand why it condenses. Thus suppose that there is such a force and that there is an energy of binding of the molecules in the liquid which is lost when they go into the vapor. That is, we are going to suppose that, in order to take a single molecule out of the liquid into the vapor, a certain amount of work W has to be done. There is a certain difference, W, in the energy of a molecule in the liquid from what it would have if it were in the vapor, because we have to pull it away from the other molecules which attract it.

Now we use the general principle that the number of atoms per unit volume in two different regions is $n_2/n_1 = e^{-(E_2-E_1)/kT}$. So the number n per unit volume in the vapor, divided by the number $1/V_a$ per unit volume in the liquid, is equal to

$$nV_a = e^{-W/kT}, \tag{42.1}$$

because that is the general rule. It is like the atmosphere in equilibrium under gravity, where the gas at the bottom is denser than that at the top because of the work mgh needed to lift the gas molecules to the height h. In the liquid, the molecules are denser than in the vapor because we have to pull them out through the energy "hill" W, and the ratio of the densities is $e^{-W/kT}$.

This is what we wanted to deduce—that the vapor density varies as e to the minus some energy or other over kT. The factors in front are not really interesting to us, because in most cases the vapor density is very much lower than the liquid density. In those circumstances, where we are not near the critical point where they are almost the same, but where the vapor density is much lower than the liquid density, then the fact that n is very much less than $1/V_a$ is occasioned by the fact that W is very much greater than kT. So formulas such as (42.1) are interesting only when W *is* very much bigger than kT, because in those circumstances, since we are raising e to minus a tremendous amount, if we change T a little bit, that tremendous power changes a bit, and the change produced in the exponential factor is very much more important than any change that might occur in the factors out in front. Why should there be any changes in such factors as V_a? Because ours was an approximate analysis. After all, there is not really a definite volume for each molecule; as we change the temperature, the volume V_a does not stay constant —the liquid expands. There are other little features like that, and so the actual situation is more complicated. There are slowly varying temperature-dependent factors all over the place. In fact, we might say that W itself varies slightly with temperature, because at a higher temperature, at a different molecular volume, there would be different average attractions, and so on. So, while we might think that if we have a formula in which everything varies in an unknown way with temperature then we have no formula at all, if we realize that the exponent W/kT is, in general, very large, we see that in the curve of the vapor density as a function of temperature most of the variation is occasioned by the exponential factor, and if we take W as a constant and the coefficient $1/V_a$ as nearly constant, it is a good approximation for short intervals along the curve. Most of the variation, in other words, is of the general nature $e^{-W/kT}$.

It turns out that there are many, many phenomena in nature which are characterized by having to borrow an energy from somewhere, and in which the central feature of the temperature variation is e to the minus the energy over kT. This is a useful fact only when the energy is large compared with kT, so that most of the variation is contained in the variation of the kT and not in the constant and in other factors.

Now let us consider another way of obtaining a somewhat similar result for the evaporation, but looking at it in more detail. To arrive at (42.1), we simply applied a rule which is valid at equilibrium, but in order to understand things better, there is no harm in trying to look at the details of what is going on. We may also describe what is going on in the following way: the molecules that are in the vapor continually bombard the surface of the liquid; when they hit it, they may bounce off or they may get stuck. There is an unknown factor for that—maybe 50–50, maybe 10 to 90—we do not know. Let us say they always get stuck—we can analyze it over again later on the assumption that they do not always get stuck. Then at a given moment there will be a certain number of atoms which are condensing onto the surface of the liquid. The number of condensing molecules, the number that arrive on a unit area, is the number n per unit volume times the velocity v. This velocity of the molecules is related to the temperature, because we know that $\frac{1}{2}mv^2$ is equal to $\frac{3}{2}kT$ on the average. So v is some kind of a mean velocity. Of course we should integrate over the angles and get some kind of an average, but it is roughly proportional to the root-mean-square velocity, within some factor. Thus

$$N_c = nv \tag{42.2}$$

is the number which arrive per unit area and are condensing.

At the same time, however, the atoms in the liquid are jiggling about, and from time to time one of them gets kicked out. Now we have to estimate how fast they get kicked out. The idea will be that at equilibrium the number that are kicked out per second and the number that arrive per second are equal.

How many get kicked out? In order to get kicked out, a particular molecule has to have acquired by accident an excess energy over its neighbors—a considerable excess energy, because it is attracted very strongly by the other molecules in the liquid. Ordinarily it does not leave because it is so strongly attracted, but in the collisions sometimes one of them gets an extra energy by accident. And the chance that it gets the extra energy W which it needs in our case is very small if $W \gg kT$. In fact, $e^{-W/kT}$ is the chance that an atom has picked up more than this much energy. That is the general principle in kinetic theory: in order to borrow an excess energy W over the average, the odds are e to the minus the energy that we have to borrow, over kT. Now suppose that some molecules have borrowed this energy. We now have to estimate how many leave the surface per second. Of course, just because a molecule has the necessary energy does not mean that it will actually evaporate, since it may be buried too deeply inside the liquid or, even if it is near the surface, it may be travelling in the wrong direction. The number that are going to leave a unit area per second is going to be something like this: the number of atoms there are near the surface, per unit area, divided by the time it takes one to escape, multiplied by the probability $e^{-W/kT}$ that they are ready to escape in the sense that they have enough energy.

We shall suppose that each molecule at the surface of the liquid occupies a certain cross-sectional area A. Then the number of molecules per unit area of liquid surface will be $1/A$. And now, how long does it take a molecule to escape? If the molecules have a certain average speed v, and have to move, say, one molecular diameter D, the thickness of the first layer, then the time it takes to get across that thickness is the time needed to escape, if the molecule has enough energy. The time will be D/v. Thus the number evaporating should be approximately

$$N_e = (1/A)(v/D)e^{-W/kT}. \tag{42.3}$$

Now the area of each atom times the thickness of the layer is approximately the same as the volume V_a occupied by a single atom. And so, in order to get equilibrium, we must have $N_c = N_e$, or

$$nv = (v/V_a)e^{-W/kT}. \tag{42.4}$$

We may cancel the v's, since they are equal; even though one is the velocity of a molecule in the vapor and the other is the velocity of an evaporating molecule, these are the same, because we know their mean kinetic energy (in one direction) is $\frac{1}{2}kT$. But one may object, "No! No! These are the especially fast-moving ones; these are the ones that have picked up excess energy." Not really, because the moment they start to pull away from the liquid, they have to *lose* that excess energy against the potential energy. So, as they come to the surface they are slowed down to the velocity v! It is the same as it was in our discussion of the distribution of molecular velocities in the atmosphere—at the bottom, the molecules had a certain distribution of energy. The ones that arrive at the top have the *same* distribution of energy, because the slow ones did not arrive at all, and the fast ones were slowed down. The molecules that are evaporating have the same distribution of energy as the ones inside—a rather remarkable fact. Anyway, it is useless to try to argue so closely about our formula because of other inaccuracies, such as the probability of bouncing back rather than entering the liquid, and so on. Thus we have a rough idea of the rate of evaporation and condensation, and we see, of course, that the vapor density n varies in the same way as before, but now we have understood it in some detail rather than just as an arbitrary formula.

This deeper understanding permits us to analyze some things. For example, suppose that we were to pump away the vapor at such a great rate that we removed the vapor as fast as it formed (if we had very good pumps and the liquid was evaporating very slowly), how fast would evaporation occur if we maintained a liquid temperature T? Suppose that we have already experimentally measured the equilibrium vapor density, so that we know, at the given temperature, how many molecules per unit volume are in equilibrium with the liquid. Now we would like to know *how fast* it will evaporate. Even though we have used only a rough analysis so far as the evaporation part of it is concerned, the number of vapor

molecules *arriving* was not done so badly, aside from the unknown factor of reflection coefficient. So therefore we may use the fact that the number that are leaving, at equilibrium, is the same as the number that arrive. True, the vapor is being swept away and so the molecules are only coming out, but if the vapor were left alone, it would attain the equilibrium density at which the number that come back would equal the number that are evaporating. Therefore, we can easily see that the number that are coming off the surface per second is equal to the unknown reflection coefficient R times the number that would come down to the surface per second were the vapor still there, because that is how many would balance the evaporation at equilibrium:

$$N_e = nvR = (vR/V_a)e^{-W/kT}. \tag{42.5}$$

Of course, the number of molecules that hit the liquid from the vapor is easy to calculate, since we do not need to know as much about the forces as we do when we are worrying about how they get to escape through the liquid surface; it is much easier to make the argument the other way.

42–2 Thermionic emission

We may give another example of a very practical situation that is similar to the evaporation of a liquid—so similar that it is not worth making a separate analysis. It is essentially the same problem. In a radio tube there is a source of electrons, namely a heated tungsten filament, and a positively charged plate to attract the electrons. Any electron that escapes from the surface of the tungsten is immediately swept away to the plate. That is our ideal "pump," which is "pumping" the electrons away all the time. Now the question is: How many electrons per second can we get out of a piece of tungsten, and how does that number vary with temperature? The answer to that problem is the same as (42.5), because it turns out that in a piece of metal, electrons are attracted to the ions, or to atoms, of the metal. They are attracted, if we may say it crudely, to the metal. In order to get an electron out of a piece of metal, it takes a certain amount of energy or work to pull it out. This work varies with the different kinds of metal. In fact, it varies even with the character of the surface of a given kind of metal, but the total work may be a few electron volts, which, incidentally, is typical of the energy involved in chemical reactions. We can remember the latter fact by remembering that the voltage in a chemical cell like a flashlight battery, which is produced by chemical reactions, is about one volt.

How can we find out how many electrons come out per second? It would be quite difficult to analyze the effects on the electrons going out; it is easier to analyze the situation the other way. So, we could start out by imagining that we did not draw the electrons away, and that the electrons were like a gas, and could come back to the metal. Then there would be a certain density of electrons at equilibrium which would, of course, be given by exactly the same formula as (42.1), where V_a is the volume per electron in the metal, roughly, and W is equal to $q_e\phi$, where ϕ is the so-called *work function*, or the voltage needed to pull an electron off the surface. This would tell us how many electrons would have to be in the surrounding space and striking the metal in order to balance the ones that are coming out. And thus it is easy to calculate how many are coming out if we sweep away all of them, because the number that are coming out is exactly equal to the number that would be going in with the above density of electron "vapor." In other words, the answer is that the current of electricity that comes in per unit area is equal to the charge on each times the number that arrive per second per unit area, which is the number per unit volume times the velocity, as we have seen many times:

$$I = q_e nv = (q_e v/V_a)e^{-q_e\phi/kT}. \tag{42.6}$$

Now one electron volt corresponds to kT at a temperature of 11,600 degrees. The filament of the tube may be operating at a temperature of, say, 1100 degrees, so the exponential factor is something like e^{-10}; when we change the temperature a

little bit, the exponential factor changes a lot. Thus, again, the central feature of the formula is the $e^{-q_e\phi/kT}$. As a matter of fact, the factor in front is quite wrong—it turns out that the behavior of electrons in a metal is not correctly described by the classical theory, but by quantum mechanics, but this only changes the factor in front a little. Actually, no one has ever been able to get the thing straightened out very well, even though many people have used the high-class quantum-mechanical theory for their calculations. The big problem is, does W change slightly with temperature? If it does, one cannot distinguish a W changing slowly with temperature from a different coefficient in front. That is, if W changed linearly, say, with temperature, so that $W = W_0 + \alpha kT$, then we would have

$$e^{-W/kT} = e^{-(W_0 + \alpha kT)/kT} = e^{-\alpha}e^{-W_0/kT}.$$

Thus a linearly temperature-dependent W is equivalent to a shifted "constant." It is really quite difficult and usually fruitless to try to obtain the coefficient in the front accurately.

42–3 Thermal ionization

Now we go on to another example of the same idea; always the same idea. This has to do with ionization. Suppose that in a gas we have a whole lot of atoms which are in the neutral state, say, but the gas is hot and the atoms can become ionized. We would like to know how many ions there are in a given circumstance if we have a certain density of atoms per unit volume at a certain temperature. Again we consider a box in which there are N atoms which can hold electrons. (If an electron has come off an atom, it is called an *ion*, and if the atom is neutral, we simply call it an atom.) Then suppose that, at any given moment, the number of neutral atoms is n_a, the number of ions is n_i, and the number of electrons is n_e, all per unit volume. The problem is: What is the relationship of these three numbers?

In the first place, we have two conditions or constraints on the numbers. For instance, as we vary different conditions, like the temperature and so on, $n_a + n_i$ would remain constant, because this would be simply the number N of atomic nuclei that are in the box. If we keep the number of nuclei per unit volume fixed, and change, say, the temperature, then as the ionization proceeded some atoms would turn to ions, but the total number of atoms plus ions would be unchanged. That is, $n_a + n_i = N$. Another condition is that if the entire gas is to be electrically neutral (and if we neglect double or triple ionization), that means that the number of ions is equal to the number of electrons at all times, or $n_i = n_e$. These are subsidiary equations that simply express the conservation of charge and the conservation of atoms.

These equations are true, and we ultimately will use them when we consider a real problem. But we want to obtain another relationship between the quantities. We can do this as follows. We again use the idea that it takes a certain amount of energy to lift the electron out of the atom, which we call the *ionization energy*, and we would write it as W, in order to make all of the formulas look the same. So we let W equal the energy needed to pull an electron out of an atom and make an ion. Now we again say that the number of free electrons per unit volume in the "vapor" is equal to the number of bound electrons per unit volume in the atoms, times e to the minus the energy difference between being bound and being free, over kT. That is the basic equation again. How can we write it? The number of free electrons per unit volume would, of course, be n_e, because that is the definition of n_e. Now what about the number of electrons per unit volume that are bound to atoms? The total number of places that we could put the electrons is apparently $n_a + n_i$, and we will suppose that when they are bound each one is bound within a certain volume V_a. So the total amount of volume which is available to electrons which would be bound is $(n_a + n_i)V_a$, so we might want to write our formula as

$$n_e = \frac{n_a}{(n_a + n_i)V_a}e^{-W/kT}.$$

The formula is wrong, however, in one essential feature, which is the following: when an electron is already on an atom, another electron cannot come to that volume anymore! In other words, all the volumes of all the possible sites are not really available for the one electron which is trying to make up its mind whether or not to be in the vapor or in the condensed position, because in this problem there is an extra feature that when one electron is where another electron is, it is not allowed to go—it is repelled. For that reason, it comes out that we should count only that part of the volume which is available for an electron to sit on or not. That is, those which are already occupied do not count in the total available volume, but the only volume which is allowed is that of the *ions*, where there are vacant places for the electron to go. Then, in those circumstances, we find that a nicer way to write our formula is

$$\frac{n_e n_i}{n_a} = \frac{1}{V_a} e^{-W/kT}. \tag{42.7}$$

This formula is called the *Saha ionization equation*. Now let us see if we can understand qualitatively why a formula like this is right, by arguing about the kinetic things that are happening.

First, every once in a while an electron comes to an ion and they combine to make an atom. And also, every once in a while, an atom gets into a collision and breaks up into an ion and an electron. Now those two rates must be equal. How fast do electrons and ions find each other? The rate is certainly increased if the number of electrons per unit volume is increased. It is also increased if the number of ions per unit volume is increased. That is, the total rate at which recombination is occurring is certainly proportional to the number of electrons times the number of ions. Now the total rate at which ionization is occurring due to collisions must be dependent linearly on how many atoms there are to ionize. And so the rates will balance when there is some relationship between the product $n_e n_i$ and the number of atoms, n_a. The fact that this relationship happens to be given by this particular formula, where W is the ionization energy, is of course a little bit more information, but we can easily understand that the formula would necessarily involve the concentrations of the electrons, ions, and atoms in the combination $n_e n_i/n_a$ to produce a constant independent of the n's, and dependent only on temperature, the atomic cross sections, and other constant factors.

We may also note that, since the equation involves the numbers *per unit volume*, if we were to do two experiments with a given total number N of atoms plus ions, that is, a certain fixed number of nuclei, but using boxes with different volumes, the n's would all be smaller in the larger box. But since the ratio $n_e n_i/n_a$ stays the same, the *total number* of electrons and ions must be greater in the larger box. To see this, suppose that there are N nuclei inside a box of volume V, and that a fraction f of them are ionized. Then $n_e = fN/V = n_i$, and $n_a = (1-f)N/V$. Then our equation becomes

$$\frac{f^2}{1-f}\frac{N}{V} = \frac{e^{-W/kT}}{V_a}. \tag{42.8}$$

In other words, if we take a smaller and smaller density of atoms, or make the volume of the container bigger and bigger, the fraction f of electrons and ions must increase. That ionization, just from "expansion" as the density goes down, is the reason why we believe that at very low densities, such as in the cold space between the stars, there may be ions present, even though we might not understand it from the point of view of the available energy. Although it takes many, many kT of energy to make them, there are ions present.

Why can there be ions present when there is so much space around, while if we increase the density, the ions tend to disappear? *Answer*: Consider an atom. Every once in a while, light, or another atom, or an ion, or whatever it is that maintains thermal equilibrium, strikes it. Very rarely, because it takes such a terrific amount of excess energy, an electron comes off and an ion is left. Now that electron, if the space is enormous, wanders and wanders and does not come near anything for years, perhaps. But once in a very great while, it does come back

to an ion and they combine to make an atom. So the rate at which electrons are coming out from the atoms is very slow. But if the volume is enormous, an electron which has escaped takes so long to find another ion to recombine with that its probability of recombination is very, very small; thus, in spite of the large excess energy needed, there may be a reasonable number of electrons.

42–4 Chemical kinetics

The same situation that we have just called "ionization" is also found in a chemical reaction. For instance, if two objects A and B combine into a compound AB, then if we think about it for a while we see that AB is what we have called an atom, B is what we call an electron, and A is what we call an ion. With these substitutions the equations of equilibrium are exactly the same in form:

$$\frac{n_A n_B}{n_{AB}} = ce^{-W/kT}. \tag{42.9}$$

This formula, of course, is not exact, since the "constant" c depends on how much volume is allowed for the A and B to combine, and so on, but by thermodynamic arguments one can identify what the meaning of the W in the exponential factor is, and it turns out that it is very close to the energy needed in the reaction.

Suppose that we tried to understand this formula as a result of collisions, much in the way that we understood the evaporation formula, by arguing about how many electrons came off and how many of them came back per unit time. Suppose that A and B combine in a collision every once in a while to form a compound AB. And suppose that the compound AB is a complicated molecule which jiggles around and is hit by other molecules, and from time to time it gets enough energy to explode and break up again into A and B.

Now it actually turns out, in chemical reactions, that if the atoms come together with too small an energy, even though energy may be released in the reaction $A + B \to AB$, the fact that A and B may touch each other does not necessarily make the reaction start. It usually is required that the collision be rather hard, in fact, to get the reaction to go at all—a "soft" collision between A and B may not do it, even though energy may be released in the process. So let us suppose that it is very common in chemical reactions that, in order for A and B to form AB, they cannot just hit each other, but they have to hit each other *with sufficient energy*. This energy is called the *activation energy*—the energy needed to "activate" the reaction. Call A^* the activation energy, the excess energy needed in a collision in order that the reaction may really occur. Then the rate R_f at which A and B produce AB would involve the number of atoms of A times the number of atoms of B, times the rate at which a single atom would strike a certain cross section σ_{AB}, times a factor $e^{-A^*/kT}$, which is the probability that they have enough energy:

$$R_f = n_A n_B v \sigma_{AB} e^{-A^*/kT}. \tag{42.10}$$

Now we have to find the opposite rate, R_r. There is a certain chance that AB will fly apart. In order to fly apart, it not only must have the energy W which it needs in order to get apart at all but, just as it was hard for A and B to come together, so there is a kind of hill that A and B have to climb over to get apart again; they must have not only enough energy just to get ready to pull apart, but a certain excess. It is like climbing a hill to get into a deep valley; they have to climb the hill coming in and they have to climb out of the valley and then over the hill coming back (Fig. 42–1). Thus the rate at which AB goes to A and B will be proportional to the number n_{AB} that are present, times $e^{-(W+A^*)/kT}$:

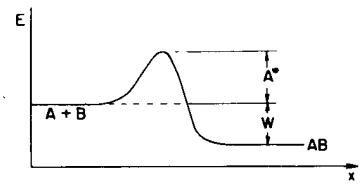


Fig. 42–1. The energy relationship for the reaction $A + B \to AB$.

$$R_r = c' n_{AB} e^{-(W+A^*)/kT}. \tag{42.11}$$

The c' will involve the volume of atoms and the rate of collisions, which we can work out, as we did the case of evaporation, with areas and times and thicknesses; but we shall not do this. The main feature of interest to us is that when these two

rates are equal, the ratio of them is equal to unity. This tells us that $n_A n_B/n_{AB} = ce^{-W/kT}$, as before, where c involves the cross sections, velocities, and other factors independent of the n's.

The interesting thing is that the rate of the reaction also varies as $e^{-\text{const}/kT}$, although the constant is not the same as that which governs the concentrations; the activation energy A^* is quite different from the energy W. *W governs the proportions of A, B, and AB that we have in equilibrium*, but if we want to know how fast $A + B$ goes to AB, that is not a question of equilibrium, and here a different energy, *the activation energy*, governs the rate of reaction through an exponential factor.

Furthermore, A^* is not a fundamental constant like W. Suppose that at the surface of the wall—or at some other place—A and B could temporarily stick there in such a way that they could combine more easily. In other words, we might find a "tunnel" through the hill, or perhaps a lower hill. By the conservation of energy, when we are all finished we have still made AB out of A and B, so the energy difference W will be quite independent of the way the reaction occurred, but the activation energy A^* will depend *very much* on the way the reaction occurs. This is why the rates of chemical reactions are very sensitive to outside conditions. We can change the rate by putting in a surface of a different kind, we can put it in a "different barrel" and it will go at a different rate, if it depends on the nature of the surface. Or if we put in a third kind of object it may change the rate very much; some things produce enormous changes in rate simply by changing the A^* a little bit—they are called *catalysts*. A reaction might practically not occur at all because A^* is too big at the given temperature, but when we put in this special stuff, the catalyst, then the reaction goes very fast indeed, because A^* is reduced.

Incidentally, there is some trouble with such a reaction, A plus B, making AB, because we cannot conserve both energy and momentum when we try to put two objects together to make one that is more stable. Therefore, we need at least a third object C, so the actual reaction is much more complicated. The forward rate would involve the product $n_A n_B n_C$, and it might seem that our formula is going wrong, but no! When we look at the rate at which AB goes the other way, we find that it also needs to collide with C, so there is an $n_{AB}n_C$ in the reverse rate; the n_C's cancel out in the formula for the equilibrium concentrations. The law of equilibrium, (42.9), which we first wrote down is absolutely guaranteed to be true, no matter what the mechanism of the reaction may be!

42–5 Einstein's laws of radiation

We now turn to an interesting analogous situation having to do with the blackbody radiation law. In the last chapter we worked out the distribution law for the radiation in a cavity the way Planck did, considering the radiation from an oscillator. The oscillator had to have a certain mean energy, and since it was oscillating, it would radiate and would keep pumping radiation into the cavity until it piled up enough radiation to balance the absorption and emission. In that way we found that the intensity of radiation at frequency ω was given by the formula

$$I(\omega)\,d\omega = \frac{\hbar\omega^3 d\omega}{\pi^2 c^2 (e^{\hbar\omega/kT} - 1)}. \tag{42.12}$$

This result involved the assumption that the oscillator which was generating the radiation had definite, equally spaced energy levels. We did not say that light had to be a photon or anything like that. There was no discussion about how, when an atom goes from one level to another, the energy must come out in one unit of energy, $\hbar\omega$, in the form of light. Planck's original idea was that the matter was quantized but not the light: material oscillators cannot take up just any energy, but have to take it in lumps. Furthermore, the trouble with the derivation is that it was partially classical. We calculated the rate of radiation from an oscillator according to classical physics; then we turned around and said, "No, this oscillator has a lot of energy levels." So gradually, in order to find the right result, the completely quantum-mechanical result, there was a slow develop-

ment which culminated in the quantum mechanics of 1927. But in the meantime, there was an attempt by Einstein to convert Planck's viewpoint that only oscillators of matter were quantized, to the idea that light was really photons and could be considered in a certain way as particles with energy $\hbar\omega$. Furthermore, Bohr had pointed out that *any* system of atoms has energy levels, but they are not necessarily equally spaced like Planck's oscillator. And so it became necessary to rederive or at least rediscuss the radiation law from a more completely quantum-mechanical viewpoint.

Einstein assumed that Planck's final formula was right, and he used that formula to obtain some new information, previously unknown, about the interaction of radiation with matter. His discussion went as follows: Consider any two of the many energy levels of an atom, say the mth level and the nth level (Fig. 42–2). Now Einstein proposed that when such an atom has light of the right frequency shining on it, it can absorb that photon of light and make a transition from state n to state m, and that the probability that this occurs per second depends upon the two levels, of course, but is *proportional to how intense the light is* that is shining on it. Let us call the proportionality constant B_{nm}, merely to remind us that this is not a universal constant of nature, but depends on the particular pair of levels: some levels are easy to excite; some levels are hard to excite. Now what is the formula going to be for the rate of emission from m to n? Einstein proposed that this must have two parts to it. First, even if there were no light present, there would be some chance that an atom in an excited state would fall to a lower state, emitting a photon; this we call *spontaneous emission.* It is analogous to the idea that an oscillator with a certain amount of energy, even in classical physics, does not keep that energy, but loses it by radiation. Thus the analog of spontaneous radiation of a classical system is that if the atom is in an excited state there is a certain probability A_{mn}, which depends on the levels again, for it to go down from m to n, and this probability is independent of whether light is shining on the atom or not. But then Einstein went further, and by comparison with the classical theory and by other arguments, concluded that emission was also influenced by the presence of light—that when light of the right frequency is shining on an atom, it has an increased rate of emitting a photon that is proportional to the intensity of the light, with a proportionality constant B_{mn}. Later, if we deduce that this coefficient is zero, then we will have found that Einstein was wrong. Of course we will find he was right.

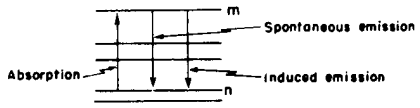


Fig. 42–2. Transitions between two energy levels of an atom.

Thus Einstein assumed that there are three kinds of processes: an absorption proportional to the intensity of light, an emission proportional to the intensity of light, called *induced emission* or sometimes *stimulated emission*, and a spontaneous emission independent of light.

Now suppose that we have, in equilibrium at temperature T, a certain number of atoms N_n in the state n and another number N_m in the state m. Then the total number of atoms that are going from n to m is the number that are in the state n times the rate per second that, if one is in n, it goes up to m. So we have a formula for the number that are going from n to m per second:

$$R_{n\to m} = N_n B_{nm} I(\omega). \tag{42.13}$$

The number that will go from m to n is expressed in the same manner, as the number N_m that are in m, times the chance per second that each one goes down to n. This time our expression is

$$R_{m\to n} = N_m[A_{mn} + B_{mn} I(\omega)]. \tag{42.14}$$

Now we shall suppose that in thermal equilibrium the number of atoms going up must equal the number coming down. That is one way, at least, in which the number will be sure to stay constant in each level.* So we take these two rates

* This is not the only way one can arrange to keep the numbers of atoms in the various levels constant, but it is the way it actually works. That every process must, in thermal equilibrium, be balanced by its exact opposite is called the *principle of detailed balancing.*

to be equal at equilibrium. But we have one other piece of information: we know how large N_m is compared with N_n—the ratio of those two is $e^{-(E_m-E_n)/kT}$. Now Einstein assumed that the only light which is effective in making the transition from n to m is the light which has the frequency corresponding to the energy difference, so $E_m - E_n = \hbar\omega$ in all our formulas. Thus

$$N_m = N_n e^{-\hbar\omega/kT}. \tag{42.15}$$

Thus if we set the two rates equal: $N_n B_{nm} I(\omega) = N_m[A_{mn} + B_{mn} I(\omega)]$, and divide by N_m, we get

$$B_{nm} I(\omega) e^{\hbar\omega/kT} = A_{mn} + B_{mn} I(\omega). \tag{42.16}$$

From this equation, we can calculate $I(\omega)$. It is simply

$$I(\omega) = \frac{A_{mn}}{B_{nm} e^{\hbar\omega/kT} - B_{mn}}. \tag{42.17}$$

But Planck has already told us that the formula must be (42.12). Therefore we can deduce something: First, that B_{nm} must equal B_{mn}, since otherwise we cannot get the $(e^{\hbar\omega/kT} - 1)$. So Einstein discovered some things that he did not know how to calculate, namely *that the induced emission probability and the absorption probability must be equal.* This is interesting. And furthermore, in order for (42.17) and (42.12) to agree,

$$A_{mn}/B_{mn} \quad \text{must be} \quad \hbar\omega^3/\pi^2c^2. \tag{42.18}$$

So if we know, for instance, the absorption rate for a given level, we can deduce the spontaneous emission rate and the induced emission rate, or any combination.

This is as far as Einstein or anyone else could go using such arguments. To actually compute the absolute spontaneous emission rate or the other rates for any specific atomic transition, of course, requires a knowledge of the machinery of the atom, called quantum electrodynamics, which was not discovered until eleven years later. This work of Einstein was done in 1916.

The possibility of induced emission has, today, found interesting applications. If there is light present, it will tend to induce the downward transition. The transition then adds its $\hbar\omega$ to the available light energy, if there were some atoms sitting in the upper state. Now we can arrange, by some nonthermal method, to have a gas in which the number in the state m is very much greater than the number in the state n. This is far out of equilibrium, and so is not given by the formula $e^{-\hbar\omega/kT}$, which is for equilibrium. We can even arrange it so that the number in the upper state is very large, while the number in the lower state is practically zero. Then light which has the frequency corresponding to the energy difference $E_m - E_n$ will not be strongly absorbed, because there are not many atoms in state n to absorb it. On the other hand, when that light is present, it will induce the emission from this upper state! So, if we had a lot of atoms in the upper state, there would be a sort of chain reaction, in which, the moment the atoms began to emit, more would be caused to emit, and the whole lot of them would dump down together. This is what is called a *laser*, or, in the case of the far infrared, a *maser*.

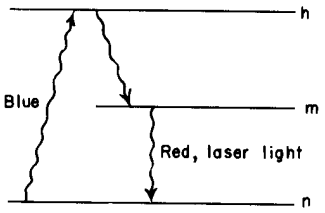


Fig. 42–3. By exciting, say by blue light, a higher state h, which may emit a photon leaving atoms in state m, the number in this state m becomes sufficiently large to start laser action.

Various tricks can be used to obtain the atoms in state m. There may be higher levels to which the atoms can get if we shine in a strong beam of light of high frequency. From these high levels, they may trickle down, emitting various photons, until they all get stuck in the state m. If they tend to stay in the state m without emitting, the state is called *metastable*. And then they are all dumped down together by induced emissions. One more technical point—if we put this system in an ordinary box, it would radiate in so many different directions spontaneously, compared with the induced effect, that we would still be in trouble. But we can enhance the induced effect, increase its efficiency, by putting nearly perfect mirrors on each side of the box, so that the light which is emitted gets another chance, and another chance, and another chance, to induce more emission. Although the mirrors are almost one hundred percent reflecting, there is a slight amount of transmission of the mirror, and a little light gets out. In the end, of course, from the conservation of energy, all the light goes out in a nice uniform straight direction which makes the strong light beams that are possible today with lasers.

43

Diffusion

43–1 Collisions between molecules

43–2 The mean free path

43–3 The drift speed

43–4 Ionic conductivity

43–5 Molecular diffusion

43–6 Thermal conductivity

43–1 Collisions between molecules

We have considered so far only the molecular motions in a gas which is in thermal equilibrium. We want now to discuss what happens when things are near, but not exactly in, equilibrium. In a situation far from equilibrium, things are extremely complicated, but in a situation very close to equilibrium we can easily work out what happens. To see what happens, we must, however, return to the kinetic theory. Statistical mechanics and thermodynamics deal with the equilibrium situation, but away from equilibrium we can only analyze what occurs atom by atom, so to speak.

As a simple example of a nonequilibrium circumstance, we shall consider the diffusion of ions in a gas. Suppose that in a gas there is a relatively small concentration of ions—electrically charged molecules. If we put an electric field on the gas, then each ion will have a force on it which is different from the forces on the neutral molecules of the gas. If there were no other molecules present, an ion would have a constant acceleration until it reached the wall of the container. But because of the presence of the other molecules, it cannot do that; its velocity increases only until it collides with a molecule and loses its momentum. It starts again to pick up more speed, but then it loses its momentum again. The net effect is that an ion works its way along an erratic path, but with a net motion in the direction of the electric force. We shall see that the ion has an average "drift" with a mean speed which is proportional to the electric field—the stronger the field, the faster it goes. While the field is on, and while the ion is moving along, it is, of course, *not* in thermal equilibrium, it is trying to get to equilibrium, which is to be sitting at the end of the container. By means of the kinetic theory we can compute the drift velocity.

It turns out that with our present mathematical abilities we cannot really compute *precisely* what will happen, but we can obtain approximate results which exhibit all the essential features. We can find out how things will vary with pressure, with temperature, and so on, but it will not be possible to get precisely the correct numerical factors in front of all the terms. We shall, therefore, in our derivations, not worry about the precise value of numerical factors. They can be obtained only by a very much more sophisticated mathematical treatment.

Before we consider what happens in nonequilibrium situations, we shall need to look a little closer at what goes on in a gas in thermal equilibrium. We shall need to know, for example, what the average time between successive collisions of a molecule is.

Any molecule experiences a sequence of collisions with other molecules—in a random way, of course. A particular molecule will, in a long period of time T, have a certain number, N, of hits. If we double the length of time, there will be twice as many hits. So the number of collisions is proportional to the time T. We would like to write it this way:

$$N = T/\tau \tag{43.1}$$

We have written the constant of proportionality as $1/\tau$, where τ will have the dimensions of a time. The constant τ is the average time between collisions. Suppose, for example, that in an hour there are 60 collisions; then τ is one minute. We would say that τ (one minute) is the *average time* between the collisions.

We may often wish to ask the following question: "What is the *chance* that a molecule will experience a collision during the next *small interval* of time dt?" The answer, we may intuitively understand, is dt/τ. But let us try to make a more

convincing argument. Suppose that there were a very large number N of molecules. How many will have collisions in the next interval of time dt? If there is equilibrium, nothing is changing *on the average* with time. So N molecules waiting the time dt will have the same number of collisions as *one* molecule waiting for the time $N\,dt$. That number we know is $N\,dt/\tau$. So the number of hits of N molecules is $N\,dt/\tau$ in a time dt, and the chance, or probability, of a hit for any one molecule is just $1/N$ as large, or $(1/N)(N\,dt/\tau) = dt/\tau$, as we guessed above. That is to say, the fraction of the molecules which will suffer a collision in the time dt is dt/τ. To take an example, if τ is one minute, then in one second the fraction of particles which will suffer collisions is 1/60. What this means, of course, is that 1/60 of the molecules happen to be close enough to what they are going to hit next that *their* collisions will occur in the next minute.

When we say that τ, the mean time between collisions, is one minute, we do not mean that all the collisions will occur at times separated by exactly one minute. A particular particle does not have a collision, wait one minute, and then have another collision. The times between successive collisions are quite variable. We will not need it for our later work here, but we may make a small diversion to answer the question: "What *are* the times between collisions?" We know that for the case above, the *average* time is one minute, but we might like to know, for example, what is the chance that we get no collision for *two* minutes?

We shall find the answer to the general question: "What is the probability that a molecule will go for a time t without having a collision?" At some arbitrary instant—that we call $t = 0$—we begin to watch a particular molecule. What is the chance that it gets by until t without colliding with another molecule? To compute the probability, we observe what is happening to all N_0 molecules in a container. After we have waited a time t, some of them will have had collisions. We let $N(t)$ be the number that have *not* had collisions up to the time t. $N(t)$ is, of course, less than N_0. We can find $N(t)$ because we know how it changes with time. If we know that $N(t)$ molecules have got by until t, then $N(t + dt)$, the number which get by until $t + dt$, is *less* than $N(t)$ by the number that have collisions in dt. The number that collide in dt we have written above in terms of the mean time τ as $dN = N(t)\,dt/\tau$. We have the equation

$$N(t + dt) = N(t) - N(t)\frac{dt}{\tau}. \tag{43.2}$$

The quantity on the left-hand side, $N(t + dt)$, can be written, according to the definitions of calculus, as $N(t) + (dN/dt)\,dt$. Making this substitution, Eq. (43.2) yields

$$\frac{dN(t)}{dt} = -\frac{N(t)}{\tau}. \tag{43.3}$$

The number that are being lost in the interval dt is proportional to the number that are present, and inversely proportional to the mean life τ. Equation (43.3) is easily integrated if we rewrite it as

$$\frac{dN(t)}{N(t)} = -\frac{dt}{\tau}. \tag{43.4}$$

Each side is a perfect differential, so the integral is

$$\ln N(t) = -t/\tau + (\text{a constant}), \tag{43.5}$$

which says the same thing as

$$N(t) = (\text{constant})e^{-t/\tau}. \tag{43.6}$$

We know that the constant must be just N_0, the total number of molecules present, since all of them start at $t = 0$ to wait for their "next" collision. We can write our result as

$$N(t) = N_0 e^{-t/\tau}. \tag{43.7}$$

If we wish the *probability* of no collision, $P(t)$, we can get it by dividing $N(t)$ by N_0, so

$$P(t) = e^{-t/\tau}. \tag{43.8}$$

Our result is: the probability that a particular molecule survives a time t without a collision is $e^{-t/\tau}$, where τ is the mean time between collisions. The probability starts out at 1 (or certainty) for $t = 0$, and gets less as t gets bigger and bigger. The probability that the molecule avoids a collision for a time equal to τ is $e^{-1} = 0.37\ldots$ The chance is less than one-half that it will have a greater than average time between collisions. That is all right, because there are enough molecules which go collision-free for times much *longer* than the mean time before colliding, so that the average time can still be τ.

We originally defined τ as the average time *between* collisions. The result we have obtained in Eq. (43.7) also says that the mean time from an *arbitrary* starting instant to the *next* collision is *also* τ. We can demonstrate this somewhat surprising fact in the following way. The number of molecules which experience their *next* collision in the interval dt at the time t after an arbitrarily chosen starting time is $N(t)\,dt/\tau$. Their "time until the next collision" is, of course, just t. The "average time until the next collision" is obtained in the usual way:

$$\text{Average time until the next collision} = \frac{1}{N_0}\int_0^\infty t\,\frac{N(t)\,dt}{\tau}.$$

Using $N(t)$ obtained in (43.7) and evaluating the integral, we find indeed that τ is the average time from *any* instant until the next collision.

43–2 The mean free path

Another way of describing the molecular collisions is to talk not about the *time* between collisions, but about *how far* the particle moves between collisions. If we say that the average time between collisions is τ, and that the molecules have a mean velocity v, we can expect that the average *distance* between collisions, which we shall call l, is just the product of τ and v. This distance between collisions is usually called the *mean free path:*

$$\text{Mean free path } l = \tau v. \tag{43.9}$$

In this chapter we shall be a little careless about *what kind of average* we mean in any particular case. The various possible averages—the mean, the root-mean-square, etc.—are all nearly equal and differ by factors which are near to one. Since a detailed analysis is required to obtain the correct numerical factors anyway, we need not worry about which average is required at any particular point. We may also warn the reader that the algebraic symbols we are using for some of the physical quantities (e.g., l for the mean free path) do not follow a generally accepted convention, mainly because there is no general agreement.

Just as the chance that a molecule will have a collision in a short time dt is equal to dt/τ, the chance that it will have a collision in going a distance dx is dx/l. Following the same line of argument used above, the reader can show that the probability that a molecule will go at least the distance x before having its next collision is $e^{-x/l}$.

The average distance a molecule goes before colliding with another molecule—the mean free path l—will depend on how many molecules there are around and on the "size" of the molecules, i.e., how big a target they represent. The effective "size" of a target in a collision we usually describe by a "collision cross section," the same idea that is used in nuclear physics, or in light-scattering problems.

Consider a moving particle which travels a distance dx through a gas which has n_0 scatterers (molecules) per unit volume (Fig. 43–1). If we look at each unit of area perpendicular to the direction of motion of our selected particle, we will find there $n_0\,dx$ molecules. If each one presents an effective collision area or, as it is usually called, "collision cross section," σ_c, then the total area covered by the scatterers is $\sigma_c n_0\,dx$.

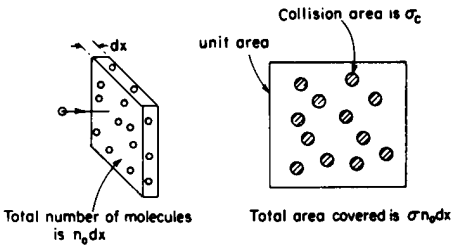


Fig. 43–1. Collision cross section.

By "collision cross section" we mean the area within which the center of our particle must be located if it is to collide with a particular molecule. If molecules were little spheres (a classical picture) we would expect that $\sigma_c = \pi(r_1 + r_2)^2$, where r_1 and r_2 are the radii of the two colliding objects. The chance that our particle will have a collision is the ratio of the area covered by scattering molecules to the total area, which we have taken to be one. So the probability of a collision in going a distance dx is just $\sigma_c n_0\, dx$:

$$\text{Chance of a collision in } dx = \sigma_c n_0\, dx. \tag{43.10}$$

We have seen above that the chance of a collision in dx can also be written in terms of the mean free path l as dx/l. Comparing this with (43.10), we can relate the mean free path to the collision cross section:

$$\frac{1}{l} = \sigma_c n_0, \tag{43.11}$$

which is easier to remember if we write it as

$$\sigma_c n_0 l = 1. \tag{43.12}$$

This formula can be thought of as saying that there should be one collision, on the average, when the particle goes through a distance l in which the scattering molecules *could* just cover the total area. In a cylindrical volume of length l and a base of unit area, there are $n_0 l$ scatterers; if each one has an area σ_c the total area covered is $n_0 l \sigma_c$, which is just one unit of area. The whole area is *not* covered, of course, because some molecules are partly hidden behind others. That is why some molecules go farther than l before having a collision. It is only *on the average* that the molecules have a collision by the time they go the distance l. From measurements of the mean free path l we can determine the scattering cross section σ_c, and compare the result with calculations based on a detailed theory of atomic structure. But that is a different subject! So we return to the problem of nonequilibrium states.

43–3 The drift speed

We want to describe what happens to a molecule, or several molecules, which are different in some way from the large majority of the molecules in a gas. We shall refer to the "majority" molecules as the "background" molecules, and we shall call the molecules which are different from the background molecules "special" molecules or, for short, the *S*-molecules. A molecule could be special for any number of reasons: It might be heavier than the background molecules. It might be a different chemical. It might have an electric charge—i.e., be an ion in a background of uncharged molecules. Because of their different masses or charges the *S*-molecules may have forces on them which are different from the forces on the background molecules. By considering what happens to these *S*-molecules we can understand the basic effects which come into play in a similar way in many different phenomena. To list a few: the diffusion of gases, electric currents in batteries, sedimentation, centrifugal separation, etc.

We begin by concentrating on the basic process: an *S*-molecule in a background gas is acted on by some specific force **F** (which might be, e.g., gravitational or electrical) and *in addition* by the not-so-specific forces due to collisions with the background molecules. We would like to describe the *general* behavior of the *S*-molecule. What happens to it, *in detail*, is that it darts around hither and yon as it collides over and over again with other molecules. But if we watch it carefully we see that it does make some net progress in the direction of the force **F**. We say that there is a *drift* superposed on its random motion. We would like to know what the speed of its drift is—its *drift velocity*—due to the force **F**.

If we start to observe an *S*-molecule at some instant we may expect that it is somewhere between two collisions. In addition to the velocity it was left with after its last collision it is picking up some velocity component due to the force **F**. In a

short time (on the average, in a time τ) it will experience a collision and start out on a new piece of its trajectory. It will have a new starting velocity, but the same acceleration from **F**.

To keep things simple for the moment, we shall suppose that after each collision our *S*-molecule gets a completely "fresh" start. That is, that it keeps no remembrance of its past acceleration by **F**. This might be a reasonable assumption if our *S*-molecule were much lighter than the background molecules, but it is certainly not valid in general. We shall discuss later an improved assumption.

For the moment, then, our assumption is that the *S*-molecule leaves each collision with a velocity which may be in any direction with equal likelihood. The starting velocity will take it equally in all directions and will not contribute to any net motion, so we shall not worry further about its initial velocity after a collision. In addition to its random motion, each *S*-molecule will have, at any moment, an additional velocity in the direction of the force **F**, which it has picked up *since* its last collision. What is the *average* value of *this* part of the velocity? It is just the acceleration $\mathbf{F}/m$ (where m is the mass of the *S*-molecule) times the *average* time *since* the last collision. Now the average time *since* the *last* collision must be the same as the average time *until* the *next* collision, which we have called τ, above. The *average* velocity from **F**, of course, is just what is called the drift velocity, so we have the relation

$$v_{\text{drift}} = \frac{F\tau}{m}. \tag{43.13}$$

This basic relation is the heart of our subject. There may be some complication in determining what τ is, but the basic process is defined by Eq. (43.13).

You will notice that the drift velocity is *proportional* to the force. There is, unfortunately, no generally used name for the constant of proportionality. Different names have been used for each different kind of force. If in an electrical problem the force is written as the charge times the electric field, $\mathbf{F} = q\mathbf{E}$, then the constant of proportionality between the velocity and the electric field **E** is usually called the "mobility." In spite of the possibility of some confusion, *we* shall use the term *mobility* for the ratio of the drift velocity to the force for *any* force. We write

$$v_{\text{drift}} = \mu F \tag{43.14}$$

in general, and we shall call μ the mobility. We have from Eq. (43.13) that

$$\mu = \tau/m. \tag{43.15}$$

The mobility is proportional to the mean time between collisions (there are fewer collisions to slow it down) and inversely proportional to the mass (more inertia means less speed picked up between collisions).

To get the correct numerical coefficient in Eq. (43.13), which is correct as given, takes some care. Without intending to confuse, we should still point out that the arguments have a subtlety which can be appreciated only by a careful and detailed study. To illustrate that there are difficulties, in spite of appearances, we shall make over again the argument which led to Eq. (43.13) in a reasonable *but erroneous* way (and the way one will find in many textbooks!).

We might have said: The mean time between collisions is τ. After a collision the particle starts out with a random velocity, but it picks up an additional velocity between collisions, which is equal to the acceleration times the time. Since it takes the time τ to arrive at the *next* collision it gets there with the velocity $(F/m)\tau$. At the beginning of the collision it had zero velocity. So between the two collisions it has, on the average, a velocity one-half of the final velocity, so the mean drift velocity is $\frac{1}{2}F\tau/m$. (Wrong!) This result is wrong and the result in Eq. (43.13) is right, although the arguments may sound equally satisfactory. The reason the second result is wrong is somewhat subtle, and has to do with the following: The argument is made as though all collisions were separated by the mean time τ. The fact is that some times are shorter and others are longer than the mean. Short times occur *more often* but make *less* contribution to the drift velocity because they have less

chance "to really get going." If one takes proper account of the *distribution* of free times between collisions, one can show that there should not be the factor $\frac{1}{2}$ that was obtained from the second argument. The error was made in trying to relate by a simple argument the *average final* velocity to the average velocity itself. This relationship is not simple, so it is best to concentrate on what is wanted: the average velocity itself. The first argument we gave determines the average velocity directly—and correctly! But we can perhaps see now why we shall not in general try to get all of the correct numerical coefficients in our elementary derivations!

We return now to our simplifying assumption that each collision knocks out all memory of the past motion—that a fresh start is made after each collision. Suppose our S-molecule is a heavy object in a background of lighter molecules. Then our S-molecule will not lose its "forward" momentum in each collision. It would take several collisions before its motion was "randomized" again. We should assume, instead, that at each collision—in each time τ on the average —it loses a certain fraction of its momentum. We shall not work out the details, but just state that the result is equivalent to replacing τ, the average collision time, by a new—and longer—τ which corresponds to the average "forgetting time," i.e., the average time to forget its forward momentum. With such an interpretation of τ we can use our formula (43.15) for situations which are not quite as simple as we first assumed.

43–4 Ionic conductivity

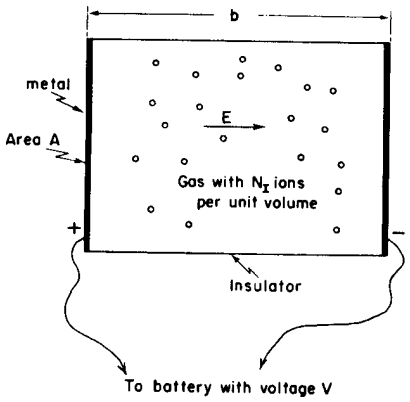


Fig. 43–2. Electric current from an ionized gas.

We now apply our results to a special case. Suppose we have a gas in a vessel in which there are also some ions—atoms or molecules with a net electric charge. We show the situation schematically in Fig. 43–2. If two opposite walls of the container are metallic plates, we can connect them to the terminals of a battery and thereby produce an electric field in the gas. The electric field will result in a force on the ions, so they will begin to drift toward one or the other of the plates. An electric current will be induced, and the gas with its ions will behave like a resistor. By computing the ion flow from the drift velocity we can compute the resistance. We ask, specifically: How does the flow of electric current depend on the voltage difference V that we apply across the two plates?

We consider the case that our container is a rectangular box of length b and cross-sectional area A (Fig. 43–2). If the potential difference, or voltage, from one plate to the other is V, the electric field E between the plates is V/b. (The electric potential is the work done in carrying a unit charge from one plate to the other. The force on a unit charge is $\mathbf{E}$. If $\mathbf{E}$ is the same everywhere between the plates, which is a good enough approximation for now, the work done on a unit charge is just Eb, so $V = Eb$.) The special force on an ion of the gas is $q\mathbf{E}$, where q is the charge on the ion. The drift velocity of the ion is then μ times this force, or

$$v_{\text{drift}} = \mu F = \mu qE = \mu q \frac{V}{b}. \tag{43.16}$$

An electric current I is the flow of charge in a unit time. The electric current to one of the plates is given by the total charge of the ions which arrive at the plate in a unit of time. If the ions drift toward the plate with the velocity v_{drift}, then those which are within a distance $(v_{\text{drift}} \cdot T)$ will arrive at the plate in the time T. If there are n_i ions per unit volume, the number which reach the plate in the time T is $(n_i \cdot A \cdot v_{\text{drift}} \cdot T)$. Each ion carries the charge q, so we have that

$$\text{Charge collected in } T = qn_iAv_{\text{drift}}T. \tag{43.17}$$

The current I is the charge collected in T divided by T, so

$$I = qn_iAv_{\text{drift}}. \tag{43.18}$$

Substituting v_{drift} from (43.16), we have

$$I = \mu q^2 n_i \frac{A}{b} V. \tag{43.19}$$

We find that the current is proportional to the voltage, which is just the form of Ohm's law, and the resistance R is the inverse of the proportionality constant:

$$\frac{1}{R} = \mu q^2 n_i \frac{A}{b}. \tag{43.20}$$

We have a relation between the resistance and the molecular properties n_i, q, and μ, which depends in turn on m and τ. If we know n_i and q from atomic measurements, a measurement of R could be used to determine μ, and from μ also τ.

43–5 Molecular diffusion

We turn now to a different kind of problem, and a different kind of analysis: the theory of diffusion. Suppose that we have a container of gas in thermal equilibrium, and that we introduce a small amount of a different kind of gas at some place in the container. We shall call the original gas the "background" gas and the new one the "special" gas. The special gas will start to spread out through the whole container, but it will spread slowly because of the presence of the background gas. This slow spreading-out process is called *diffusion*. The diffusion is controlled mainly by the molecules of the special gas getting knocked about by the molecules of the background gas. After a large number of collisions, the special molecules end up spread out more or less evenly throughout the whole volume. We must be careful *not* to confuse diffusion of a gas with the gross transport that may occur due to convection currents. Most commonly, the mixing of two gases occurs by a combination of convection and diffusion. We are interested now only in the case that there are *no "wind" currents*. The gas is spreading only by molecular motions, by diffusion. We wish to compute how fast diffusion takes place.

We now compute the *net flow* of molecules of the "special" gas due to the molecular motions. There will be a net flow only when there is some nonuniform distribution of the molecules, otherwise all of the molecular motions would average to give no net flow. Let us consider first the flow in the x-direction. To find the flow, we consider an imaginary plane surface perpendicular to the x-axis and count the number of special molecules that cross this plane. To obtain the net flow, we must count as positive those molecules which cross in the direction of positive x and *subtract* from this number the number which cross in the negative x-direction. As we have seen many times, the number which cross a surface area in a time ΔT is given by the number which start the interval ΔT in a volume which extends the distance $v\,\Delta T$ from the plane. (Note that v, here, is the actual molecular velocity, not the drift velocity.)

We shall simplify our algebra by giving our surface one unit of area. Then the number of special molecules which pass from left to right (taking the $+x$-direction to the right) is $n_- v\,\Delta T$, where n_- is the number of special molecules per unit volume to the left (within a factor of 2 or so, but we are ignoring such factors!). The number which cross from right to left is, similarly, $n_+ v\,\Delta T$, where n_+ is the number density of special molecules on the right-hand side of the plane. If we call the molecular current J, by which we mean the net flow of molecules per unit area per unit time, we have

$$J = \frac{n_- v\,\Delta T - n_+ v\,\Delta T}{\Delta T}, \tag{43.21}$$

or

$$J = (n_- - n_+)v. \tag{43.22}$$

What shall we use for n_- and n_+? When we say "the density on the left," how *far* to the left do we mean? We should choose the density at the place from which the molecules started their "flight," because the number which *start* such

trips is determined by the number present at that place. So by n_- we should mean the density a distance to the left equal to the mean free path l, and by n_+, the density at the distance l to the right of our imaginary surface.

It is convenient to consider that the distribution of our special molecules in space is described by a continuous function of x, y, and z which we shall call n_a. By $n_a(x, y, z)$ we mean the number density of special molecules in a small volume element centered on (x, y, z). In terms of n_a we can express the difference $(n_+ - n_-)$ as

$$(n_+ - n_-) = \frac{dn_a}{dx}\,\Delta x = \frac{dn_a}{dx}\cdot 2l. \tag{43.23}$$

Substituting this result in Eq. (43.22) and neglecting the factor of 2, we get

$$J_x = -lv\,\frac{dn_a}{dx}. \tag{43.24}$$

We have found that the flow of special molecules is proportional to the derivative of the density, or to what is sometimes called the "gradient" of the density.

It is clear that we have made several rough approximations. Besides various factors of two we have left out, we have used v where we should have used v_x, and we have assumed that n_+ and n_- refer to places at the perpendicular distance l from our surface, whereas for those molecules which do not travel perpendicular to the surface element, l should correspond to the *slant* distance from the surface. All of these refinements can be made; the result of a more careful analysis shows that the right-hand side of Eq. (43.24) should be multiplied by 1/3. So a better answer is

$$J_x = -\frac{lv}{3}\,\frac{dn_a}{dx}. \tag{43.25}$$

Similar equations can be written for the currents in the y- and z-directions.

The current J_x and the density gradient dn_a/dx can be measured by macroscopic observations. Their experimentally determined ratio is called the "diffusion coefficient," D. That is,

$$J_x = -D\,\frac{dn_a}{dx}. \tag{43.26}$$

We have been able to show that for a gas we expect

$$D = \tfrac{1}{3}lv. \tag{43.27}$$

So far in this chapter we have considered two distinct processes: *mobility*, the drift of molecules due to "outside" forces; and *diffusion*, the spreading determined only by the internal forces, the random collisions. There is, however, a relation between them, since they both depend basically on the thermal motions, and the mean free path l appears in both calculations.

If, in Eq. (43.25), we substitute $l = v\tau$ and $\tau = \mu m$, we have

$$J_x = -\tfrac{1}{3}mv^2\mu\,\frac{dn_a}{dx}. \tag{43.28}$$

But mv^2 depends only on the temperature. We recall that

$$\tfrac{1}{2}mv^2 = \tfrac{3}{2}kT, \tag{43.29}$$

so

$$J_x = -\mu kT\,\frac{dn_a}{dx}. \tag{43.30}$$

We find that D, the *diffusion* coefficient, is just kT times μ, the *mobility* coefficient:

$$D = \mu kT. \tag{43.31}$$

And it turns out that the numerical coefficient in (43.31) is exactly right—no extra factors have to be thrown in to adjust for our rough assumptions. We can show,

in fact, that (43.31) must *always* be correct—even in complicated situations (for example, the case of a suspension in a liquid) where the details of our simple calculations would not apply at all.

To show that (43.31) must be correct in general, we shall derive it in a different way, using only our basic principles of statistical mechanics. Imagine a situation in which there is a gradient of "special" molecules, and we have a diffusion current proportional to the density gradient, according to Eq. (43.26). We now apply a force field in the x-direction, so that each special molecule feels the force F. According to the *definition* of the mobility μ there will be a drift velocity given by

$$v_{\text{drift}} = \mu F. \tag{43.32}$$

By our usual arguments, the *drift current* (the *net* number of molecules which pass a unit of area in a unit of time) will be

$$J_{\text{drift}} = n_a v_{\text{drift}}, \tag{43.33}$$

or

$$J_{\text{drift}} = n_a \mu F. \tag{43.34}$$

We now *adjust* the force F so that the drift current due to F just *balances* the diffusion, so that there is *no net flow* of our special molecules. We have $J_x + J_{\text{drift}} = 0$, or

$$D \frac{dn_a}{dx} = n_a \mu F. \tag{43.35}$$

Under the "balance" conditions we find a steady (with time) gradient of density given by

$$\frac{dn_a}{dx} = \frac{n_a \mu F}{D}. \tag{43.36}$$

But notice! We are describing an *equilibrium* condition, so our *equilibrium* laws of statistical mechanics apply. According to these laws the probability of finding a molecule at the coordinate x is proportional to $e^{-U/kT}$, where U is the potential energy. In terms of the number density n_a, this means that

$$n_a = n_0 e^{-U/kT}. \tag{43.37}$$

If we differentiate (43.37) with respect to x, we find

$$\frac{dn_a}{dx} = -n_0 e^{-U/kT} \cdot \frac{1}{kT} \frac{dU}{dx}, \tag{43.38}$$

or

$$\frac{dn_a}{dx} = -\frac{n_a}{kT} \frac{dU}{dx}. \tag{43.39}$$

In our situation, since the force F is in the x-direction, the potential energy U is just $-Fx$, and $-dU/dx = F$. Equation (43.39) then gives

$$\frac{dn_a}{dx} = \frac{n_a F}{kT}. \tag{43.40}$$

[This is just exactly Eq. (40.2), from which we deduced $e^{-U/kT}$ in the first place, so we have come in a circle]. Comparing (43.40) with (43.36), we get exactly Eq. (43.31). We have shown that Eq. (43.31), which gives the diffusion current in terms of the mobility, has the correct coefficient and is very generally true. Mobility and diffusion are intimately connected. This relation was first deduced by Einstein.

43–6 Thermal conductivity

The methods of the kinetic theory that we have been using above can be used also to compute the *thermal conductivity* of a gas. If the gas at the top of a container is hotter than the gas at the bottom, heat will flow from the top to the bottom. (We think of the top being hotter because otherwise convection currents would be

set up and the problem would no longer be one of heat *conduction*.) The transfer of heat from the hotter gas to the colder gas is by the diffusion of the "hot" molecules—those with more energy—downward and the diffusion of the "cold" molecules upward. To compute the flow of thermal energy we can ask about the energy carried downward across an element of area by the downward-moving molecules, and about the energy carried upward across the surface by the upward-moving molecules. The difference will give us the net downward flow of energy.

The thermal conductivity κ is defined as the ratio of the rate at which thermal energy is carried across a unit surface area, to the temperature gradient:

$$\frac{1}{A}\frac{dQ}{dt} = -\kappa\frac{dT}{dz}. \tag{43.41}$$

Since the details of the calculations are quite similar to those we have done above in considering the flow of electric current in an ionized gas, we shall leave it as an exercise for the reader to show that

$$\kappa = \frac{knlv}{\gamma - 1}, \tag{43.42}$$

where $(\gamma - 1)kT$ is the average energy of a molecule at the temperature T.

If we use our relation $nl\sigma_c = 1$, the heat conductivity can be written as

$$\kappa = \frac{1}{\gamma - 1}\frac{kv}{\sigma_c}. \tag{43.43}$$

We have a rather surprising result. We know that the average velocity of gas molecules depends on the temperature but *not on the density*. We expect σ_c to depend only on the *size* of the molecules. So our simple result says that the thermal conductivity κ (and therefore the *rate* of flow of heat in any particular circumstance) is independent of the *density* of the gas! The change in the number of "carriers" of energy with a change in density is just compensated by the larger distance the "carriers" can go between collisions.

One may ask: "Is the heat flow independent of the gas density in the limit as the density goes to zero? When there is no gas at all?" Certainly not! The formula (43.43) was derived, as were all the others in this chapter, under the assumption that the mean free path between collisions is much smaller than any of the dimensions of the container. Whenever the gas density is so low that a molecule has a fair chance of crossing from one wall of its container to the other without having a collision, none of the calculations of this chapter apply. We must in such cases go back to kinetic theory and calculate again the details of what will occur.

44

The Laws of Thermodynamics

44–1 Heat engines; the first law

44–2 The second law

44–3 Reversible engines

44–4 The efficiency of an ideal engine

44–5 The thermodynamic temperature

44–6 Entropy

44–1 Heat engines; the first law

So far we have been discussing the properties of matter from the atomic point of view, trying to understand roughly what will happen if we suppose that things are made of atoms obeying certain laws. However, there are a number of relationships among the properties of substances which can be worked out without consideration of the detailed structure of the materials. The determination of the relationships among the various properties of materials, without knowing their internal structure, is the subject of *thermodynamics*. Historically, thermodynamics was developed before an understanding of the internal structure of matter was achieved.

To give an example: we know from the kinetic theory that the pressure of a gas is caused by molecular bombardment, and we know that if we heat a gas, so that the bombardment increases, the pressure must increase. Conversely, if the piston in a container of the gas is moved inward against the force of bombardment, the energy of the molecules bombarding the piston will increase, and consequently the temperature will increase. So, on the one hand, if we increase the temperature at a given volume, we increase the pressure. On the other hand, if we compress the gas, we will find that the temperature will rise. From the kinetic theory, one can derive a quantitative relationship between these two effects, but instinctively one might guess that they are related in some necessary fashion which is independent of the details of the collisions.

Let us consider another example. Many people are familiar with this interesting property of rubber: If we take a rubber band and pull it, it gets warm. If one puts it between his lips, for example, and pulls it out, he can feel a distinct warming, and this warming is reversible in the sense that if he relaxes the rubber band quickly while it is between his lips, it is distinctly cooled. That means that when we stretch a rubber band it heats, and when we release the tension of the band it cools. Now our instincts might suggest that if we heated a band, it might pull: that the fact that pulling a band heats it might imply that heating a band should cause it to contract. And, in fact, if we apply a gas flame to a rubber band holding a weight, we will see that the band contracts abruptly (Fig. 44–1). So it is true that when we heat a rubber band it pulls, and this fact is definitely related to the fact that when we release the tension of it, it cools.

The internal machinery of rubber that causes these effects is quite complicated. We will describe it from a molecular point of view to some extent, although our main purpose in this chapter is to understand the relationship of these effects independently of the molecular model. Nevertheless, we can show from the molecular model that the effects are closely related. One way to understand the behavior of rubber is to recognize that this substance consists of an enormous tangle of long chains of molecules, a kind of "molecular spaghetti," with one extra complication: between the chains there are cross-links—like spaghetti that is sometimes welded together where it crosses another piece of spaghetti—a grand tangle. When we pull out such a tangle, some of the chains tend to line up along the direction of the pull. At the same time, the chains are in thermal motion, so they hit each other continually. It follows that such a chain, if stretched, would not by itself remain stretched, because it would be hit from the sides by the other chains and other molecules, and would tend to kink up again. So the real reason why a rubber band tends to contract is this: when one pulls it out, the chains are lengthwise, and the thermal agitations of the molecules on the sides of the chains tend

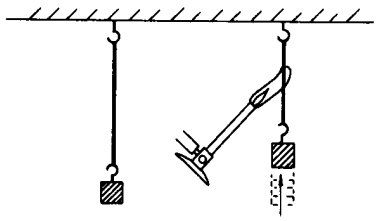


Fig. 44–1. The heated rubber band.

to kink the chains up, and make them shorten. One can then appreciate that if the chains are held stretched and the temperature is increased, so that the vigor of the bombardment on the sides of the chains is also increased, the chains tend to pull in, and they are able to pull a stronger weight when heated. If, after being stretched for a time, a rubber band is allowed to relax, each chain becomes soft, and the molecules striking it lose energy as they pound into the relaxing chain. So the temperature falls.

We have seen how these two process, contraction when heated and cooling during relaxation, can be related by the kinetic theory, but it would be a tremendous challenge to determine from the theory the precise relationship between the two. We would have to know how many collisions there were each second and what the chains look like, and we would have to take account of all kinds of other complications. The detailed mechanism is so complex that we cannot, by kinetic theory, really determine exactly what happens; still, a definite relation between the two effects we observe can be worked out without knowing anything about the internal machinery!

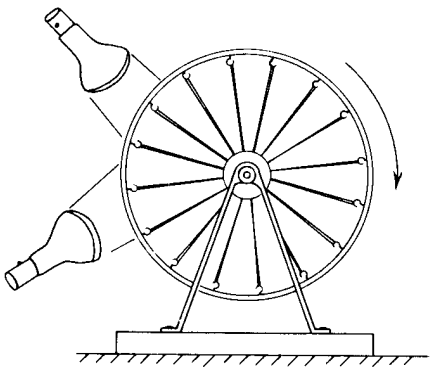


Fig. 44–2. The rubber-band heat engine.

The whole subject of thermodynamics depends essentially upon the following kind of consideration: because a rubber band is "stronger" at higher temperatures than it is at lower temperatures, it ought to be possible to lift weights, and to move them around, and thus to do work with heat. In fact, we have already seen experimentally that a heated rubber band can lift a weight. The study of the way that one does work with heat is the beginning of the science of thermodynamics. Can we make an engine which uses the heating effect on a rubber band to do work? One can make a silly looking engine that does just this. It consists of a bicycle wheel in which all the spokes are rubber bands (Fig. 44–2). If one heats the rubber bands on one side of the wheel with a pair of heat lamps, they become "stronger" than the rubber bands on the other side. The center of gravity of the wheel will be pulled to one side, away from the bearing, so that the wheel turns. As it turns, cool rubber bands move toward the heat, and the heated bands move away from the heat and cool, so that the wheel turns slowly so long as the heat is applied. The efficiency of this engine is extremely low. Four hundred watts of power pour into the two lamps, but it is just possible to lift a fly with such an engine! An interesting question, however, is whether we can get heat to do the work in more efficient ways.

In fact, the science of thermodynamics began with an analysis, by the great engineer Sadi Carnot, of the problem of how to build the best and most efficient engine, and this constitutes one of the few famous cases in which engineering has contributed fundamentally to physical theory. Another example that comes to mind is the more recent analysis of information theory by Claude Shannon. These two analyses, incidentally, turn out to be closely related.

Now the way a steam engine ordinarily operates is that heat from a fire boils some water, and the steam so formed expands and pushes on a piston which makes a wheel go around. So the steam pushes the piston—what then? One has to finish the job: a stupid way to complete the cycle would be to let the steam escape into the air, for then one has to keep supplying water. It is cheaper—more efficient—to let the steam go into another box, where it is condensed by cool water, and then pump the water back into the boiler, so that it circulates continuously. Heat is thus supplied to the engine and converted into work. Now would it be better to use alcohol? What property should a substance have so that it makes the best possible engine? That was the question to which Carnot addressed himself, and one of the by-products was the discovery of the type of relationship that we have just explained above.

The results of thermodynamics are all contained implicitly in certain apparently simple statements called the *laws of thermodynamics*. At the time when Carnot lived, the first law of thermodynamics, the conservation of energy, was not known. Carnot's arguments were so carefully drawn, however, that they are valid even though the first law was not known in his time! Some time afterwards, Clausius made a simpler derivation that could be understood more easily than Carnot's very subtle reasoning. But it turned out that Clausius assumed, not the conserva-

tion of energy in general, but that *heat* was conserved according to the caloric theory, which was later shown to be false. So it has often been said that Carnot's logic was wrong. But his logic was quite correct. Only Clausius' simplified version, that everybody read, was incorrect.

The so-called second law of thermodynamics was thus discovered by Carnot before the first law! It would be interesting to give Carnot's argument that did not use the first law, but we shall not do so because we want to learn physics, not history. We shall use the first law from the start, in spite of the fact that a great deal can be done without it.

Let us begin by stating the first law, the conservation of energy: if one has a system and puts heat into it, and does work on it, then its energy is increased by the heat put in and the work done. We can write this as follows: The heat Q put into the system, plus the work W done on the system, is the increase in the energy U of the system; the latter energy is sometimes called the internal energy:

$$\text{Change in } U = Q + W. \tag{44.1}$$

The change in U can be represented as adding a little heat ΔQ and adding a little work ΔW:

$$\Delta U = \Delta Q + \Delta W, \tag{44.2}$$

which is a differential form of the same law. We know that very well, from an earlier chapter.

44–2 The second law

Now, what about the second law of thermodynamics? We know that if we do work against friction, say, the work lost to us is equal to the heat produced. If we do work in a room at temperature T, and we do the work slowly enough, the room temperature does not change much, and we have converted work into heat at a given temperature. What about the reverse possibility? Is it possible to convert the heat back into work at a given temperature? The second law of thermodynamics asserts that it is not. It would be very convenient to be able to convert heat into work merely by reversing a process like friction. If we consider only the conservation of energy, we might think that heat energy, such as that in the vibrational motions of molecules, might provide a goodly supply of useful energy. But Carnot assumed that it is impossible to extract the energy of heat at a single temperature. In other words, if the whole world were at the same temperature, one could not convert any of its heat energy into work: while the process of making work go into heat can take place at a given temperature, one cannot reverse it to get the work back again. Specifically, Carnot assumed that heat cannot be taken in at a certain temperature and converted into work *with no other change* in the system or the surroundings.

That last phrase is very important. Suppose we have a can of compressed air at a certain temperature, and we let the air expand. It can do work; it can make hammers go, for example. It cools off a little in the expansion, but if we had a big sea, like the ocean, at a given temperature—a heat reservoir—we could warm it up again. So we have taken the heat out of the sea, and we have done work with the compressed air. But Carnot was not wrong, because *we did not leave everything as it was.* If we recompress the air that we let expand, we will find we are doing extra work, and when we are finished we will discover that we not only got no work out of the system at temperature T, but we actually put some in. We must talk only about situations in which the *net result* of the whole process is to take heat away and convert it into work, just as the net result of the process of doing work against friction is to take work and convert it into heat. If we move in a circle, we can bring the system back precisely to its starting point, with the net result that we did work against friction and produced heat. Can we reverse the process? Turn a switch, so that everything goes backwards, so the friction does work against us, and cools the sea? According to Carnot: no! So let us suppose that this is impossible.

If it were possible it would mean, among other things, that we could take heat out of a cold body and put it into a hot body at no cost, as it were. Now we know it is natural that a hot thing can warm up a cool thing; if we simply put a hot body and a cold one together, and change nothing else, our experience assures us that it is not going to happen that the hot one gets hotter, and the cold one gets colder! But if we could obtain work by extracting the heat out of the ocean, say, or from anything else at a single temperature, then that work could be converted back into heat by friction at some other temperature. For instance, the other arm of a working machine could be rubbing something that is already hot. The net result would be to take heat from a "cold" body, the ocean, and to put it into a hot body. Now, the hypothesis of Carnot, the second law of thermodynamics, is sometimes stated as follows: heat cannot, of itself, flow from a cold to a hot object. But, as we have just seen, these two statements are equivalent: first, that one cannot devise a process whose only result is to convert heat to work at a single temperature, and second, that one cannot make heat flow by itself from a cold to a hot place. We shall mostly use the first form.

Carnot's analysis of heat engines is quite similar to the argument that we gave about weight-lifting engines in our discussion of the conservation of energy in Chapter 4. In fact, that argument was patterned after Carnot's argument about heat engines, and so the present treatment will sound very much the same.

Suppose we build a heat engine that has a "boiler" somewhere at a temperature T_1. A certain heat Q_1 is taken from the boiler, the steam engine does some work W, and it then delivers some heat Q_2 into a "condenser" at another temperature T_2 (Fig. 44–3). Carnot did not say how much heat, because he did not know the first law, and he did not use the law that Q_2 was equal to Q_1 because he did not believe it. Although everybody thought that, according to the caloric theory, the heats Q_1 and Q_2 would have to be the same, Carnot did not say they were the same—that is part of the cleverness of his argument. If we do use the first law, we find that the heat delivered, Q_2, is the heat Q_1 that was put in minus the work W that was done:

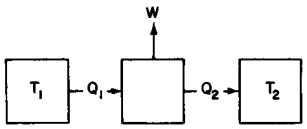


Fig. 44–3. Heat engine.

$$Q_2 = Q_1 - W. \tag{44.3}$$

(If we have some kind of cyclic process where water is pumped back into the boiler after it is condensed, we will say that we have heat Q_1 absorbed and work W done, during each cycle, for a certain amount of water that goes around the cycle.)

Now we shall build another engine, and see if we cannot get more work from the same amount of heat being delivered at the temperature T_1, with the condenser still at the temperature T_2. We shall use the same amount of heat Q_1 from the boiler, and we shall try to get more work than we did out of the steam engine, perhaps by using another fluid, such as alcohol.

44–3 Reversible engines

Now we must analyze our engines. One thing is clear: we will lose something if the engines contain devices in which there is friction. The best engine will be a frictionless engine. We assume, then, the same idealization that we did when we studied the conservation of energy; that is, a perfectly frictionless engine.

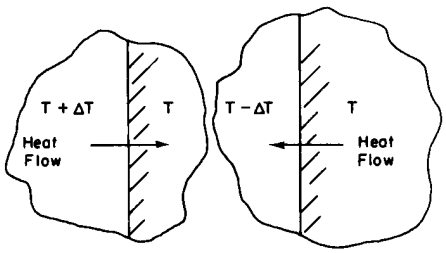


Fig. 44–4. Reversible heat transfer.

We must also consider the analog of frictionless motion, "frictionless" heat transfer. If we put a hot object at a high temperature against a cold object, so that the heat flows, then it is not possible to make that heat flow in a reverse direction by a very small change in the temperature of either object. But when we have a practically frictionless machine, if we push it with a little force one way, it goes that way, and if we push it with a little force the other way, it goes the other way. We need to find the analog of frictionless motion: heat transfer whose direction we can reverse with only a tiny change. If the difference in temperature is finite, that is impossible, but if one makes sure that heat flows always between two things at essentially the same temperature, with just an infinitesimal difference to make it flow in the desired direction, the flow is said to be reversible (Fig. 44–4). If we heat

the object on the left a little, heat will flow to the right; if we cool it a little, heat will flow to the left. So we find that the ideal engine is a so-called *reversible* engine, in which every process is reversible in the sense that, by minor changes, infinitesimal changes, we can make the engine go in the opposite direction. That means that nowhere in the machine must there be any appreciable friction, and nowhere in the machine must there be any place where the heat of the reservoirs, or the flame of the boiler, is in direct contact with something definitely cooler or warmer.

Let us now consider an idealized engine in which all the processes are reversible. To show that such a thing is possible in principle, we will give an example of an engine cycle which may or may not be practical, but which is at least reversible, in the sense of Carnot's idea. Suppose that we have a gas in a cylinder equipped with a frictionless piston. The gas is not necessarily a perfect gas. The fluid does not even have to be a gas, but to be specific let us say we do have a perfect gas. Also, suppose that we have two heat pads, T_1 and T_2—great big things that have definite temperatures, T_1 and T_2. We will suppose in this case that T_1 is higher than T_2. Let us first heat the gas and at the same time expand it, while it is in contact with the heat pad at T_1. As we do this, pulling the piston out very slowly as the heat flows into the gas, we will make sure that the temperature of the gas never gets very far from T_1. If we pull the piston out too fast, the temperature of the gas will fall too much below T_1 and then the process will not be quite reversible, but if we pull it out slowly enough, the temperature of the gas will never depart much from T_1. On the other hand, if we push the piston back slowly, the temperature would be only infinitesimally higher than T_1, and the heat would pour back. We see that such an isothermal (constant-temperature) expansion, done slowly and gently enough, is a reversible process.

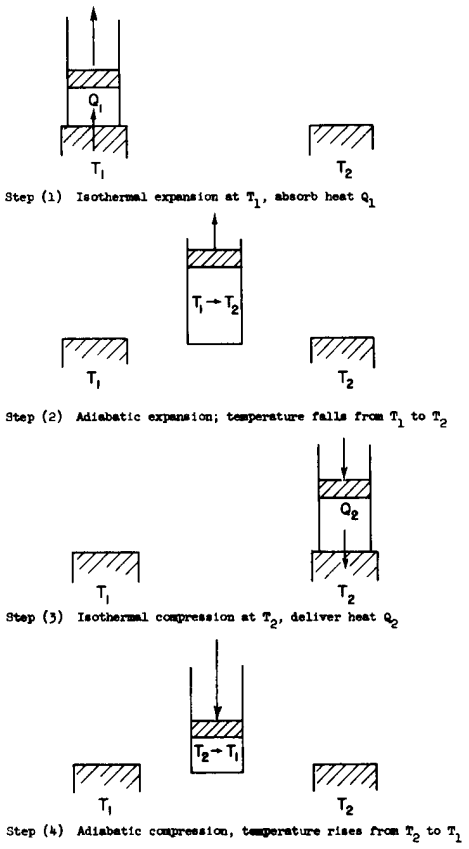


Fig. 44–5. Steps in Carnot cycle.

To understand what we are doing, we shall use a plot (Fig. 44–6) of the pressure of the gas against its volume. As the gas expands, the pressure falls. The curve marked (1) tells us how the pressure and volume change if the temperature is kept fixed at the value T_1. For an ideal gas this curve would be $PV = NkT_1$. During an isothermal expansion the pressure falls as the volume increases until we stop at the point b. At the same time, a certain heat Q_1 must flow into the gas from the reservoir, for if the gas were expanded without being in contact with the reservoir it would cool off, as we already know. Having completed the isothermal expansion, stopping at the point b, let us take the cylinder away from the reservoir and continue the expansion. This time we permit no heat to enter the cylinder. Again we perform the expansion slowly, so there is no reason why we cannot reverse it, and we again assume there is no friction. The gas continues to expand and the temperature falls, since there is no longer any heat entering the cylinder.

We let the gas expand, following the curve marked (2), until the temperature falls to T_2, at the point marked c. This kind of expansion, made without adding heat, is called an *adiabatic* expansion. For an ideal gas, we already know that curve (2) has the form PV^γ = constant, where γ is a constant greater than 1, so that the adiabatic curve has a more negative slope than the isothermal curve. The gas cylinder has now reached the temperature T_2, so that if we put it on the heat pad at temperature T_2 there will be no irreversible changes. Now we slowly compress the gas while it is in contact with the reservoir at T_2, following the curve marked (3) (Fig. 44–5, Step 2). Because the cylinder is in contact with the reservoir, the temperature does not rise, but heat Q_2 flows from the cylinder into the reservoir at the temperature T_2. Having compressed the gas isothermally along curve (3) to the point d, we remove the cylinder from the heat pad at temperature T_2 and compress it still further, without letting any heat flow out. The temperature will rise, and the pressure will follow the curve marked (4). If we carry out each step properly, we can return to the point a at temperature T_1 where we started, and repeat the cycle.

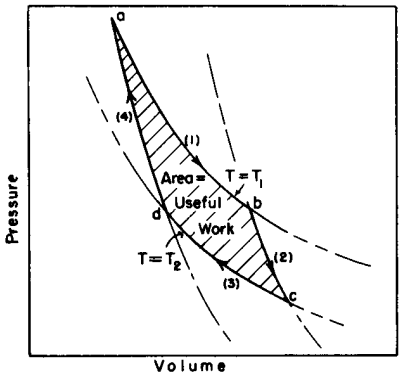


Fig. 44–6. The Carnot cycle.

We see that on this diagram we have carried the gas around a complete cycle, and during one cycle we have put Q_1 in at temperature T_1, and have removed Q_2 at temperature T_2. Now the point is that this cycle is reversible, so that we could represent all the steps the other way around. We could have gone backwards instead of forwards: we could have started at point a, at temperature T_1, expanded

along the curve (4), expanded further at the temperature T_2, absorbing heat Q_2, and so on, going around the cycle backward. If we go around the cycle in one direction, we must do work on the gas; if we go in the other direction, the gas does work on us.

Incidentally, it is easy to find out what the total amount of work is, because the work during any expansion is the pressure times the change in volume, $\int P\,dV$. On this particular diagram, we have plotted P vertically and V horizontally. So if we call the vertical distance y and the horizontal distance x, this is $\int y\,dx$—in other words, the area under the curve. So the area under each of the numbered curves is a measure of the work done by or on the gas in the corresponding step. It is easy to see that the net work done is the shaded area of the picture.

Now that we have given a single example of a reversible machine, we shall suppose that other such engines are also possible. Let us assume that we have a reversible engine A which takes Q_1 at T_1, does work W, and delivers some heat at T_2. Now let us assume we have any other engine B, made by man, already designed or not yet invented, made of rubber bands, steam, or whatever, reversible or not, which is designed so that it takes in the same amount of heat Q_1 at T_1, and rejects the heat at the lower temperature T_2 (Fig. 44–7). Assume that engine B does some work, W'. Now we shall show that W' is not greater than W—that no engine can do more work than a reversible one. Why? Suppose that, indeed, W' were bigger than W. Then we could take the heat Q_1 out of the reservoir at T_1, and with engine B we could do work W' and deliver some heat to the reservoir at T_2; we do not care how much. That done, we could save some of the work W', which is supposed to be greater than W; we could use a part of it, W, and save the remainder, $W' - W$, for useful work. With the work W we could run engine A backwards *because it is a reversible engine*. It will absorb some heat from the reservoir at T_2 and deliver Q_1 back to the reservoir at T_1. After this double cycle, the net result would be that we would have put everything back the way it was before, and we would have done some excess work, namely $W' - W$, and *all* we would have done would be to extract energy from the reservoir at T_2! We were careful to restore the heat Q_1 to the reservoir at T_1. So that reservoir can be small and "inside" our combined machine $A + B$, whose net effect is therefore to extract a net heat $W' - W$ from the reservoir at T_2 and convert it into work. But to obtain useful work from a reservoir at a single temperature *with no other changes* is impossible according to Carnot's postulate; it cannot be done. Therefore no engine which absorbs a given amount of heat from a higher temperature T_1 and delivers it at the temperature T_2 can do more work than a reversible engine operating under the same temperature conditions.

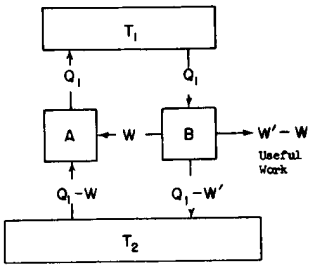


Fig. 44–7. Reversible engine A being driven backwards by engine B.

Now suppose that engine B is also reversible. Then, of course, not only must W' be not greater than W, but now we can reverse the argument and show that W cannot be greater than W'. So, if both engines are reversible they must both do the same amount of work, and we thus come to Carnot's brilliant conclusion: that if an engine is reversible, it makes no difference how it is designed, because the amount of work one will obtain if the engine absorbs a given amount of heat at temperature T_1 and delivers heat at some other temperature T_2 *does not depend on the design of the engine*. It is a property of the world, not a property of a particular engine.

If we could find out what the law is that determines how much work we obtain when we absorb the heat Q_1 at T_1 and deliver heat at T_2, this quantity would be a universal thing, independent of the substance. Of course if we knew the properties of a particular substance, we could work it out and then say that all other substances must give the same amount of work in a reversible engine. That is the key idea, the clue by which we can find the relationship between how much, for instance, a rubber band contracts when we heat it, and how much it cools when we let it contract. Imagine that we put that rubber band in a reversible machine, and that we make it go around a reversible cycle. The net result, the total amount of work done, is that universal function, that great function which is independent of substance. So we see that a substance's properties must be limited in a certain way; one

cannot make up anything he wants, or he would be able to invent a substance which he could use to produce more than the maximum allowable work when he carried it around a reversible cycle. This principle, this limitation, is the only real rule that comes out of the thermodynamics.

44–4 The efficiency of an ideal engine

Now we shall try to find the law which determines the work W as a function of Q_1, T_1, and T_2. It is clear that W is proportional to Q_1, for if we consider two reversible engines in parallel, both working together and both double engines, the combination is also a reversible engine. If each one absorbed heat Q_1, the two together absorb $2Q_1$ and the work done is $2W$, and so on. So it is not unreasonable that W is proportional to Q_1.

Now the next important step is to find this universal law. We can, and will, do so by studying a reversible engine with the one particular substance whose laws we know, a perfect gas. It is also possible to obtain the rule by a purely logical argument, using no particular substance at all. This is one of the very beautiful pieces of reasoning in physics and we are reluctant not to show it to you, so for those who would like to see it we shall discuss it in just a moment. But first we shall use the much less abstract and simpler method of direct calculation for a perfect gas.

We need only obtain formulas for Q_1 and Q_2 (for W is just $Q_1 - Q_2$), the heats exchanged with the reservoirs during the isothermal expansion or contraction. For example, how much heat Q_1 is absorbed from the reservoir at temperature T_1 during the isothermal expansion [marked (1) in Fig. 44–6] from point a, at pressure p_a, volume V_a, temperature T_1, to point b with pressure p_b, volume V_b, and the same temperature T_1? For a perfect gas each molecule has an energy that depends only on the temperature, and since the temperature and the number of molecules are the same at a and at b, the internal energy is the same. *There is no change in U;* all the work done by the gas,

$$W = \int_a^b p\,dV\,,$$

during the expansion is energy Q_1 taken from the reservoir. During the expansion, $pV = NkT_1$, or

$$p = \frac{NkT_1}{V}$$

or

$$Q_1 = \int_a^b p\,dV = \int_a^b NkT_1\,\frac{dV}{V} \tag{44.4}$$

or

$$Q_1 = NkT_1 \ln \frac{V_b}{V_a}$$

is the heat taken from the reservoir at T_1. In the same way, for the compression at T_2 [curve (3) of Fig. 44–6] the heat delivered to the reservoir at T_2 is

$$Q_2 = NkT_2 \ln \frac{V_c}{V_d}\cdot \tag{44.5}$$

To finish our analysis we need only find a relation between V_c/V_d and V_b/V_a. This we do by noting that (2) is an adiabatic expansion from b to c, during which pV^γ is a constant. Since $pV = NkT$, we can write this as $(pV)V^{\gamma-1} = \text{const}$ or, in terms of T and V, as $TV^{\gamma-1} = \text{const}$, or

$$T_1V_b^{\gamma-1} = T_2V_c^{\gamma-1}. \tag{44.6}$$

Likewise, since (4), the expansion from d to a, is also adiabatic, we find

$$T_1V_a^{\gamma-1} = T_2V_d^{\gamma-1}. \tag{44.6a}$$

If we divide this equation by the previous one, we find that V_b/V_a must equal V_c/V_d, so the ln's in (44.4) and (44.5) are equal, and that

$$\frac{Q_1}{T_1} = \frac{Q_2}{T_2}. \tag{44.7}$$

This is the relation we were seeking. Although proved for a perfect gas engine, we know it must be true *for any reversible engine at all.*

Now we shall see how this universal law could also be obtained by logical argument, without knowing the properties of any specific substances, as follows. Suppose that we have three engines and three temperatures, let us say T_1, T_2, and T_3. Let one engine absorb heat Q_1 from the temperature T_1 and do a certain amount of work W_{13}, and let it deliver heat Q_3 to the temperature T_3 (Fig. 44–8). Let another engine run backwards between T_2 and T_3. Suppose that we let the second engine be of such a size that it will absorb the same heat Q_3, and deliver the heat Q_2. We will have to put a certain amount of work, W_{32}, into it—negative because the engine is running backwards. When the first machine goes through a cycle, it absorbs heat Q_1 and delivers Q_3 at the temperature T_3; then the second machine takes the same heat Q_3 out of the reservoir at the temperature T_3 and delivers it into the reservoir at temperature T_2. Therefore the net result of the two machines in tandem is to take the heat Q_1 from T_1, and deliver Q_2 at T_2. The two machines are thus equivalent to a third one, which absorbs Q_1 at T_1, does work W_{12}, and delivers heat Q_2 at T_2, because $W_{12} = W_{13} - W_{32}$, as one can immediately show from the first law, as follows:

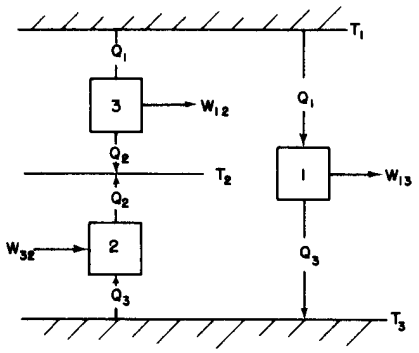


Fig. 44–8. Engines 1 and 2 together are equivalent to engine 3.

$$W_{13} - W_{32} = (Q_1 - Q_3) - (Q_2 - Q_3) = Q_1 - Q_2 = W_{12}. \tag{44.8}$$

We can now obtain the laws which relate the efficiencies of the engines, because there clearly must be some kind of relationship between the efficiencies of engines running between the temperatures T_1 and T_3, and between T_2 and T_3, and between T_1 and T_2.

We can make the argument very clear in the following way: We have just seen that we can always relate the heat absorbed at T_1 to the heat delivered at T_2 by finding the heat delivered at some other temperature T_3. Therefore we can get all the engines' properties if we introduce a standard temperature, analyzing everything with that standard temperature. In other words, if we knew the efficiency of an engine running between a certain temperature T and a certain arbitrary standard temperature, then we could work out the efficiency for any other difference in temperature. Because we assume we are using only reversible engines, we can work from the initial temperature down to the standard temperature and back up to the final temperature again. We shall define the standard temperature arbitrarily as *one degree.* We shall also adopt a special symbol for the heat which is delivered at this standard temperature: we shall call it Q_S. In other words, when a reversible engine absorbs the heat Q at temperature T, it will deliver, at the unit temperature, a heat Q_S. *If one engine, absorbing heat Q_1 at T_1, delivers the heat Q_S at one degree, and if an engine absorbing heat Q_2 at temperature T_2 will also deliver the same heat Q_S at one degree, then it follows that an engine which absorbs heat Q_1 at the temperature T_1 will deliver heat Q_2 if it runs between T_1 and T_2,* as we have already proved by considering engines running between three temperatures. So all we really have to do is to find how much heat Q_1 we need to put in at the temperature T_1 in order to deliver a certain amount of heat Q_S at the unit temperature. If we discover that, we have everything. The heat Q, of course, is a function of the temperature T. It is easy to see that the heat must increase as the temperature increases, for we know that it takes work to run an engine backwards and deliver heat at a higher temperature. It is also easy to see that the heat Q_1 must be proportional to Q_S. So the great law is something like this: for a given amount of heat Q_S delivered at one degree from an engine running at temperature T degrees, the heat Q absorbed must be that amount Q_S times some increasing function of the temperature:

$$Q = Q_S f(T). \tag{44.9}$$

44–5 The thermodynamic temperature

At this stage we are not going to try to find the formula for the above increasing function of the temperature in terms of our familiar mercury temperature scale, but instead *we shall define temperature by a new scale.* At one time "the temperature" was defined arbitrarily by dividing the expansion of water into even degrees of a certain size. But when one then measures temperature with a mercury thermometer, one finds that the degrees are no longer even. But *now we can make a definition of temperature which is independent of any particular substance.* We can use that function $f(T)$, which does not depend on what device we use, because the efficiency of these reversible engines is independent of their working substances. Since the function we found is rising with temperature, we will define the *function itself* as the temperature, measured in units of the standard one-degree temperature, as follows:

$$Q = ST, \tag{44.10}$$

where

$$Q_S = S \cdot 1°. \tag{44.11}$$

This means that we can tell how hot an object is by finding out how much heat is absorbed by a reversible engine working between the temperature of the object and the unit temperature (Fig. 44–9). If seven times more heat is taken out of a boiler than is delivered at a one-degree condenser, the temperature of the boiler will be called seven degrees, and so forth. So, by measuring how much heat is absorbed at different temperatures, we determine the temperature. The temperature defined in this way is called the *absolute thermodynamic temperature*, and it is independent of the substance. We shall use this definition exclusively from now on.*

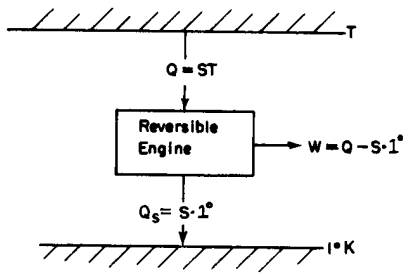


Fig. 44–9. Absolute thermodynamic temperature.

Now we see that when we have two engines, one working between T_1 and one degree, the other working between T_2 and one degree, delivering the same heat at unit temperature, then the heats absorbed must be related by

$$\frac{Q_1}{T_1} = S = \frac{Q_2}{T_2}. \tag{44.12}$$

But that means that if we have a single engine running between T_1 and T_2, then the result of the whole analysis, the grand finale, is that Q_1 is to T_1 as Q_2 is to T_2, if the engine absorbs energy Q_1 at temperature T_1 and delivers heat Q_2 at temperature T_2. Whenever the engine is reversible, this relationship between the heats must follow. That is all there is to it: that is the center of the universe of thermodynamics.

If this is all there is to thermodynamics, why is it considered such a difficult subject? In doing a problem involving a given mass of some substance, the condition of the substance at any moment can be described by telling what its temperature is and what its volume is. If we know the temperature and volume of a substance, and that the pressure is some function of the temperature and volume, then we know the internal energy. One could say, "I do not want to do it that way. Tell me the temperature and the pressure, and I will tell you the volume. I can think of the volume as a function of temperature and pressure, and the internal energy as a function of temperature and pressure, and so on." That is why thermodynamics is hard, because everyone uses a different approach. If we could only sit down once and decide on our variables, and stick to them, it would be fairly easy.

Now we start to make deductions. Just as $F = ma$ is the center of the universe in mechanics, and it goes on and on and on after that, in the same way the principle just found is all there is to thermodynamics. But can one make conclusions out of it?

---

* We have previously defined our scale of temperature in a different way, namely by stating that the mean kinetic energy of a molecule in a perfect gas is proportional to the temperature, or that the perfect gas law says pV is proportional to T. Is this new definition equivalent? Yes, since the final result (44.7) derived from the gas law is the same as that derived here. We shall discuss this point again in the next chapter.

We begin. To obtain our first conclusion, we shall combine both laws, the law of conservation of energy and this law which relates the heats Q_2 and Q_1, and we can easily obtain *the efficiency of a reversible engine*. From the first law, we have $W = Q_1 - Q_2$. According to our new principle,

$$Q_2 = \frac{T_2}{T_1} Q_1,$$

so the work becomes

$$W = Q_1\left(1 - \frac{T_2}{T_1}\right) = Q_1 \frac{T_1 - T_2}{T_1}, \tag{44.13}$$

which tells us the efficiency of the engine—how much work we get out of so much heat. The efficiency of an engine is proportional to the difference in the temperatures between which the engine runs, divided by the higher temperature:

$$\text{Efficiency} = \frac{W}{Q_1} = \frac{T_1 - T_2}{T_1}. \tag{44.14}$$

The efficiency cannot be greater than unity and the absolute temperature cannot be less than zero, absolute zero. So, since T_2 must be positive, the efficiency is always less than unity. That is our first conclusion.

44–6 Entropy

Equation (44.7) or (44.12) can be interpreted in a special way. Working always with reversible engines, a heat Q_1 at temperature T_1 is "equivalent" to Q_2 at T_2 if $Q_1/T_1 = Q_2/T_2$, in the sense that as one is absorbed the other is delivered. This suggests that if we call Q/T something, we can say: in a reversible process as much Q/T is absorbed as is liberated; there is no gain or loss of Q/T. This Q/T is called *entropy*, and we say "there is no net change in entropy in a reversible cycle." If T is 1°, then the entropy is $Q/1°$ or, as we symbolized it, $Q_S/1° = S$. Actually, S *is* the letter usually used for entropy, and it is numerically equal to the heat (which we have called Q_S) delivered to a 1°-reservoir (entropy is not itself a heat, it is heat divided by a temperature, hence it is measured in *joules per degree*).

Now it is interesting that besides the pressure, which is a function of the temperature and the volume, and the internal energy, which is a function of temperature and volume, we have found another quantity which is a function of the condition, i.e., the entropy of the substance. Let us try to explain how we compute it, and what we mean when we call it a "function of the condition." Consider the system in two different conditions, much as we had in the experiment where we did the adiabatic and isothermal expansions. (Incidentally, there is no need that a heat engine have only two reservoirs, it could have three or four different temperatures at which it takes in and delivers heats, and so on.) We can move around on a pV diagram all over the place, and go from one condition to another. In other words, we could say the gas is in a certain condition a, and then it goes over to some other condition, b, and we will require that this transition, made from a to b, be reversible. Now suppose that all along the path from a to b we have little reservoirs at different temperatures, so that the heat dQ removed from the substance at each little step is delivered to each reservoir at the temperature corresponding to that point on the path. Then let us connect all these reservoirs, by reversible heat engines, to a single reservoir at the unit temperature. When we are finished carrying the substance from a to b, we shall bring all the reservoirs back to their original condition. Any heat dQ that has been absorbed from the substance at temperature T has now been converted by a reversible machine, and a certain amount of entropy dS has been delivered at the unit temperature as follows:

$$dS = dQ/T. \tag{44.15}$$

Let us compute the total amount of entropy which has been delivered. The entropy difference, or the entropy needed to go from a to b by this particular

reversible transformation, is the total entropy, the total of the entropy taken out of the little reservoirs, and delivered at the unit temperature:

$$S_b - S_a = \int_a^b \frac{dQ}{T}. \tag{44.16}$$

The question is, does the entropy difference depend upon the path taken? There is more than one way to go from a to b. Remember that in the Carnot cycle we could go from a to c in Fig. 44–6 by first expanding isothermally and then adiabatically; or we could first expand adiabatically and then isothermally. So the question is whether the entropy change which occurs when we go from a to b in Fig. 44–10 is the same on one route as it is on another. *It must be the same,* because if we went all the way around the cycle, going forward on one path and backward on another, we would have a reversible engine, and there would be no loss of heat to the reservoir at unit temperature. In a totally reversible cycle, no heat must be taken from the reservoir at the unit temperature, so the entropy needed to go from a to b is the same over one path as it is over another. It is *independent of path,* and depends only on the endpoints. We can, therefore, say that there is a certain function, which we call the entropy of the substance, that depends only on the condition, i.e., only on the volume and temperature.

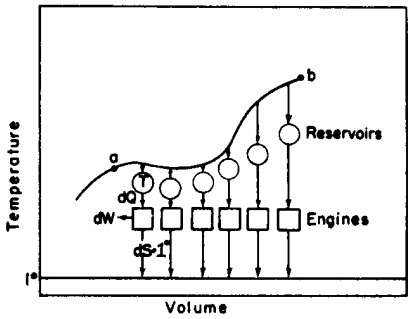


Fig. 44–10. Change in entropy during a reversible transformation.

We can find a function $S(V, T)$ which has the property that if we compute the change in entropy, as the substance is moved along any reversible path, in terms of the heat rejected at unit temperature, then

$$\Delta S = \int \frac{dQ}{T}, \tag{44.17}$$

where dQ is the heat removed from the substance at temperature T. This total entropy change is the difference between the entropy calculated at the initial and final points:

$$\Delta S = S(V_b, T_b) - S(V_a, T_a) = \int_a^b \frac{dQ}{T}. \tag{44.18}$$

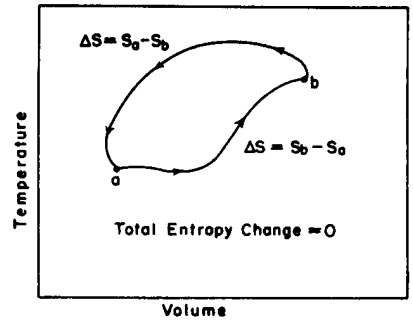


Fig. 44–11. Change in entropy in a completely reversible cycle.

This expression does not completely define the entropy, but rather only the *difference* of entropy between two different conditions. Only if we can evaluate the entropy for one special condition can we really define S absolutely.

For a long time it was believed that absolute entropy meant nothing—that only differences could be defined—but finally Nernst proposed what he called the *heat theorem,* which is also called the third law of thermodynamics. It is very simple. We will say what it is, but we will not explain why it is true. Nernst's postulate states simply that the entropy of any object at absolute zero is zero. We know of one case of T and V, namely $T = 0$, where S is zero; and so we can get the entropy at any other point.

To give an illustration of these ideas, let us calculate the entropy of a perfect gas. In an isothermal (and therefore reversible) expansion, $\int dQ/T$ is Q/T, since T is constant. Therefore (from 44.4) the change in entropy is

$$S(V_a, T) - S(V_b, T) = Nk \ln \frac{V_a}{V_b},$$

so $S(V, T) = Nk \ln V$ plus some function of T only. How does S depend on T? We know that for a reversible adiabatic expansion, *no heat is exchanged.* Thus the entropy does not change even though V changes, provided that T changes also, such that $TV^{\gamma-1} = \text{constant}$. Can you see that this implies that

$$S(V, T) = Nk\left[\ln V + \frac{1}{\gamma - 1} \ln T\right] + a,$$

where a is some constant independent of *both* V and T? [a is called the chemical constant. It depends on the gas in question, and may be determined experimentally from the Nernst theorem by measuring the heat liberated in cooling and condensing

the gas until it is brought to a solid (or for helium, a liquid) at 0°, by integrating $\int dQ/T$. It can also be determined theoretically by means of Planck's constant and quantum mechanics, but we shall not study it in this course.]

Now we shall remark on some of the properties of the entropy of things. We first remember that if we go along a reversible cycle from a to b, then the entropy of the substance will change by $S_b - S_a$. And we remember that as we go along the path, the entropy—the heat delivered at unit temperature—increases according to the rule $dS = dQ/T$, where dQ is the heat we remove from the substance when its temperature is T.

We already know that if we have a reversible *cycle*, the total entropy of everything is not changed, because the heat Q_1 absorbed at T_1 and the heat Q_2 delivered at T_2 correspond to equal and opposite changes in entropy, so that the net change in the entropy is zero. So for a reversible cycle there is no change in the entropy of anything, including the reservoirs. This rule may look like the conservation of energy again, but it is not; it applies only to reversible cycles. If we include irreversible cycles there is no law of conservation of entropy.

We shall give two examples. First, suppose that we do irreversible work on an object by friction, generating a heat Q on some object at temperature T. The entropy is increased by Q/T. The heat Q is equal to the work, and thus when we do a certain amount of work by friction against an object whose temperature is T, the entropy of the whole world increases by W/T.

Another example of irreversibility is this: If we put together two objects that are at different temperatures, say T_1 and T_2, a certain amount of heat will flow from one to the other by itself. Suppose, for instance, we put a hot stone in cold water. Then when a certain heat ΔQ is transferred from T_1 to T_2, how much does the entropy of the hot stone change? It decreases by $\Delta Q/T_1$. How much does the water entropy change? It increases by $\Delta Q/T_2$. The heat will, of course, flow only from the higher temperature T_1 to the lower temperature T_2, so that ΔQ is positive if T_1 is greater than T_2. So the change in entropy of the whole world is positive, and it is the difference of the two fractions:

$$\Delta S = \frac{\Delta Q}{T_2} - \frac{\Delta Q}{T_1}. \tag{44.19}$$

So the following proposition is true: in any process that is irreversible, the entropy of the whole world is increased. Only in reversible processes does the entropy remain constant. Since no process is absolutely reversible, there is always at least a small gain in the entropy; a reversible process is an idealization in which we have made the gain of entropy minimal.

Unfortunately, we are not going to enter into the field of thermodynamics very far. Our purpose is only to illustrate the principal ideas involved and the reasons why it is possible to make such arguments, but we will not use thermodynamics very much in this course. Thermodynamics is used very often by engineers and, particularly, by chemists. So we must learn our thermodynamics in practice in chemistry or engineering. Because it is not worth while duplicating everything, we shall just give some discussion of the origin of the theory, rather than much detail for special applications.

The two laws of thermodynamics are often stated this way:

First law: the energy of the universe is always constant.

Second law: the entropy of the universe is always increasing.

That is not a very good statement of the second law; it does not say, for example, that in a reversible cycle the entropy stays the same, and it does not say exactly what the entropy is. It is just a clever way of remembering the two laws, but it does not really tell us exactly where we stand. We have summarized the laws discussed in this chapter in Table 44–1. In the next chapter we shall apply these laws to discover the relationship between the heat generated in the expansion of a rubber band, and the extra tension when it is heated.

Table 44–1

Summary of the laws of thermodynamics

First law:

Heat put into a system + Work done on a system = Increase in internal energy of the system:

$$dQ + dW = dU.$$

Second law:

A process whose *only* net result is to take heat from a reservoir and convert it to work is impossible.

No heat engine taking heat Q_1 from T_1 and delivering heat Q_2 at T_2 can do more work than a reversible engine, for which

$$W = Q_1 - Q_2 = Q_1\left(\frac{T_1 - T_2}{T_1}\right).$$

The entropy of a system is defined this way:

(a) If heat ΔQ is added reversibly to a system at temperature T, the increase in entropy of the system is $\Delta S = \Delta Q/T$.

(b) At $T = 0$, $S = 0$ (*third law*).

In a *reversible change*, the total entropy of all parts of the system (including reservoirs) does not change.

In *irreversible change*, the total entropy of the system always increases.

45

Illustrations of Thermodynamics

45–1 Internal energy
45–2 Applications
45–3 The Clausius-Clapeyron equation

45–1 Internal energy

Thermodynamics is a rather difficult and complex subject when we come to apply it, and it is not appropriate for us to go very far into the applications in this course. The subject is of very great importance, of course, to engineers and chemists, and those who are interested in the subject can learn about the applications in physical chemistry or in engineering thermodynamics. There are also good reference books, such as Zemansky's *Heat and Thermodynamics*, where one can learn more about the subject. In the Encyclopedia Britannica, fourteenth edition, one can find excellent articles on thermodynamics and thermochemistry, and in the article on chemistry, the sections on physical chemistry, vaporization, liquefication of gases, and so on.

The subject of thermodynamics is complicated because there are so many different ways of describing the same thing. If we wish to describe the behavior of a gas, we can say that the pressure depends on the temperature and on the volume, or we can say that the volume depends on the temperature and the pressure. Or with respect to the internal energy U, we might say that it depends on the temperature and volume, if those are the variables we have chosen—but we might also say that it depends on the temperature and the pressure, or the pressure and the volume, and so on. In the last chapter we discussed another function of temperature and volume, called the entropy S, and we can of course construct as many other functions of these variables as we like: $U - TS$ is a function of temperature and volume. So we have a large number of different quantities which can be functions of many different combinations of variables.

To keep the subject simple in this chapter, we shall decide at the start to use *temperature* and *volume* as the independent variables. Chemists use temperature and pressure, because they are easier to measure and control in chemical experiments, but we shall use temperature and volume throughout this chapter, except in one place where we shall see how to make the transformation into the chemists' system of variables.

We shall first, then, consider only one system of independent variables: temperature and volume. Secondly, we shall discuss only two dependent functions: the internal energy and the pressure. All the other functions can be derived from these, so it is not necessary to discuss them. With these limitations, thermodynamics is still a fairly difficult subject, but it is not quite so impossible!

First we shall review some mathematics. If a quantity is a function of two variables, the idea of the derivative of the quantity requires a little more careful thought than for the case where there is only one variable. What do we mean by the derivative of the pressure with respect to the temperature? The pressure change accompanying a change in the temperature depends partly, of course, on what happens to the *volume* while T is changing. We must specify the change in V before the concept of a derivative with respect to T has a precise meaning. We might ask, for example, for the rate of change of P with respect to T if V is held constant. This ratio is just the ordinary derivative that we usually write as dP/dT. We customarily use a special symbol, $\partial P/\partial T$, to remind us that P depends on another variable V as well as on T, and that this other variable is held constant. We shall not only use the symbol ∂ to call attention to the fact that the other variable is held constant, but we shall also write the variable that is held constant as a subscript, $(\partial P/\partial T)_V$. Since we have only two independent variables, this notation is redundant, but it will help us keep our wits about us in the thermodynamic jungle of partial derivatives.

Let us suppose that the function $f(x, y)$ depends on the two independent variables x and y. By $(\partial f/\partial x)_y$ we mean simply the ordinary derivative, obtained in the usual way, if we treat y as a constant:

$$\left(\frac{\partial f}{\partial x}\right)_y = \lim_{\Delta x \to 0} \frac{f(x + \Delta x, y) - f(x, y)}{\Delta x}.$$

Similarly, we define

$$\left(\frac{\partial f}{\partial y}\right)_x = \lim_{\Delta y \to 0} \frac{f(x, y + \Delta y) - f(x, y)}{\Delta y}.$$

For example, if $f(x, y) = x^2 + yx$, then $(\partial f/\partial x)_y = 2x + y$, and $(\partial f/\partial y)_x = x$. We can extend this idea to higher derivatives: $\partial^2 f/\partial y^2$ or $\partial^2 f/\partial y \partial x$. The latter symbol indicates that we first differentiate f with respect to x, treating y as a constant, then differentiate the result with respect to y, treating x as a constant. The actual order of differentiation is immaterial: $\partial^2 f/\partial x \partial y = \partial^2 f/\partial y \partial x$.

We will need to compute the change Δf in $f(x, y)$ when x changes to $x + \Delta x$ *and* y changes to $y + \Delta y$. We assume throughout the following that Δx and Δy are infinitesimally small:

$$\begin{aligned}\Delta f &= f(x + \Delta x, y + \Delta y) - f(x, y) \\ &= \underbrace{f(x + \Delta x, y + \Delta y) - f(x, y + \Delta y)}_{\Delta x \left(\frac{\partial f}{\partial x}\right)_y} + \underbrace{f(x, y + \Delta y) - f(x, y)}_{\Delta y \left(\frac{\partial f}{\partial y}\right)_x}.\end{aligned} \tag{45.1}$$

The last equation is the fundamental relation that expresses Δf in terms of Δx and Δy.

As an example of the use of this relation, let us calculate the change in the internal energy $U(T, V)$ when the temperature changes from T to $T + \Delta T$ and the volume changes from V to $V + \Delta V$. Using Eq. (45.1), we write

$$\Delta U = \Delta T \left(\frac{\partial U}{\partial T}\right)_V + \Delta V \left(\frac{\partial U}{\partial V}\right)_T. \tag{45.2}$$

In our last chapter we found another expression for the change ΔU in the internal energy when a quantity of heat ΔQ was added to the gas:

$$\Delta U = \Delta Q - P \Delta V. \tag{45.3}$$

In comparing Eqs. (45.2) and (45.3) one might at first be inclined to think that $P = (\partial U/\partial V)_T$, but this is not correct. To obtain the correct relation, let us first suppose that we add a quantity of heat ΔQ to the gas while keeping the volume constant, so that $\Delta V = 0$. With $\Delta V = 0$, Eq. (45.3) tells us that $\Delta U = \Delta Q$, and Eq. (45.2) tells us that $\Delta U = (\partial U/\partial T)_V \Delta T$, so that $(\partial U/\partial T)_V = \Delta Q/\Delta T$. The ratio $\Delta Q/\Delta T$, the amount of heat one must put into a substance in order to change its temperature by one degree with the volume held constant, is called the *specific heat at constant volume* and is designated by the symbol C_V. By this argument we have shown that

$$\left(\frac{\partial U}{\partial T}\right)_V = C_V. \tag{45.4}$$

Now let us again add a quantity of heat ΔQ to the gas, but this time we will hold T constant and allow the volume to change by ΔV. The analysis in this case is more complex, but we can calculate ΔU by the argument of Carnot, making use of the Carnot cycle we introduced in the last chapter.

The pressure-volume diagram for the Carnot cycle is shown in Fig. 45–1. As we have already shown, the total amount of work done by the gas in a reversible cycle is $\Delta Q(\Delta T/T)$, where ΔQ is the amount of heat energy added to the gas as it expands isothermally at temperature T from volume V to $V + \Delta V$, and $T - \Delta T$ is the final temperature reached by the gas as it expands adiabatically on the second leg of the cycle. Now we will show that this work done is also given by

the shaded area in Fig. 45–1. In any circumstances, the work done by the gas is $\int P\,dV$, and is positive when the gas expands and negative when the gas is compressed. If we plot P vs. V, the variation of P and V is represented by a curve which gives the value of P corresponding to a particular value of V. As the volume changes from one value to another, the work done by the gas, the integral $\int P\,dV$, is the area under the curve connecting the initial and final values of V. When we apply this idea to the Carnot cycle, we see that as we go around the cycle, paying attention to the sign of the work done by the gas, the net work done by the gas is just the shaded area in Fig. 45–1.

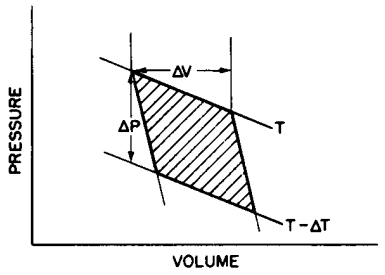


Fig. 45–1. Pressure-volume diagram for a Carnot cycle. The curves marked T and $T - \Delta T$ are isothermal lines; the steeper curves are adiabatic lines. ΔV is the volume change as heat ΔQ is added to the gas at constant temperature T. ΔP is the pressure change at constant volume as the gas temperature is changed from T to $T - \Delta T$.

Now we want to evaluate the shaded area geometrically. The cycle we have used in Fig. 45–1 differs from that used in the previous chapter in that we now suppose that ΔT and ΔQ are infinitesimally small. We are working between adiabatic lines and isothermal lines that are very close together, and the figure described by the heavy lines in Fig. 45–1 will approach a parallelogram as the increments ΔT and ΔQ approach zero. The area of this parallelogram is just $\Delta V\,\Delta P$, where ΔV is the change in volume as energy ΔQ is added to the gas at constant temperature, and ΔP is the change in pressure as the temperature changes by ΔT at constant volume. One can easily show that the shaded area in Fig. 45–1 is given by $\Delta V\,\Delta P$ by recognizing that the shaded area is equal to the area enclosed by the dotted lines in Fig. 45–2, which in turn differs from the rectangle bounded by ΔP and ΔV only by the addition and subtraction of the equal triangular areas in Fig. 45–2.

Now let us summarize the results of the arguments we have developed so far:

$$\left.\begin{aligned} &\text{Work done by the gas} = \text{shaded area} = \Delta V\,\Delta P = \Delta Q\left(\frac{\Delta T}{T}\right) \\ &\text{or} \\ &\frac{\Delta T}{T}\cdot(\text{heat needed to change } V \text{ by } \Delta V)_{\text{constant}\,T} \\ &= \Delta V\cdot(\text{change in } P \text{ when } T \text{ changes by } \Delta T)_{\text{constant}\,V} \\ &\text{or} \\ &\frac{1}{\Delta V}\cdot(\text{heat needed to change } V \text{ by } \Delta V)_T = T(\partial P/\partial T)_V. \end{aligned}\right\} \quad (45.5)$$

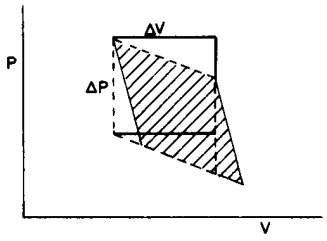


Fig. 45–2. Shaded area = area enclosed by dashed lines = area of rectangle = $\Delta P\,\Delta V$.

Equation (45.5) expresses the essential result of Carnot's argument. The whole of thermodynamics can be deduced from Eq. (45.5) and the First Law, which is stated in Eq. (45.3). Equation (45.5) is essentially the Second Law, although it was originally deduced by Carnot in a slightly different form, since he did not use our definition of temperature.

Now we can proceed to calculate $(\partial U/\partial V)_T$. By how much would the internal energy U change if we changed the volume by ΔV? First, U changes because heat is put in, and second, U changes because work is done. The heat put in is

$$\Delta Q = T\left(\frac{\partial P}{\partial T}\right)_V \Delta V,$$

according to Eq. (45.5), and the work done on the substance is $-P\,\Delta V$. Therefore the change ΔU in internal energy has two pieces:

$$\Delta U = T\left(\frac{\partial P}{\partial T}\right)_V \Delta V - P\,\Delta V. \quad (45.6)$$

Dividing both sides by ΔV, we find for the rate of change of U with V at constant T

$$\left(\frac{\partial U}{\partial V}\right)_T = T\left(\frac{\partial P}{\partial T}\right)_V - P. \quad (45.7)$$

In our thermodynamics, in which T and V are the only variables and P and U are the only functions, Eqs. (45.3) and (45.7) are the basic equations from which all the result of the subject can be deduced.

45–2 Applications

Now let us discuss the meaning of Eq. (45.7) and see why it answers the questions which we proposed in our last chapter. We considered the following problem: in kinetic theory it is obvious that an increase in temperature leads to an increase in pressure, because of the bombardments of the atoms on a piston. For the same physical reason, when we let the piston move back, heat is taken out of the gas and, in order to keep the temperature constant, heat will have to be put back in. The gas cools when it expands, and the pressure rises when it is heated. There must be some connection between these two phenomena, and this connection is given explicitly in Eq. (45.7). If we hold the volume fixed and increase the temperature, the pressure rises at a rate $(\partial P/\partial T)_V$. Related to that fact is this: if we increase the volume, the gas will cool unless we pour some heat in to maintain the temperature constant, and $(\partial U/\partial V)_T$ tells us the amount of heat needed to maintain the temperature. Equation (45.7) expresses the fundamental interrelationship between these two effects. That is what we promised we would find when we came to the laws of thermodynamics. Without knowing the internal mechanism of the gas, and knowing only that we cannot make perpetual motion of the second type, we can deduce the relationship between the amount of heat needed to maintain a constant temperature when the gas expands, and the pressure change when the gas is heated!

Now that we have the result we wanted for a gas, let us consider the rubber band. When we stretch a rubber band, we find that its temperature falls, and when we heat a rubber band, we find that it pulls itself in. What is the equation that gives the same relation for a rubber band as Eq. (45.3) gives for gas? For a rubber band the situation will be something like this: when heat ΔQ is put in, the internal energy is changed by ΔU and some work is done. The only difference will be that the work done by the rubber band is $-F\,\Delta L$ instead of $P\,\Delta V$, where F is the force on the band, and L is the length of the band. The force F is a function of temperature and of length of the band. Replacing $P\Delta V$ in Eq. (45.3) by $-F\,\Delta L$, we get

$$\Delta U = \Delta Q + F\Delta L \tag{45.8}$$

Comparing Eqs. (45.3) and (45.8), we see that the rubber band equation is obtained by a mere substitution of one letter for another. Furthermore, if we substitute L for V, and $-F$ for P, all of our discussion of the Carnot cycle applies to the rubber band. We can immediately deduce, for instance, that the heat ΔQ needed to change the length by ΔL is given by the analog to Eq. (45.5): $\Delta Q = -T(\partial F/\partial T)_L\,\Delta L$. This equation tells us that if we keep the length of a rubber band fixed and heat the band, we can calculate how much the force will increase in terms of the heat needed to keep the temperature constant when the band is stretched a little bit. So we see that the same equation applies to both gas and a rubber band. In fact, if one can write $\Delta U = \Delta Q + A\Delta B$, where A and B represent different quantities, force and length, pressure and volume, etc., one can apply the results obtained for a gas by substituting A and B for P and V. For example, consider the electric potential difference, or "voltage," E in a battery and the charge ΔZ that moves through the battery. We know that the work done in a reversible electric cell, like a storage battery, is $E\Delta Z$. (Since we include no $P\Delta V$ term in the work, we require that our battery maintain a constant volume.) Let us see what thermodynamics can tell us about the performance of a battery. If we substitute E for P and Z for V in Eq. (45.6), we obtain

$$\frac{\Delta U}{\Delta Z} = -T\left(\frac{\partial E}{\partial T}\right)_Z - E. \tag{45.9}$$

Equation (45.9) says that the internal energy U is changed when a charge ΔZ moves through the cell. Why is $\Delta U/\Delta Z$ not simply the voltage E of the battery? The answer is that a real battery gets warm when charge moves through the cell. The internal energy of the battery is changed, first, because the battery did some work on the outside circuit, and second, because the battery is heated. The re-

markable thing is that the second part can again be expressed in terms of the way in which the battery voltage changes with temperature. Incidentally, when the charge moves through the cell, chemical reactions occur, and Eq. (45.9) suggests a nifty way of measuring the amount of energy required to produce a chemical reaction. All we need do is construct a cell that works on the reaction, measure the voltage, and measure how much the voltage changes with temperature when we draw no charge from the battery!

Now we have assumed that the volume of the battery can be maintained constant, since we have omitted the $P\Delta V$ term when we set the work done by the battery equal to $E\Delta Z$. It turns out that it is technically quite difficult to keep the volume constant. It is much easier to keep the cell at constant atmospheric pressure. For that reason, the chemists do not like any of the equations we have written above: they prefer equations which describe performance under constant *pressure*. We chose at the beginning of this chapter to use V and T as independent variables. The chemists prefer P and T, and we will now consider how the results we have obtained so far can be transformed into the chemists' system of variables. Remember that in the following treatment confusion can easily set in because we are shifting gears from T and V to T and P.

We started in Eq. (45.3) with $\Delta U = \Delta Q - P\Delta V$; $P\Delta V$ may be replaced by $E\Delta Z$ or $A\Delta B$. If we could somehow replace the last term, $P\Delta V$, by $V\Delta P$, then we would have interchanged V and P, and the chemists would be happy. Well, a clever man noticed that the differential of the product PV is $d(PV) = PdV + V\,dP$, and if he added this equality to Eq. (45.3), he obtained

$$\begin{aligned} \Delta(PV) &= P\Delta V + V\Delta P \\ \Delta U &= \Delta Q - P\Delta V \\ \hline \Delta(U + PV) &= \Delta Q + V\Delta P \end{aligned}$$

In order that the result look like Eq. (45.3), we define $U + PV$ to be something new, called the *enthalpy*, H, and we write $\Delta H = \Delta Q + V\Delta P$.

Now we are ready to transform our results into chemists' language with the following rules: $U \to H$, $P \to -V$, $V \to P$. For example, the fundamental relationship that chemists would use instead of Eq. (45.7) is

$$\left(\frac{\partial H}{\partial P}\right)_T = T\left(\frac{\partial V}{\partial T}\right)_P - V.$$

It should now be clear how one transforms to the chemists' variables T and P. We now go back to our original variables: for the remainder of this chapter, T and V are the independent variables.

Now let us apply the results we have obtained to a number of physical situations. Consider first the ideal gas. From kinetic theory we know that the internal energy of a gas depends only on the motion of the molecules and the number of molecules. The internal energy depends on T, but not on V. If we change V, but keep T constant, U is not changed. Therefore $(\partial U/\partial V)_T = 0$, and Eq. (45.7) tells us that for an ideal gas

$$T\left(\frac{\partial P}{\partial T}\right)_V - P = 0. \tag{45.10}$$

Equation (45.10) is a differential equation that can tell us something about P. We take account of the partial derivatives in the following way: Since the partial derivative is at constant V, we will replace the partial derivative by an ordinary derivative and write explicitly, to remind us, "constant V." Equation (45.10) then becomes

$$T\frac{\Delta P}{\Delta T} - P = 0; \qquad \text{const } V, \tag{45.11}$$

which we can integrate to get

$$\begin{aligned} \ln P &= \ln T + \text{const}; \qquad \text{const } V, \\ P &= \text{const} \times T; \qquad \text{const } V. \end{aligned} \tag{45.12}$$

We know that for an ideal gas the pressure is equal to

$$P = \frac{RT}{V}, \tag{45.13}$$

which is consistent with (45.12), since V and R are constants. Why did we bother to go through this calculation if we already knew the results? Because we have been using *two independent definitions of temperature!* At one stage we assumed that the kinetic energy of the molecules was proportional to the temperature, an assumption that defines one scale of temperature which we will call the ideal gas scale. The T in Eq. (45.13) is based on the gas scale. We also call temperatures measured on the gas scale *kinetic* temperatures. Later, we defined the temperature in a second way which was completely independent of any substance. From arguments based on the Second Law we defined what we might call the "grand thermodynamic absolute temperature" T, the T that appears in Eq. (45.12). What we proved here is that the pressure of an ideal gas (defined as one for which the internal energy does not depend on the volume) is proportional to the grand thermodynamic absolute temperature. We also know that the pressure is proportional to the temperature measured on the gas scale. Therefore we can deduce that the kinetic temperature is proportional to the "grand thermodynamic absolute temperature." That means, of course, that if we were sensible we could make two scales agree. In this instance, at least, the two scales *have* been chosen so that they coincide; the proportionality constant has been chosen to be 1. Most of the time man chooses trouble for himself, but in this case he made them equal!

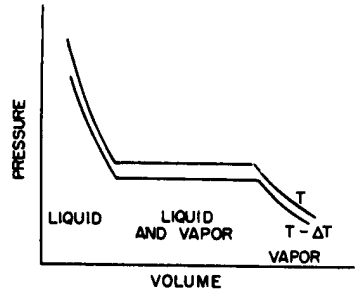


Fig. 45–3. Isothermal lines for a condensable vapor compressed in a cylinder. At the left, the substance is in the liquid phase. At the right, the substance is vaporized. In the center, both liquid and vapor are present in the cylinder.

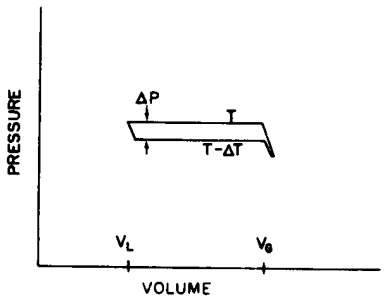


Fig. 45–4. Pressure-volume diagram for a Carnot cycle with a condensable vapor in the cylinder. At the left, the substance is in the liquid state. A quantity of heat L is added at temperature T to vaporize the liquid. The vapor expands adiabatically as T changes to $T - \Delta T$.

45–3 The Clausius-Clapeyron equation

The vaporization of a liquid is another application of the results we have derived. Suppose we have some liquid in a cylinder, such that we can compress it by pushing on the piston, and we ask ourselves, "If we keep the temperature constant, how does the pressure vary with volume?" In other words, we want to draw an isothermal line on the P-V diagram. The substance in the cylinder is not the ideal gas that we considered earlier; now it may be in the liquid or the vapor phase, or both may be present. If we apply sufficient pressure, the substance will condense to a liquid. Now if we squeeze still harder, the volume changes very little, and our isothermal line rises rapidly with decreasing volume, as shown at the left in Fig. 45–3.

If we increase the volume by pulling the piston out, the pressure drops until we reach the point at which the liquid starts to boil, and then vapor starts to form. If we pull the piston out farther, all that happens is that more liquid vaporizes. When there is part liquid and part vapor in the cylinder, the two phases are in equilibrium—liquid is evaporating and vapor is condensing at the same rate. If we make more room for the vapor, more vapor is needed to maintain the pressure, so a little more liquid evaporates, but the pressure remains constant. On the flat part of the curve in Fig. 45–3 the pressure does not change, and the value of the pressure here is called the *vapor pressure at temperature T*. As we continue to increase the volume, there comes a time when there is no more liquid to evaporate. At this juncture, if we expand the volume further, the pressure will fall as for an ordinary gas, as shown at the right of the P-V diagram. The lower curve in Fig. 45–3 is the isothermal line at a slightly lower temperature $T - \Delta T$. The pressure in the liquid phase is slightly reduced because liquid expands with an increase in temperature (for most substances, but not for water near the freezing point) and, of course, the vapor pressure is lower at the lower temperature.

We will now make a cycle out of the two isothermal lines by connecting them (say by adiabatic lines) at the ends of the flat sections, as shown in Fig. 45–4. The little jiggle in the lower right-hand corner of the figure will make little difference and we will neglect it. We are going to use the argument of Carnot, which tells us that the heat added to the substance in changing it from a liquid to a vapor is related to the work done by the substance as it goes around the cycle. Let

us call L the heat needed to vaporize the substance in the cylinder. As in the argument immediately preceding Eq. (45.5), we know that $L(\Delta T/T)$ = work done by the substance. As before, the work done by the substance is the shaded area, which is approximately $\Delta P(V_G - V_L)$, where ΔP is the difference in vapor pressure at the two temperatures T and $T - \Delta T$, V_G is the volume of the gas, and V_L is the volume of the liquid, both volumes measured at the vapor pressure. Setting these two expressions for the area equal, we get $L\Delta T/T = \Delta P(V_G - V_L)$, or

$$\frac{L}{T(V_G - V_L)} = (\partial P_{\text{vap}}/\partial T). \tag{45.14}$$

Equation (45.14) gives the relationship between the rate of change of vapor pressure with temperature and the amount of heat required to evaporate the liquid. This relationship was deduced by Carnot, but it is called the Clausius-Clapeyron equation.

Now let us compare Eq. (45.14) with the results deduced from kinetic theory. Usually V_G is very much larger than V_L. So $V_G - V_L \approx V_G = RT/P$ per mole. If we further assume that L is a constant, independent of temperature—not a very good approximation—then we would have $\partial P/\partial T = L/(RT^2P)$. The solution of this differential equation is

$$P = \text{const}\, e^{-L/RT}. \tag{45.15}$$

Let us compare this with the pressure variation with temperature that we deduced earlier from kinetic theory. Kinetic theory indicated the possibility, at least roughly, that the number of molecules of vapor above a liquid would be

$$n = \left(\frac{1}{V_A}\right) e^{-(U_G - U_L)/RT}, \tag{45.16}$$

where $U_G - U_L$ is the internal energy per mole in the liquid minus the internal energy per mole in the gas, i.e., the energy needed to vaporize a mole of liquid. Equation (45.15) from thermodynamics and Eq. (45.16) from kinetic theory are very closely related because the pressure is nkT, but they are not exactly the same. However, they will turn out to be exactly the same if we assume $L - U_G$ = const, instead of L = const. If we assume $L - U_G$ = const, independent of temperature, then the argument leading to Eq. (45.15) will produce Eq. (45.16).

This comparison shows the advantages and disadvantages of thermodynamics over kinetic theory: First of all, Eq. (45.14) obtained by thermodynamics is exact, while Eq. (45.16) can only be approximated, for instance, if U is nearly constant, and if the model is right. Second, we may not understand correctly how the gas goes into the liquid; nevertheless, Eq. (45.14) is right, while (45.16) is only approximate. Third, although our treatment applies to a gas condensing into a liquid, the argument is true for any other change of state. For instance, the solid-to-liquid transition has the same kind of curve as that shown in Figs. 45–3 and 45–4. Introducing the latent heat for melting, M/mole, the formula analogous to Eq. (45.14) then is $(\partial P_{\text{melt}}/\partial T)_V = M/[T(V_{\text{liq}} - V_{\text{solid}})]$. Although we may not understand the kinetic theory of the melting process, we nevertheless have a correct equation. However, when we *can* understand the kinetic theory, we have another advantage. Equation (45.14) is only a differential relationship, and we have no way of obtaining the constants of integration. In the kinetic theory we can obtain the constants also if we have a good model that describes the phenomenon completely. So there are advantages and disadvantages to each. When knowledge is weak and the situation is complicated, thermodynamic relations are really the most powerful. When the situation is very simple and a theoretical analysis can be made, then it is better to try to get more information from theoretical analysis.

One more example: blackbody radiation. We have discussed a box containing radiation and nothing else. We have talked about the equilibrium between the oscillator and the radiation. We also found that the photons hitting the wall of the box would exert the pressure P, and we found $PV = U/3$, where U is the

total energy of all the photons and V is the volume of the box. If we substitute $U = 3PV$ in the basic Eq. (45.7), we find

$$\left(\frac{\partial U}{\partial V}\right)_T = 3P = T\left(\frac{\partial P}{\partial T}\right)_V - P. \tag{45.17}$$

Since the volume of our box is constant, we can replace $(\partial P/\partial T)_V$ by dP/dT to obtain an ordinary differential equation we can integrate: $\ln P = 4 \ln T + \text{const}$, or $P = \text{const} \times T^4$. The pressure of radiation varies as the fourth power of the temperature, and the energy content of the radiation, $U/V = P/3$, also varies as T^4. It is usual to write $U/V = (4\sigma/c)T^4$, where c is the speed of light and σ is a constant. It is not possible to get σ from thermodynamics alone. Here is a good example of its power, and its limitations. To know that U/V goes as T^4 is a great deal, but to know how big U/V actually is at any temperature requires that we go into the kind of detail that only a complete theory can supply. For blackbody radiation we have such a theory and we can derive an expression for the constant σ in the following manner.

Let $I(\omega)\,d\omega$ be the intensity distribution, the energy flow through 1 m^2 in one second with frequency between ω and $\omega + d\omega$. The energy density distribution = energy/volume = $I(\omega)\,d\omega/c$ is

$$\begin{aligned}\frac{U}{V} &= \text{total energy density} \\ &= \int_{\omega=0}^{\infty} \text{energy density } \omega \text{ and } \omega + d\omega \\ &= \int_0^{\infty} \frac{I(\omega)\,d\omega}{c}.\end{aligned}$$

From our earlier discussions, we know that

$$I(\omega) = \frac{\hbar\omega^3}{\pi^2c^2(e^{\hbar\omega/kT} - 1)}.$$

Substituting this expression for $I(\omega)$ in our equation for U/V, we get

$$\frac{U}{V} = \frac{1}{\pi^2c^3}\int_0^{\infty} \frac{\hbar\omega^3\,d\omega}{e^{\hbar\omega/kT} - 1}.$$

If we substitute $x = \hbar\omega/kT$, the expression becomes

$$\frac{U}{V} = \frac{(kT)^4}{\hbar^3\pi^2c^3}\int_0^{\infty} \frac{x^3\,dx}{e^x - 1}.$$

This integral is just some number that we can get, approximately, by drawing a curve and taking the area by counting squares. It is roughly 6.5. The mathematicians among us can show that the integral is exactly $\pi^4/15$.* Comparing this expression with $U/V = (4\sigma/c)T^4$, we find

$$\sigma = \frac{k^4\pi^2}{60\hbar^3c^2} = 5.67 \times 10^{-8}\,\frac{\text{watts}}{(\text{meter})^2\,(\text{degree})^4}.$$

---

* Since $(e^x - 1)^{-1} = e^{-x} + e^{-2x} + \ldots$, the integral is

$$\sum_{n=1}^{\infty}\int_0^{\infty} e^{-nx}x^3\,dx.$$

But $\int_0^{\infty} e^{-nx}\,dx = 1/n$, and differentiating with respect to n three times gives $\int_0^{\infty} x^3e^{-nx}\,dx = 6/n^4$, so the integral is $6(1 + \frac{1}{16} + \frac{1}{81} + \ldots)$ and a good estimate comes from adding the first few terms. In Chapter 50 we will find a way to show that the sum of the reciprocal fourth powers of the integers is, in fact, $\pi^4/90$.

If we make a small hole in our box, how much energy will flow per second through the hole of unit area? To go from energy density to energy flow, we multiply the energy density U/V by c. We also multiply by $\frac{1}{4}$, which arises as follows: first, a factor of $\frac{1}{2}$, because only the energy which is flowing *out* escapes; and second, another factor $\frac{1}{2}$, because energy which approaches the hole at an angle to the normal is less effective in getting through the hole by a cosine factor. The average value of the cosine is $\frac{1}{2}$. It is clear now why we write $U/V = (4\sigma/c)T^4$: so that we can ultimately say that the flux from a small hole is σT^4 per unit area.

46

Ratchet and pawl

46–1 How a ratchet works

46–2 The ratchet as an engine

46–3 Reversibility in mechanics

46–4 Irreversibility

46–5 Order and entropy

46–1 How a ratchet works

In this chapter we discuss the ratchet and pawl, a very simple device which allows a shaft to turn only one way. The possibility of having something turn only one way requires some detailed and careful analysis, and there are some very interesting consequences.

The plan of the discussion came about in attempting to devise an elementary explanation, from the molecular or kinetic point of view, for the fact that there is a maximum amount of work which can be extracted from a heat engine. Of course we have seen the essence of Carnot's argument, but it would be nice to find an explanation which is elementary in the sense that we can see what is happening physically. Now, there are complicated mathematical demonstrations which follow from Newton's laws to demonstrate that we can get only a certain amount of work out when heat flows from one place to another, but there is great difficulty in converting this into an elementary demonstration. In short, we do not understand it, although we can follow the mathematics.

In Carnot's argument, the fact that more than a certain amount of work cannot be extracted in going from one temperature to another is deduced from another axiom, which is that if everything is at the same temperature, heat cannot be converted to work by means of a cyclic process. First, let us back up and try to see, in at least one elementary example, why this simpler statement is true.

Let us try to invent a device which will violate the Second Law of Thermodynamics, that is, a gadget which will generate work from a heat reservoir with everything at the same temperature. Let us say we have a box of gas at a certain temperature, and inside there is an axle with vanes in it. (See Fig. 46–1 but take $T_1 = T_2 = T$, say.) Because of the bombardments of gas molecules on the vane, the vane oscillates and jiggles. All we have to do is to hook onto the other end of the axle a wheel which can turn only one way—the ratchet and pawl. Then when the shaft tries to jiggle one way, it will not turn, and when it jiggles the other, it will turn. Then the wheel will slowly turn, and perhaps we might even tie a flea onto a string hanging from a drum on the shaft, and lift the flea! Now let us ask if this is possible. According to Carnot's hypothesis, it is impossible. But if we just look at it, we see, *prima facie*, that it seems quite possible. So we must look more closely. Indeed, if we look at the ratchet and pawl, we see a number of complications.

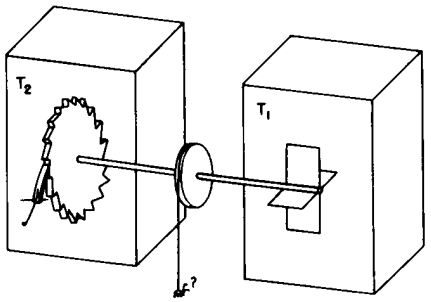


Fig. 46–1. The ratchet and pawl machine.

First, our idealized ratchet is as simple as possible, but even so, there is a pawl, and there must be a spring in the pawl. The pawl must return after coming off a tooth, so the spring is necessary.

Another feature of this ratchet and pawl, not shown in the figure, is quite essential. Suppose the device were made of perfectly elastic parts. After the pawl is lifted off the end of the tooth and is pushed back by the spring, it will bounce against the wheel and continue to bounce. Then, when another fluctuation came, the wheel could turn the other way, because the tooth could get underneath during the moment when the pawl was up! Therefore an essential part of the irreversibility of our wheel is a damping or deadening mechanism which stops the bouncing. When the damping happens, of course, the energy that was in the pawl goes into the wheel and shows up as heat. So, as it turns, the wheel will get hotter and hotter. To make the thing simpler, we can put a gas around the wheel to take up some of the heat. Anyway, let us say the gas keeps rising in temperature, along with the wheel. Will it go on forever? No! The pawl and wheel, both at some temperature

T, also have Brownian motion. This motion is such that, every once in a while, by accident, the pawl lifts itself up and over a tooth just at the moment when the Brownian motion on the vanes is trying to turn the axle backwards. And as things get hotter, this happens more often.

So, this is the reason this device does not work in perpetual motion. When the vanes get kicked, sometimes the pawl lifts up and goes over the end. But sometimes, when it tries to turn the other way, the pawl has already lifted due to the fluctuations of the motions on the wheel side, and the wheel goes back the other way! The net result is nothing. It is not hard to demonstrate that when the temperature on both sides is equal, there will be no net average motion of the wheel. Of course the wheel will do a lot of jiggling this way and that way, but it will not do what we would like, which is to turn just one way.

Let us look at the reason. It is necessary to do work against the spring in order to lift the pawl to the top of a tooth. Let us call this energy ϵ, and let θ be the angle between the teeth. The chance that the system can accumulate enough energy, ϵ, to get the pawl over the top of the tooth, is $e^{-\epsilon/kT}$. But the probability that the pawl will accidentally be up is also $e^{-\epsilon/kT}$. So the number of times that the pawl is up and the wheel can turn backwards freely is equal to the number of times that we have enough energy to turn it forward when the pawl is down. We thus get a "balance," and the wheel will not go around.

46–2 The ratchet as an engine

Let us now go further. Take the example where the temperature of the vanes is T_1 and the temperature of the wheel, or ratchet, is T_2, and T_2 is less than T_1. Because the wheel is cold and the fluctuations of the pawl are relatively infrequent, it will be very hard for the pawl to attain an energy ϵ. Because of the high temperature T_1, the vanes will often attain the energy ϵ, so our gadget will go in one direction, as designed.

We would now like to see if it can lift weights. Onto the drum in the middle we tie a string, and put a weight, such as our flea, on the string. We let L be the torque due to the weight. If L is not too great, our machine will lift the weight because the Brownian fluctuations make it more likely to move in one direction than the other. We want to find how much weight it can lift, how fast it goes around, and so on.

First we consider a forward motion, the usual way one designs a ratchet to run. In order to make one step forward, how much energy has to be borrowed from the vane end? We must borrow an energy ϵ to lift the pawl. The wheel turns through an angle θ against a torque L, so we also need the energy $L\theta$. The total amount of energy that we have to borrow is thus $\epsilon + L\theta$. The probability that we get this energy is proportional to $e^{-(\epsilon+L\theta)/kT_1}$. Actually, it is not only a question of getting the energy, but we also would like to know the number of times per second it has this energy. The probability per second is proportional to $e^{-(\epsilon+L\theta)/kT_1}$, and we shall call the proportionality constant $1/\tau$. It will cancel out in the end anyway. When a forward step happens, the work done on the weight is $L\theta$. The energy taken from the vane is $\epsilon + L\theta$. The spring gets wound up with energy ϵ, then it goes clatter, clatter, bang, and this energy goes into heat. All the energy taken out goes to lift the weight and to drive the pawl, which then falls back and gives heat to the other side.

Now we look at the opposite case, which is backward motion. What happens here? To get the wheel to go backwards all we have to do is supply the energy to lift the pawl high enough so that the ratchet will slip. This is still energy ϵ. Our probability per second for the pawl to lift this high is now $(1/\tau)e^{-\epsilon/kT_2}$. Our proportionality constant is the same, but this time kT_2 shows up because of the different temperature. When this happens, the work is released because the wheel slips backward. It loses one notch, so it releases work $L\theta$. The energy taken from the ratchet system is ϵ, and the energy given to the gas at T_1 on the vane side is $L\theta + \epsilon$. It takes a little thinking to see the reason for that. Suppose the pawl has lifted itself up accidentally by a fluctuation. Then when it falls back and the spring

Table 46–1

Summary of operation of ratchet and pawl.

Forward:	Need energy	$\epsilon + L\theta$	from vane. $\therefore$ Rate $= \frac{1}{\tau} e^{-(L\theta+\epsilon)/kT_1}$
	Takes from vane	$L\theta + \epsilon$	
	Does work	$L\theta$	
	Gives to ratchet	ϵ	
Backward:	Needs energy	ϵ	for pawl. $\therefore$ Rate $= \frac{1}{\tau} e^{-\epsilon/kT_2}$
	Takes from ratchet	ϵ	} same as above with sign reversed.
	Releases work	$L\theta$	
	Gives to vane	$L\theta + \epsilon$	

If system is reversible, rates are equal, hence $\frac{\epsilon + L\theta}{T_1} = \frac{\epsilon}{T_2}$.

$\frac{\text{Heat to ratchet}}{\text{Heat from vane}} = \frac{\epsilon}{L\theta + \epsilon}$. Hence $\frac{Q_2}{Q_1} = \frac{T_2}{T_1}$.

pushes it down against the tooth, there is a force trying to turn the wheel, because the tooth is pushing on an inclined plane. This force is doing work, and so is the force due to the weights. So both together make up the total force, and all the energy which is slowly released appears at the vane end as heat. (Of course it must, by conservation of energy, but one must be careful to think the thing through!) We notice that all these energies are exactly the same, but reversed. So, depending upon which of these two rates is greater, the weight is either slowly lifted or slowly released. Of course, it is constantly jiggling around, going up for a while and down for a while, but we are talking about the average behavior.

Suppose that for a particular weight the rates happen to be equal. Then we add an infinitesimal weight to the string. The weight will slowly go down, and work will be done on the machine. Energy will be taken from the wheel and given to the vanes. If instead we take off a little bit of weight, then the imbalance is the other way. The weight is lifted, and heat is taken from the vane and put into the wheel. So we have the conditions of Carnot's reversible cycle, provided that the weight is just such that these two are equal. This condition is evidently that $(\epsilon + L\theta)/T_1 = \epsilon/T_2$. Let us say that the machine is slowly lifting the weight. Energy Q_1 is taken from the vanes and energy Q_2 is delivered to the wheel, and these energies are in the ratio $(\epsilon + L\theta)/\epsilon$. If we are lowering the weight, we also have $Q_1/Q_2 = (\epsilon + L\theta)/\epsilon$. Thus (Table 46–1) we have

$$Q_1/Q_2 = T_1/T_2.$$

Furthermore, the work we get out is to the energy taken from the vane as $L\theta$ is to $L\theta + \epsilon$, hence as $(T_1 - T_2)/T_1$. We see that our device cannot extract more work than this, operating reversibly. This is the result that we expected from Carnot's argument, and the main result of this lecture. However, we can use our device to understand a few other phenomena, even out of equilibrium, and therefore beyond the range of thermodynamics.

Let us now calculate *how fast* our one-way device would turn if everything were at the same temperature and we hung a weight on the drum. If we pull very, very hard, of course, there are all kinds of complications. The pawl slips over the ratchet, or the spring breaks, or something. But suppose we pull gently enough that everything works nicely. In those circumstances, the above analysis is right for the probability of the wheel going forward and backward, if we remember

that the two temperatures are equal. In each step an angle θ is obtained, so the angular velocity is θ times the probability of one of these jumps per second. It goes forward with probability $(1/\tau)e^{-(\epsilon+L\theta)/kT}$ and backward with probability $(1/\tau)e^{-\epsilon/kT}$, so that for the angular velocity we have

$$\begin{aligned}\omega &= (\theta/\tau)e^{-(\epsilon+L\theta)/kT} - e^{-\epsilon/kT} \\ &= (\theta/\tau)e^{-\epsilon/kT}(e^{-L\theta/kT} - 1).\end{aligned} \tag{46.1}$$

If we plot ω against L, we get the curve shown in Fig. 46–2. We see that it makes a great difference whether L is positive or negative. If L increases in the positive range, which happens when we try to drive the wheel backward, the backward velocity approaches a constant. As L becomes negative, ω really "takes off" forward, since e to a tremendous power is very great!

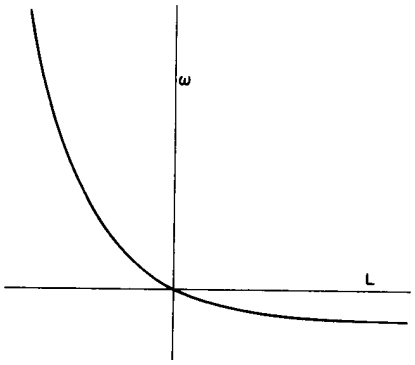


Fig. 46–2. Angular velocity of the ratchet as a function of torque.

The angular velocity that was obtained from different forces is thus very unsymmetrical. Going one way it is easy: we get a lot of angular velocity for a little force. Going the other way, we can put on a lot of force, and yet the wheel hardly goes around.

We find the same thing in an *electrical rectifier*. Instead of the force, we have the electric field, and instead of the angular velocity, we have the electric current. In the case of a rectifier, the voltage is not proportional to resistance, and the situation is unsymmetrical. The same analysis that we made for the mechanical rectifier will also work for an electrical rectifier. In fact, the kind of formula we obtained above is typical of the current-carrying capacities of rectifiers as a function of their voltages.

Now let us take all the weights away, and look at the original machine. If T_2 were less than T_1, the ratchet would go forward, as anybody would believe. But what is hard to believe, at first sight, is the opposite. If T_2 is greater than T_1, the ratchet goes around the opposite way! A dynamic ratchet with lots of heat in it runs itself backwards, because the ratchet pawl is bouncing. If the pawl, for a moment, is on the incline somewhere, it pushes the inclined plane sideways. But it is *always* pushing on an inclined plane, because if it happens to lift up high enough to get past the point of a tooth, then the inclined plane slides by, and it comes down again on an inclined plane. So a hot ratchet and pawl is ideally built to go around in a direction exactly opposite to that for which it was originally designed!

In spite of all our cleverness of lopsided design, if the two temperatures are exactly equal there is no more propensity to turn one way than the other. The moment we look at it, it may be turning one way or the other, but in the long run, it gets nowhere. The fact that it gets nowhere is really the fundamental deep principle on which all of thermodynamics is based.

46–3 Reversibility in mechanics

What deeper mechanical principle tells us that, in the long run, if the temperature is kept the same everywhere, our gadget will turn neither to the right nor to the left? We evidently have a fundamental proposition that there is no way to design a machine which, left to itself, will be more likely to be turning one way than the other after a long enough time. We must try to see how this follows from the laws of mechanics.

The laws of mechanics go something like this: the mass times the acceleration is the force, and the force on each particle is some complicated function of the positions of all the other particles. There are other situations in which forces depend on velocity, such as in magnetism, but let us not consider that now. We take a simpler case, such as gravity, where forces depend only on position. Now suppose that we have solved our set of equations and we have a certain motion $x(t)$ for each particle. In a complicated enough system, the solutions are very complicated, and what happens with time turns out to be very surprising. If we write down any arrangement we please for the particles, we will see this arrangement actually occur if we wait long enough! If we follow our solution for a long

enough time, it tries everything that it can do, so to speak. This is not absolutely necessary in the simplest devices, but when systems get complicated enough, with enough atoms, it happens. Now there is something else the solution can do. If we solve the equations of motion, we may get certain functions such as $t + t^2 + t^3$. We claim that another solution would be $-t + t^2 - t^3$. In other words, if we substitute $-t$ everywhere for t throughout the entire solution, we will once again get a solution of the same equation. This follows from the fact that if we substitute $-t$ for t in the original differential equation, nothing is changed, since only second derivatives with respect to t appear. This means that if we have a certain motion, then the exact opposite motion is also possible. In the complete confusion which comes if we wait long enough, it finds itself going one way sometimes, and it finds itself going the other way sometimes. There is nothing more beautiful about one of the motions than about the other. So it is impossible to design a machine which, in the long run, is more likely to be going one way than the other, if the machine is sufficiently complicated.

One might think up an example for which this is obviously untrue. If we take a wheel, for instance, and spin it in empty space, it will go the same way forever. So there are some conditions, like the conservation of angular momentum, which violate the above argument. This just requires that the argument be made with a little more care. Perhaps the walls take up the angular momentum, or something like that, so that we have no special conservation laws. Then, if the system is complicated enough, the argument is true. It is based on the fact that the laws of mechanics are reversible.

For historical interest, we would like to remark on a device invented by Maxwell, who first worked out the dynamical theory of gases. He supposed the following situation: We have two boxes of gas at the same temperature, with a little hole between them. At the hole sits a little demon (who may be a machine of course!). There is a door on the hole, which can be opened or closed by the demon. He watches the molecules coming from the left. Whenever he sees a fast molecule, he opens the door. When he sees a slow one, he leaves it closed. If we want him to be an extra special demon, he can have eyes at the back of his head, and do the opposite to the molecules from the other side. He lets the slow ones through to the left, and the fast through to the right. Pretty soon the left side will get cold and the right side hot. Then, are the ideas of thermodynamics violated because we could have such a demon?

It turns out, if we build a finite-sized demon, that the demon himself gets so warm that he cannot see very well after a while. The simplest possible demon, as an example, would be a trap door held over the hole by a spring. A fast molecule comes through, because it is able to lift the trap door. The slow molecule cannot get through, and bounces back. But this thing is nothing but our ratchet and pawl in another form, and ultimately the mechanism will heat up. If we assume that the specific heat of the demon is not infinite, it must heat up. It has but a finite number of internal gears and wheels, so it cannot get rid of the extra heat that it gets from observing the molecules. Soon it is shaking from Brownian motion so much that it cannot tell whether it is coming or going, much less whether the molecules are coming or going, so it does not work.

46–4 Irreversibility

Are all the laws of physics reversible? Evidently not! Just try to unscramble an egg! Run a moving picture backwards, and it takes only a few minutes for everybody to start laughing. The most natural characteristic of all phenomena is their obvious irreversibility.

Where does irreversibility come from? It does not come from Newton's laws. If we claim that the behavior of everything is ultimately to be understood in terms of the laws of physics, and if it also turns out that all the equations have the fantastic property that if we put $t = -t$ we have another solution, then every phenomenon is reversible. How then does it come about in nature on a large scale that things are not reversible? Obviously there must be some law, some

obscure but fundamental equation, perhaps in electricity, maybe in neutrino physics, in which it *does* matter which way time goes.

Let us discuss that question now. We already know one of those laws, which says that the entropy is always increasing. If we have a hot thing and a cold thing, the heat goes from hot to cold. So the law of entropy is one such law. But we expect to understand the law of entropy from the point of view of mechanics. In fact, we have just been successful in obtaining all the consequences of the argument that heat cannot flow backwards by itself from just mechanical arguments, and we thereby obtained an understanding of the Second Law. Apparently we can get irreversibility from reversible equations. But *was* it *only* a mechanical argument that we used? Let us look into it more closely.

Since our question has to do with the entropy, our problem is to try to find a microscopic description of entropy. If we say we have a certain amount of energy in something, like a gas, then we can get a microscopic picture of it, and say that every atom has a certain energy. All these energies added together give us the total energy. Similarly, maybe every atom has a certain entropy. If we add everything up, we would have the total entropy. It does not work so well, but let us see what happens.

As an example, we calculate the entropy difference between a gas at a certain temperature at one volume, and a gas at the same temperature at another volume. We remember, from Chapter 44, that we had, for the change in entropy,

$$\Delta S = \int \frac{dQ}{T}.$$

In the present case, the energy of the gas is the same before and after expansion, since the temperature does not change. So we have to add enough heat to equal the work done by the gas or, for each little change in volume,

$$dQ = P\,dV.$$

Putting this in for dQ, we get

$$\Delta S = \int_{V_1}^{V_2} P\,\frac{dV}{T} = \int_{V_1}^{V_2} \frac{NkT}{V}\,\frac{dV}{T}$$

$$= Nk \ln \frac{V_2}{V_1},$$

as we obtained in Chapter 44. For instance, if we expand the volume by a factor of 2, the entropy change is $Nk \ln 2$.

Let us now consider another interesting example. Suppose we have a box with a barrier in the middle. On one side is neon ("black" molecules), and on the other, argon ("white" molecules). Now we take out the barrier, and let them mix. How much has the entropy changed? It is possible to imagine that instead of the barrier we have a piston, with holes in it that let the whites through but not the blacks, and another kind of piston which is the other way around. If we move one piston to each end, we see that, for each gas, the problem is like the one we just solved. So we get an entropy change of $Nk \ln 2$, which means that the entropy has increased by $k \ln 2$ per molecule. The 2 has to do with the extra room that the molecule has, which is rather peculiar. It is not a property of the molecule itself, but of *how much room* the molecule has to run around in. This is a strange situation, where entropy increases but where everything has the same temperature and the same energy! The only thing that is changed is that the molecules are distributed differently.

We well know that if we just pull the barrier out, everything will get mixed up after a long time, due to the collisions, the jiggling, the banging, and so on. Every once in a while a white molecule goes toward a black, and a black one goes toward a white, and maybe they pass. Gradually the whites worm their way, by accident, across into the space of blacks, and the blacks worm their way, by accident, into the space of whites. If we wait long enough we get a mixture.

Clearly, this is an irreversible process in the real world, and ought to involve an increase in the entropy.

Here we have a simple example of an irreversible process which is completely composed of reversible events. Every time there is a collision between any two molecules, they go off in certain directions. If we took a moving picture of a collision in reverse, there would be nothing wrong with the picture. In fact, one kind of collision is just as likely as another. So the mixing is completely reversible, and yet it is irreversible. Everyone knows that if we started with white and with black, separated, we would get a mixture within a few minutes. If we sat and looked at it for several more minutes, it would not separate again but would stay mixed. So we have an irreversibility which is based on reversible situations. But we also see the *reason* now. We started with an arrangement which is, in some sense, *ordered.* Due to the chaos of the collisions, it becomes disordered. *It is the change from an ordered arrangement to a disordered arrangement which is the source of the irreversibility.*

It is true that if we took a motion picture of this, and showed it backwards, we would see it gradually become ordered. Someone would say, "That is against the laws of physics!" So we would run the film over again, and we would look at every collision. Every one would be perfect, and every one would be obeying the laws of physics. The reason, of course, is that every molecule's velocities are just right, so if the paths are all followed back, they get back to their original condition. But that is a very unlikely circumstance to have. If we start with the gas in no special arrangement, just whites and blacks, it will never get back.

46–5 Order and entropy

So we now have to talk about what we mean by disorder and what we mean by order. It is not a question of pleasant order or unpleasant disorder. What is different in our mixed and unmixed cases is the following. Suppose we divide the space into little volume elements. If we have white and black molecules, how many ways could we distribute them among the volume elements so that white is on one side, and black on the other? On the other hand, how many ways could we distribute them with no restriction on which goes where? Clearly, there are many more ways to arrange them in the latter case. We measure "disorder" by the number of ways that the insides can be arranged, so that from the outside it looks the same. *The logarithm of that number of ways is the entropy.* The number of ways in the separated case is less, so the entropy is less, or the "disorder" is less.

So with the above technical definition of disorder we can understand the proposition. First, the entropy measures the disorder. Second, the universe always goes from "order" to "disorder," so entropy always increases. Order is not order in the sense that we like the arrangement, but in the sense that the number of different ways we can hook it up, and still have it look the same from the outside, is relatively restricted. In the case where we reversed our motion picture of the gas mixing, there was not as much disorder as we thought. Every single atom had exactly the correct speed and direction to come out right! The entropy was not high after all, even though it appeared so.

What about the reversibility of the other physical laws? When we talked about the electric field which comes from an accelerating charge, it was said that we must take the retarded field. At a time t and at a distance r from the charge, we take the field due to the acceleration at a time $t - r/c$, not $t + r/c$. So it looks, at first, as if the law of electricity is not reversible. Very strangely, however, the laws we used come from a set of equations called Maxwell's equations, which are, in fact, reversible. Furthermore, it is possible to argue that if we were to use only the advanced field, the field due to the state of affairs at $t + r/c$, and do it absolutely consistently in a completely enclosed space, everything happens exactly the same way as if we use retarded fields! This apparent irreversibility in electricity, at least in an enclosure, is thus not an irreversibility at all. We have some feeling for that already, because we know that when we have an oscillating charge which generates fields which are bounced from the walls of an enclosure we ultimately

get to an equilibrium in which there is no one-sidedness. The retarded field approach is only a convenience in the method of solution.

So far as we know, all the fundamental laws of physics, like Newton's equations, are reversible. Then where does irreversibility come from? It comes from order going to disorder, but we do not understand this until we know the origin of the order. Why is it that the situations we find ourselves in every day are always out of equilibrium? One possible explanation is the following. Look again at our box of mixed white and black molecules. Now it is possible, if we wait long enough, by sheer, grossly improbable, but possible, accident, that the distribution of molecules gets to be mostly white on one side and mostly black on the other. After that, as times goes on and accidents continue, they get more mixed up again.

Thus one possible explanation of the high degree of order in the present-day world is that it is just a question of luck. Perhaps our universe happened to have had a fluctuation of some kind in the past, in which things got somewhat separated, and now they are running back together again. This kind of theory is not unsymmetrical, because we can ask what the separated gas looks like either a little in the future or a little in the past. In either case, we see a grey smear at the interface, because the molecules are mixing again. No matter which way we run time, the gas mixes. So this theory would say the irreversibility is just one of the accidents of life.

We would like to argue that this is not the case. Suppose we do not look at the whole box at once, but only at a piece of the box. Then, at a certain moment, suppose we discover a certain amount of order. In this little piece, white and black are separate. What should we deduce about the condition in places where we have not yet looked? If we really believe that the order arose from complete disorder by a fluctuation, we must surely take the most likely fluctuation which could produce it, and the most likely condition is *not* that the rest of it has also become disentangled! Therefore, from the hypothesis that the world is a fluctuation, all of the predictions are that if we look at a part of the world we have never seen before, we will find it mixed up, and not like the piece we just looked at. If our order were due to a fluctuation, we would not expect order anywhere but where we have just noticed it.

Now we assume the separation is because the past of the universe was really ordered. It is not due to a fluctuation, but the whole thing used to be white and black. This theory now predicts that there will be order in other places—the order is not due to a fluctuation, but due to a much higher ordering at the beginning of time. Then we would expect to find order in places where we have not yet looked.

The astronomers, for example, have only looked at some of the stars. Every day they turn their telescopes to other stars, and the new stars are doing the same thing as the other stars. We therefore conclude that the universe is *not* a fluctuation, and that the order is a memory of conditions when things started. This is not to say that we understand the logic of it. For some reason, the universe at one time had a very low entropy for its energy content, and since then the entropy has increased. So that is the way toward the future. That is the origin of all irreversibility, that is what makes the processes of growth and decay, that makes us remember the past and not the future, remember the things which are closer to that moment in the history of the universe when the order was higher than now, and why we are not able to remember things where the disorder is higher than now, which we call the future. So, as we commented in an earlier chapter, the entire universe is in a glass of wine, if we look at it closely enough. In this case the glass of wine is complex, because there is water and glass and light and everything else.

Another delight of our subject of physics is that even simple and idealized things, like the ratchet and pawl, work only because they are part of the universe. The ratchet and pawl works in only one direction because it has some ultimate contact with the rest of the universe. If the ratchet and pawl were in a box and isolated for some sufficient time, the wheel would be no more likely to go one way than the other. But because we pull up the shades and let the light out, because

we cool off on the earth and get heat from the sun, the ratchets and pawls that we make can turn one way. This one-wayness is interrelated with the fact that the ratchet is part of the universe. It is part of the universe not only in the sense that it obeys the physical laws of the universe, but its one-way behavior is tied to the one-way behavior of the entire universe. It cannot be completely understood until the mystery of the beginnings of the history of the universe are reduced still further from speculation to scientific understanding.

47

Sound. The wave equation

47–1 Waves

47–2 The propagation of sound

47–3 The wave equation

47–4 Solutions of the wave equation

47–5 The speed of sound

47–1 Waves

In this chapter we shall discuss the phenomenon of *waves*. This is a phenomenon which appears in many contexts throughout physics, and therefore our attention should be concentrated on it not only because of the particular example considered here, which is sound, but also because of the much wider application of the ideas in all branches of physics.

It was pointed out when we studied the harmonic oscillator that there are not only mechanical examples of oscillating systems but electrical ones as well. Waves are related to oscillating systems, except that wave oscillations appear not only as time-oscillations at one place, but propagate in space as well.

We have really already studied waves. When we studied light, in learning about the properties of waves in that subject, we paid particular attention to the interference in space of waves from several sources at different locations and all at the same frequency. There are two important wave phenomena that we have not yet discussed which occur in light, i.e., electromagnetic waves, as well as in any other form of waves. The first of these is the phenomenon of *interference in time* rather than interference in space. If we have two sources of sound which have slightly different frequencies and if we listen to both at the same time, then sometimes the waves come with the crests together and sometimes with the crest and trough together (see Fig. 47–1). The rising and falling of the sound that results is the phenomenon of *beats* or, in other words, of interference in time. The second phenomenon involves the wave patterns which result when the waves are confined within a given volume and reflect back and forth from walls.

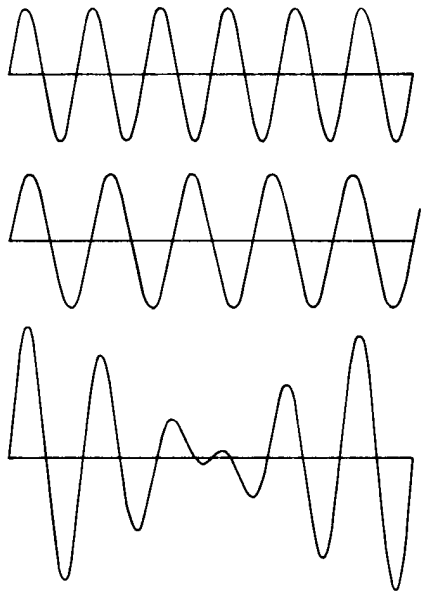


Fig. 47–1. Interference in time of two sound sources with slightly different frequencies, resulting in beats.

These effects could have been discussed, of course, for the case of electromagnetic waves. The reason for not having done this is that by using one example we would not generate the feeling that we are actually learning about many different subjects at the same time. In order to emphasize the general applicability of waves beyond electrodynamics, we consider here a different example, in particular sound waves.

Other examples of waves are water waves consisting of long swells that we see coming in to the shore, or the smaller water waves consisting of surface tension ripples. As another example, there are two kinds of elastic waves in solids; a compressional (or longitudinal) wave in which the particles of the solid oscillate back and forth along the direction of propagation of the wave (sound waves in a gas are of this kind), and a transverse wave in which the particles of the solid oscillate in a direction perpendicular to the direction of propagation. Earthquake waves contain elastic waves of both kinds, generated by a motion at some place in the earth's crust.

Still another example of waves is found in modern physics. These are waves which give the probability amplitude of finding a particle at a given place—the "matter waves" which we have already discussed. Their frequency is proportional to the energy and their wave number is proportional to the momentum. They are the waves of quantum mechanics.

In this chapter we shall consider only waves for which the velocity is independent of the wavelength. This is, for example, the case for light in a vacuum. The speed of light is then the same for radiowaves, blue light, green light, or for any other wavelength. Because of this behavior, when we began to describe the wave phenomenon we did not notice at first that we had wave propagation. Instead, we said that if a charge is moved at one place, the electric field at a distance x was

proportional to the acceleration, not at the time t, but at the earlier time $t - x/c$. Therefore if we were to picture the electric field in space at some instant of time, as in Fig. 47–2, the electric field at a time t later would have moved the distance ct, as indicated in the figure. Mathematically, we can say that in the one-dimensional example we are taking, the electric field is a function of $x - ct$. We see that at $t = 0$, it is some function of x. If we consider a later time, we need only to increase x somewhat to get the same value of the electric field. For example, if the maximum field occurred at $x = 3$ at time zero, then to find the new position of the maximum field at time t we need

$$x - ct = 3 \quad \text{or} \quad x = 3 + ct.$$

We see that this kind of function represents the propagation of a wave.

Such a function, $f(x - ct)$, then represents a wave. We may summarize this description of a wave by saying simply that

$$f(x - ct) = f(x + \Delta x - c(t + \Delta t)),$$

when $\Delta x = c\,\Delta t$. There is, of course, another possibility, i.e., that instead of a source to the left as indicated in Fig. 47–2, we have a source on the right, so that the wave propagates toward negative x. Then the wave would be described by $g(x + ct)$.

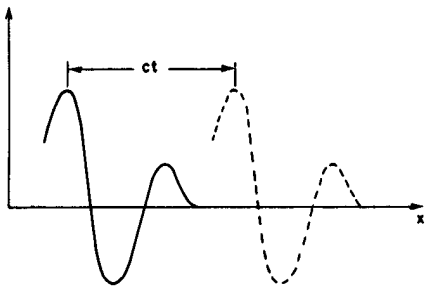


Fig. 47–2. The solid curve shows what the electric field might be like at some instant of time and the dashed curve shows what the electric field is at a time t later.

There is the additional possibility that more than one wave exists in space at the same time, and so the electric field is the sum of the two fields, each one propagating independently. This behavior of electric fields may be described by saying that if $f_1(x - ct)$ is a wave, and if $f_2(x - ct)$ is another wave, then their sum is also a wave. This is called the principle of superposition. The same principle is valid in sound.

We are familiar with the fact that if a sound is produced, we hear with complete fidelity the same sequence of sounds as was generated. If we had high frequencies travelling faster than low frequencies, a short, sharp noise would be heard as a succession of musical sounds. Similarly, if red light travelled faster than blue light, a flash of white light would be seen first as red, then as white, and finally as blue. We are familiar with the fact that this is not the case. Both sound and light travel with a speed in air which is very nearly independent of frequency. Examples of wave propagation for which this independence is not true will be considered in Chapter 48.

In the case of light (electromagnetic waves) we gave a rule which determined the electric field at a point as a result of the acceleration of a charge. One might expect now that what we should do is give a rule whereby some quality of the air, say the pressure, is determined at a given distance from a source in terms of the source motion, delayed by the travel time of the sound. In the case of light this procedure was acceptable because all that we knew was that a charge at one place exerts a force on another charge at another place. The details of propagation from the one place to the other were not absolutely essential. In the case of sound, however, we know that it propagates through the air between the source and the hearer, and it is certainly a natural question to ask what, at any given moment, the pressure of the air is. We would like, in addition, to know exactly how the air moves. In the case of electricity we could accept a rule, since we could say that we do not yet know the laws of electricity, but we cannot make the same remark with regard to sound. We would not be satisfied with a rule stating how the sound pressure moves through the air, because the process ought to be understandable as a consequence of the laws of mechanics. In short, sound is a branch of mechanics, and so it is to be understood in terms of Newton's laws. The propagation of sound from one place to another is merely a consequence of mechanics and the properties of gases, if it propagates in a gas, or of the properties of liquids or solids, if it propagates through such mediums. Later we shall derive the properties of light and its wave propagation in a similar way from the laws of electrodynamics.

47–2 The propagation of sound

We shall give a derivation of the properties of the propagation of sound *between* the source and the receiver as a consequence of Newton's laws, and we shall not consider the interaction with the source and the receiver. Ordinarily we emphasize a result rather than a particular derivation of it. In this chapter we take the opposite view. The point here, in a certain sense, is the derivation itself. This problem of explaining new phenomena in terms of old ones, when we know the laws of the old ones, is perhaps the greatest art of mathematical physics. The mathematical physicist has two problems: one is to find solutions, given the equations, and the other is to find the equations which describe a new phenomenon. The derivation here is an example of the second kind of problem.

We shall take the simplest example here—the propagation of sound in one dimension. To carry out such a derivation it is necessary first to have some kind of understanding of what is going on. Fundamentally what is involved is that if an object is moved at one place in the air, we observe that there is a disturbance which travels through the air. If we ask what kind of disturbance, we would say that we would expect that the motion of the object produces a change of pressure. Of course, if the object is moved gently, the air merely flows around it, but what we are concerned with is a rapid motion, so that there is not sufficient time for such a flow. Then, with the motion, the air is compressed and a change of pressure is produced which pushes on additional air. This air is in turn compressed, which leads again to an extra pressure, and a wave is propagated.

We now want to formulate such a process. We have to decide what variables we need. In our particular problem we would need to know how much the air has moved, so that the air *displacement* in the sound wave is certainly one relevant variable. In addition we would like to describe how the air *density* changes as it is displaced. The air *pressure* also changes, so this is another variable of interest. Then, of course, the air has a *velocity*, so that we shall have to describe the velocity of the air particles. The air particles also have *accelerations*—but as we list these many variables we soon realize that the velocity and acceleration would be known if we knew how the air *displacement* varies with time.

As we said, we shall consider the wave in one dimension. We can do this if we are sufficiently far from the source that what we call the *wavefronts* are very nearly planes. We thus make our argument simpler by taking the least complicated example. We shall then be able to say that the displacement, χ, depends only on x and t, and not on y and z. Therefore the description of the air is given by $\chi(x, t)$.

Is this description complete? It would appear to be far from complete, for we know none of the details of how the air molecules are moving. They are moving in all directions, and this state of affairs is certainly not described by means of this function $\chi(x, t)$. From the point of view of kinetic theory, if we have a higher density of molecules at one place and a lower density adjacent to that place, the molecules would move away from the region of higher density to the one of lower density, so as to equalize this difference. Apparently we would not get an oscillation and there would be no sound. What is necessary to get the sound wave is this situation: as the molecules rush out of the region of higher density and higher pressure, they give momentum to the molecules in the adjacent region of lower density. For sound to be generated, the regions over which the density and pressure change must be much larger than the distance the molecules travel before colliding with other molecules. This distance is the mean free path, and the distance between pressure crests and troughs must be much larger than this. Otherwise the molecules would move freely from the crest to the trough and immediately smear out the wave.

It is clear that we are going to describe the gas behavior on a scale large compared with the mean free path, and so the properties of the gas will not be described in terms of the individual molecules. The displacement, for example, will be the displacement of the center of mass of a small element of the gas, and the pressure or density will be the pressure or density in this region. We shall call the pressure P and the density ρ, and they will be functions of x and t. We must keep in mind that this description is an approximation which is valid only when these gas properties do not vary too rapidly with distance.

47–3 The wave equation

The physics of the phenomenon of sound waves thus involves three features:

I. The gas moves and changes the density.

II. The change in density corresponds to a change in pressure.

III. Pressure inequalities generate gas motion.

Let us consider II first. For a gas, a liquid, or a solid, the pressure is some function of the density. Before the sound wave arrives, we have equilibrium, with a pressure P_0 and a corresponding density ρ_0. A pressure P in the medium is connected to the density by some characteristic relation $P = f(\rho)$ and, in particular, the equilibrium pressure P_0 is given by $P_0 = f(\rho_0)$. The changes of pressure in sound from the equilibrium value are extremely small. A convenient unit for measuring pressure is the *bar*, where 1 bar $= 10^5$ n/m^2. The pressure of 1 standard atmosphere is very nearly 1 bar: 1 atm $= 1.0133$ bars. In sound we use a logarithmic scale of intensities since the sensitivity of the ear is roughly logarithmic. This scale is the decibel scale, in which the acoustic pressure level for the pressure amplitude P is defined as

$$I \text{ (acoustic pressure level)} = 20 \log_{10}(P/P_{\text{ref}}) \text{ in db}, \tag{47.1}$$

where the reference pressure $P_{\text{ref}} = 2 \times 10^{-10}$ bar. A pressure amplitude of $P = 10^3 P_{\text{ref}} = 2 \times 10^{-7}$ bar* corresponds to a moderately intense sound of 60 decibels. We see that the pressure changes in sound are extremely small compared with the equilibrium, or mean, pressure of 1 atm. The displacements and the density changes are correspondingly extremely small. In explosions we do not have such small changes; the excess pressures produced can be greater than 1 atm. These large pressure changes lead to new effects which we shall consider later. In sound we do not often consider acoustic intensity levels over 100 db; 120 db is a level which is painful to the ear. Therefore, for sound, if we write

$$P = P_0 + P_e, \qquad \rho = \rho_0 + \rho_e, \tag{47.2}$$

we shall always have the pressure change P_e very small compared with P_0 and the density change ρ_e very small compared with ρ_0. Then

$$P_0 + P_e = f(\rho_0 + \rho_e) = f(\rho_0) + \rho_e f'(\rho_0), \tag{47.3}$$

where $P_0 = f(\rho_0)$ and $f'(\rho_0)$ stands for the derivative of $f(\rho)$ evaluated at $\rho = \rho_0$. We can take the second step in this equality only because ρ_e is very small. We find in this way that the excess pressure P_e is proportional to the excess density ρ_e, and we may call the proportionality factor κ:

$$P_e = \kappa \rho_e, \quad \text{where } \kappa = f'(\rho_0) = (dP/d\rho)_0. \qquad \text{(II)} \tag{47.4}$$

The relation we needed for II is this very simple one.

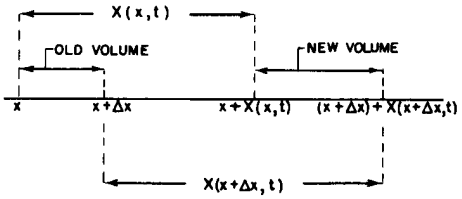


Fig. 47–3. The displacement of the air at x is $\chi(x, t)$, and at $x + \Delta x$ it is $\chi(x + \Delta x, t)$. The original volume of the air for a unit area of the plane wave is Δx; the new volume is $\Delta x + \chi(x + \Delta x, t) - \chi(x, t)$.

Let us now consider I. We shall suppose that the position of a portion of air undisturbed by the sound wave is x and the displacement at the time t due to the sound is $\chi(x, t)$, so that its new position is $x + \chi(x, t)$, as in Fig. 47–3. Now the undisturbed position of a nearby portion of air is $x + \Delta x$, and its new position is $x + \Delta x + \chi(x + \Delta x, t)$. We can now find the density changes in the following way. Since we are limiting ourselves to plane waves, we can take a unit area perpendicular to the x-direction, which is the direction of propagation of the sound wave. The amount of air, per unit area, in Δx is then $\rho_0 \Delta x$, where ρ_0 is the undisturbed, or equilibrium, air density. This air, when displaced by the sound wave, now lies between $x + \chi(x, t)$ and $x + \Delta x + \chi(x + \Delta x, t)$, so that we have the same matter in this interval that was in Δx when undisturbed. If ρ is the new density, then

$$\rho_0 \Delta x = \rho[x + \Delta x + \chi(x + \Delta x, t) - x - \chi(x, t)]. \tag{47.5}$$

---

* With this choice of P_{ref}, the P is not the peak pressure in the sound wave but the "root-mean-square" pressure, which is $1/(2)^{1/2}$ times the peak pressure.

Since Δx is small, we can write $\chi(x + \Delta x, t) - \chi(x, t) = (\partial\chi/\partial x)\,\Delta x$. This derivative is a partial derivative, since χ depends on the time as well as on x. Our equation then is

$$\rho_0\,\Delta x = \rho\left(\frac{\partial\chi}{\partial x}\,\Delta x + \Delta x\right) \tag{47.6}$$

or

$$\rho_0 = (\rho_0 + \rho_e)\,\frac{\partial\chi}{\partial x} + \rho_0 + \rho_e. \tag{47.7}$$

Now in sound waves all changes are small so that ρ_e is small, χ is small, and $\partial\chi/\partial x$ is also small. Therefore in the relation that we have just found,

$$\rho_e = -\rho_0\,\frac{\partial\chi}{\partial x} - \rho_e\,\frac{\partial\chi}{\partial x}, \tag{47.8}$$

we can neglect $\rho_e\,\partial\chi/\partial x$ compared with $\rho_0\,\partial\chi/\partial x$. Thus we get the relation we needed for I:

$$\rho_e = -\rho_0\,\frac{\partial\chi}{\partial x}. \qquad \text{(I)} \tag{47.9}$$

This equation is what we would expect physically. If the displacements vary with x, then there will be density changes. The sign is also right: if the displacement χ increases with x, so that the air is stretched out, the density must go down.

We now need the third equation, which is the equation of the motion produced by the pressure. If we know the relation between the force and the pressure, we can then get the equation of motion. If we take a thin slab of air of length Δx and of unit area perpendicular to x, then the mass of air in this slab is $\rho_0\,\Delta x$ and it has the acceleration $\partial^2\chi/\partial t^2$, so the mass times the acceleration for this slab of matter is $\rho_0\,\Delta x\,(\partial^2\chi/\partial t^2)$. (It makes no difference for small Δx whether the acceleration $\partial^2\chi/\partial t^2$ is evaluated at an edge of the slab or at some intermediate position.) If now we find the force on this matter for a unit area perpendicular to x, it will then be equal to $\rho_0\,\Delta x\,(\partial^2\chi/\partial t^2)$. We have the force in the $+x$-direction, at x, of amount $P(x, t)$ per unit area, and we have the force in the opposite direction, at $x + \Delta x$, of amount $P(x + \Delta x, t)$ per unit area (Fig. 47–4):

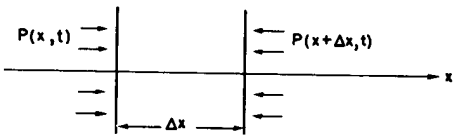


Fig. 47–4. The net force in the positive x-direction produced by the pressure acting on unit area perpendicular to x is $-(\partial P/\partial x)\,\Delta x$.

$$P(x, t) - P(x + \Delta x, t) = -\frac{\partial P}{\partial x}\,\Delta x = -\frac{\partial P_e}{\partial x}\,\Delta x, \tag{47.10}$$

since Δx is small and since the only part of P which changes is the excess pressure P_e. We now have III:

$$\rho_0\,\frac{\partial^2\chi}{\partial t^2} = -\frac{\partial P_e}{\partial x}, \qquad \text{(III)} \tag{47.11}$$

and so we have enough equations to interconnect things and reduce down to one variable, say to χ. We can eliminate P_e from III by using II, so that we get

$$\rho_0\,\frac{\partial^2\chi}{\partial t^2} = -\kappa\,\frac{\partial\rho_e}{\partial x}, \tag{47.12}$$

and then we can use I to eliminate ρ_e. In this way we find that ρ_0 cancels out and that we are left with

$$\frac{\partial^2\chi}{\partial t^2} = \kappa\,\frac{\partial^2\chi}{\partial x^2}. \tag{47.13}$$

We shall call $c_s^2 = \kappa$, so that we can write

$$\frac{\partial^2\chi}{\partial x^2} = \frac{1}{c_s^2}\,\frac{\partial^2\chi}{\partial t^2}. \tag{47.14}$$

This is the wave equation which describes the behavior of sound in matter.

47–4 Solutions of the wave equation

We now can see whether this equation really does describe the essential properties of sound waves in matter. We want to deduce that a sound pulse, or disturbance, will move with a constant speed. We want to verify that two different pulses can move through each other—the principle of superposition. We also want to verify that sound can go either to the right or to the left. All these properties should be contained in this one equation.

We have remarked that any plane-wave disturbance which moves with a constant velocity v has the form $f(x - vt)$. Now we have to see whether $\chi(x, t) = f(x - vt)$ is a solution of the wave equation. When we calculate $\partial\chi/\partial x$, we get the derivative of the function, $\partial\chi/\partial x = f'(x - vt)$. Differentiating once more, we find

$$\frac{\partial^2\chi}{\partial x^2} = f''(x - vt). \tag{47.15}$$

The differentiation of this same function with respect to t gives $-v$ times the derivative of the function, or $\partial\chi/\partial t = -vf'(x - vt)$, and the second time derivative is

$$\frac{\partial^2\chi}{\partial t^2} = v^2 f''(x - vt). \tag{47.16}$$

It is evident that $f(x - vt)$ will satisfy the wave equation provided the wave velocity v is equal to c_s.

We find, therefore, from the *laws of mechanics* that any sound disturbance propagates with the velocity c_s, and in addition we find that

$$c_s = \kappa^{1/2} = (dP/d\rho)_0^{1/2},$$

and so *we have related the wave velocity to a property of the medium.*

If we consider a wave travelling in the opposite direction, so that $\chi(x, t) = g(x + vt)$, it is easy to see that such a disturbance also satisfies the wave equation. The only difference between such a wave and one travelling from left to right is in the sign of v, but whether we have $x + vt$ or $x - vt$ as the variable in the function does not affect the sign of $\partial^2\chi/\partial t^2$, since it involves only v^2. It follows that we have a solution for waves propagating in either direction with speed c_s.

An extremely interesting question is that of superposition. Suppose one solution of the wave equation has been found, say χ_1. This means that the second derivative of χ_1 with respect to x is equal to $1/c_s^2$ times the second derivative of χ_1 with respect to t. Now any other solution χ_2 has this same property. If we superpose these two solutions, we have

$$\chi(x, t) = \chi_1(x, t) + \chi_2(x, t), \tag{47.17}$$

and we wish to verify that $\chi(x, t)$ is also a wave, i.e., that χ satisfies the wave equation. We can easily prove this result, since we have

$$\frac{\partial^2\chi}{\partial x^2} = \frac{\partial^2\chi_1}{\partial x^2} + \frac{\partial^2\chi_2}{\partial x^2} \tag{47.18}$$

and, in addition,

$$\frac{\partial^2\chi}{\partial t^2} = \frac{\partial^2\chi_1}{\partial t^2} + \frac{\partial^2\chi_2}{\partial t^2}. \tag{47.19}$$

It follows that $\partial^2\chi/\partial x^2 = (1/c_s^2)\,\partial^2\chi/\partial t^2$, so we have verified the principle of superposition. The proof of the principle of superposition follows from the fact that the wave equation is *linear* in χ.

We can now expect that a plane light wave propagating in the x-direction, polarized so that the electric field is in the y-direction, will satisfy the wave equation

$$\frac{\partial^2 E_y}{\partial x^2} = \frac{1}{c^2}\frac{\partial^2 E_y}{\partial t^2}, \tag{47.20}$$

where c is the speed of light. This wave equation is one of the consequences of Maxwell's equations. The equations of electrodynamics will lead to the wave equation for light just as the equations of mechanics lead to the wave equation for sound.

47–5 The speed of sound

Our deduction of the wave equation for sound has given us a *formula* which connects the wave speed with the rate of change of pressure with the density at the normal pressure:

$$c_s^2 = \left(\frac{dP}{d\rho}\right)_0. \tag{47.21}$$

In evaluating this rate of change, it is essential to know how the temperature varies. In a sound wave, we would expect that in the region of compression the temperature would be raised, and that in the region of rarefaction the temperature would be lowered. Newton was the first to calculate the rate of change of pressure with density, and he supposed that the temperature remained unchanged. He argued that the heat was conducted from one region to the other so rapidly that the temperature could not rise or fall. This argument gives the isothermal speed of sound, and it is wrong. The correct deduction was given later by Laplace, who put forward the opposite idea—that the pressure and temperature change adiabatically in a sound wave. The heat flow from the compressed region to the rarefied region is negligible so long as the wavelength is long compared with the mean free path. Under this condition the slight amount of heat flow in a sound wave does not affect the speed, although it gives a small absorption of the sound energy. We can expect correctly that this absorption increases as the wavelength approaches the mean free path, but these wavelengths are smaller by factors of about a million than the wavelengths of audible sound.

The actual variation of pressure with density in a sound wave is the one that allows no heat flow. This corresponds to the adiabatic variation, which we found to be $PV^\gamma = \text{const}$, where V was the volume. Since the density ρ varies inversely with V, the adiabatic connection between P and ρ is

$$P = \text{const}\,\rho^\gamma, \tag{47.22}$$

from which we get $dP/d\rho = \gamma P/\rho$. We then have for the speed of sound the relation

$$c_s^2 = \frac{\gamma P}{\rho}. \tag{47.23}$$

We can also write $c_s^2 = \gamma PV/\rho V$ and make use of the relation $PV = NkT$. Further, we see that ρV is the mass of gas, which can also be expressed as Nm, or as μ, where m is the mass of a molecule aud μ is the molecular weight. In this way we find that

$$c_s^2 = \frac{\gamma kT}{m} = \frac{\gamma RT}{\mu}, \tag{47.24}$$

from which it is evident that the speed of sound depends only on the gas temperature and not on the pressure or the density. We also have observed that

$$kT = \tfrac{1}{3}m\langle v^2\rangle, \tag{47.25}$$

where $\langle v^2\rangle$ is the mean square of the speed of the molecules. It follows that $c_s^2 = (\gamma/3)\langle v^2\rangle$, or

$$c_s = \left(\frac{\gamma}{3}\right)^{1/2} v_{\text{av}}. \tag{47.26}$$

This equation states that the speed of sound is some number which is roughly $1/(3)^{1/2}$ times some average speed, v_{av}, of the molecules (the square root of the

mean square velocity). In other words, the speed of sound is of the same order of magnitude as the speed of the molecules, and is actually somewhat less than this average speed.

Of course we could expect such a result, because a disturbance like a change in pressure is, after all, propagated by the motion of the molecules. However, such an argument does not tell us the precise propagation speed; it could have turned out that sound was carried primarily by the fastest molecules, or by the slowest molecules. It is reasonable and satisfying that the speed of sound is roughly $\frac{1}{2}$ of the average molecular speed v_{av}.

48

Beats

48–1 Adding two waves

48–2 Beat notes and modulation

48–3 Side bands

48–4 Localized wave trains

48–5 Probability amplitudes for particles

48–6 Waves in three dimensions

48–7 Normal modes

48–1 Adding two waves

Some time ago we discussed in considerable detail the properties of light waves and their interference—that is, the effects of the superposition of two waves from different sources. In all these analyses we assumed that the frequencies of the sources were all the same. In this chapter we shall discuss some of the phenomena which result from the interference of two sources which have *different* frequencies.

It is easy to guess what is going to happen. Proceeding in the same way as we have done previously, suppose we have two equal oscillating sources of the same frequency whose phases are so adjusted, say, that the signals arrive in phase at some point P. At that point, if it is light, the light is very strong; if it is sound, it is very loud; or if it is electrons, many of them arrive. On the other hand, if the arriving signals were 180° out of phase, we would get no signal at P, because the net amplitude there is then a minimum. Now suppose that someone twists the "phase knob" of one of the sources and changes the phase at P back and forth, say, first making it 0° and then 180°, and so on. Of course, we would then find variations in the net signal strength. Now we also see that if the phase of one source is slowly changing relative to that of the other in a gradual, uniform manner, starting at zero, going up to ten, twenty, thirty, forty degrees, and so on, then what we would measure at P would be a series of strong and weak "pulsations," because when the phase shifts through 360° the amplitude returns to a maximum. Of course, to say that one source is shifting its phase relative to another at a uniform rate is the same as saying that the number of oscillations per second is slightly different for the two.

So we know the answer: if we have two sources at slightly different frequencies we should find, as a net result, an oscillation with a slowly pulsating intensity. That is all there really is to the subject!

It is very easy to formulate this result mathematically also. Suppose, for example, that we have two waves, and that we do not worry for the moment about all the spatial relations, but simply analyze what arrives at P. From one source, let us say, we would have $\cos \omega_1 t$, and from the other source, $\cos \omega_2 t$, where the two ω's are not exactly the same. Of course the amplitudes may not be the same, either, but we can solve the general problem later; let us first take the case where the amplitudes are equal. Then the total amplitude at P is the sum of these two cosines. If we plot the amplitudes of the waves against the time, as in Fig. 48–1, we see that

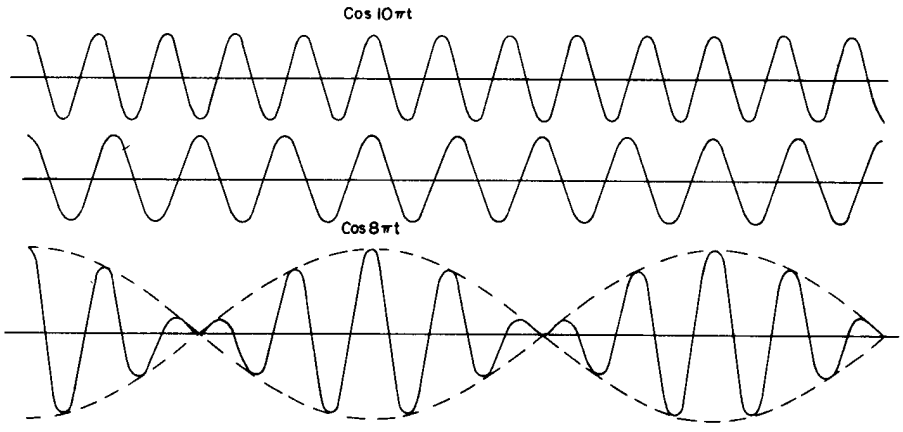


Fig. 48–1. The superposition of two cosine waves with frequencies in the ratio 8:10. The precise repetition of the pattern within each "beat" is not typical of the general case.

where the crests coincide we get a strong wave, and where a trough and crest coincide we get practically zero, and then when the crests coincide again we get a strong wave again.

Mathematically, we need only to add two cosines and rearrange the result somehow. There exist a number of useful relations among cosines which are not difficult to derive. Of course we know that

$$e^{i(a+b)} = e^{ia}e^{ib}, \tag{48.1}$$

and that e^{ia} has a real part, $\cos a$, and an imaginary part, $\sin a$. If we take the real part of $e^{i(a+b)}$, we get $\cos(a + b)$. If we multiply out:

$$e^{ia}e^{ib} = (\cos a + i \sin a)(\cos b + i \sin b),$$

we get $\cos a \cos b - \sin a \sin b$, plus some imaginary parts. But we now need only the real part, so we have

$$\cos(a + b) = \cos a \cos b - \sin a \sin b. \tag{48.2}$$

Now if we change the sign of b, since the cosine does not change sign while the sine does, the same equation, for negative b, is

$$\cos(a - b) = \cos a \cos b + \sin a \sin b. \tag{48.3}$$

If we add these two equations together, we lose the sines and we learn that the product of two cosines is half the cosine of the sum, plus half the cosine of the difference:

$$\cos a \cos b = \tfrac{1}{2}\cos(a + b) + \tfrac{1}{2}\cos(a - b). \tag{48.4}$$

Now we can also reverse the formula and find a formula for $\cos\alpha + \cos\beta$ if we simply let $\alpha = a + b$ and $\beta = a - b$. That is, $a = \frac{1}{2}(\alpha + \beta)$ and $b = \frac{1}{2}(\alpha - \beta)$, so that

$$\cos\alpha + \cos\beta = 2\cos\tfrac{1}{2}(\alpha + \beta)\cos\tfrac{1}{2}(\alpha - \beta). \tag{48.5}$$

Now we can analyze our problem. The sum of $\cos\omega_1 t$ and $\cos\omega_2 t$ is

$$\cos\omega_1 t + \cos\omega_2 t = 2\cos\tfrac{1}{2}(\omega_1 + \omega_2)t\cos\tfrac{1}{2}(\omega_1 - \omega_2)t. \tag{48.6}$$

Now let us suppose that the two frequencies are nearly the same, so that $\frac{1}{2}(\omega_1 + \omega_2)$ is the average frequency, and is more or less the same as either. But $\omega_1 - \omega_2$ is *much smaller* than ω_1 or ω_2 because, as we suppose, ω_1 and ω_2 are nearly equal. That means that we can represent the solution by saying that there is a high-frequency cosine wave more or less like the ones we started with, but that its "size" is slowly changing—its "size" is pulsating with a frequency which appears to be $\frac{1}{2}(\omega_1 - \omega_2)$. But is this the frequency at which the beats are heard? Although (48.6) says that the amplitude goes as $\cos\frac{1}{2}(\omega_1 - \omega_2)$, what it is really telling us is that the high-frequency oscillations are contained between two opposed cosine curves (shown dotted in Fig. 48–1). On this basis one could say that the amplitude varies at the frequency $\frac{1}{2}(\omega_1 - \omega_2)$, but if we are talking about the *intensity* of the wave we must think of it as having twice this frequency. That is, the modulation of the amplitude, in the sense of the strength of its intensity, is at frequency $\omega_1 - \omega_2$, although the formula tells us that we multiply by a cosine wave at half that frequency. The technical basis for the difference is that the high frequency-wave has a little different phase relationship in the second half-cycle.

Ignoring this small complication, we may conclude that if we add two waves of frequency ω_1 and ω_2, we will get a net resulting wave of average frequency $\frac{1}{2}(\omega_1 + \omega_2)$ which oscillates in strength with a frequency $\omega_1 - \omega_2$.

If the two amplitudes are different, we can do it all over again by multiplying the cosines by different amplitudes A_1 and A_2, and do a lot of mathematics, rearranging, and so on, using equations like (48.2)–(48.5). However, there are other, easier ways of doing the same analysis. For example, we know that it is much

easier to work with exponentials than with sines and cosines and that we can represent $A_1 \cos \omega_1 t$ as the real part of $A_1 e^{i\omega_1 t}$. The other wave would similarly be the real part of $A_2 e^{i\omega_2 t}$. If we add the two, we get $A_1 e^{i\omega_1 t} + A_2 e^{i\omega_2 t}$. If we then factor out the average frequency, we have

$$A_1 e^{i\omega_1 t} + A_2 e^{i\omega_2 t} = e^{1/2 i(\omega_1+\omega_2)t}[A_1 e^{1/2 i(\omega_1-\omega_2)t} + A_2 e^{-1/2 i(\omega_1-\omega_2)t}]. \quad (48.7)$$

Again we have the high-frequency wave with a modulation at the lower frequency.

48–2 Beat notes and modulation

If we are now asked for the intensity of the wave of Eq. (48.7), we can either take the absolute square of the left side, or of the right side. Let us take the left side. The intensity then is

$$I = A_1^2 + A_2^2 + 2A_1A_2 \cos (\omega_1 - \omega_2)t. \quad (48.8)$$

We see that the intensity swells and falls at a frequency $\omega_1 - \omega_2$, varying between the limits $(A_1 + A_2)^2$ and $(A_1 - A_2)^2$. If $A_1 \neq A_2$, the minimum intensity is not zero.

One more way to represent this idea is by means of a drawing, like Fig. 48–2. We draw a vector of length A_1, rotating at a frequency ω_1, to represent one of the waves in the complex plane. We draw another vector of length A_2, going around at a frequency ω_2, to represent the second wave. If the two frequencies are exactly equal, their resultant is of fixed length as it keeps revolving, and we get a definite, fixed intensity from the two. But if the frequencies are slightly different, the two complex vectors go around at different speeds. Figure 48–3 shows what the situation looks like relative to the vector $A_1 e^{i\omega_1 t}$. We see that A_2 is turning slowly away from A_1, and so the amplitude that we get by adding the two is first strong, and then, as it opens out, when it gets to the 180° relative position the resultant gets particularly weak, and so on. As the vectors go around, the amplitude of the sum vector gets bigger and smaller, and the intensity thus pulsates. It is a relatively simple idea, and there are many different ways of representing the same thing.

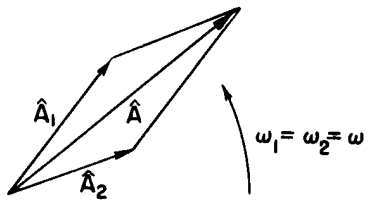


Fig. 48–2. The resultant of two complex vectors of equal frequency.

The effect is very easy to observe experimentally. In the case of acoustics, we may arrange two loudspeakers driven by two separate oscillators, one for each loudspeaker, so that they each make a tone. We thus receive one note from one source and a different note from the other source. If we make the frequencies exactly the same, the resulting effect will have a definite strength at a given space location. If we then de-tune them a little bit, we hear some variations in the intensity. The farther they are de-tuned, the more rapid are the variations of sound. The ear has some trouble following variations more rapid than ten or so per second.

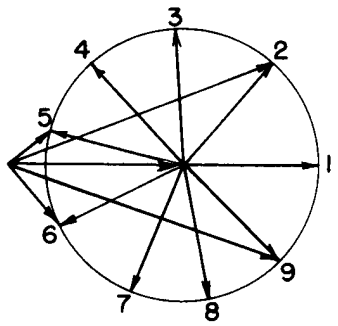


Fig. 48–3. The resultant of two complex vectors of unequal frequency, as seen in the rotating frame of reference of one vector. Nine successive positions of the slowly rotating vector are shown.

We may also see the effect on an oscilloscope which simply displays the sum of the currents to the two speakers. If the frequency of pulsing is relatively low, we simply see a sinusoidal wave train whose amplitude pulsates, but as we make the pulsations more rapid we see the kind of wave shown in Fig. 48–1. As we go to greater frequency differences, the "bumps" move closer together. Also, if the amplitudes are not equal and we make one signal stronger than the other, then we get a wave whose amplitude does not ever become zero, just as we expect. Everything works the way it should, both acoustically and electrically.

The opposite phenomenon occurs too! In radio transmission using so-called *amplitude modulation* (AM), the sound is broadcast by the radio station as follows: the radio transmitter has an AC electric oscillation which is at a very high frequency, for example 800 kilocycles per second, in the broadcast band. If this *carrier signal* is turned on, the radio station emits a wave which is of uniform amplitude at 800,000 oscillations a second. The way the "information" is transmitted, the useless kind of information about what kind of car to buy, is that when somebody talks into a microphone the amplitude of the carrier signal is changed in step with the vibrations of sound entering the microphone.

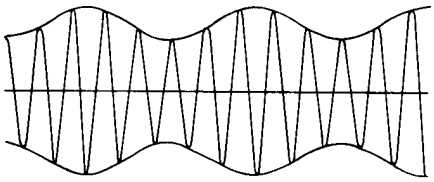


Fig. 48–4. A modulated carrier wave. In this schematic sketch, $\omega_c/\omega_m = 5$. In an actual radiowave, $\omega_c/\omega_m \sim 100$.

If we take as the simplest mathematical case the situation where a soprano is singing a perfect note, with perfect sinusoidal oscillations of her vocal cords, then we get a signal whose strength is alternating as shown in Fig. 48–4. The audiofrequency alternation is then recovered in the receiver; we get rid of the carrier wave and just look at the envelope which represents the oscillations of the vocal cords, or the sound of the singer. The loudspeaker then makes corresponding vibrations at the same frequency in the air, and the listener is then essentially unable to tell the difference, so they say. Because of a number of distortions and other subtle effects, it is, in fact, possible to tell whether we are listening to a radio or to a real soprano; otherwise the idea is as indicated above.

48–3 Side bands

Mathematically, the modulated wave described above would be expressed as

$$S = (1 + b\cos\omega_m t)\cos\omega_c t, \tag{48.9}$$

where ω_c represents the frequency of the carrier and ω_m is the frequency of the audio tone. Again we use all those theorems about the cosines, or we can use $e^{i\theta}$; it makes no difference—it is easier with $e^{i\theta}$, but it is the same thing. We then get

$$S = \cos\omega_c t + \tfrac{1}{2}b\cos(\omega_c + \omega_m)t + \tfrac{1}{2}b\cos(\omega_c - \omega_m)t. \tag{48.10}$$

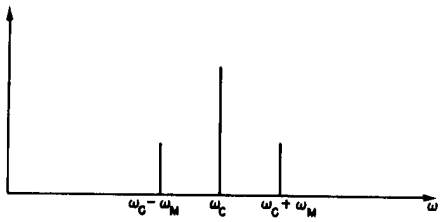


Fig. 48–5. The frequency spectrum of a carrier wave ω_c modulated by a single cosine wave ω_m.

So, from another point of view, we can say that the output wave of the system consists of three waves added in superposition: first, the regular wave at the frequency ω_c, that is, at the carrier frequency, and then two new waves at two new frequencies. One is the carrier frequency plus the modulation frequency, and the other is the carrier frequency minus the modulation frequency. If, therefore, we make some kind of plot of the intensity being generated by the generator as a function of frequency, we would find a lot of intensity at the frequency of the carrier, naturally, but when a singer started to sing, we would suddenly also find intensity proportional to the strength of the singer, b^2, at frequency $\omega_c + \omega_m$ and $\omega_c - \omega_m$, as shown in Fig. 48–5. These are called *side bands;* when there is a modulated signal from the transmitter, there are side bands. If there is more than one note at the same time, say ω_m and $\omega_{m'}$, there are two instruments playing; or if there is any other complicated cosine wave, then, of course, we can see from the mathematics that we get some more waves that correspond to the frequencies $\omega_c \pm \omega_{m'}$.

Therefore, when there is a complicated modulation that can be represented as the sum of many cosines,* we find that the actual transmitter is transmitting over a range of frequencies, namely the carrier frequency plus or minus the maximum frequency that the modulation signal contains.

Although at first we might believe that a radio transmitter transmits only at the nominal frequency of the carrier, since there are big, superstable crystal oscillators in there, and everything is adjusted to be at precisely 800 kilocycles, the moment someone *announces* that they are at 800 kilocycles, he modulates the 800 kilocycles, and so they are no longer precisely at 800 kilocycles! Suppose that the amplifiers are so built that they are able to transmit over a good range of the ear's sensitivity (the ear can hear up to 20,000 cycles per second, but usually radio transmitters and receivers do not work beyond 10,000, so we do not hear the highest parts), then, when the man speaks, his voice may contain frequencies ranging up, say, to 10,000 cycles, so the transmitter is transmitting frequencies

* A slight side remark: In what circumstances can a curve be represented as a sum of a lot of cosines? *Answer:* In all ordinary circumstances, except for certain cases the mathematicians can dream up. Of course, the curve must have only one value at a given point, and it must not be a crazy curve which jumps an infinite number of times in an infinitesimal distance, or something like that. But aside from such restrictions any reasonable curve (one that a singer is going to be able to make by shaking her vocal cords) can always be compounded by adding cosine waves together.

which may range from 790 to 810 kilocycles per second. Now if there were another station at 795 kc/sec, there would be a lot of confusion. Also, if we made our receiver so sensitive that it picked up only 800, and did not pick up the 10 kilocycles on either side, we would not hear what the man was saying, because the information would be on these other frequencies! Therefore it is absolutely essential to keep the stations a certain distance apart, so that their side bands do not overlap and, also, the receiver must not be so selective that it does not permit reception of the side bands as well as of the main nominal frequency. In the case of sound, this problem does not really cause much trouble. We can hear over a ± 20 kc/sec range, and we have usually from 500 to 1500 kc/sec in the broadcast band, so there is plenty of room for lots of stations.

The television problem is more difficult. As the electron beam goes across the face of the picture tube, there are various little spots of light and dark. That "light" and "dark" is the "signal." Now ordinarily the beam scans over the whole picture, 500 lines, approximately, in a thirtieth of a second. Let us consider that the resolution of the picture vertically and horizontally is more or less the same, so that there are the same number of spots per inch along a scan line. We want to be able to distinguish dark from light, dark from light, dark from light, over, say, 500 lines. In order to be able to do this with cosine waves, the shortest wavelength needed thus corresponds to a wavelength, from maximum to maximum, of one 250th of the screen size. So we have $250 \times 500 \times 30$ pieces of information per second. The highest frequency that we are going to carry, therefore, is close to 4 megacycles per second. Actually, to keep the television stations apart, we have to use a little bit more than this, about 6 mc/sec; part of it is used to carry the sound signal, and other information. So, television channels are 6 megacycles per second wide. It certainly would not be possible to transmit TV on an 800 kc/sec carrier, since we cannot modulate at a higher frequency than the carrier.

At any rate, the television band starts at 54 megacycles. The first transmission channel, which is channel 2 (!), has a frequency range from 54 to 60 mc/sec, which is 6 mc/sec wide. "But," one might say, "we have just proved that there were side bands on both sides, and therefore it should be twice that wide." It turns out that the radio engineers are rather clever. If we analyze the modulation signal using not just cosine terms, but cosine and sine terms, to allow for phase differences, we then see that there is a definite, invariant relationship between the side band on the high-frequency side and the side band on the low-frequency side. What we mean is that there is no new information on that other side band. So what is done is to suppress one side band, and the receiver is wired inside such that the information which is missing is reconstituted by looking at the single side band and the carrier. Single side-band transmission is a clever scheme for decreasing the band widths needed to transmit information.

48–4 Localized wave trains

The next subject we shall discuss is the interference of waves in both space and time. Suppose that we have two waves travelling in space. We know, of course, that we can represent a wave travelling in space by $e^{i(\omega t-kx)}$. This might be, for example, the displacement in a sound wave. This is a solution of the wave equation provided that $\omega^2 = k^2c^2$, where c is the speed of propagation of the wave. In this case we can write it as $e^{ik(x-ct)}$, which is of the general form $f(x - ct)$. Therefore this must be a wave which is travelling at this velocity, ω/k, and that is c and everything is all right.

Now we want to add two such waves together. Suppose we have a wave that is travelling with one frequency, and another wave travelling with another frequency. We leave to the reader to consider the case where the amplitudes are different; it makes no real difference. Thus we want to add $e^{i(\omega_1 t-k_1 x)} + e^{i(\omega_2 t-k_2 x)}$. We can add these by the same kind of mathematics we used when we added signal waves. Of course, if c is the same for both, this is easy, since it is the same as what we did before:

$$e^{i\omega_1(t-x/c)} + e^{i\omega_2(t-x/c)} = e^{i\omega_1 t'} + e^{i\omega_2 t'}, \tag{48.11}$$

except that $t' = t - x/c$ is the variable instead of t. So we get the same kind of modulations, naturally, but we see, of course, that those modulations are moving along with the wave. In other words, if we added two waves, but these waves were not just oscillating, but also moving in space, then the resultant wave would move along also, at the same speed.

Now we would like to generalize this to the case of waves in which the relationship between the frequency and the wave number k is not so simple. Example: material having an index of refraction. We have already studied the theory of the index of refraction in Chapter 31, where we found that we could write $k = n\omega/c$, where n is the index of refraction. As an interesting example, for x-rays we found that the index n is

$$n = 1 - \frac{Nq_e^2}{2\epsilon_0 m\omega^2}. \tag{48.12}$$

We actually derived a more complicated formula in Chapter 31, but this one is as good as any, as an example.

Incidentally, we know that even when ω and k are not linearly proportional, the ratio ω/k is certainly the speed of propagation for the particular frequency and wave number. We call this ratio the *phase velocity;* it is the speed at which the phase, or the nodes of a single wave, would move along:

$$v_p = \frac{\omega}{k}. \tag{48.13}$$

This phase velocity, for the case of x-rays in glass, is greater than the speed of light in vacuum (since n in 48.12 is less than 1), and that is a bit bothersome, because we do not think we can send signals faster than the speed of light!

What we are going to discuss now is the interference of two waves in which ω and k have a definite formula relating them. The above formula for n says that k is given as a definite function of ω. To be specific, in this particular problem, the formula for k in terms of ω is

$$k = \frac{\omega}{c} - \frac{a}{\omega c}, \tag{48.14}$$

where $a = Nq_e^2/2\epsilon_0 m$, a constant. At any rate, for each frequency there is a definite wave number, and we want to add two such waves together.

Let us do it just as we did in Eq. (48.7):

$$e^{i(\omega_1 t - k_1 x)} + e^{i(\omega_2 t - k_2 x)} = e^{1/2\,i[(\omega_1+\omega_2)t-(k_1+k_2)x]} \times \{e^{1/2\,i[(\omega_1-\omega_2)t-(k_1-k_2)x]} + e^{-1/2\,i[(\omega_1-\omega_2)t-(k_1-k_2)x]}\}. \tag{48.15}$$

So we have a modulated wave again, a wave which travels with the mean frequency and the mean wave number, but whose strength is varying with a form which depends on the difference frequency and the difference wave number.

Now let us take the case that the difference between the two waves is relatively small. Let us suppose that we are adding two waves whose frequencies are nearly equal; then $(\omega_1 + \omega_2)/2$ is practically the same as either one of the ω's, and similarly for $(k_1 + k_2)/2$. Thus the speed of the wave, the fast oscillations, the nodes, is still essentially ω/k. But look, the speed of propagation of the modulation is not the same! How much do we have to change x to account for a certain amount of t? The speed of this modulation wave is the ratio

$$v_M = \frac{\omega_1 - \omega_2}{k_1 - k_2}. \tag{48.16}$$

The speed of modulation is sometimes called the *group velocity.* If we take the case that the difference in frequency is relatively small, and the difference in wave number is then also relatively small, then this expression approaches, in the limit,

$$v_g = \frac{d\omega}{dk}. \tag{48.17}$$

In other words, for the slowest modulation, the slowest beats, there is a definite speed at which they travel which is not the same as the phase speed of the waves—what a mysterious thing!

The group velocity is the derivative of ω with respect to k, and the phase velocity is ω/k.

Let us see if we can understand why. Consider two waves, again of slightly different wavelength, as in Fig. 48–1. They are out of phase, in phase, out of phase, and so on. Now these waves represent, really, the waves in space travelling with slightly different frequencies also. Now because the phase velocity, the velocity of the nodes of these two waves, is not precisely the same, something new happens. Suppose we ride along with one of the waves and look at the other one; if they both went at the same speed, then the other wave would stay right where it was relative to us, as we ride along on this crest. We ride on that crest and right opposite us we see a crest; if the two velocities are equal the crests stay on top of each other. But it is *not so* that the two velocities are really equal. There is only a small difference in frequency and therefore only a small difference in velocity, but because of that difference in velocity, as we ride along the other wave moves slowly forward, say, or behind, relative to our wave. So as time goes on, what happens to the node? If we move one wave train just a shade forward, the node moves forward (or backward) a considerable distance. That is, the sum of these two waves has an envelope, and as the waves travel along, the envelope rides on them at a different speed. The *group velocity* is the speed at which modulated signals would be transmitted.

If we made a signal, i.e., some kind of change in the wave that one could recognize when he listened to it, a kind of modulation, then that modulation would travel at the group velocity, provided that the modulations were relatively slow. (When they are fast, it is much more difficult to analyze.)

Now we may show (at long last), that the speed of propagation of x-rays in a block of carbon is *not* greater than the speed of light, although the phase velocity *is* greater than the speed of light. In order to do that, we must find $d\omega/dk$, which we get by differentiating (48.14): $dk/d\omega = 1/c + a/\omega^2 c$. The group velocity, therefore, is the reciprocal of this, namely,

$$v_g = \frac{c}{1 + a/\omega^2}, \tag{48.18}$$

which is smaller than c! So although the phases can travel faster than the speed of light, the modulation signals travel slower, and that is the resolution of the apparent paradox! Of course, if we have the simple case that $\omega = kc$, then $d\omega/dk$ is also c. So when all the phases have the same velocity, naturally the group has the same velocity.

48–5 Probability amplitudes for particles

Let us now consider one more example of the phase velocity which is extremely interesting. It has to do with quantum mechanics. We know that the amplitude to find a particle at a place can, in some circumstances, vary in space and time, let us say in one dimension, in this manner:

$$\psi = Ae^{i(\omega t - kx)}, \tag{48.19}$$

where ω is the frequency, which is related to the classical idea of the energy through $E = \hbar\omega$, and k is the wave number, which is related to the momentum through $p = \hbar k$. We would say the particle had a definite momentum p if the wave number were exactly k, that is, a perfect wave which goes on with the same amplitude everywhere. Equation (48.19) gives the amplitude, and if we take the absolute square, we get the relative probability for finding the particle as a function of position and time. This is a *constant*, which means that the probability is the same to find a particle anywhere. Now suppose, instead, that we have a situation where we know that the particle is more likely to be at one place than at another. We would

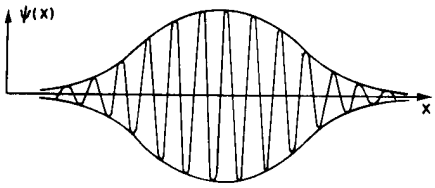


Fig. 48–6. A localized wave train.

represent such a situation by a wave which has a maximum and dies out on either side (Fig. 48–6). (It is not quite the same as a wave like (48.1) which has a series of maxima, but it is possible, by adding several waves of nearly the same ω and k together, to get rid of all but one maximum.)

Now in those circumstances, since the square of (48.19) represents the chance of finding a particle somewhere, we know that at a given instant the particle is most likely to be near the center of the "lump," where the amplitude of the wave is maximum. If now we wait a few moments, the waves will move, and after some time the "lump" will be somewhere else. If we knew that the particle originally was situated somewhere, classically, we would *expect* that it would later be elsewhere as a matter of fact, because it has a *speed*, after all, and a momentum. The quantum theory, then, will go into the correct classical theory for the relationship of momentum, energy, and velocity only if the group velocity, the velocity of the modulation, is equal to the velocity that we would obtain classically for a particle of the same momentum.

It is now necessary to demonstrate that this is, or is not, the case. According to the classical theory, the energy is related to the velocity through an equation like

$$E = \frac{mc^2}{\sqrt{1 - v^2/c^2}}. \tag{48.20}$$

Similarly, the momentum is

$$p = \frac{mv}{\sqrt{1 - v^2/c^2}}. \tag{48.21}$$

That is the classical theory, and as a consequence of the classical theory, by eliminating v, we can show that

$$E^2 - p^2c^2 = m^2c^4.$$

That is the four-dimensional grand result that we have talked and talked about, that $p_\mu p_\mu = m^2$; that is the relation between energy and momentum in the classical theory. Now that means, since these E's and p's are going to become ω's and k's, by substitution of $E = \hbar\omega$ and $p = \hbar k$, that for quantum mechanics it is necessary that

$$\frac{\hbar^2\omega^2}{c^2} - \hbar^2k^2 = m^2c^2. \tag{48.22}$$

This, then, is the relationship between the frequency and the wave number of a quantum-mechanical amplitude wave representing a particle of mass m. From this equation we can deduce that ω is

$$\omega = c\sqrt{k^2 + m^2c^2/\hbar^2}.$$

The phase velocity, ω/k, is here again faster than the speed of light!

Now let us look at the group velocity. The group velocity should be $d\omega/dk$, the speed at which the modulations move. We have to differentiate a square root, which is not very difficult. The derivative is

$$\frac{d\omega}{dk} = \frac{kc}{\sqrt{k^2 + m^2c^2/\hbar^2}}.$$

Now the square root is, after all, ω, so we could write this as $d\omega/dk = c^2k/\omega$. Further, k/ω is p/E, so

$$v_g = \frac{c^2p}{E}.$$

But from (48.20) and (48.21), $c^2p/E = v$, the velocity of the particle, according to classical mechanics. So we see that whereas the fundamental quantum-mechanical relationship $E = \hbar\omega$ and $p = \hbar k$, for the identification of ω and k with the classical E and p, only produces the equation $\omega^2 - k^2c^2 = m^2c^4/\hbar^2$, now we also understand the relationships (48.20) and (48.21) which connected E and p

to the velocity. Of course the group velocity must be the velocity of the particle if the interpretation is going to make any sense. If we think the particle is over here at one time, and then ten minutes later we think it is over there, as the quantum mechanics said, the distance traversed by the "lump," divided by the time interval, must be, classically, the velocity of the particle.

48–6 Waves in three dimensions

We shall now bring our discussion of waves to a close with a few general remarks about the wave equation. These remarks are intended to give some view of the future—not that we can understand everything exactly just now, but rather to see what things are going to look like when we study waves a little more. First of all, the wave equation for sound in one dimension was

$$\frac{\partial^2\chi}{\partial x^2} = \frac{1}{c^2}\frac{\partial^2\chi}{\partial t^2},$$

where c is the speed of whatever the wave is—in the case of sound, it is the sound speed; in the case of light, it is the speed of light. We showed that for a sound wave the displacements would propagate themselves at a certain speed. But the excess pressure also propagates at a certain speed, and so does the excess density. So we should expect that the pressure would satisfy the same equation, as indeed it does. We shall leave it to the reader to prove that it does. *Hint:* ρ_e is proportional to the rate of change of χ with respect to x. Therefore if we differentiate the wave equation with respect to x, we will immediately discover that $\partial\chi/\partial x$ satisfies the same equation. That is to say, ρ_e satisfies the same equation. But P_e is proportional to ρ_e, and therefore P_e does too. So the pressure, the displacements, everything, satisfy the same wave equation.

Usually one sees the wave equation for sound written in terms of pressure instead of in terms of displacement, because the pressure is a scalar and has no direction. But the displacement is a vector and has direction, and it is thus easier to analyze the pressure.

The next matter we discuss has to do with the wave equation in three dimensions. We know that the sound wave solution in one dimension is $e^{i(\omega t-kx)}$, with $\omega = kc_s$, but we also know that in three dimensions a wave would be represented by $e^{i(\omega t-k_x x-k_y y-k_z z)}$, where, in this case, $\omega^2 = k^2c_s^2$, which is, of course, $(k_x^2+k_y^2+k_z^2)c_s^2$. Now what we want to do is to guess what the correct wave equation in three dimensions is. Naturally, for the case of sound this can be deduced by going through the same dynamic argument in three dimensions that we made in one dimension. But we shall not do that; instead we just write down what comes out: the equation for the pressure (or displacement, or anything) is

$$\frac{\partial^2 P_e}{\partial x^2}+\frac{\partial^2 P_e}{\partial y^2}+\frac{\partial^2 P_e}{\partial z^2}=\frac{1}{c_s^2}\frac{\partial^2 P_e}{\partial t^2}. \tag{48.23}$$

That this is true can be verified by substituting in $e^{i(\omega t-\mathbf{k}\cdot\mathbf{r})}$. Clearly, every time we differentiate with respect to x, we multiply by $-ik_x$. If we differentiate twice, it is equivalent to multiplying by $-k_x^2$, so the first term would become $-k_x^2P_e$, for that wave. Similarly, the second term becomes $-k_y^2P_e$, and the third term becomes $-k_z^2P_e$. On the right, we get $-(\omega^2/c_s^2)P_e$. Then, if we take away the P_e's and change the sign, we see that the relationship between k and ω is the one that we want.

Working backwards again, we cannot resist writing down the grand equation which corresponds to the dispersion equation (48.22) for quantum-mechanical waves. If ϕ represents the amplitude for finding a particle at position x, y, z, at the time t, then the great equation of quantum mechanics for free particles is this:

$$\frac{\partial^2\phi}{\partial x^2}+\frac{\partial^2\phi}{\partial y^2}+\frac{\partial^2\phi}{\partial z^2}-\frac{1}{c^2}\frac{\partial^2\phi}{\partial t^2}=\frac{m^2c^2}{h^2}\phi. \tag{48.24}$$

First of all, the relativity character of this expression is suggested by the appearance

of x, y, z and t in the nice combination relativity usually involves. Second, it is a wave equation which, if we try a plane wave, would produce as a consequence that $-k^2 + \omega^2/c^2 = m^2c^2/h^2$, which is the right relationship for quantum mechanics. There is still another great thing contained in the wave equation: the fact that any superposition of waves is also a solution. So this equation contains all of the quantum mechanics and the relativity that we have been discussing so far, at least so long as it deals with a single particle in empty space with no external potentials or forces on it!

48–7 Normal modes

Now we turn to another example of the phenomenon of beats which is rather curious and a little different. Imagine two equal pendulums which have, between them, a rather weak spring connection. They are made as nearly as possible the same length. If we pull one aside and let go, it moves back and forth, and it pulls on the connecting spring as it moves back and forth, and so it really is a machine for generating a force which has the natural frequency of the other pendulum. Therefore, as a consequence of the theory of resonance, which we studied before, when we put a force on something at just the right frequency, it will drive it. So, sure enough, one pendulum moving back and forth drives the other. However, in this circumstance there is a new thing happening, because the total energy of the system is finite, so when one pendulum pours its energy into the other to drive it, it finds itself gradually losing energy, until, if the timing is just right along with the speed, it loses all its energy and is reduced to a stationary condition! Then, of course, it is the other pendulum ball that has all the energy and the first one which has none, and as time goes on we see that it works also in the opposite direction, and that the energy is passed back into the first ball; this is a very interesting and amusing phenomenon. We said, however, that this is related to the theory of beats, and we must now explain how we can analyze this motion from the point of view of the theory of beats.

We note that the motion of either of the two balls is an oscillation which has an amplitude which changes cyclically. Therefore the motion of one of the balls is presumably analyzable in a different way, in that it is the sum of two oscillations, present at the same time but having two slightly different frequencies. Therefore it ought to be possible to find two other motions in this system, and to claim that what we saw was a superposition of the two solutions, because this is of course a linear system. Indeed, it is easy to find two ways that we could start the motion, each one of which is a perfect, single-frequency motion—absolutely periodic. The motion that we started with before was not strictly periodic, since it did not last; soon one ball was passing energy to the other and so changing its amplitude; but there are ways of starting the motion so that nothing changes and, of course, as soon as we see it we understand why. For example, if we made both pendulums go together, then, since they are of the same length and the spring is not then doing anything, they will of course continue to swing like that for all time, assuming no friction and that everything is perfect. On the other hand, there is another possible motion which also has a definite frequency: that is, if we move the pendulums oppositely, pulling them aside exactly equal distances, then again they would be in absolutely periodic motion. We can appreciate that the spring just adds a little to the restoring force that the gravity supplies, that is all, and the system just keeps oscillating at a slightly higher frequency than in the first case. Why higher? Because the spring is pulling, in addition to the gravitation, and it makes the system a little "stiffer," so that the frequency of this motion is just a shade higher than that of the other.

Thus this system has two ways in which it can oscillate with unchanging amplitude: it can either oscillate in a manner in which both pendulums go the same way and oscillate all the time at one frequency, or they could go in opposite directions at a slightly higher frequency.

Now the actual motion of the thing, because the system is linear, can be represented as a superposition of the two. (The subject of this chapter, remember,

is the effects of adding two motions with different frequencies.) So think what would happen if we combined these two solutions. If at $t = 0$ the two motions are started with equal amplitude and in the same phase, the sum of the two motions means that one ball, having been impressed one way by the first motion and the other way by the second motion, is at zero, while the other ball, having been displaced the same way in both motions, has a large amplitude. As time goes on, however, the two basic *motions* proceed independently, so the phase of one relative to the other is slowly shifting. That means, then, that after a sufficiently long time, when the time is enough that one motion could have gone "900½" oscillations, while the other went only "900," the relative phase would be just reversed with respect to what it was before. That is, the large-amplitude motion will have fallen to zero, and in the meantime, of course, the initially motionless ball will have attained full strength!

So we see that we could analyze this complicated motion either by the idea that there is a resonance and that one passes energy to the other, or else by the superposition of two constant-amplitude motions at two different frequencies.

49

Modes

49–1 The reflection of waves

49–2 Confined waves, with natural frequencies

49–3 Modes in two dimensions

49–4 Coupled pendulums

49–5 Linear systems

49–1 The reflection of waves

This chapter will consider some of the remarkable phenomena which are a result of confining waves in some finite region. We will be led first to discover a few particular facts about vibrating strings, for example, and then the generalization of these facts will give us a principle which is probably the most far-reaching principle of mathematical physics.

Our first example of confining waves will be to confine a wave at one boundary. Let us take the simple example of a one-dimensional wave on a string. One could equally well consider sound in one dimension against a wall, or other situations of a similar nature, but the example of a string will be sufficient for our present purposes. Suppose that the string is held at one end, for example by fastening it to an "infinitely solid" wall. This can be expressed mathematically by saying that the displacement y of the string at the position $x = 0$ must be zero, because the end does not move. Now if it were not for the wall, we know that the general solution for the motion is the sum of two functions, $F(x - ct)$ and $G(x + ct)$, the first representing a wave travelling one way in the string, and the second a wave travelling the other way in the string:

$$y = F(x - ct) + G(x + ct) \tag{49.1}$$

is the general solution for any string. But we have next to satisfy the condition that the string does not move at one end. If we put $x = 0$ in Eq. (49.1) and examine y for any value of t, we get $y = F(-ct) + G(+ct)$. Now if this is to be zero for all times, it means that the function $G(ct)$ must be $-F(-ct)$. In other words, G of anything must be $-F$ of minus that same thing. If this result is put back into Eq. (49.1), we find that the solution for the problem is

$$y = F(x - ct) - F(-x - ct). \tag{49.2}$$

It is easy to check that we will get $y = 0$ if we set $x = 0$.

Figure 49–1 shows a wave travelling in the negative x-direction near $x = 0$, and a hypothetical wave travelling in the other direction reversed in sign and on the other side of the origin. We say hypothetical because, of course, there is no string to vibrate on that side of the origin. The total motion of the string is to be regarded as the sum of these two waves in the region of positive x. As they reach the origin, they will always cancel at $x = 0$, and finally the second (reflected) wave will be the only one to exist for positive x and it will, of course, be travelling in the opposite direction. These results are equivalent to the following statement: if a wave reaches the clamped end of a string, it will be reflected with a change in sign. Such a reflection can always be understood by imagining that what is coming to the end of the string comes out upside down from behind the wall. In short, if we assume that the string is infinite and that whenever we have a wave going one way we have another one going the other way with the stated symmetry, the displacement at $x = 0$ will always be zero and it would make no difference if we clamped the string there.

The next point to be discussed is the reflection of a periodic wave. Suppose that the wave represented by $F(x - ct)$ is a sine wave and has been reflected; then the reflected wave $-F(-x - ct)$ is also a sine wave of the same frequency, but travelling in the opposite direction. This situation can be most simply described by using the complex function notation: $F(x - ct) = e^{i\omega(t-x/c)}$ and

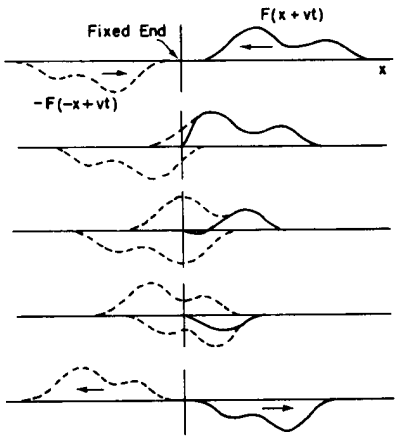


Fig. 49–1. Reflection of a wave as a superposition of two travelling waves.

$F(-x - ct) = e^{i\omega(t+x/c)}$. It can be seen that if these are substituted in (49.2) and if x is set equal to 0, then $y = 0$ for all values of t, so it satisfies the necessary condition. Because of the properties of exponentials, this can be written in a simpler form:

$$y = e^{i\omega t}(e^{-i\omega x/c} - e^{i\omega x/c}) = -2ie^{i\omega t}\sin(\omega x/c). \tag{49.3}$$

There is something interesting and new here, in that this solution tells us that if we look at any fixed x, the string oscillates at frequency ω. No matter where this point is, the frequency is the same! But there are some places, in particular wherever $\sin(\omega x/c) = 0$, where there is no displacement at all. Furthermore, if at any time t we take a snapshot of the vibrating string, the picture will be a sine wave. However, the displacement of this sine wave will depend upon the time t. From inspection of Eq. (49.3) we can see that the length of one cycle of the sine wave is equal to the wavelength of either of the superimposed waves:

$$\lambda = 2\pi c/\omega. \tag{49.4}$$

The points where there is no motion satisfy the condition $\sin(\omega x/c) = 0$, which means that $(\omega x/c) = 0, \pi, 2\pi, \ldots, n\pi \ldots$ These points are called *nodes*. Between any two successive nodes, every point moves up and down sinusoidally, but the pattern of motion stays fixed in space. This is the fundamental characteristic of what we call a *mode*. If one can find a pattern of motion which has the property that at any point the object moves perfectly sinusoidally, and that all points move at the same frequency (though some will move more than others), then we have what is called a mode.

49–2 Confined waves, with natural frequencies

The next interesting problem is to consider what happens if the string is held at both ends, say at $x = 0$ and $x = L$. We can begin with the idea of the reflection of waves, starting with some kind of a bump moving in one direction. As time goes on, we would expect the bump to get near one end, and as time goes still further it will become a kind of little wobble, because it is combining with the reversed-image bump which is coming from the other side. Finally the original bump will disappear and the image bump will move in the other direction to repeat the process at the other end. This problem has an easy solution, but an interesting question is whether we can have a sinusoidal motion (the solution just described is *periodic*, but of course it is not *sinusoidally* periodic). Let us try to put a sinusoidally periodic wave on a string. If the string is tied at one end, we know it must look like our earlier solution (49.3). If it is tied at the other end, it has to look the same at the other end. So the only possibility for periodic sinusoidal motion is that the sine wave must neatly fit into the string length. If it does not fit into the string length, then it is not a natural frequency at which the string can continue to oscillate. In short, if the string is started with a sine wave shape that just fits in, then it will continue to keep that perfect shape of a sine wave and will oscillate harmonically at some frequency.

Mathematically, we can write $\sin kx$ for the shape, where k is equal to the factor (ω/c) in Eqs. (49.3) and (49.4), and this function will be zero at $x = 0$. However, it must also be zero at the other end. The significance of this is that k is no longer arbitrary, as was the case for the half-open string. With the string closed at both ends, the only possibility is that $\sin(kL) = 0$, because this is the only condition that will keep both ends fixed. Now in order for a sine to be zero, the angle must be either $0, \pi, 2\pi$, or some other integral multiple of π. The equation

$$kL = n\pi \tag{49.5}$$

will, therefore, give any one of the possible k's, depending on what integer is put in. For each of the k's there is a certain frequency ω, which, according to (49.3), is simply

$$\omega = kc = n\pi c/L. \tag{49.6}$$

So we have found the following: that a string has a property that it can have sinusoidal motions, *but only at certain frequencies.* This is the most important characteristic of confined waves. No matter how complicated the system is, it always turns out that there are some patterns of motion which have a perfect sinusoidal time dependence, but with frequencies that are a property of the particular system and the nature of its boundaries. In the case of the string we have many different possible frequencies, each one, by definition, corresponding to a mode, because a mode is a pattern of motion which repeats itself sinusoidally. Figure 49–2 shows the first three modes for a string. For the first mode the wavelength λ is $2L$. This can be seen if one continues the wave out to $x = 2L$ to obtain one complete cycle of the sine wave. The angular frequency ω is $2\pi c$ divided by the wavelength, in general, and in this case, since λ is $2L$, the frequency is $\pi c/L$, which is in agreement with (49.6) with $n = 1$. Let us call the first mode frequency ω_1. Now the next mode shows two loops with one node in the middle. For this mode the wavelength, then, is simply L. The corresponding value of k is twice as great and the frequency is twice as large; it is $2\omega_1$. For the third mode it is $3\omega_1$, and so on. So all the different frequencies of the string are multiples, 1, 2, 3, 4, and so on, of the lowest frequency ω_1.

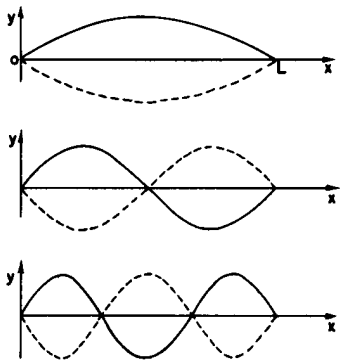


Fig. 49–2. The first three modes of a vibrating string.

Returning now to the general motion of the string, it turns out that any possible motion can always be analyzed by asserting that more than one mode is operating at the same time. In fact, for general motion an infinite number of modes must be excited at the same time. To get some idea of this, let us illustrate what happens when there are two modes oscillating at the same time: Suppose that we have the first mode oscillating as shown by the sequence of pictures in Fig. 49–3, which illustrates the deflection of the string for equally spaced time intervals extending through half a cycle of the lowest frequency.

Now, at the same time, we suppose that there is an oscillation of the second mode also. Figure 49–3 also shows a sequence of pictures of this mode, which at the start is 90° out of phase with the first mode. This means that at the start it has no displacement, but the two halves of the string have oppositely directed velocities. Now we recall a general principle relating to linear systems: if there are any two solutions, then their sum is also a solution. Therefore a third possible motion of the string would be a displacement obtained by adding the two solutions shown in Fig. 49–3. The result, also shown in the figure, begins to suggest the idea of a bump running back and forth between the ends of the string, although with only two modes we cannot make a very good picture of it; more modes are needed. This result is, in fact, a special case of a great principle for linear systems:

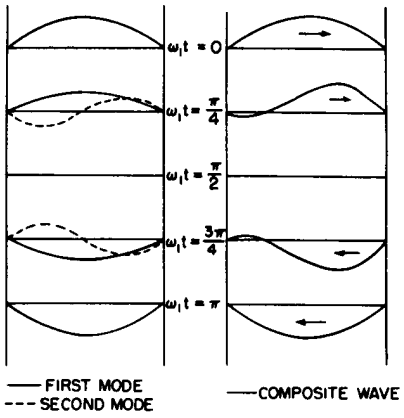


Fig. 49–3. Two modes combine to give a travelling wave.

Any motion at all can be analyzed by assuming that it is the sum of the motions of all the different modes, combined with appropriate amplitudes and phases.

The importance of the principle derives from the fact that each mode is very simple—it is nothing but a sinusoidal motion in time. It is true that even the general motion of a string is not really very complicated, but there are other systems, for example the whipping of an airplane wing, in which the motion is much more complicated. Nevertheless, even with an airplane wing, we find there is a certain particular way of twisting which has one frequency and other ways of twisting that have other frequencies. If these modes can be found, then the complete motion can always be analyzed as a superposition of harmonic oscillations (except when the whipping is of such degree that the system can no longer be considered as linear).

49–3 Modes in two dimensions

The next example to be considered is the interesting situation of modes in two dimensions. Up to this point we have talked only about one-dimensional situations —a stretched string or sound waves in a tube. Ultimately we should consider three dimensions, but an easier step will be that to two dimensions. Consider for definiteness a rectangular rubber drumhead which is confined so as to have no displacement anywhere on the rectangular edge, and let the dimensions of the rectangle

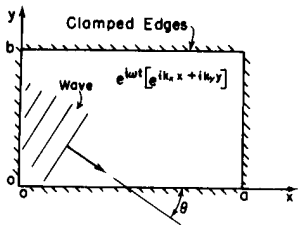


Fig. 49–4. Vibrating rectangular plate.

be a and b, as shown in Fig. 49–4. Now the question is, what are the characteristics of the possible motion? We can start with the same procedure used for the string. If we had no confinement at all, we would expect waves travelling along with some kind of wave motion. For example, $(e^{i\omega t})(e^{-ik_x x+ik_y y})$ would represent a sine wave travelling in some direction which depends on the relative values of k_x and k_y. Now how can we make the x-axis, that is, the line $y = 0$, a node? Using the ideas developed for the one-dimensional string, we can imagine another wave represented by the complex function $(-e^{i\omega t})(e^{-ik_x x-ik_y y})$. The superposition of these waves will give zero displacement at $y = 0$ regardless of the values of x and t. (Although these functions are defined for negative y where there is no drumhead to vibrate, this can be ignored, since the displacement is truly zero at $y = 0$.) In this case we can look upon the second function as the reflected wave.

However, we want a nodal line at $y = b$ as well as at $y = 0$. How do we do that? The solution is related to something we did when studying reflection from crystals. These waves which cancel each other at $y = 0$ will do the same at $y = b$ only if $2b \sin\theta$ is an integral multiple of λ, where θ is the angle shown in Fig. 49–4:

$$m\lambda = 2b \sin\theta, \qquad m = 0, 1, 2, \ldots \tag{49.7}$$

Now in the same way we can make the y-axis a nodal line by adding two more functions $-(e^{i\omega t})(e^{+ik_x x+ik_y y})$ and $+(e^{i\omega t})(e^{+ik_x x-ik_y y})$, each representing a reflection of one of the other two waves from the $x = 0$ line. The condition for a nodal line at $x = a$ is similar to the one for $y = b$. It is that $2a \cos\theta$ must also be an integral multiple of λ:

$$n\lambda = 2a \cos\theta. \tag{49.8}$$

Then the final result is that the waves bouncing about in the box produce a standing-wave pattern, that is, a definite mode.

So we must satisfy the above two conditions if we are to have a mode. Let us first find the wavelength. This can be obtained by eliminating the angle θ from (49.7) and (49.8) to obtain the wavelength in terms of a, b, n and m. The easiest way to do that is to divide both sides of the respective equations by $2b$ and $2a$, square them, and add the two equations together. The result is $\sin^2\theta + \cos^2\theta = 1 = (n\lambda/2a)^2 + (m\lambda/2b)^2$, which can be solved for λ:

$$\frac{1}{\lambda^2} = \frac{n^2}{4a^2} + \frac{m^2}{4b^2}. \tag{49.9}$$

In this way we have determined the wavelength in terms of two integers, and from the wavelength we immediately get the frequency ω, because, as we know, the frequency is equal to $2\pi c$ divided by the wavelength.

This result is interesting and important enough that we should deduce it by a purely mathematical analysis instead of by an argument about the reflections. Let us represent the vibration by a superposition of four waves chosen so that the four lines $x = 0$, $x = a$, $y = 0$, and $y = b$ are all nodes. In addition we shall require that all waves have the same frequency, so that the resulting motion will represent a mode. From our earlier treatment of light reflection we know that $(e^{i\omega t})(e^{-ik_x x+ik_y y})$ represents a wave travelling in the direction indicated in Fig. 49–4. Equation (49.6), that is, $k = \omega/c$, still holds, provided

$$k^2 = k_x^2 + k_y^2. \tag{49.10}$$

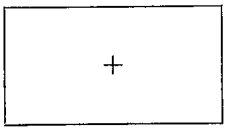
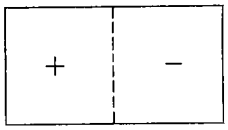
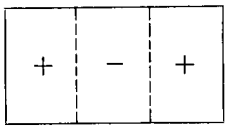
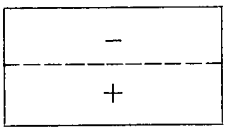
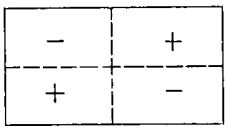
It is clear from the figure that $k_x = k \cos\theta$ and $k_y = k \sin\theta$.

Now our equation for the displacement, say ϕ, of the rectangular drumhead takes on the grand form

$$\phi = [e^{i\omega t}][e^{(-ik_x x+ik_y y)} - e^{(+ik_x x+ik_y y)} - e^{(-ik_x x-ik_y y)} + e^{(+ik_x x-ik_y y)}]. \tag{49.11a}$$

Although this looks rather a mess, the sum of these things now is not very hard.

Table 49–1

Mode shape	m	n	$(\omega/\omega_0)^2$	ω/ω_0
	1	1	1.25	1.12
	1	2	2.00	1.41
	1	3	3.25	1.80
	2	1	4.25	2.06
	2	2	5.00	2.24

The exponentials can be combined to give sine functions, so that the displacement turns out to be

$$\phi = [-4 \sin k_x x \sin k_y y][e^{i\omega t}]. \tag{49.11b}$$

In other words, it is a sinusoidal oscillation, all right, with a pattern that is also sinusoidal in both the x- and the y-direction. Our boundary conditions are of course satisfied at $x = 0$ and $y = 0$. We also want ϕ to be zero when $x = a$ and when $y = b$. Therefore we have to put in two other conditions: $k_x a$ must be an integral multiple of π, and $k_x b$ must be another integral multiple of π. Since we have seen that $k_x = k \cos\theta$ and $k_y = k \sin\theta$, we immediately get equations (49.7) and (49.8) and from these the final result (49.9).

Now let us take as an example a rectangle whose width is twice the height. If we take $a = 2b$ and use Eqs. (49.4) and (49.9), we can calculate the frequencies of all of the modes:

$$\omega^2 = \left(\frac{\pi c}{b}\right)^2 \frac{4m^2 + n^2}{4}. \tag{49.12}$$

Table 49–1 lists a few of the simple modes and also shows their shape in a qualitative way.

The most important point to be emphasized about this particular case is that the frequencies are not multiples of each other, nor are they multiples of any number. The idea that the natural frequencies are harmonically related is not generally true. It is not true for a system of more than one dimension, nor is it true for one-dimensional systems which are more complicated than a string with uniform density and tension. A simple example of the latter is a hanging chain in which the tension is higher at the top than at the bottom. If such a chain is set in harmonic oscillation, there are various modes and frequencies, but the frequencies are not simple multiples of any number, nor are the mode shapes sinusoidal.

The modes of more complicated systems are still more elaborate. For example, inside the mouth we have a cavity above the vocal cords, and by moving the tongue and the lips, and so forth, we make an open-ended pipe or a closed-ended pipe of different diameters and shapes; it is a terribly complicated resonator, but

it is a resonator nevertheless. Now when one talks with the vocal cords, they are made to produce some kind of tone. The tone is rather complicated and there are many sounds coming out, but the cavity of the mouth further modifies that tone because of the various resonant frequencies of the cavity. For instance, a singer can sing various vowels, a, or o, or oo, and so forth, at the same pitch, but they sound different because the various harmonics are in resonance in this cavity to different degrees. The very great importance of the resonant frequencies of a cavity in modifying the voice sounds can be demonstrated by a simple experiment. Since the speed of sound goes as the reciprocal of the square root of the density, the speed of sound may be varied by using different gases. If one uses helium instead of air, so that the density is lower, the speed of sound is much higher, and all the frequencies of a cavity will be raised. Consequently if one fills one's lungs with helium before speaking, the character of his voice will be drastically altered even though the vocal cords may still be vibrating at the same frequency.

49–4 Coupled pendulums

Finally we should emphasize that not only do modes exist for complicated continuous systems, but also for very simple mechanical systems. A good example is the system of two coupled pendulums discussed in the preceding chapter. In that chapter it was shown that the motion could be analyzed as a superposition of two harmonic motions with different frequencies. So even this system can be analyzed in terms of harmonic motions or modes. The string has an infinite number of modes and the two-dimensional surface also has an infinite number of modes. In a sense it is a double infinity, if we know how to count infinities. But a simple mechanical thing which has only two degrees of freedom, and requires only two variables to describe it, has only two modes.

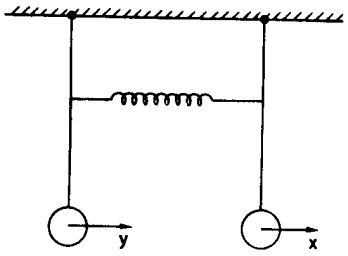


Fig. 49–5. Two coupled pendulums.

Let us make a mathematical analysis of these two modes for the case where the pendulums are of equal length. Let the displacement of one be x, and the displacement of the other be y, as shown in Fig. 49–5. Without a spring, the force on the first mass is proportional to the displacement of that mass, because of gravity. There would be, if there were no spring, a certain natural frequency ω_0 for this one alone. The equation of motion without a spring would be

$$m\frac{d^2x}{dt^2} = -m\omega_0^2 x. \tag{49.13}$$

The other pendulum would swing in the same way if there were no spring. In addition to the force of restoration due to gravitation, there is an additional force pulling the first mass. That force depends upon the excess distance of x over y and is proportional to that difference, so it is some constant which depends on the geometry, times $(x - y)$. The same force in reverse sense acts on the second mass. The equations of motion that have to be solved are therefore

$$m\frac{d^2x}{dt^2} = -m\omega_0^2 x - k(x - y), \qquad m\frac{d^2y}{dx^2} = -m\omega_0^2 y - k(y - x). \tag{49.14}$$

In order to find a motion in which both of the masses move at the same frequency, we must determine how much each mass moves. In other words, pendulum x and pendulum y will oscillate at the same frequency, but their amplitudes must have certain values, A and B, whose relation is fixed. Let us try this solution:

$$x = Ae^{i\omega t}, \qquad y = Be^{i\omega t}. \tag{49.15}$$

If these are substituted in Eqs. (49.14) and similar terms are collected, the results are

$$\begin{aligned}\left(\omega^2 - \omega_0^2 - \frac{k}{m}\right)A &= -\frac{k}{m}B,\\ \left(\omega^2 - \omega_0^2 - \frac{k}{m}\right)B &= -\frac{k}{m}A.\end{aligned} \tag{49.16}$$

The equations as written have had the common factor $e^{i\omega t}$ removed and have been divided by m.

Now we see that we have two equations for what looks like two unknowns. But there really are not *two* unknowns, because the whole size of the motion is something that we cannot determine from these equations. The above equations can determine only the *ratio* of A to B, *but they must both give the same ratio.* The necessity for both of these equations to be consistent is a requirement that the frequency be something very special.

In this particular case this can be worked out rather easily. If the two equations are multiplied together, the result is

$$\left(\omega^2 - \omega_0^2 - \frac{k}{m}\right)^2 AB = \left(\frac{k}{m}\right)^2 AB. \tag{49.17}$$

The term AB can be removed from both sides unless A and B are zero, which means there is no motion at all. If there is motion, then the other terms must be equal, giving a quadratic equation to solve. The result is that there are two possible frequencies:

$$\omega_1^2 = \omega_0^2, \qquad \omega_2^2 = \omega_0^2 + \frac{2k}{m}. \tag{49.18}$$

Furthermore, if these values of frequency are substituted back into Eq. (49.16), we find that for the first frequency $A = B$, and for the second frequency $A = -B$. These are the "mode shapes," as can be readily verified by experiment.

It is clear that in the first mode, where $A = B$, the spring is never stretched, and both masses oscillate at the frequency ω_0, as though the spring were absent. In the other solution, where $A = -B$, the spring contributes a restoring force and raises the frequency. A more interesting case results if the pendulums have different lengths. The analysis is very similar to that given above, and is left as an exercise for the reader.

49–5 Linear systems

Now let us summarize the ideas discussed above, which are all aspects of what is probably the most general and wonderful principle of mathematical physics. If we have a linear system whose character is independent of the time, then the motion does not have to have any particular simplicity, and in fact may be exceedingly complex, but there are very special motions, usually a series of special motions, in which the whole pattern of motion varies exponentially with the time. For the vibrating systems that we are talking about now, the exponential is imaginary, and instead of saying "exponentially" we might prefer to say "sinusoidally" with time. However, one can be more general and say that the motions will vary exponentially with the time in very special modes, with very special shapes. The most general motion of the system can always be represented as a superposition of motions involving each of the different exponentials.

This is worth stating again for the case of sinusoidal motion: a linear system need not be moving in a purely sinusoidal motion, i.e., at a definite single frequency, but no matter how it does move, this motion can be represented as a superposition of pure sinusoidal motions. The frequency of each of these motions is a characteristic of the system, and the pattern or waveform of each motion is also a characteristic of the system. The general motion in any such system can be characterized by giving the strength and the phase of each of these modes, and adding them all together. Another way of saying this is that any linear vibrating system is equivalent to a set of independent harmonic oscillators, with the natural frequencies corresponding to the modes.

We conclude this chapter by remarking on the connection of modes with quantum mechanics. In quantum mechanics the vibrating object, or the thing that varies in space, is the amplitude of a probability function that gives the probability of finding an electron, or system of electrons, in a given configuration. This amplitude function can vary in space and time, and satisfies, in fact, a linear equation.

But in quantum mechanics there is a transformation, in that what we call frequency of the probability amplitude is equal, in the classical idea, to energy. Therefore we can translate the principle stated above to this case by taking the word *frequency* and replacing it with *energy*. It becomes something like this: a quantum-mechanical system, for example an atom, need not have a definite energy, just as a simple mechanical system does not have to have a definite frequency; but no matter how the system behaves, its behavior can always be represented as a superposition of states of definite energy. The energy of each state is a characteristic of the atom, and so is the pattern of amplitude which determines the probability of finding particles in different places. The general motion can be described by giving the amplitude of each of these different energy states. This is the origin of energy levels in quantum mechanics. Since quantum mechanics is represented by waves, in the circumstance in which the electron does not have enough energy to ultimately escape from the proton, they are *confined waves*. Like the confined waves of a string, there are definite frequencies for the solution of the wave equation for quantum mechanics. The quantum-mechanical interpretation is that these are definite *energies*. Therefore a quantum-mechanical system, because it is represented by waves, can have definite states of fixed energy; examples are the energy levels of various atoms.

50

Harmonics

50–1 Musical tones

50–2 The Fourier series

50–3 Quality and consonance

50–4 The Fourier coefficients

50–5 The energy theorem

50–6 Nonlinear responses

50–1 Musical tones

Pythagoras is said to have discovered the fact that two similar strings under the same tension and differing only in length, when sounded together give an effect that is pleasant to the ear *if* the lengths of the strings are in the ratio of two small integers. If the lengths are as one is to two, they then correspond to the octave in music. If the lengths are as two is to three, they correspond to the interval between C and G, which is called a fifth. These intervals are generally accepted as "pleasant" sounding chords.

Pythagoras was so impressed by this discovery that he made it the basis of a school—Pythagoreans they were called—which held mystic beliefs in the great powers of numbers. It was believed that something similar would be found out about the planets—or "spheres." We sometimes hear the expression: "the music of the spheres." The idea was that there would be some numerical relationships between the orbits of the planets or between other things in nature. People usually think that this is just a kind of superstition held by the Greeks. But is it so different from our own scientific interest in quantitative relationships? Pythagoras' discovery was the first example, outside geometry, of any numerical relationship in nature. It must have been very surprising to suddenly discover that there was a *fact* of nature that involved a simple numerical relationship. Simple measurements of lengths gave a prediction about something which had no apparent connection to geometry—the production of pleasant sounds. This discovery led to the extension that perhaps a good tool for understanding nature would be arithmetic and mathematical analysis. The results of modern science justify that point of view.

Pythagoras could only have made his discovery by making an experimental observation. Yet this important aspect does not seem to have impressed him. If it had, physics might have had a much earlier start. (It is always easy to look back at what someone else has done and to decide what he *should* have done!)

We might remark on a third aspect of this very interesting discovery: that the discovery had to do with two notes that *sound pleasant* to the ear. We may question whether *we* are any better off than Pythagoras in understanding *why* only certain sounds are pleasant to our ear. The general theory of aesthetics is probably no further advanced now than in the time of Pythagoras. In this one discovery of the Greeks, there are the three aspects: experiment, mathematical relationships, and aesthetics. Physics has made great progress on only the first two parts. This chapter will deal with our present-day understanding of the discovery of Pythagoras.

Among the sounds that we hear, there is one kind that we call *noise*. Noise corresponds to a sort of irregular vibration of the eardrum that is produced by the irregular vibration of some object in the neighborhood. If we make a diagram to indicate the pressure of the air on the eardrum (and, therefore, the displacement of the drum) as a function of time, the graph which corresponds to a noise might look like that shown in Fig. 50–1(a). (Such a noise might correspond roughly to the sound of a stamped foot.) The sound of *music* has a different character. Music is characterized by the presence of more-or-less *sustained tones*—or musical "notes." (Musical instruments may make noises as well!) The tone may last for a relatively short time, as when a key is pressed on a piano, or it may be sustained almost indefinitely, as when a flute player holds a long note.

What is the special character of a musical note from the point of view of the pressure in the air? A musical note differs from a noise in that there is a periodicity in its graph. There is some uneven shape to the variation of the air pressure with

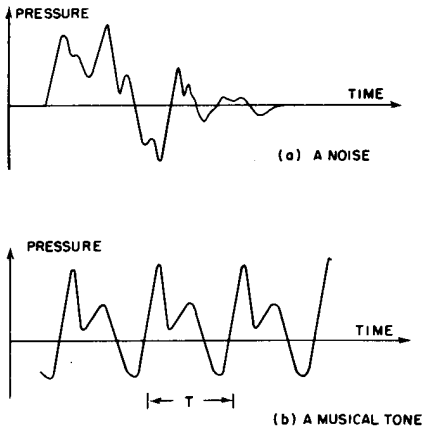


Fig. 50–1. Pressure as a function of time for (a) a noise, and (b) a musical tone.

time, and the shape repeats itself over and over again. An example of a pressure-time function that would correspond to a musical note is shown in Fig. 50–1(b).

Musicians will usually speak of a musical tone in terms of three characteristics: the loudness, the pitch, and the "quality." The "loudness" is found to correspond to the magnitude of the pressure changes. The "pitch" corresponds to the period of time for one repetition of the basic pressure function. ("Low" notes have longer periods than "high" notes.) The "quality" of a tone has to do with the differences we may still be able to hear between two notes of the same loudness and pitch. An oboe, a violin, or a soprano are still distinguishable even when they sound notes of the same pitch. The quality has to do with the structure of the repeating pattern.

Let us consider, for a moment, the sound produced by a vibrating string. If we pluck the string, by pulling it to one side and releasing it, the subsequent motion will be determined by the motions of the waves we have produced. We know that these waves will travel in both directions, and will be reflected at the ends. They will slosh back and forth for a long time. No matter how complicated the wave is, however, it will repeat itself. The period of repetition is just the time T required for the wave to travel two full lengths of the string. For that is just the time required for any wave, once started, to reflect off each end and return to its starting position, and be proceeding in the original direction. The time is the same for waves which start out in either direction. Each point on the string will, then, return to its starting position after one period, and again one period later, etc. The sound wave produced must also have the same repetition. We see why a plucked string produces a musical tone.

50–2 The Fourier series

We have discussed in the preceding chapter another way of looking at the motion of a vibrating system. We have seen that a string has various natural modes of oscillation, and that any particular kind of vibration that may be set up by the starting conditions can be thought of as a combination—in suitable proportions—of several of the natural modes, oscillating together. For a string we found that the normal modes of oscillation had the frequencies ω_0, $2\omega_0$, $3\omega_0, \ldots$ The most general motion of a plucked string, therefore, is composed of the sum of a sinusoidal oscillation at the fundamental frequency ω_0, another at the second harmonic frequency $2\omega_0$, another at the third harmonic $3\omega_0$, etc. Now the fundamental mode repeats itself every period $T_1 = 2\pi/\omega_0$. The second harmonic mode repeats itself every $T_2 = 2\pi/2\omega_0$. It *also* repeats itself every $T_1 = 2T_2$, after *two* of its periods. Similarly, the third harmonic mode repeats itself after a time T_1 which is 3 of its periods. We see again why a plucked string repeats its whole pattern with a periodicity of T_1. It produces a musical tone.

We have been talking about the motion of the string. But the *sound*, which is the motion of the air, is produced by the motion of the string, so its vibrations too must be composed of the same harmonics—though we are no longer thinking about the normal modes of the air. Also, the relative strength of the harmonics may be different in the air than in the string, particularly if the string is "coupled" to the air via a sounding board. The efficiency of the coupling to the air is different for different harmonics.

If we let $f(t)$ represent the air pressure as a function of time for a musical tone [such as that in Fig. 50–1(b)], then we expect that $f(t)$ can be written as the sum of a number of simple harmonic functions of time—like $\cos \omega t$—for each of the various harmonic frequencies. If the period of the vibration is T, the fundamental angular frequency will be $\omega = 2\pi/T$, and the harmonics will be 2ω, 3ω, etc.

There is one slight complication. For each frequency we may expect that the starting phases will not necessarily be the same for all frequencies. We should, therefore, use functions like $\cos(\omega t + \phi)$. It is, however, simpler to use instead both the sine and cosine functions for *each* frequency. We recall that

$$\cos(\omega t + \phi) = (\cos\phi \cos\omega t - \sin\phi \sin\omega t) \tag{50.1}$$

and since ϕ is a constant, *any* sinusoidal oscillation at the frequency ω can be written as the sum of a term with cos ωt and another term with sin ωt.

We conclude, then, that *any* function $f(t)$ that is periodic with the period T can be written mathematically as

$$\begin{aligned} f(t) = a_0 & \\ & + a_1 \cos \omega t + b_1 \sin \omega t \\ & + a_2 \cos 2\omega t + b_2 \sin 2\omega t \\ & + a_3 \cos 3\omega t + b_3 \sin 3\omega t \\ & + \ldots \qquad + \ldots \end{aligned} \tag{50.2}$$

where $\omega = 2\pi/T$ and the a's and b's are numerical constants which tell us how much of each component oscillation is present in the oscillation $f(t)$. We have added the "zero-frequency" term a_0 so that our formula will be completely general, although it is usually zero for a musical tone. It represents a shift of the average value (that is, the "zero" level) of the sound pressure. With it our formula can take care of any case. The equality of Eq. (50.2) is represented schematically in Fig. 50–2. (The amplitudes, a_n and b_n, of the harmonic functions must be suitably chosen. They are shown schematically and without any particular scale in the figure.) The series (50.2) is called the *Fourier series* for $f(t)$.

We have said that *any* periodic function can be made up in this way. We should correct that and say that any sound wave, or any function we ordinarily encounter in physics, can be made up of such a sum. The mathematicians can invent functions which cannot be made up of simple harmonic functions—for instance, a function that has a "reverse twist" so that it has two values for some values of t! We need not worry about such functions here.

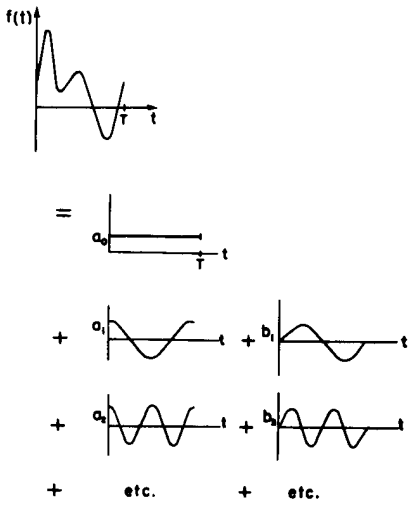


Fig. 50–2. Any periodic function $f(t)$ is equal to a sum of simple harmonic functions.

50–3 Quality and consonance

Now we are able to describe what it is that determines the "quality" of a musical tone. It is the relative amounts of the various harmonics—the values of the a's and b's. A tone with only the first harmonic is a "pure" tone. A tone with many strong harmonics is a "rich" tone. A violin produces a different proportion of harmonics than does an oboe.

We can "manufacture" various musical tones if we connect several "oscillators" to a loudspeaker. (An oscillator usually produces a nearly pure simple harmonic function.) We should choose the frequencies of the oscillators to be ω, 2ω, 3ω, etc. Then by adjusting the volume control on each oscillator, we can add in any amount we wish of each harmonic—thereby producing tones of different quality. An electric organ works in much this way. The "keys" select the frequency of the fundamental oscillator and the "stops" are switches that control the relative proportions of the harmonics. By throwing these switches, the organ can be made to sound like a flute, or an oboe, or a violin.

It is interesting that to produce such "artificial" tones we need only one oscillator for each frequency—we do not need separate oscillators for the sine and cosine components. The ear is not very sensitive to the relative phases of the harmonics. It pays attention mainly to the *total* of the sine and cosine parts of each frequency. Our analysis is more accurate than is necessary to explain the *subjective* aspect of music. The response of a microphone or other physical instrument does depend on the phases, however, and our complete analysis may be needed to treat such cases.

The "quality" of a spoken sound also determines the vowel sounds that we recognize in speech. The shape of the mouth determines the frequencies of the natural modes of vibration of the air in the mouth. Some of these modes are set into vibration by the sound waves from the vocal chords. In this way, the amplitudes of some of the harmonics of the sound are increased with respect to others. When we change the shape of our mouth, harmonics of different frequencies are given preference. These effects account for the difference between an "e–e–e" sound and an "a–a–a" sound.

We all know that a particular vowel sound—say "e–e–e"—still "sounds like" the same vowel whether we say (or sing) it at a high or a low pitch. From the mechanism we describe, we would expect that *particular* frequencies are emphasized when we shape our mouth for an "e–e–e," and that they *do not* change as we change the pitch of our voice. So the relation of the important harmonics to the fundamental—that is, the "quality"—changes as we change pitch. Apparently the mechanism by which we recognize speech is not based on specific harmonic relationships.

What should we say now about Pythagoras' discovery? We understand that two similar strings with lengths in the ratio of 2 to 3 will have fundamental frequencies in the ratio 3 to 2. But why should they "sound pleasant" together? Perhaps we should take our clue from the frequencies of the harmonics. The second harmonic of the lower shorter string will have the *same* frequency as the third harmonic of the longer string. (It is easy to show—or to believe—that a plucked string produces strongly the several lowest harmonics.)

Perhaps we should make the following rules. Notes sound consonant when they have harmonics with the same frequency. Notes sound dissonant if their upper harmonics have frequencies near to each other but far enough apart that there are rapid beats between the two. Why beats do not sound pleasant, and why unison of the upper harmonics does sound pleasant, is something that we do not know how to define or describe. We cannot say from this knowledge of what *sounds* good, what ought, for example, to *smell* good. In other words, our understanding of it is not anything more general than the statement that when they are in unison they sound good. It does not permit us to deduce anything more than the properties of concordance in music.

It is easy to check on the harmonic relationships we have described by some simple experiments with a piano. Let us label the 3 successive C's near the middle of the keyboard by C, C′, and C″, and the G's just above by G, G′, and G″. Then the fundamentals will have relative frequencies as follows:

C –2	G – 3
C′ –4	G′ – 6
C″–8	G″–12

These harmonic relationships can be demonstrated in the following way: Suppose we press C′ *slowly*—so that it does not sound but we cause the damper to be lifted. If we then sound C, it will produce its own fundamental *and* some second harmonic. The second harmonic will set the strings of C′ into vibration. If we now release C (keeping C′ pressed) the damper will stop the vibration of the C strings, and we can hear (softly) the note C′ as it dies away. In a similar way, the third harmonic of C can cause a vibration of G′. Or the sixth of C (now getting much weaker) can set up a vibration in the fundamental of G″.

A somewhat different result is obtained if we press G quietly and then sound C′. The third harmonic of C′ will correspond to the fourth harmonic of G, so *only* the fourth harmonic of G will be excited. We can hear (if we listen closely) the sound of G″, which is two octaves above the G we have pressed! It is easy to think up many more combinations for this game.

We may remark in passing that the major scale can be defined just by the condition that the three major chords (F–A–C); (C–E–G); and (G–B–D) *each* represent tone sequences with the frequency ratio (4: 5: 6). These ratios—plus the fact that an octave (C–C′, B–B′, etc.) has the ratio 1: 2—determine the whole scale for the "ideal" case, or for what is called "just intonation." Keyboard instruments like the piano are *not* usually tuned in this manner, but a little "fudging" is done so that the frequencies are *approximately* correct for all possible starting tones. For this tuning, which is called "tempered," the octave (still 1: 2) is divided into 12 equal intervals for which the frequency ratio is $(2)^{1/12}$. A fifth no longer has the frequency ratio 3/2, but $2^{7/12} = 1.499$, which is apparently close enough for most ears.

We have stated a rule for consonance in terms of the coincidence of harmonics. Is this coincidence perhaps the *reason* that two notes are consonant? One worker has claimed that two *pure* tones—tones carefully manufactured to be free of harmonics—do not give the *sensations* of consonance or dissonance as the relative frequencies are placed at or near the expected ratios. (Such experiments are difficult because it is difficult to manufacture pure tones, for reasons that we shall see later.) We cannot still be certain whether the ear is matching harmonics or doing arithmetic when we decide that we like a sound.

50–4 The Fourier coefficients

Let us return now to the idea that any note—that is, a *periodic* sound—can be represented by a suitable combination of harmonics. We would like to show how we can find out what amount of each harmonic is required. It is, of course, easy to compute $f(t)$, using Eq. (50.2), if we are *given* all the coefficients a and b. The question now is, if we are given $f(t)$ how can we know what the coefficients of the various harmonic terms should be? (It is easy to make a cake from a recipe; but can we write down the recipe if we are given a cake?)

Fourier discovered that it was not really very difficult. The term a_0 is certainly easy. We have already said that it is just the average value of $f(t)$ over one period (from $t = 0$ to $t = T$). We can easily see that this is indeed so. The average value of a sine or cosine function over one period is zero. Over two, or three, or any whole number of periods, it is also zero. So the average value of all of the terms on the right-hand side of Eq. (50.2) is zero, except for a_0. (Recall that we must choose $\omega = 2\pi/T$.)

Now the average of a sum is the sum of the averages. So the average of $f(t)$ is just the average of a_0. But a_0 is a *constant*, so its average is just the same as its value. Recalling the definition of an average, we have

$$a_0 = \frac{1}{T}\int_0^T f(t)\,dt. \tag{50.3}$$

The other coefficients are only a little more difficult. To find them we can use a trick discovered by Fourier. Suppose we multiply both sides of Eq. (50.2) by some harmonic function—say by $\cos 7\omega t$. We have then

$$\begin{aligned} f(t)\cdot\cos 7\omega t = {} & a_0\cdot\cos 7\omega t \\ & + a_1\cos\omega t\cdot\cos 7\omega t + b_1\sin\omega t\cdot\cos 7\omega t \\ & + a_2\cos 2\omega t\cdot\cos 7\omega t + b_2\sin 2\omega t\cdot\cos 7\omega t \\ & + \cdots \qquad\qquad + \cdots \\ & + a_7\cos 7\omega t\cdot\cos 7\omega t + b_7\sin 7\omega t\cdot\cos 7\omega t \\ & + \cdots \qquad\qquad + \cdots \end{aligned} \tag{50.4}$$

Now let us average both sides. The average of $a_0\cos 7\omega t$ over the time T is proportional to the average of a cosine over 7 whole periods. But that is just zero. The average of *almost all* of the rest of the terms is *also* zero. Let us look at the a_1 term. We know, in general, that

$$\cos A\cos B = \tfrac{1}{2}\cos(A + B) + \tfrac{1}{2}\cos(A - B). \tag{50.5}$$

The a_1 term becomes

$$\tfrac{1}{2}a_1(\cos 8\omega t + \cos 6\omega t). \tag{50.6}$$

We thus have two cosine terms, one with 8 full periods in T and the other with 6. *They both average to zero.* The average of the a_1 term is therefore zero.

For the a_2 term, we would find $a_2\cos 9\omega t$ and $a_2\cos 5\omega t$, each of which also averages to zero. For the a_9 term, we would find $\cos 16\omega t$ and $\cos(-2\omega t)$. But $\cos(-2\omega t)$ is the same as $\cos 2\omega t$, so both of these have zero averages. It is clear

that *all* of the a terms will have a zero average *except* one. And that one is the a_7 term. For this one we have

$$\tfrac{1}{2}a_7(\cos 14\omega t + \cos 0). \tag{50.7}$$

The cosine of zero is one, and its average, of course, is one. So we have the result that the average of all of the a terms of Eq. (50.4) equals $\frac{1}{2}a_7$.

The b terms are even easier. When we multiply by any cosine term like $\cos n\omega t$, we can show by the same method that *all* of the b terms have the average value zero.

We see that Fourier's "trick" has acted like a sieve. When we multiply by $\cos 7\omega t$ and average, all terms drop out except a_7, and we find that

$$\text{Average } [f(t) \cdot \cos 7\omega t] = a_7/2, \tag{50.8}$$

or

$$a_7 = \frac{2}{T}\int_0^T f(t) \cdot \cos 7\omega t \, dt. \tag{50.9}$$

We shall leave it for the reader to show that the coefficient b_7 can be obtained by multiplying Eq. (50.2) by $\sin 7\omega t$ and averaging both sides. The result is

$$b_7 = \frac{2}{T}\int_0^T f(t) \cdot \sin 7\omega t \, dt. \tag{50.10}$$

Now what is true for 7 we expect is true for any integer. So we can summarize our proof and result in the following more elegant mathematical form. If m and n are integers other than zero, and if $\omega = 2\pi/T$, then

$$\text{I.} \quad \int_0^T \sin n\omega t \cos m\omega t \, dt = 0. \tag{50.11}$$

$$\left.\begin{aligned} \text{II.} \quad & \int_0^T \cos n\omega t \cos m\omega t \, dt = \\ \text{III.} \quad & \int_0^T \sin n\omega t \sin m\omega t \, dt = \end{aligned}\right\} \begin{cases} 0 & \text{if } n \neq m. \\ T/2 & \text{if } n = m. \end{cases} \tag{50.12}$$

$$\text{IV.} \quad f(t) = a_0 + \sum_{n=1}^{\infty} a_n \cos n\omega t + \sum_{n=1}^{\infty} b_n \sin n\omega t. \tag{50.13}$$

$$\text{V.} \quad a_0 = \frac{1}{T}\int_0^T f(t) \cdot dt. \tag{50.14}$$

$$a_n = \frac{2}{T}\int_0^T f(t) \cdot \cos n\omega t \, dt. \tag{50.15}$$

$$b_n = \frac{2}{T}\int_0^T f(t) \cdot \sin n\omega t \, dt. \tag{50.16}$$

In earlier chapters it was convenient to use the exponential notation for representing simple harmonic motion. Instead of $\cos \omega t$ we used $\text{Re } e^{i\omega t}$, the real part of the exponential function. We have used cosine and sine functions in this chapter because it made the derivations perhaps a little clearer. Our final result of Eq. (50.13) can, however, be written in the compact form

$$f(t) = \text{Re} \sum_{n=0}^{\infty} \hat{a}_n e^{in\omega t}, \tag{50.17}$$

where $\hat{a}_n$ is the complex number $a_n - ib_n$ (with $b_0 = 0$). If we wish to use the same notation throughout, we can write also

$$a_n = \frac{2}{T}\int_0^T f(t)e^{-in\omega t} \, dt \qquad (n \geq 1). \tag{50.18}$$

We now know how to "analyze" a periodic wave into its harmonic components. The procedure is called *Fourier analysis*, and the separate terms are called Fourier components. We have *not* shown, however, that once we find all of the Fourier components and add them together, we do indeed get back our $f(t)$. The mathematicians have shown, for a wide class of functions, in fact for all that are of interest to physicists, that if we can do the integrals we will get back $f(t)$. There is one minor exception. If the function $f(t)$ is discontinuous, i.e., if it jumps suddenly from one value to another, the Fourier sum will give a value at the breakpoint halfway between the upper and lower values at the discontinuity. So if we have the strange function $f(t) = 0, 0 \leq t < t_0$, and $f(t) = 1$ for $t_0 \leq t \leq T$, the Fourier sum will give the right value everywhere *except at* t_0, where it will have the value $\frac{1}{2}$ instead of 1. It is rather unphysical anyway to insist that a function should be zero *up to* t_0, but 1 *right at* t_0. So perhaps we should make the "rule" for physicists that any discontinuous function (which can only be a simplification of a *real* physical function) should be defined with halfway values at the discontinuities. Then any such function—with any finite number of such jumps—as well as all other physically interesting functions, are given correctly by the Fourier sum.

As an exercise, we suggest that the reader determine the Fourier series for the function shown in Fig. 50–3. Since the function cannot be written in an explicit algebraic form, you will not be able to do the integrals from zero to T in the usual way. The integrals are easy, however, if we separate them into two parts: the integral from zero to $T/2$ (over which $f(t) = 1$) and the integral from $T/2$ to T (over which $f(t) = -1$). The result should be

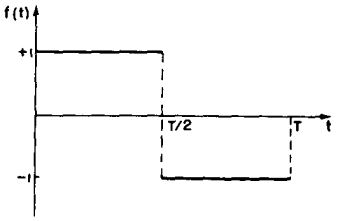


Fig. 50–3. Square-wave function.
$f(t) = +1$ for $0 < t < t/2$,
$f(t) = -1$ for $T/2 < t < T$.

$$f(t) = \frac{4}{\pi}(\sin \omega t + \tfrac{1}{3}\sin 3\omega t + \tfrac{1}{5}\sin 5\omega t + \cdots), \tag{50.19}$$

where $\omega = 2\pi/T$. We thus find that our square wave (with the particular phase chosen) has only odd harmonics, and their amplitudes are in inverse proportion to their frequencies.

Let us check that Eq. (50.19) does indeed give us back $f(t)$ for some value of t. Let us choose $t = T/4$, or $\omega t = \pi/2$. We have

$$f(t) = \frac{4}{\pi}\left(\sin\frac{\pi}{2} + \frac{1}{3}\sin\frac{3\pi}{2}\,\frac{1}{5}\sin\frac{5\pi}{2} + \cdots\right) \tag{50.20}$$

$$= \frac{4}{\pi}\left(1 - \frac{1}{3} + \frac{1}{5} - \frac{1}{7} + \cdots\right) \tag{50.21}$$

The series* has the value $\pi/4$, and we find that $f(t) = 1$.

50–5 The energy theorem

The energy in a wave is proportional to the square of its amplitude. For a wave of complex shape, the energy in one period will be proportional to $\int_0^T f^2(t)\,dt$. We can also relate this energy to the Fourier coefficients. We write

$$\int_0^T f^2(t)\,dt = \int_0^T \left[a_0 + \sum_{n=1}^{\infty} a_n \cos n\omega t + \sum_{n=1}^{\infty} b_n \sin n\omega t\right]^2 dt. \tag{50.22}$$

When we expand the square of the bracketed term we will get all possible cross terms, such as $a_5 \cos 5\omega t \cdot b_7 \cos 7\omega t$. We have shown above, however, [Eqs. (50.11) and (50.12)] that the integrals of all such terms over one period is zero.

---

* The series can be evaluated in the following way. First we remark that $\int_0^x [dx/(1 + x^2)] = \tan^{-1} x$. Second, we expand the integrand in a series $1/(1 + x^2) = 1 - x^2 + x^4 - x^6 + \ldots$ We integrate the series term by term (from zero to x) to obtain $\tan^{-1} x = 1 - x^3/3 + x^5/5 - x^7/7 + \ldots$ Setting $x = 1$, we have the stated result, since $\tan^{-1} 1 = \pi/4$.

We have left only the square terms like $a_5^2 \cos^2 5\omega t$. The integral of any cosine squared or sine squared over one period is equal to $T/2$, so we get

$$\int_0^T f^2(t)\,dt = Ta_0^2 + \frac{T}{2}(a_1^2 + a_2^2 + \cdots + b_1^2 + b_2^2 + \cdots)$$

$$= Ta_0^2 + \frac{T}{2}\sum_{n=1}^{\infty}(a_n^2 + b_n^2). \tag{50.23}$$

This equation is called the "energy theorem," and says that the total energy in a wave is just the sum of the energies in all of the Fourier components. For example, applying this theorem to the series (50.19), since $[f(t)]^2 = 1$ we get

$$T = \frac{T}{2}\cdot\left(\frac{4}{\pi}\right)^2\left(1 + \frac{1}{3^2} + \frac{1}{5^2} + \frac{1}{7^2}\cdots\right),$$

so we learn that the sum of the squares of the reciprocals of the odd integers is $\pi^2/8$. In a similar way, by first obtaining the Fourier series for the function and using the energy theorem, we can prove that $1 + 1/2^4 + 1/3^4 + \cdots$ is $\pi^4/90$, a result we needed in Chapter 45.

50–6 Nonlinear responses

Finally, in the theory of harmonics there is an important phenomenon which should be remarked upon because of its practical importance—that of nonlinear effects. In all the systems that we have been considering so far, we have supposed that everything was linear, that the responses to forces, say the displacements or the accelerations, were always proportional to the forces. Or that the currents in the circuits were proportional to the voltages, and so on. We now wish to consider cases where there is not a strict proportionality. We think, at the moment, of some device in which the response, which we will call x_{out} at the time t, is determined by the input x_{in} at the time t. For example, x_{in} might be the force and x_{out} might be the displacement. Or x_{in} might be the current and x_{out} the voltage. If the device is linear, we would have

$$x_{\text{out}}(t) = Kx_{\text{in}}(t), \tag{50.24}$$

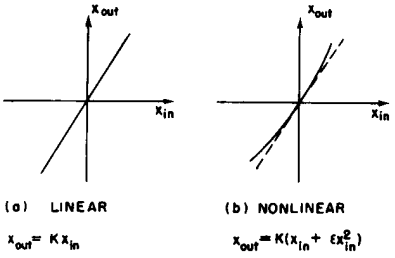


Fig. 50–4. Linear and nonlinear responses.

where K is a constant independent of t and of x_{in}. Suppose, however, that the device is nearly, but not exactly, linear, so that we can write

$$x_{\text{out}}(t) = K[x_{\text{in}}(t) + \epsilon x_{\text{in}}^2(t)], \tag{50.25}$$

where ϵ is small in comparison with unity. Such linear and nonlinear responses are shown in the graphs of Fig. 50–4.

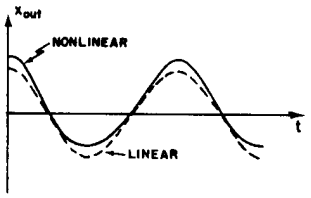


Fig. 50–5. The response of a nonlinear device to the input cos ωt. A linear response is shown for comparison.

Nonlinear responses have several important practical consequences. We shall discuss some of them now. First we consider what happens if we apply a pure tone at the input. We let $x_{\text{in}} = \cos \omega t$. If we plot x_{out} as a function of time we get the solid curve shown in Fig. 50–5. The dashed curve gives, for comparison, the response of a linear system. We see that the output is no longer a cosine function. It is more peaked at the top and flatter at the bottom. We say that the output is *distorted.* We know, however, that such a wave is no longer a pure tone, that it will have harmonics. We can find what the harmonics are. Using $x_{\text{in}} = \cos \omega t$ with Eq. (50.25), we have

$$x_{\text{out}} = K(\cos \omega t + \epsilon \cos^2 \omega t). \tag{50.26}$$

From the equality $\cos^2 \theta = \frac{1}{2}(1 - \cos 2\theta)$, we have

$$x_{\text{out}} = K\left(\cos \omega t + \frac{\epsilon}{2} - \frac{\epsilon}{2}\cos 2\omega t\right). \tag{50.27}$$

The output has not only a component at the fundamental frequency, that was present at the input, but also has some of its second harmonic. There has also

appeared at the output a constant term $K(\epsilon/2)$, which corresponds to the shift of the average value, shown in Fig. 50–5. The process of producing a shift of the average value is called *rectification.*

A nonlinear response will rectify and will produce harmonics of the frequencies at its input. Although the nonlinearity we assumed produced only second harmonics, nonlinearities of higher order—those which have terms like x_{in}^3 and x_{in}^4, for example—will produce harmonics higher than the second.

Another effect which results from a nonlinear response is *modulation.* If our input function contains two (or more) pure tones, the output will have not only their harmonics, but still other frequency components. Let $x_{\text{in}} = A\cos\omega_1 t + B\cos\omega_2 t$, where now ω_1 and ω_2 are *not* intended to be in a harmonic relation. In addition to the linear term (which is K times the input) we shall have a component in the output given by

$$x_{\text{out}} = K\epsilon(A\cos\omega_1 t + B\cos\omega_2 t)^2 \tag{50.28}$$

$$= K\epsilon(A^2\cos^2\omega_1 t + B^2\cos^2\omega_2 t + 2AB\cos\omega_1 t\cos\omega_2 t). \tag{50.29}$$

The first two terms in the parentheses of Eq. (50.29) are just those which gave the constant terms and second harmonic terms we found above. The last term is new.

We can look at this new "cross term" $AB\cos\omega_1 t\cos\omega_2 t$ in two ways. First, if the two frequencies are widely different (for example, if ω_1 is much greater than ω_2) we can consider that the cross term represents a cosine oscillation of varying amplitude. That is, we can think of the factors in this way:

$$AB\cos\omega_1 t\cos\omega_2 t = C(t)\cos\omega_1 t, \tag{50.30}$$

with

$$C(t) = AB\cos\omega_2 t. \tag{50.31}$$

We say that the amplitude of $\cos\omega_1$ is *modulated* with the frequency ω_2.

Alternatively, we can write the cross term in another way:

$$AB\cos\omega_1 t\cos\omega_2 t = \frac{AB}{2}[\cos(\omega_1 + \omega_2)t + \cos(\omega_1 - \omega_2)t]. \tag{50.32}$$

We would now say that two *new* components have been produced, one at the *sum* frequency $(\omega_1 + \omega_2)$, another at the *difference* frequency $(\omega_1 - \omega_2)$.

We have two different, but equivalent, ways of looking at the same result. In the special case that $\omega_1 \gg \omega_2$, we can relate these two different views by remarking that since $(\omega_1 + \omega_2)$ and $(\omega_1 - \omega_2)$ are near to each other we would expect to observe beats between them. But these beats have just the effect of *modulating* the amplitude of the *average* frequency ω_1 by one-half the difference frequency $2\omega_2$. We see, then, why the two descriptions are equivalent.

In summary, we have found that a nonlinear response produces several effects: rectification, generation of harmonics, and modulation, or the generation of components with sum and difference frequencies.

We should notice that all these effects (Eq. 50.29) are proportional not only to the nonlinearity coefficient ϵ, but also to the product of two amplitudes—either A^2, B^2, or AB. We expect these effects to be much more important for *strong* signals than for weak ones.

The effects we have been describing have many practical applications. First, with regard to sound, it is believed that the ear is nonlinear. This is believed to account for the fact that with loud sounds we have the sensation that we *hear* harmonics and also sum and difference frequencies even if the sound waves contain only pure tones.

The components which are used in sound-reproducing equipment—amplifiers, loudspeakers, etc.—always have some nonlinearity. They produce distortions in the sound—they generate harmonics, etc.—which were not present in the original sound. These new components are heard by the ear and are apparently objectionable. It is for this reason that "Hi-Fi" equipment is designed to be as linear as

possible. (Why the nonlinearities of the *ear* are *not* "objectionable" in the same way, or how we even know that the nonlinearity is in the *loudspeaker* rather than in the *ear* is not clear!)

Nonlinearities are quite *necessary*, and are, in fact, intentionally made large in certain parts of radio transmitting and receiving equipment. In an AM transmitter the "voice" signal (with frequencies of some kilocycles per second) is combined with the "carrier" signal (with a frequency of some megacycles per second) in a nonlinear circuit called a *modulator*, to produce the modulated oscillation that is transmitted. In the receiver, the components of the received signal are fed to a nonlinear circuit which combines the sum and difference frequencies of the modulated carrier to generate again the voice signal.

When we discussed the transmission of light, we assumed that the induced oscillations of charges were proportional to the electric field of the light—that the response was linear. That is indeed a very good approximation. It is only within the last few years that light sources have been devised (lasers) which produce an intensity of light strong enough so that nonlinear effects can be observed. It is now possible to generate harmonics of light frequencies. When a strong red light passes through a piece of glass, a little bit of blue light—second harmonic—comes out!

51

Waves

51–1 Bow waves

51–2 Shock waves

51–3 Waves in solids

51–4 Surface waves

51–1 Bow waves

Although we have finished our quantitative analyses of waves, this added chapter on the subject is intended to give some appreciation, qualitatively, for various phenomena that are associated with waves, which are too complicated to analyze in detail here. Since we have been dealing with waves for several chapters, more properly the subject might be called "some of the more complex phenomena associated with waves."

The first topic to be discussed concerns the effects that are produced by a source of waves which is moving faster than the wave velocity, or the phase velocity. Let us first consider waves that have a definite velocity, like sound and light. If we have a source of sound which is moving faster than the speed of sound, then something like this happens: Suppose at a given moment a sound wave is generated from the source at point x_1 in Fig. 51–1; then, in the next moment, as the source moves to x_2, the wave from x_1 expands by a radius r_1 smaller than the distance that the source moves; and, of course, another wave starts from x_2. When the sound source has moved still farther, to x_3, and a wave is starting there, the wave from x_2 has now expanded to r_2, and the one from x_1 has expanded to r_3. Of course the thing is done continuously, not in steps, and therefore, we have a series of wave circles with a common tangent line which goes through the center of the source. We see that instead of a source generating spherical waves, as it would if it were standing still, it generates a wavefront which forms a cone in three dimensions, or a pair of lines in two dimensions. The angle of the cone is very easy to figure out. In a given amount of time the source moves a distance, say $x_3 - x_1$, proportional to v, the velocity of the source. In the meantime the wavefront has moved out a distance r_3, proportional to c_w, the speed of the wave. Therefore it is clear that the half-angle of opening has a sine equal to the ratio of the speed of the waves, divided by the speed of the source, and this sine has a solution only if c_w is less than v, or the speed of the object is faster than the speed of the wave:

$$\sin\theta = \frac{c_w}{v}. \tag{51.1}$$

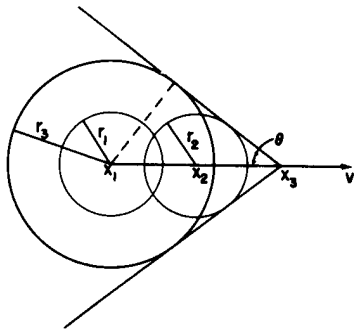


Fig. 51–1. The shock wave front lies on a cone with apex at the source and half-angle $\theta = \sin^{-1} v/c_w$.

Incidentally, although we implied that it is necessary to have a *source* of sound, it turns out, very interestingly, that once the object is moving faster than the speed of sound, it will *make* sound. That is, it is not necessary that it have a certain tone vibrational character. Any object moving through a medium faster than the speed at which the medium carries waves will generate waves on each side, automatically, just from the motion itself. This is simple in the case of sound, but it also occurs in the case of light. At first one might think nothing can move faster than the speed of light. However, light in glass has a phase velocity less than the speed of light in a vacuum, and it is possible to shoot a charged particle of very high energy through a block of glass such that the particle velocity is close to the speed of light in a vacuum, while the speed of light in the glass may be only $\frac{2}{3}$ the speed of light in the vacuum. A particle moving faster than the speed of light in the medium will produce a conical wave of light with its apex at the source, like the wave wake from a boat (which is from the same effect, as a matter of fact). By measuring the cone angle, we can determine the speed of the particle. This is used technically to determine the speeds of particles as one of the methods of determining their energy in high-energy research. The direction of the light is all that needs to be measured.

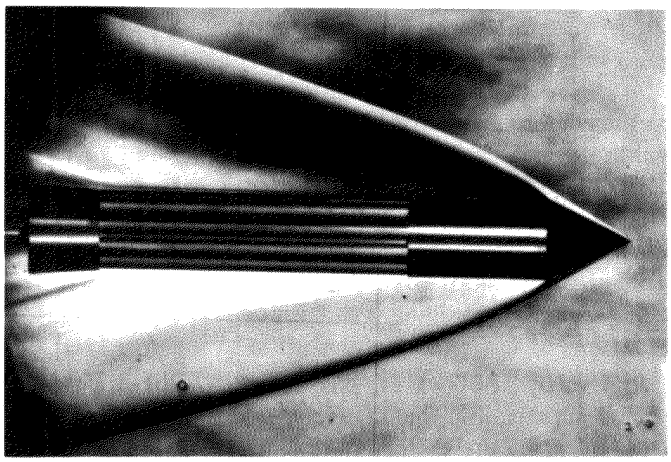


Fig. 51–2. A shock wave induced in a gas by a projectile moving faster than sound.

This light is sometimes called Çerenkov radiation, because it was first observed by Çerenkov. How intense this light should be was analyzed theoretically by Frank and Tamm. The 1958 Nobel Prize for physics was awarded jointly to all three for this work.

The corresponding circumstances in the case of sound are illustrated in Fig. 51–2, which is a photograph of an object moving through a gas at a speed greater than the speed of sound. The changes in pressure produce a change in refractive index, and with a suitable optical system the edges of the waves can be made visible. We see that the object moving faster than the speed of sound does, indeed, produce a conical wave. But closer inspection reveals that the surface is actually curved. It is straight asymptotically, but it is curved near the apex, and we have now to discuss how that can be, which brings us to the second topic of this chapter.

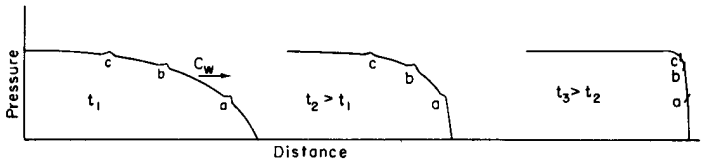


Fig. 51–3. Wavefront "snapshots" at successive instants in time.

51–2 Shock waves

Wave speed often depends on the amplitude, and in the case of sound the speed depends upon the amplitude in the following way. An object moving through the air has to move the air out of the way, so the disturbance produced in this case is some kind of a pressure step, with the pressure higher behind the wavefront than in the undisturbed region not yet reached by the wave (running along at the normal speed, say). But the air that is left behind, after the wavefront passes, has been compressed adiabatically, and therefore the temperature is increased. Now the speed of sound increases with the temperature, so the speed in the region behind the jump is faster than in the air in front. That means that any other disturbance that is made behind this step, say by a continuous pushing of the body, or any other disturbance, will ride faster than the front, the speed increasing with higher pressure. Figure 51–3 illustrates the situation, with some little bumps of pressure added to the pressure contour to aid visualization. We see that the higher pressure regions at the rear overtake the front as time goes on, until ultimately the compressional wave develops a sharp front. If the strength is very high, "ultimately" means right away; if it is rather weak, it takes a long time; it may be, in fact, that the sound is spreading and dying out before it has time to do this.

The sounds we make in talking are extremely weak relative to the atmospheric pressure—only 1 part in a million or so. But for pressure changes of the order of 1 atmosphere, the wave velocity increases by about twenty percent, and the wavefront sharpens up at a correspondingly high rate. In nature nothing happens *infinitely* rapidly, presumably, and what we call a "sharp" front has, actually, a very slight thickness; it is not infinitely steep. The distances over which it is varying are of the

order of one mean free path, in which the theory of the wave equation begins to fail because we did not consider the structure of the gas.

Now, referring again to Fig. 51–2, we see that the curvature can be understood if we appreciate that the pressures near the apex are higher than they are farther back, and so the angle θ is greater. That is, the curve is the result of the fact that the speed depends upon the strength of the wave. Therefore the wave from an atomic bomb explosion travels much faster than the speed of sound for a while, until it gets so far out that it is weakened to such an extent from spreading that the pressure bump is small compared with atmospheric pressure. The speed of the bump then approaches the speed of sound in the gas into which it is going. (Incidentally, it always turns out that the speed of the shock is higher than the speed of sound in the gas ahead, but is lower than the speed of sound in the gas behind. That is, impulses from the back will arrive at the front, but the front rides into the medium in which it is going faster than the normal speed of signals. So one cannot tell, acoustically, that the shock is coming until it is too late. The light from the bomb arrives first, but one cannot tell that the shock is coming until it arrives, because there is no sound signal coming ahead of it.)

This is a very interesting phenomenon, this piling up of waves, and the main point on which it depends is that after a wave is present, the speed of the resulting wave should be higher. Another example of the same phenomenon is the following. Consider water flowing in a long channel with finite width and finite depth. If a piston, or a wall across the channel, is moved along the channel fast enough, water piles up, like snow before a snow plow. Now suppose the situation is as shown in Fig. 51–4, with a sudden step in water height somewhere in the channel. It can be demonstrated that long waves in a channel travel faster in deeper water than they do in shallow water. Therefore any new bumps or irregularities in energy supplied by the piston run off forward and pile up at the front. Again, ultimately what we have is just water with a sharp front, theoretically. However, as Fig. 51–4 shows, there are complications. Pictured is a wave coming up a channel; the piston is at the far right end of the channel. At first it might have appeared like a well-behaved wave, as one might expect, but farther along the channel, it has become sharper and sharper until the events pictured occurred. There is a terrible churning at the surface, as the pieces of water fall down, but it is essentially a very sharp rise with no disturbance of the water ahead.



Figure 51–4

Actually water is much more complicated than sound. However, just to illustrate a point, we will try to analyze the speed of such a so-called *bore, in a channel.* The point here is not that this is of any basic importance for our purposes—it is not a great generalization—it is only to illustrate that the laws of mechanics that we already know are capable of explaining the phenomenon.

Imagine, for a moment, that the water does look something like Fig. 51–5(a), that water at the higher height h_2 is moving with a velocity v, and that the front is moving with velocity u into undisturbed water which is at height h_1. We would like to determine the speed at which the front moves. In a time Δt a vertical plane initially at x_1 moves a distance $v\,\Delta t$ to x_2, while the front of the wave has moved $u\,\Delta t$.

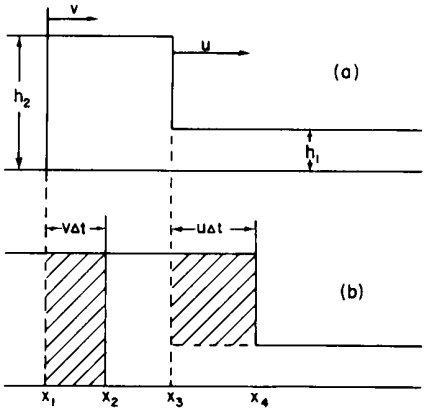


Fig. 51–5. Two cross sections of a bore in a channel, with (b) at an interval Δt later than (a).

Now we apply the equations of conservation of matter and momentum. First, the former: Per unit channel width, we see that the amount $h_2 v\,\Delta t$ of matter that has moved past x_1 (shown shaded) is compensated by the other shaded region, which amounts to $(h_2 - h_1)u\,\Delta t$. So, dividing by Δt, $vh_2 = u(h_2 - h_1)$. That does not yet give us enough, because although we have h_2 and h_1, we do not know either u or v; we are trying to get both of them.

Now the next step is to use conservation of momentum. We have not discussed the problems of water pressure, or anything in hydrodynamics, but it is clear anyway that the pressure of water at a given depth is just enough to hold up the column of water above it. Therefore the pressure of water is equal to ρ, the density of water, times g, times the depth below the surface. Since the pressure increases linearly with depth, the average pressure over the plane at x_1, say, is $\frac{1}{2}\rho g h_2$, which is also the average force per unit width and per unit height pushing the plane toward x_2. So we multiply by another h_2 to get the total force which is acting on the water

pushing from the left. On the other hand, there is pressure in the water on the right also, exerting an opposite force on the region in question, which is, by the same kind of analysis, $\frac{1}{2}\rho g h_1^2$. Now we must balance the forces against the rate of change of the momentum. Thus we have to figure out how much more momentum there is in situation (b) in Fig. 51–5 than there was in (a). We see that the additional mass that has acquired the speed v is just $\rho h_2 u\,\Delta t - \rho h_2 v\,\Delta t$ (per unit width), and multiplying this by v gives the additional momentum to be equated to the impulse $F\,\Delta t$:

$$(\rho h_2 u\,\Delta t - \rho h_2 v\,\Delta t)v = (\tfrac{1}{2}\rho g h_2^2 - \tfrac{1}{2}\rho g h_1^2)\,\Delta t.$$

If we eliminate v from this equation by substituting $vh_2 = u(h_2 - h_1)$, already found, and simplify, we get finally that $u^2 = gh_2(h_1 + h_2)/2h_1$.

If the height difference is very small, so that h_1 and h_2 are nearly equal, this says that the velocity $= \sqrt{gh}$. As we will see later, that is only true provided the wavelength of the wave is longer than the depth of the channel.

We could also do the analogous thing for sound waves—including the conservation of internal energy, not the conservation of entropy, because the shock is irreversible. In fact, if one checks the conservation of energy in the bore problem, one finds that energy is not conserved. If the height difference is small, it is almost perfectly conserved, but as soon as the height difference becomes very appreciable, there is a net loss of energy. This is manifested as the falling water and the churning shown in Fig. 51–4.

In shock waves there is a corresponding apparent loss of energy, from the point of view of adiabatic reactions. The energy in the sound wave, behind the shock, goes into heating of the gas after shock passes, corresponding to churning of the water in the bore. In working it out, three equations for the sound case turn out to be necessary for solution, and the temperature behind the shock is not the same as the temperature in front, as we have seen.

If we try to make a bore that is upside down ($h_2 < h_1$), then we find that the energy loss per second is negative. Since energy is not available from anywhere, that bore cannot then maintain itself; it is unstable. If we were to start a wave of that sort, it would flatten out, because the speed dependence on height that resulted in sharpening in the case we discussed would now have the opposite effect.

51–3 Waves in solids

The next kind of waves to be discussed are the more complicated waves in solids. We have already discussed sound waves in gas and in liquid, and there is a direct analog to a sound wave in a solid. If a sudden push is applied to a solid, it is compressed. It resists the compression, and a wave analogous to sound is started. However there is another kind of wave that is possible in a solid, and which is not possible in a fluid. If a solid is distorted by pushing it sideways (called *shearing*), then it tries to pull itself back. That is by definition what distinguishes a solid from a liquid: if we distort a liquid (internally), hold it a minute so that it calms down, and then let go, it will stay that way, but if we take a solid and push it, like shearing a piece of "Jello," and let it go, it flies back and starts a *shear* wave, travelling in the same way the compressions travel. In all cases, the shear wave speed is less than the speed of longitudinal waves. The shear waves are somewhat more analogous, so far as their polarizations are concerned, to light waves. Sound has no polarization, it is just a pressure wave. Light has a characteristic orientation perpendicular to its direction of travel.

In a solid, the waves are of both kinds. First, there is a compression wave, analogous to sound, that runs at one speed. If the solid is not crystalline, then a shear wave polarized in any direction will propagate at a characteristic speed. (Of course all solids are crystalline, but if we use a block made up of microcrystals of all orientations, the crystal anisotropies average out.)

Another interesting question concerning sound waves is the following: What happens if the wavelength in a solid gets shorter, and shorter, and shorter? How short can it get? It is interesting that it cannot get any shorter than the space

between the atoms, because if there is supposed to be a wave in which one point goes up and the next down, etc., the shortest possible wavelength is clearly the atom spacing. In terms of the modes of oscillation, we say that there are longitudinal modes, and transverse modes, long wave modes, short wave modes. As we consider wavelengths comparable to the spacing between the atoms, then the speeds are no longer constant; there is a dispersion effect where the velocity is not independent of the wave number. But, ultimately, the highest mode of transverse waves would be that in which every atom is doing the opposite of neighboring atoms.

Now from the point of view of atoms, the situation is like the two pendulums that we were talking about, for which there are two modes, one in which they both go together, and the other in which they go apart. It is possible to analyze the solid waves another way, in terms of a system of coupled harmonic oscillators, like an enormous number of pendulums, with the highest mode such that they oscillate oppositely, and lower modes with different relationships of the timing.

The shortest wavelengths are so short that they are not usually available technically. However they are of great interest because, in the theory of thermodynamics of a solid, the heat properties of a solid, for example specific heats, can be analyzed in terms of the properties of the short sound waves. Going to the extreme of sound waves of ever shorter wavelength, one necessarily comes to the individual motions of the atoms; the two things are the same ultimately.

A very interesting example of sound waves in a solid, both longitudinal and transverse, are the waves that are in the solid earth. Who makes the noises we do not know, but inside the earth, from time to time, there are earthquakes—some rock slides past some other rock. That is like a little noise. So waves like sound waves start out from such a source very much longer in wavelength than one usually considers in sound waves, but still they are sound waves, and they travel around in the earth. The earth is not homogeneous, however, and the properties of pressure, density, compressibility, and so on, change with depth, and therefore the speed varies with depth. Then the waves do not travel in straight lines—there is a kind of index of refraction and they go in curves. The longitudinal waves and the transverse waves have different speeds, so there are different solutions for the different speeds. Therefore if we place a seismograph at some location and watch the way the thing jiggles after there has been an earthquake somewhere else, then we do not just get an irregular jiggling. We might get a jiggling, and a quieting down, and then another jiggling—what happens depends upon the location. If it were close enough, we would first receive longitudinal waves from the disturbance, and then, a few moments later, transverse waves, because they travel more slowly. By measuring the time difference between the two, we can tell how far away the earthquake is, if we know enough about the speeds and composition of the interior regions involved.

An example of the behavior pattern of waves in the earth is shown in Fig. 51–6. The two kinds of waves are represented by different symbols. If there were an earthquake at the place marked "source," the transverse waves and longitudinal waves would arrive at different times at the station by the most direct routes, and there would also be reflections at discontinuities, resulting in other paths and times. It turns out that there is a core in the earth which does not carry transverse waves. If the station is opposite the source, transverse waves still arrive, but the timing is not right. What happens is that the transverse wave comes to the core, and whenever the transverse waves come to a surface which is oblique, between two materials, two new waves are generated, one transverse and one longitudinal. But inside the core of the earth, a transverse wave is not propagated (or at least, there is no evidence for it, only for a longitudinal wave); it comes out again in both forms and comes to the station.

It is from the behavior of these earthquake waves that it has been determined that transverse waves cannot be propagated within the inner circle. This means that the center of the earth is liquid in the sense that it cannot propagate transverse waves. The only way we know what is inside the earth is by studying earthquakes. So, by using a large number of observations of many earthquakes at different

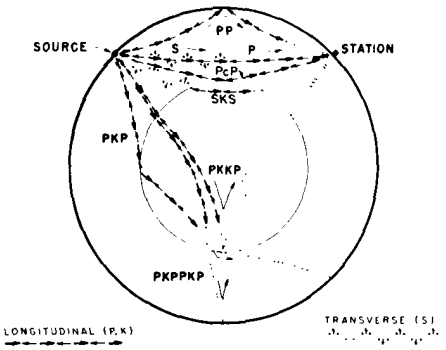


Fig. 51–6. Schematic of the earth, showing paths of longitudinal and transverse sound waves.

stations, the details have been worked out—the speed, the curves, etc. are all known. We know what the speeds of various kinds of waves are at every depth. Knowing that, therefore, it is possible to figure out what the normal modes of the earth are, because we know the speed of propagation of sound waves—in other words, the elastic properties of both kinds of waves at every depth. Suppose the earth were distorted into an ellipsoid and let go. It is just a matter of superposing waves travelling around in the ellipsoid to determine the period and shapes in a free mode. We have figured out that if there is a disturbance, there are a lot of modes, from the lowest, which is ellipsoidal, to higher modes with more structure.

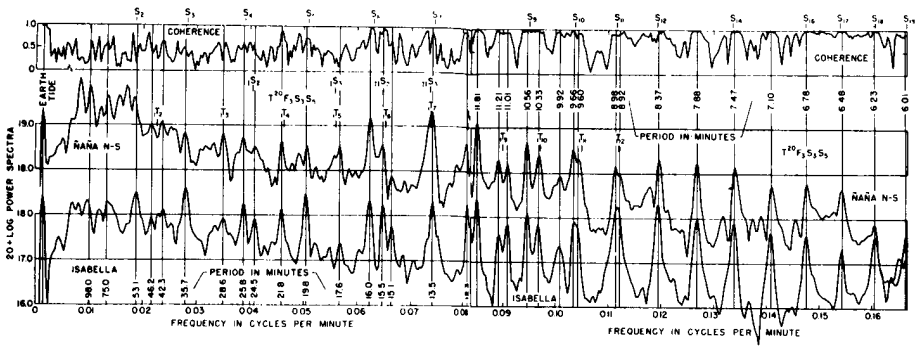


Fig. 51–7. Power versus frequency as detected at seismographs in Ñaña, Peru, and Isabella, California. The coherence is a measure of the coupling between the stations. [From Benioff, Press and Smith, *J. Geoph. Research* **66**, 605 (1961)].

The Chilean earthquake of May 1960 made a loud enough "noise" that the signals went around the earth many times, and new seismographs of great delicacy were made just in time to determine the frequencies of the fundamental modes of the earth and to compare them with the values that were calculated from the theory of sound with the known velocities, as measured from the independent earthquakes. The result of this experiment is illustrated in Fig. 51–7, which is a plot of the strength of the signal versus the frequency of its oscillation (a *Fourier analysis*). Note that at certain particular frequencies there is much more being received than at other frequencies; there are very definite maxima. These are the natural frequencies of the earth, because these are the main frequencies at which the earth can oscillate. In other words, if the entire motion of the earth is made up of many different modes, we would expect to obtain, for each station, irregular bumpings which indicate a superposition of many frequencies. If we analyze this in terms of frequencies, we should be able to find the characteristic frequencies of the earth. The vertical dark lines in the figure are the calculated frequencies, and we find a remarkable agreement, an agreement due to the fact that the theory of sound is right for the inside of the earth.

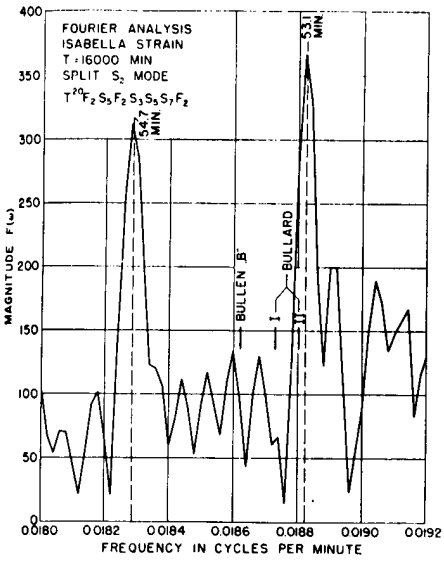


Fig. 51–8. High-resolution analysis of one of the seismograph records, showing spectral doublet.

A very curious point is revealed in Fig. 51–8, which shows a very careful measurement, with better resolution of the lowest mode, the ellipsoidal mode of the earth. Note that it is not a single maximum, but a double one, 54.7 minutes and 53.1 minutes—slightly different. The reason for the two different frequencies was not known at the time that it was measured, although it may have been found in the meantime. There are at least two possible explanations: One would be that there may be asymmetry in the earth's distribution, which would result in two similar modes. Another possibility, which is even more interesting, is this: Imagine the waves going around the earth in two directions from the source. The speeds will not be equal because of effects of the rotation of the earth in the equations of motion, which have not been taken into account in making the analysis. Motion in a rotating system is modified by Coriolis forces, and these may cause the observed splitting.

Regarding the method by which these quakes have been analyzed, what is obtained on the seismograph is not a curve of amplitude as a function of frequency, but displacement as a function of time, always a very irregular tracing. To find the amount of all the different sine waves for all different frequencies, we know that the trick is to multiply the data by a sine wave of a given frequency and integrate, i.e., average it, and in the average all other frequencies disappear. The figures were thus plots of the integrals found when the data were multiplied by sine waves of different cycles per minute, and integrated.

51–4 Surface waves

Now, the next waves of interest, that are easily seen by everyone and which are usually used as an example of waves in elementary courses, are water waves. As we shall soon see, they are the worst possible example, because they are in no respects like sound and light; they have all the complications that waves can have. Let us start with long water waves in deep water. If the ocean is considered infinitely deep and a disturbance is made on the surface, waves are generated. All kinds of irregular motions occur, but the sinusoidal type motion, with a very small disturbance, might look like the common smooth ocean waves coming in toward the shore. Now with such a wave, the water, of course, on the average, is standing still, but the wave moves. What is the motion, is it transverse or longitudinal? It must be neither; it is not transverse, nor is it longitudinal. Although the water at a given place is alternately trough or hill, it cannot simply be moving up and down, by the conservation of water. That is, if it goes down, where is the water going to go? The water is essentially incompressible. The speed of compression of waves—that is, sound in the water—is much, much higher, and we are not considering that now. Since water is incompressible on this scale, as a hill comes down the water must move away from the region. What actually happens is that particles of water near the surface move approximately in circles. When smooth swells are coming, a person floating in a tire can look at a nearby object and see it going in a circle. So it is a mixture of longitudinal and transverse, to add to the confusion. At greater depths in the water the motions are smaller circles until, reasonably far down, there is nothing left of the motion (Fig. 51–9).

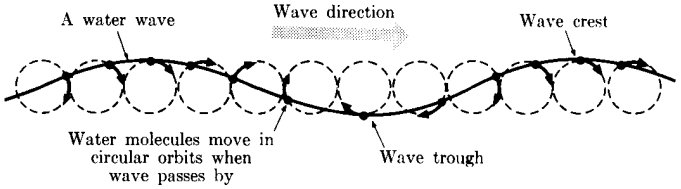


Fig. 51–9. Deep-water waves are formed from particles moving in circles. Note the systematic phase shifts from circle to circle. How would a floating object move?

To find the velocity of such waves is an interesting problem: it must be some combination of the density of the water, the acceleration of gravity, which is the restoring force that makes the waves, and possibly of the wavelength and of the depth. If we take the case where the depth goes to infinity, it will no longer depend on the depth. Whatever formula we are going to get for the velocity of the phases of the waves must combine the various factors to make the proper dimensions, and if we try this in various ways, we find only one way to combine the density, g, and λ in order to make a velocity, namely, $\sqrt{g\lambda}$, which does not include the density at all. Actually, this formula for the phase velocity is not exactly right, but a complete analysis of the dynamics, which we will not go into, shows that the factors are as we have them, except for $\sqrt{2\pi}$:

$$v_{\text{phase}} = \sqrt{g\lambda/2\pi} \text{ (for gravity waves).}$$

It is interesting that the long waves go faster than the short waves. Thus if a boat makes waves far out, because there is some sports-car driver in a motorboat travelling by, then after a while the waves come to shore with slow sloshings at first and then more and more rapid sloshings, because the first waves that come are long. The waves get shorter and shorter as the time goes on, because the velocities go as the square root of the wavelength.

One may object, "That is not right, we must look at the *group* velocity in order to figure it out!" Of course that is true. The formula for the phase velocity does not tell us what is going to arrive first; what tells us is the group velocity. So we have to work out the group velocity, and it is left as a problem to show it to be one-half of the phase velocity, assuming that the velocity goes as the square root of the wavelength, which is all that is needed. The group velocity also goes as the square root of the wavelength. How can the group velocity go half as fast as the phase? If one looks at the bunch of waves that are made by a boat travelling

along, following a particular crest, he finds that it moves forward in the group and gradually gets weaker and dies out in the front, and mystically and mysteriously a weak one in the back works its way forward and gets stronger. In short, the waves are moving through the group while the group is only moving at half the speed that the waves are moving.

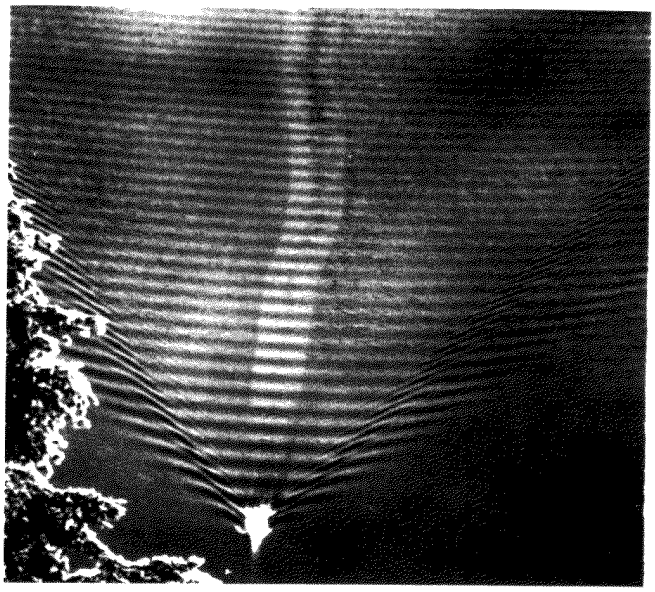


Fig. 51–10. The wake of a boat.

Because the group velocities and phase velocities are not equal, then the waves that are produced by an object moving through are no longer simply a cone, but it is much more interesting. We can see that in Fig. 51–10, which shows the waves produced by an object moving through the water. Note that it is quite different than what we would have for sound, in which the velocity is independent of wavelength, where we would have wavefronts only along the cone, travelling outward. Instead of that, we have waves in the back with fronts moving parallel to the motion of the boat, and then we have little waves on the sides at other angles. This entire pattern of waves can, with ingenuity, be analyzed by knowing only this: that the phase velocity is proportional to the square root of the wavelength. The trick is that the pattern of waves is stationary relative to the (constant-velocity) boat; any other pattern would get lost from the boat.

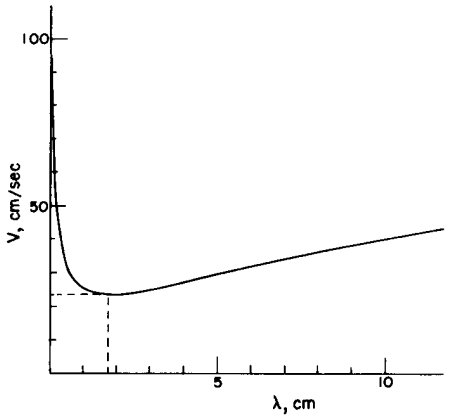


Fig. 51–11. Phase velocity vs. wavelength for water.

The water waves that we have been considering so far were long waves in which the force of restoration is due to gravitation. But when waves get very short in the water, the main restoring force is capillary attraction, i.e., the energy of the surface, the surface tension. For surface tension waves, it turns out that the phase velocity is

$$v_{\text{phase}} = \sqrt{2\pi T/\lambda\rho} \text{ (for ripples)},$$

where T is the surface tension and ρ the density. It is the exact opposite: the phase velocity is *higher*, the shorter the wavelength, when the wavelength gets very small. When we have both gravity and capillary action, as we always do, we get the combination of these two together:

$$v_{\text{phase}} = \sqrt{Tk/\rho + g/k},$$

where $k = 2\pi/\lambda$ is the wave number. So the velocity of the waves of water is really quite complicated. The phase velocity as a function of the wavelength is shown in Fig. 51–11; for very short waves it is fast, for very long waves it is fast, and there is a minimum speed at which the waves can go. The group velocity can be calculated from the formula: it goes to $\frac{3}{2}$ the phase velocity for ripples and $\frac{1}{2}$ the phase velocity for gravity waves. To the left of the minimum the group velocity is higher than the phase velocity; to the right, the group velocity is less than the

phase velocity. There are a number of interesting phenomena associated with these facts. In the first place, since the group velocity is increasing so rapidly as the wavelength goes down, if we make a disturbance there will be a slowest end of the disturbance going at the minimum speed with the corresponding wavelength, and then in front, going at higher speed, will be a short wave and a very long wave. It is very hard to see the long ones, but it is easy to see the short ones in a water tank.

So we see that the ripples often used to illustrate simple waves are quite interesting and complicated; they do not have a sharp wavefront at all, as is the case for simple waves like sound and light. The main wave has little ripples which run out ahead. A sharp disturbance in the water does not produce a sharp wave because of the dispersion. First come the very fine waves. Incidentally, if an object moves through the water at a certain speed, a rather complicated pattern results, because all the different waves are going at different speeds. One can demonstrate this with a tray of water and see that the fastest ones are the fine capillary waves. There are slowest waves, of a certain kind, which go behind. By inclining the bottom, one sees that where the depth is lower, the speed is lower. If a wave comes in at an angle to the line of maximum slope, it bends and tends to follow that line. In this way one can show various things, and we conclude that waves are more complicated in water than in air.

The speed of long waves in water with circulational motions is slower when the depth is less, faster in deep water. Thus as water comes toward a beach where the depth lessens, the waves go slower. But where the water is deeper, the waves are faster, so we get the effects of shock waves. This time, since the wave is not so simple, the shocks are much more contorted, and the wave over-curves itself, in the familiar way shown in Fig. 51–12. This is what happens when waves come into the shore, and the real complexities in nature are well revealed in such a circumstance. No one has yet been able to figure out what shape the wave should take as it breaks. It is easy enough when the waves are small, but when one gets large and breaks, then it is much more complicated.



Fig. 51–12. A water wave.

An interesting feature about capillary waves can be seen in the disturbances made by an object moving through the water. From the point of view of the object itself, the water is flowing past, and the waves which ultimately sit around it are always the waves which have just the right speed to stay still with the object in the water. Similarly, around an object in a stream, with the stream flowing by, the pattern of waves is stationary, and at just the right wavelengths to go at the same speed as the water going by. But if the group velocity is less than the phase velocity, then the disturbances propagate out backwards in the stream, because the group velocity is not quite enough to keep up with the stream. If the group velocity is faster than the velocity of the phase, the pattern of waves will appear in front of the object. If one looks closely at objects in a stream, one can see that there are little ripples in front and long "slurps" in the back.

Another interesting feature of this sort can be observed in pouring liquids. If milk is poured fast enough out of a bottle, for instance, a large number of lines can be seen crossing both ways in the outgoing stream. They are waves starting from the disturbance at the edges and running out, much like the waves about an object in a stream. There are effects from both sides which produce the crossed pattern.

We have investigated some of the interesting properties of waves and the various complications of dependence of phase velocity on wavelength, the speed of the waves on depth, and so forth, that produce the really complex, and therefore interesting, phenomena of nature.

52

Symmetry in Physical Laws

52–1 Symmetry operations

52–2 Symmetry in space and time

52–3 Symmetry and conservation laws

52–4 Mirror reflections

52–5 Polar and axial vectors

52–6 Which hand is right?

52–7 Parity is not conserved!

52–8 Antimatter

52–9 Broken symmetries

52–1 Symmetry operations

The subject of this chapter is what we may call *symmetry in physical laws.* We have already discussed certain features of symmetry in physical laws in connection with vector analysis (Chapter 11), the theory of relativity (Chapter 16), and rotation (Chapter 20).

Why should we be concerned with symmetry? In the first place, symmetry is fascinating to the human mind, and everyone likes objects or patterns that are in some way symmetrical. It is an interesting fact that nature often exhibits certain kinds of symmetry in the objects we find in the world around us. Perhaps the most symmetrical object imaginable is a sphere, and nature is full of spheres—stars, planets, water droplets in clouds. The crystals found in rocks exhibit many different kinds of symmetry, the study of which tells us some important things about the structure of solids. Even the animal and vegetable worlds show some degree of symmetry, although the symmetry of a flower or of a bee is not as perfect or as fundamental as is that of a crystal.

But our main concern here is not with the fact that the *objects* of nature are often symmetrical. Rather, we wish to examine some of the even more remarkable symmetries of the universe—the symmetries that exist in the *basic laws themselves* which govern the operation of the physical world.

First, what *is* symmetry? How can a physical *law* be "symmetrical"? The problem of defining symmetry is an interesting one and we have already noted that Weyl gave a good definition, the substance of which is that a thing is symmetrical if there is something we can do to it so that after we have done it, it looks the same as it did before. For example, a symmetrical vase is of such a kind that if we reflect or turn it, it will look the same as it did before. The question we wish to consider here is what we can do to physical phenomena, or to a physical situation in an experiment, and yet leave the result the same. A list of the known operations under which various physical phenomena remain invariant is shown in Table 52–1.

Table 52–1

Symmetry Operations

Translation in space
Translation in time
Rotation through a fixed angle
Uniform velocity in a straight line (Lorentz transformation)
Reversal of time
Reflection of space
Interchange of identical atoms or identical particles
Quantum-mechanical phase
Matter-antimatter (charge conjugation)

52–2 Symmetry in space and time

The first thing we might try to do, for example, is to *translate* the phenomenon in space. If we do an experiment in a certain region, and then build another apparatus at another place in space (or move the original one over) then, whatever went on in one apparatus, in a certain order in time, will occur in the same way if we have arranged the same condition, with all due attention to the restrictions that we mentioned before: that all of those features of the environment which make it not behave the same way have also been moved over—we talked about how to define how much we should include in those circumstances, and we shall not go into those details again.

In the same way, we also believe today that *displacement in time* will have no effect on physical laws. (That is, *as far as we know today*—all of these things are as far as we know today!) That means that if we build a certain apparatus and start it at a certain time, say on Thursday at 10:00 a.m., and then build the same apparatus and start it, say, three days later in the same condition, the two apparatuses will go through the same motions in exactly the same way as a function of time no matter what the starting time, provided again, of course, that the relevant features of the environment are also modified appropriately in *time*. That symmetry means,

of course, that if one bought General Motors stock three months ago, the same thing would happen to it if he bought it now!

We have to watch out for geographical differences too, for there are, of course, variations in the characteristics of the earth's surface. So, for example, if we measure the magnetic field in a certain region and move the apparatus to some other region, it may not work in precisely the same way because the magnetic field is different, but we say that is because the magnetic field is associated with the earth. We can imagine that if we move the whole earth and the equipment, it would make no difference in the operation of the apparatus.

Another thing that we discussed in considerable detail was rotation in space: if we turn an apparatus at an angle it works just as well, provided we turn everything else that is relevant along with it. In fact, we discussed the problem of symmetry under rotation in space in some detail in Chapter 11, and we invented a mathematical system called *vector analysis* to handle it as neatly as possible.

On a more advanced level we had another symmetry—the symmetry under uniform velocity in a straight line. That is to say—a rather remarkable effect—that if we have a piece of apparatus working a certain way and then take the same apparatus and put it in a car, and move the whole car, plus all the relevant surroundings, at a uniform velocity in a straight line, then so far as the phenomena inside the car are concerned there is no difference: all the laws of physics appear the same. We even know how to express this more technically, and that is that the mathematical equations of the physical laws must be unchanged under a *Lorentz transformation.* As a matter of fact, it was a study of the relativity problem that concentrated physicists' attention most sharply on symmetry in physical laws.

Now the above-mentioned symmetries have all been of a geometrical nature, time and space being more or less the same, but there are other symmetries of a different kind. For example, there is a symmetry which describes the fact that we can replace one atom by another of the same kind; to put it differently, there *are* atoms of the same kind. It is possible to find groups of atoms such that if we change a pair around, it makes no difference—the atoms are identical. Whatever one atom of oxygen of a certain type will do, another atom of oxygen of that type will do. One may say, "That is ridiculous, that is the *definition* of equal types!" That may be merely the definition, but then we still do not know whether there *are* any "atoms of the same type"; the *fact* is that there are many, many atoms of the same type. Thus it does mean something to say that it makes no difference if we replace one atom by another of the same type. The so-called elementary particles of which the atoms are made are also identical particles in the above sense—all electrons are the same; all protons are the same; all positive pions are the same; and so on.

After such a long list of things that can be done without changing the phenomena, one might think we could do practically anything; so let us give some examples to the contrary, just to see the difference. Suppose that we ask: "Are the physical laws symmetrical under a change of scale?" Suppose we build a certain piece of apparatus, and then build another apparatus five times bigger in every part, will it work exactly the same way? The answer is, in this case, *no*! The wavelength of light emitted, for example, by the atoms inside one box of sodium atoms and the wavelength of light emitted by a gas of sodium atoms five times in volume is not five times longer, but is. in fact exactly the same as the other. So the ratio of the wavelength to the size of the emitter will change.

Another example: we see in the newspaper, every once in a while pictures of a great cathedral made with little matchsticks—a tremendous work of art by some retired fellow who keeps gluing matchsticks together. It is much more elaborate and wonderful than any real cathedral. If we imagine that this wooden cathedral were actually built on the scale of a real cathedral, we see where the trouble is; it would not last—the whole thing would collapse because of the fact that scaled-up matchsticks are just not strong enough. "Yes," one might say, "but we also know that when there is an influence from the outside, it also must be changed in proportion!" We are talking about the ability of the object to withstand gravitation. So what we should do is first to take the model cathedral of real matchsticks and

the real earth, and then we know it is stable. Then we should take the larger cathedral and take a bigger earth. But then it is even worse, because the gravitation is increased still more!

Today, of course, we understand the fact that phenomena depend on the scale on the grounds that matter is atomic in nature, and certainly if we built an apparatus that was so small there were only five atoms in it, it would clearly be something we could not scale up and down arbitrarily. The scale of an individual atom is not at all arbitrary—it is quite definite.

The fact that the laws of physics are not unchanged under a change of scale was discovered by Galileo. He realized that the strengths of materials were not in exactly the right proportion to their sizes, and he illustrated this property that we were just discussing, about the cathedral of matchsticks, by drawing two bones, the bone of one dog, in the right proportion for holding up his weight, and the imaginary bone of a "super dog" that would be, say, ten or a hundred times bigger—that bone was a big, solid thing with quite different proportions. We do not know whether he ever carried the argument quite to the conclusion that the laws of nature must have a definite scale, but he was so impressed with this discovery that he considered it to be as important as the discovery of the laws of motion, because he published them both in the same volume, called "On Two New Sciences."

Another example in which the laws are not symmetrical, that we know quite well, is this: a system in rotation at a uniform angular velocity does not give the same apparent laws as one that is not rotating. If we make an experiment and then put everything in a space ship and have the space ship spinning in empty space, all alone at a constant angular velocity, the apparatus will not work the same way because, as we know, things inside the equipment will be thrown to the outside, and so on, by the centrifugal or coriolis forces, etc. In fact, we can tell that the earth is rotating by using a so-called Foucault pendulum, without looking outside.

Next we mention a very interesting symmetry which is obviously false, i.e., *reversibility in time*. The physical laws apparently cannot be reversible in time, because, as we know, all obvious phenomena are irreversible on a large scale: "The moving finger writes, and having writ, moves on." So far as we can tell, this irreversibility is due to the very large number of particles involved, and if we could see the individual molecules, we would not be able to discern whether the machinery was working forward or backwards. To make it more precise: we build a small apparatus in which we know what all the atoms are doing, in which we can watch them jiggling. Now we build another apparatus like it, but which starts its motion in the final condition of the other one, with all the velocities precisely reversed. *It will then go through the same motions, but exactly in reverse.* Putting it another way: if we take a motion picture, with sufficient detail, of all the inner works of a piece of material and shine it on a screen and run it backwards, no physicist will be able to say, "That is against the laws of physics, that is doing something wrong!" If we do not see all the details, of course, the situation will be perfectly clear. If we see the egg splattering on the sidewalk and the shell cracking open, and so on, then we will surely say, "That is irreversible, because if we run the moving picture backwards the egg will all collect together and the shell will go back together, and that is obviously ridiculous!" But if we look at the individual atoms themselves, the laws look completely reversible. This is, of course, a much harder discovery to have made, but apparently it is true that the fundamental physical laws, on a microscopic and fundamental level, are completely reversible in time!

52–3 Symmetry and conservation laws

The symmetries of the physical laws are very interesting at this level, but they turn out, in the end, to be even more interesting and exciting when we come to quantum mechanics. For a reason which we cannot make clear at the level of the present discussion—a fact that most physicists still find somewhat staggering, a most profound and beautiful thing, is that, in quantum mechanics, *for each of the rules of symmetry there is a corresponding conservation law;* there is a definite

connection between the laws of conservation and the symmetries of physical laws. We can only state this at present, without any attempt at explanation.

The fact, for example, that the laws are symmetrical for translation in space when we add the principles of quantum mechanics, turns out to mean that *momentum is conserved.*

That the laws are symmetrical under translation in time means, in quantum mechanics, that *energy is conserved.*

Invariance under rotation through a fixed angle in space corresponds to the *conservation of angular momentum.* These connections are very interesting and beautiful things, among the most beautiful and profound things in physics.

Incidentally, there are a number of symmetries which appear in quantum mechanics which have no classical analog, which have no method of description in classical physics. One of these is as follows: If ψ is the amplitude for some process or other, we know that the absolute square of ψ is the probability that the process will occur. Now if someone else were to make his calculations, not with this ψ, but with a ψ' which differs merely by a change in phase (let Δ be some constant, and multiply $e^{i\Delta}$ times the old ψ), the absolute square of ψ', which is the probability of the event, is then equal to the absolute square of ψ:

$$\psi' = \psi e^{i\Delta}; \qquad |\psi'|^2 = |\psi|^2. \tag{52.1}$$

Therefore the physical laws are unchanged if the phase of the wave function is shifted by an arbitrary constant. That is another symmetry. Physical laws must be of such a nature that a shift in the quantum-mechanical phase makes no difference. As we have just mentioned, in quantum mechanics there is a conservation law for every symmetry. The conservation law which is connected with the quantum-mechanical phase seems to be the *conservation of electrical charge.* This is altogether a very interesting business!

52–4 Mirror reflections

Now the next question, which is going to concern us for most of the rest of this chapter, is the question of symmetry under *reflection in space.* The problem is this: Are the physical laws symmetrical under reflection? We may put it this way: Suppose we build a piece of equipment, let us say a clock, with lots of wheels and hands and numbers; it ticks, it works, and it has things wound up inside. We look at the clock in the mirror. How it *looks* in the mirror is not the question. But let us actually *build* another clock which is exactly the same as the first clock looks in the mirror—every time there is a screw with a right-hand thread in one, we use a screw with a left-hand thread in the corresponding place of the other; where one is marked "2" on the face, we mark a "Ƨ" on the face of the other; each coiled spring is twisted one way in one clock and the other way in the mirror-image clock; when we are all finished, we have two clocks, both physical, which bear to each other the relation of an object and its mirror image, although they are both actual, material objects, we emphasize. Now the question is: If the two clocks are started in the same condition, the springs wound to corresponding tightnesses, will the two clocks tick and go around, forever after, as exact mirror images? (This is a physical question, not a philosophical question.) Our intuition about the laws of physics would suggest that they *would.*

We would suspect that, at least in the case of these clocks, reflection in space is one of the symmetries of physical laws, that if we change everything from "right" to "left" and leave it otherwise the same, we cannot tell the difference. Let us, then, suppose for a moment that this is true. If it is true, then it would be impossible to distinguish "right" and "left" by any physical phenomenon, just as it is, for example, impossible to define a particular absolute velocity by a physical phenomenon. So it should be impossible, by any physical phenomenon, to define absolutely what we mean by "right" as opposed to "left," because the physical laws should be symmetrical.

Of course, the world does not *have* to be symmetrical. For example, using what we may call "geography," surely "right" can be defined. For instance, we stand

in New Orleans and look at Chicago, and Florida is to our right (when our feet are on the ground!). So we can define "right" and "left" by geography. Of course, the actual situation in any system does not have to have the symmetry that we are talking about; it is a question of whether the *laws* are symmetrical—in other words, whether it is *against the physical laws* to have a sphere like the earth with "left-handed dirt" on it and a person like ourselves standing looking at a city like Chicago from a place like New Orleans, but with everything the other way around, so Florida is on the other side. It clearly seems not impossible, not against the physical laws, to have everything changed left for right.

Another point is that our definition of "right" should not depend on history. An easy way to distinguish right from left is to go to a machine shop and pick up a screw at random. The odds are it has a right-hand thread—not necessarily, but it is much more likely to have a right-hand thread than a left-hand one. This is a question of history or convention, or the way things happen to be, and is again not a question of fundamental laws. As we can well appreciate, everyone could have started out making left-handed screws!

So we must try to find some phenomenon in which "right hand" is involved fundamentally. The next possibility we discuss is the fact that polarized light rotates its plane of polarization as it goes through, say, sugar water. As we saw in Chapter 33, it rotates, let us say, to the right in a certain sugar solution. That is a way of defining "right-hand," because we may dissolve some sugar in the water and then the polarization goes to the right. But sugar has come from living things, and if we try to make the sugar artificially, then we discover that it *does not* rotate the plane of polarization! But if we then take that same sugar which is made artificially and which does not rotate the plane of polarization, and put bacteria in it (they eat some of the sugar) and then filter out the bacteria, we find that we still have sugar left (almost half as much as we had before), and this time it does rotate the plane of polarization, but *the other way*! It seems very confusing, but is easily explained.

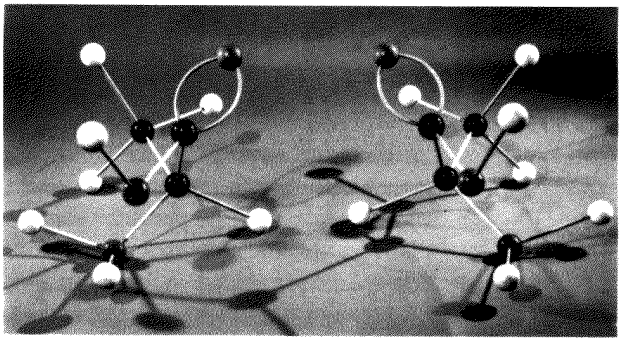


Fig. 52–1. (a) L-alanine (left), and (b) D-alanine (right).

Take another example: One of the substances which is common to all living creatures and that is fundamental to life is protein. Proteins consist of chains of amino acids. Figure 52–1 shows a model of an amino acid that comes out of a protein. This amino acid is called alanine, and the molecular arrangement would look like that in Fig. 52–1(a) if it came out of a protein of a real living thing. On the other hand, if we try to make alanine from carbon dioxide, ethane, and ammonia (and we *can* make it, it is not a complicated molecule), we discover that we are making equal amounts of this molecule and the one shown in Fig. 52–1(b)! The first molecule, the one that comes from the living thing, is called *L-alanine*. The other one, which is the same chemically, in that it has the same kinds of atoms and the same connections of the atoms, is a "right-hand" molecule, compared with the "left-hand" L-alanine, and it is called *D-alanine*. The interesting thing is that when we make alanine at home in a laboratory from simple gases, we get an equal mixture of both kinds. However, the only thing that life uses is L-alanine. (This is not exactly true. Here and there in living creatures there is a special use for D-alanine, but it is very rare. All proteins use L-alanine exclusively.) Now if we make both kinds, and we feed the mixture to some animal which likes to "eat," or use up, alanine, it cannot use D-alanine, so it only uses the L-alanine; that is what happened to our sugar—after the bacteria eat the sugar that works

well for them, only the "wrong" kind is left! (Left-handed sugar tastes sweet, but not the same as right-handed sugar.)

So it looks as though the phenomena of life permit a distinction between "right" and "left," or chemistry permits a distinction, because the two molecules are chemically different. But no, it does not! So far as physical measurements can be made, such as of energy, the rates of chemical reactions, and so on, the two kinds work exactly the same way if we make everything else in a mirror image too. One molecule will rotate light to the right, and the other will rotate it to the left in precisely the same amount, through the same amount of fluid. Thus, so far as physics is concerned, these two amino acids are equally satisfactory. So far as we understand things today, the fundamentals of the Schrödinger equation have it that the two molecules should behave in exactly corresponding ways, so that one is to the right as the other is to the left. Nevertheless, in life it is all one way!

It is presumed that the reason for this is the following. Let us suppose, for example, that life is somehow at one moment in a certain condition in which all the proteins in some creatures have left-handed amino acids, and all the enzymes are lopsided—every substance in the living creature is lopsided—it is not symmetrical. So when the digestive enzymes try to change the chemicals in the food from one kind to another, one kind of chemical "fits" into the enzyme, but the other kind does not (like Cinderella and the slipper, except that it is a "left foot" that we are testing). So far as we know, in principle, we could build a frog, for example, in which every molecule is reversed, everything is like the "left-hand" mirror image of a real frog; we have a left-hand frog. This left-hand frog would go on all right for a while, but he would find nothing to eat, because if he swallows a fly, his enzymes are not built to digest it. The fly has the wrong "kind" of amino acids (unless we give him a left-hand fly). So as far as we know, the chemical and life processes would continue in the same manner if everything were reversed.

If life is entirely a physical and chemical phenomenon, then we can understand that the proteins are all made in the same corkscrew only from the idea that at the very beginning some living molecules, by accident, got started and a few won. Somewhere, once, one organic molecule was lopsided in a certain way, and from this particular thing the "right" happened to evolve in our particular geography; a particular historical accident was one-sided, and ever since then the lopsidedness has propagated itself. Once having arrived at the state that it is in now, of course, it will always continue—all the enzymes digest the right things, manufacture the right things: when the carbon dioxide and the water vapor, and so on, go in the plant leaves, the enzymes that make the sugars make them lopsided because the enzymes are lopsided. If any new kind of virus or living thing were to originate at a later time, it would survive only if it could "eat" the kind of living matter already present. Thus it, too, must be of the same kind.

There is no conservation of the number of right-handed molecules. Once started, we could keep increasing the number of right-handed molecules. So the presumption is, then, that the phenomena in the case of life do not show a lack of symmetry in physical laws, but do show, on the contrary, the universal nature and the commonness of ultimate origin of all creatures on earth, in the sense described above.

52–5 Polar and axial vectors

Now we go further. We observe that in physics there are a lot of other places where we have "right" and "left" hand rules. As a matter of fact, when we learned about vector analysis we learned about the right-hand rules we have to use in order to get the angular momentum, torque, magnetic field, and so on, to come out right. The force on a charge moving in a magnetic field, for example, is $\mathbf{F} = q\mathbf{v} \times \mathbf{B}$. In a given situation, in which we know $\mathbf{F}$, $\mathbf{v}$, and $\mathbf{B}$, isn't that equation enough to define right-handedness? As a matter of fact, if we go back and look at where the vectors came from, we know that the "right-hand rule" was merely a convention; it was a trick. The original quantities, like the angular momenta and the angular velocities, and things of this kind, were not really vectors at all! They are all

somehow associated with a certain plane, and it is just because there are three dimensions in space that we can associate the quantity with a direction perpendicular to that plane. Of the two possible directions, we chose the "right-hand" direction.

So if the laws of physics are symmetrical, we should find that if some demon were to sneak into all the physics laboratories and replace the word "right" for "left" in every book in which "right-hand rules" are given, and instead we were to use all "left-hand rules," uniformly, then it should make no difference whatever in the physical laws.

Let us give an illustration. There are two kinds of vectors. There are "honest" vectors, for example a step $\Delta\mathbf{r}$ in space. If in our apparatus there is a piece here and something else there, then in a mirror apparatus there will be the image piece and the image something else, and if we draw a vector from the "piece" to the "something else," one vector is the mirror image of the other (Fig. 52–2). The vector arrow changes its head, just as the whole space turns inside out; such a vector we call a *polar vector.*

But the other kind of vector, which has to do with rotations, is of a different nature. For example, suppose that in three dimensions something is rotating as shown in Fig. 52–3. Then if we look at it in a mirror, it will be rotating as indicated, namely, as the mirror image of the original rotation. Now we have agreed to represent the mirror rotation by the same rule, it is a "vector" which, on reflection, does *not* change about as the polar vector does, but is reversed relative to the polar vectors and to the geometry of the space; such a vector is called an *axial vector.*

Now if the law of reflection symmetry is right in physics, then it must be true that the equations must be so designed that if we change the sign of each axial vector and each cross-product of vectors, which would be what corresponds to reflection, nothing will happen. For instance, when we write a formula which says that the angular momentum is $\mathbf{L} = \mathbf{r} \times \mathbf{p}$, that equation is all right, because if we change to a left-hand coordinate system, we change the sign of $\mathbf{L}$, but $\mathbf{p}$ and $\mathbf{r}$ do not change; the cross-product sign is changed, since we must change from a right-hand rule to a left-hand rule. As another example, we know that the force on a charge moving in a magnetic field is $\mathbf{F} = q\mathbf{v} \times \mathbf{B}$, but if we change from a right- to a left-handed system, since $\mathbf{F}$ and $\mathbf{v}$ are known to be polar vectors the sign change required by the cross-product must be cancelled by a sign change in $\mathbf{B}$, which means that $\mathbf{B}$ must be an axial vector. In other words, if we make such a reflection, $\mathbf{B}$ must go to $-\mathbf{B}$. So if we change our coordinates from right to left, we must also change the poles of magnets from north to south.

Let us see how that works in an example. Suppose that we have two magnets, as in Fig. 52–4. One is a magnet with the coils going around a certain way, and with current in a given direction. The other magnet looks like the reflection of the first magnet in a mirror—the coil will wind the other way, everything that happens inside the coil is exactly reversed, and the current goes as shown. Now, from the laws for the production of magnetic fields, which we do not know yet officially, but which we most likely learned in high school, it turns out that the magnetic field is as shown in the figure. In one case the pole is a south magnetic pole, while in the other magnet the current is going the other way and the magnetic field is reversed—it is a north magnetic pole. So we see that when we go from right to left we must indeed change from north to south!

Never mind changing north to south; these too are mere conventions. Let us talk about *phenomena.* Suppose, now, that we have an electron moving through one field, going into the page. Then, if we use the formula for the force, $\mathbf{v} \times \mathbf{B}$ (remember the charge is minus), we find that the electron will deviate in the indicated direction according to the physical law. So the phenomenon is that we have a coil with a current going in a specified sense and an electron curves in a certain way—that is the physics—never mind how we label everything.

Now let us do the same experiment with a mirror: we send an electron through in a corresponding direction and now the force is reversed, if we calculate it from the same rule, and that is very good because the corresponding *motions* are then mirror images!



Fig. 52–2. A step in space and its mirror image.

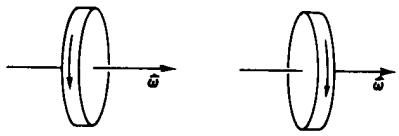


Fig. 52–3. A rotating wheel and its mirror image. Note that the angular velocity "vector" is not reversed in direction.

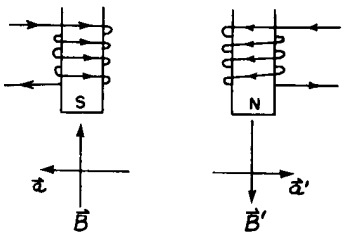


Fig. 52–4. A magnet and its mirror image.

52–6 Which hand is right?

So the fact of the matter is that in studying any phenomenon there are always two right-hand rules, or an even number of them, and the net result is that the phenomena always look symmetrical. In short, therefore, we cannot tell right from left if we also are not able to tell north from south. However, it may seem that we *can* tell the north pole of a magnet. The north pole of a compass needle, for example, is one that points to the north. But of course that is again a local property that has to do with geography of the earth; that is just like talking about in which direction is Chicago, so it does not count. If we have seen compass needles, we may have noticed that the north-seeking pole is a sort of bluish color. But that is just due to the man who painted the magnet. These are all local, conventional criteria.

However, if a magnet were to have the property that if we looked at it closely enough we would see small hairs growing on its north pole but not on its south pole, if that were the general rule, or if there were *any* unique way to distinguish the north from the south pole of a magnet, then we could tell which of the two cases we actually had, and *that would be the end of the law of reflection symmetry.*

To illustrate the whole problem still more clearly, imagine that we were talking to a Martian, or someone very far away, by telephone. We are not allowed to send him any actual samples to inspect; for instance, if we could send light, we could send him right-hand circularly polarized light and say, "That is right-hand light—just watch the way it is going." But we cannot *give* him anything, we can only talk to him. He is far away, or in some strange location, and he cannot see anything we can see. For instance, we cannot say, "Look at Ursa major; now see how those stars are arranged. What we mean by 'right' is . . ." We are only allowed to telephone him.

Now we want to tell him all about us. Of course, first we start defining numbers, and say, "Tick, tick, *two*, tick, tick, tick, *three* . . . ," so that gradually he can understand a couple of words, and so on. After a while we may become very familiar with this fellow, and he says, "What do you guys look like?" We start to describe ourselves, and say, "Well, we are six feet tall." He says, "Wait a minute, what is six feet?" Is it possible to tell him what six feet is? Certainly! We say, "You know about the diameter of hydrogen atoms—we are 17,000,000,000 hydrogen atoms high!" That is possible because physical laws are not variant under change of scale, and therefore we *can* define an absolute length. And so we define the size of the body, and tell him what the general shape is—it has prongs with five bumps sticking out on the ends, and so on, and he follows us along, and we finish describing how we look on the outside, presumably without encountering any particular difficulties. He is even making a model of us as we go along. He says, "My, you are certainly very handsome fellows; now what is on the inside?" So we start to describe the various organs on the inside, and we come to the heart, and we carefully describe the shape of it, and say, "Now put the heart on the left side." He says, "Duhhh—the left side?" Now our problem is to describe to him which side the heart goes on without his ever seeing anything that we see, and without our ever sending any sample to him of what we mean by "right"—no standard right-handed object. Can we do it?

52–7 Parity is not conserved!

It turns out that the laws of gravitation, the laws of electricity and magnetism, nuclear forces, all satisfy the principle of reflection symmetry, so these laws, or anything derived from them, cannot be used. But associated with the many particles that are found in nature there is a phenomenon called *beta decay*, or *weak decay*. One of the examples of weak decay, in connection with a particle discovered in about 1954, posed a strange puzzle. There was a certain charged particle which disintegrated into three π-mesons, as shown schematically in Fig. 52–5. This particle was called, for a while, a τ-meson. Now in Fig. 52–5 we also see another particle which disintegrates into *two* mesons; one must be neutral, from the con-

servation of charge. This particle was called a θ-meson. So on the one hand we have a particle called a τ, which disintegrates into three π-mesons, and a θ, which disintegrates into two π-mesons. Now it was soon discovered that the τ and the θ are almost equal in mass; in fact, within the experimental error, they are equal. Next, the length of time it took for them to disintegrate into three π's and two π's was found to be almost exactly the same; they live the same length of time. Next, whenever they were made, they were made in the same proportions, say, 14 percent τ's to 86 percent θ's.

Anyone in his right mind realizes immediately that they must be the same particle, that we merely produce an object which has two different ways of disintegrating—not two different particles. This object that can disintegrate in two different ways has, therefore, the same lifetime and the same production ratio (because this is simply the ratio of the odds with which it disintegrates into these two kinds).

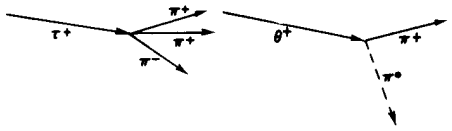


Fig. 52–5. A schematic diagram of the disintegration of a τ^+ and a θ^+ particle.

However, it was possible to prove (and we cannot here explain at all *how*), from the principle of reflection symmetry in quantum mechanics, that it was *impossible* to have these both come from the same particle—the same particle *could not* disintegrate in both of these ways. The conservation law corresponding to the principle of reflection symmetry is something which has no classical analog, and so this kind of quantum-mechanical conservation was called the *conservation of parity*. So, it was a result of the conservation of parity or, more precisely, from the symmetry of the quantum-mechanical equations of the weak decays under reflection, that the same particle could not go into both, so it must be some kind of coincidence of masses, lifetimes, and so on. But the more it was studied, the more remarkable the coincidence, and the suspicion gradually grew that possibly the deep law of the reflection symmetry of nature may be false.

As a result of this apparent failure, the physicists Lee and Yang suggested that other experiments be done in related decays to try to test whether the law was correct in other cases. The first such experiment was carried out by Miss Wu from Columbia, and was done as follows. Using a very strong magnet at a very low temperature, it turns out that a certain isotope of cobalt, which disintegrates by emitting an electron, is magnetic, and if the temperature is low enough that the thermal oscillations do not jiggle the atomic magnets about too much, they line up in the magnetic field. So the cobalt atoms will all line up in this strong field. They then disintegrate, emitting an electron, and it was discovered that when the atoms were lined up in a field whose **B** vector points upward, most of the electrons were emitted in a downward direction.

If one is not really "hep" to the world, such a remark does not sound like anything of significance, but if one appreciates the problems and interesting things in the world, then he sees that it is a most dramatic discovery: When we put cobalt atoms in an extremely strong magnetic field, more disintegration electrons go down than up. Therefore if we were to put it in a corresponding experiment in a "mirror," in which the cobalt atoms would be lined up in the opposite direction, they would spit their electrons *up*, not *down;* the action is *unsymmetrical. The magnet has grown hairs!* The south pole of a magnet is of such a kind that the electrons in a β-disintegration tend to go away from it; that distinguishes, in a physical way, the north pole from the south pole.

After this, a lot of other experiments were done: the disintegration of the π into μ and ν; μ into an electron and two neutrinos; nowadays, the Λ into proton and π; disintegration of Σ's; and many other disintegrations. In fact, in almost all cases where it could be expected, all have been found *not* to obey reflection symmetry! Fundamentally, the law of reflection symmetry, at this level in physics, is incorrect.

In short, we can tell a Martian where to put the heart: we say, "Listen, build yourself a magnet, and put the coils in, and put the current on, and then take some cobalt and lower the temperature. Arrange the experiment so the electrons go from the foot to the head, then the direction in which the current goes through the coils is the direction that goes in on what we call the right and comes out on the left." So it is possible to define right and left, now, by doing an experiment of this kind.

There are a lot of other features that were predicted. For example, it turns out that the spin, the angular momentum, of the cobalt nucleus before disintegration is 5 units of $\hbar$, and after disintegration it is 4 units. The electron carries spin angular momentum, and there is also a neutrino involved. It is easy to see from this that the electron must carry its spin angular momentum aligned along its direction of motion, the neutrino likewise. So it looks as though the electron is spinning to the left, and that was also checked. In fact, it was checked right here at Caltech by Boehm and Wapstra, that the electrons spin mostly to the left. (There were some other experiments that gave the opposite answer, but they were wrong!)

The next problem, of course, was to find the law of the failure of parity conservation. What is the rule that tells us how strong the failure is going to be? The rule is this: it occurs only in these very slow reactions, called weak decays, and when it occurs, the rule is that the particles which carry spin, like the electron, neutrino, and so on, come out with a spin tending to the left. That is a lopsided rule; it connects a polar vector velocity and an axial vector angular momentum, and says that the angular momentum is more likely to be opposite to the velocity than along it.

Now that is the rule, but today we do not really understand the whys and wherefores of it. *Why* is this the right rule, what is the fundamental reason for it, and how is it connected to anything else? At the moment we have been so shocked by the fact that this thing is unsymmetrical that we have not been able to recover enough to understand what it means with regard to all the other rules. However, the subject is interesting, modern, and still unsolved, so it seems appropriate that we discuss some of the questions associated with it.

52–8 Antimatter

The first thing to do when one of the symmetries is lost is to immediately go back over the list of known or assumed symmetries and ask whether any of the others are lost. Now we did not mention one operation on our list, which must necessarily be questioned, and that is the relation between matter and antimatter. Dirac predicted that in addition to electrons there must be another particle, called the positron (discovered at Caltech by Anderson), that is necessarily related to the electron. All the properties of these two particles obey certain rules of correspondence: the energies are equal; the masses are equal; the charges are reversed; but, more important than anything, the two of them, when they come together, can annihilate each other and liberate their entire mass in the form of energy, say γ-rays. The positron is called an *antiparticle* to the electron, and these are the characteristics of a particle and its antiparticle. It was clear from Dirac's argument that all the rest of the particles in the world should also have corresponding antiparticles. For instance, for the proton there should be an antiproton, which is now symbolized by a $\bar{p}$. The $\bar{p}$ would have a negative electrical charge and the same mass as a proton, and so on. The most important feature, however, is that a proton and an antiproton coming together can annihilate each other. The reason we emphasize this is that people do not understand it when we say there is a neutron and also an antineutron, because they say, "A neutron is neutral, so how *can* it have the opposite charge?" The rule of the "anti" is not just that it has the opposite charge, it has a certain set of properties, the whole lot of which are opposite. The antineutron is distinguished from the neutron in this way: if we bring two neutrons together, they just stay as two neutrons, but if we bring a neutron and an antineutron together, they annihilate each other with a great explosion of energy being liberated, with various π-mesons, γ-rays, and whatnot.

Now if we have antineutrons, antiprotons, and antielectrons, we can make antiatoms, in principle. They have not been made yet, but it is possible in principle. For instance, a hydrogen atom has a proton in the center with an electron going around outside. Now imagine that somewhere we can make an antiproton with a positron going around, would it go around? Well, first of all, the antiproton is electrically negative and the antielectron is electrically positive, so they attract each other in a corresponding manner—the masses are all the same; everything is the

same. It is one of the principles of the symmetry of physics, the equations seem to show, that if a clock, say, were made of matter on one hand, and then we made the same clock of antimatter, it would run in this way. (Of course, if we put the clocks together, they would annihilate each other, but that is different.)

An immediate question then arises. We can build, out of matter, two clocks, one which is "left-hand" and one which is "right-hand." For example, we could build a clock which is not built in a simple way, but has cobalt and magnets and electron detectors which detect the presence of β-decay electrons and count them. Each time one is counted, the second hand moves over. Then the mirror clock, receiving fewer electrons, will not run at the same rate. So evidently we can make two clocks such that the left-hand clock does not agree with the right-hand one. Let us make, out of matter, a clock which we call the standard or right-hand clock. Now let us make, also out of matter, a clock which we call the left-hand clock. We have just discovered that, in general, these two will *not* run the same way; prior to that famous physical discovery, it was thought that they would. Now it was also supposed that matter and antimatter were equivalent. That is, if we made an antimatter clock, right-hand, the same shape, then it would run the same as the right-hand matter clock, and if we made the same clock to the left it would run the same. In other words, in the beginning it was believed that *all four* of these clocks were the same; now of course we know that the right-hand and left-hand matter are not the same. Presumably, therefore, the right-handed antimatter and the left-handed antimatter are not the same.

So the obvious question is, which goes with which, if either? In other words, does the right-handed matter behave the same way as the right-handed antimatter? Or does the right-handed matter behave the same as the left-handed antimatter? β-decay experiments, using positron decay instead of electron decay, indicate that this is the interconnection: matter to the "right" works the same way as antimatter to the "left."

Therefore, at long last, it is really true that right and left symmetry is still maintained! If we made a left-hand clock, but made it out of the other kind of matter, antimatter instead of matter, it would run in the same way. So what has happened is that instead of having two independent rules in our list of symmetries, two of these rules go together to make a new rule, which says that matter to the right is symmetrical with antimatter to the left.

So if our Martian is made of antimatter and we give him instructions to make this "right" handed model like us, it will, of course, come out the other way around. What would happen when, after much conversation back and forth, we each have taught the other to make space ships and we meet halfway in empty space? We have instructed each other on our traditions, and so forth, and the two of us come rushing out to shake hands. Well, if he puts out his left hand, watch out!

52–9 Broken symmetries

The next question is, what can we make out of laws which are *nearly* symmetrical? The marvelous thing about it all is that for such a wide range of important, strong phenomena—nuclear forces, electrical phenomena, and even weak ones like gravitation—over a tremendous range of physics, all the laws for these seem to be symmetrical. On the other hand, this little extra piece says, "No, the laws are not symmetrical!" How is it that nature can be almost symmetrical, but not perfectly symmetrical? What shall we make of this? First, do we have any other examples? The answer is, we do, in fact, have a few other examples. For instance, the nuclear part of the force between proton and proton, between neutron and neutron, and between neutron and proton, is all exactly the same—there is a symmetry for nuclear forces, a new one, that we can interchange neutron and proton—but it evidently is not a general symmetry, for the electrical repulsion between two protons at a distance does not exist for neutrons. So it is not generally true that we can *always* replace a proton with a neutron, but only to a good approximation. Why *good?* Because the nuclear forces are much stronger than the

electrical forces. So this is an "almost" symmetry also. So we do have examples in other things.

We have, in our minds, a tendency to accept symmetry as some kind of perfection. In fact it is like the old idea of the Greeks that circles were perfect, and it was rather horrible to believe that the planetary orbits were not circles, but only nearly circles. The difference between being a circle and being nearly a circle is not a small difference, it is a fundamental change so far as the mind is concerned. There is a sign of perfection and symmetry in a circle that is not there the moment the circle is slightly off—that is the end of it—it is no longer symmetrical. Then the question is why it is only *nearly* a circle—that is a much more difficult question. The actual motion of the planets, in general, should be ellipses, but during the ages, because of tidal forces, and so on, they have been made almost symmetrical. Now the question is whether we have a similar problem here. The problem from the point of view of the circles is if they were perfect circles there would be nothing to explain, that is clearly simple. But since they are only nearly circles, there is a lot to explain, and the result turned out to be a big dynamical problem, and now our problem is to explain why they are nearly symmetrical by looking at tidal forces and so on.

So our problem is to explain where symmetry comes from. Why is nature so nearly symmetrical? No one has any idea why. The only thing we might suggest is something like this: There is a gate in Japan, a gate in Neiko, which is sometimes called by the Japanese the most beautiful gate in all Japan; it was built in a time when there was great influence from Chinese art. This gate is very elaborate, with lots of gables and beautiful carving and lots of columns and dragon heads and princes carved into the pillars, and so on. But when one looks closely he sees that in the elaborate and complex design along one of the pillars, one of the small design elements is carved upside down; otherwise the thing is completely symmetrical. If one asks why this is, the story is that it was carved upside down so that the gods will not be jealous of the perfection of man. So they purposely put an error in there, so that the gods would not be jealous and get angry with human beings.

We might like to turn the idea around and think that the true explanation of the near symmetry of nature is this: that God made the laws only nearly symmetrical so that we should not be jealous of His perfection!

Index

Aberration, 27–7, 34–10
Absolute zero, 1–5
Absorption, 31–8 ff
Acceleration, 8–8 ff
  components of, 9–3
  of gravity, 9–4
Activation energy, 42–7
Adams, J. C., 7–5
Adiabatic compression, 39–5
Adiabatic expansion, 44–5
Affective future, 17–4
Air trough, 10–5
Algebra, 22–1 ff
Amplitude modulation, 48–3
Amplitude, of oscillation, 21–3
Analog computer, 25–8
Anderson, C. D., 52–10
Angle, of incidence, 26–3
  of reflection, 26–3
Angstrom (unit), 1–3
Angular frequency, 21–3, 29–2
Angular momentum, 7–7, 18–5 f, 20–1
  conservation of, 4–7, 18–6 ff, 20–5
  of rigid body, 20–8
Anomalous refraction, 33–9 f
Antimatter 52–10 f
Antiparticle, 2–8
Aristotle, 5–1
Atom, 1–2
  metastable, 42–10
Atomic clock, 5–5
Atomic hypothesis, 1–2
Atomic particles, 2–9 f
Atomic processes, 1–5 f
Attenuation, 31–8
Avogadro, A., 39–2
Avogadro's number, 41–10
Axial vector, 52–6 f

Becquerel, A. H., 28–3
Birefringence, 33–3 ff
Blackbody radiation, 41–5 f
Boehm, 52–10
Bohr, N., 42–9
Bohr radius, 38–6
Boltzmann, 41–2
Boltzmann's law, 40–2 f
Born, M., 37–1, 38–9
Boyle's law, 40–8
Bremsstrahlung, 34–6 f
Brewster's angle, 33–6
Briggs, H., 22–6
Brown, R., 41–1
Brownian motion, 1–8, 6–5, 41–1 ff

Capacitance, 23–5
Capacitor, 14–9, 23–5
Capillary action, 51–8
Carnot, S., 4–2, 44–2 ff
Carnot cycle, 44–5 f, 45–2
Carrier signal, 48–3
Catalyst, 42–8
Cavendish, H., 7–9
Cavendish's experiment, 7–9
Center of mass, 18–1 f, 19–1 ff
Centrifugal force, 7–5, 12–11
Çerenkov, P. A., 51–2
Çerenkov radiation, 51–2
Charge, conservation of, 4–7
  on electron, 12–7
Chemical energy, 4–2
Chemical kinetics, 42–7 f
Chemical reaction, 1–6 ff
Chromaticity, 35–6 f
Circular motion, 21–4
Clausius, R., 44–2, 44–3
Clausius-Clapeyron equation, 45–6 ff
Coefficient, of friction, 12–4
  gravitational, 7–9
Collision, 16–6
  elastic, 10–7
Color vision, 35–1 ff
  physiochemistry of, 35–9 f
Complex impedance, 23–7
Complex numbers, 22–7 ff, 23–1 ff
Compound eye, 36–6 ff
Compression, adiabatic, 39–5
  isothermal, 44–5
Cones, 35–1
Conservation, of angular momentum, 4–7, 18–6 ff, 20–5
  of charge, 4–7
  of energy, 3–2, 4–1 ff
  of linear momentum, 4–7, 10–1 ff
Contraction hypothesis, 15–3
Copernicus, 7–1
Coriolis force, 19–8 f
Cornea, 35–1
Coulomb's law, 28–2
Cross section, for scattering, 32–7
Crystal diffraction, 38–4 f

Dedekind, R., 22–4
Degrees of freedom, 25–2, 39–12
Density, 1–4
Derivative, 8–5 ff
  partial, 14–9
Dicke, R. H., 7–11
Differential calculus, 8–4
Diffraction, 30–1 ff
  by screen, 31–10 f
Diffraction grating, 29–5, 30–3 ff
Diffusion, 43–1 ff
Dipole moment, 12–6
Dipole radiator, 28–5 f, 29–3 ff
Dirac, P., 52–10
Dirac equation, 20–6
Dispersion, 31–6 ff
Distance, 5–5 ff
Distance measurement, color brightness, 5–6
  triangulation, 5–6
Doppler effect, 17–8, 23–9, 34–7 f, 38–6
Double stars, 7–6

Dynamics, 7–2 f, 9–1 ff
  relativistic, 15–9 f

Efficiency, of ideal engine, 44–7 f
Einstein, A., 2–6, 7–11, 12–12, 15–1, 16–1, 41–8, 42–8, 42–9
Elastic collision, 10–7
Elastic energy, 4–2, 4–6
Electrical energy, 4–2
Electric field, 2–4, 12–7 f
Electromagnetic energy, 29–2
Electromagnetic field, 2–2, 2–5, 10–9
Electromagnetic radiation, 26–1, 28–1 ff
Electromagnetic waves, cosmic rays, 2–5
  gamma rays, 2–5
  infrared, 2–5, 23–8, 26–1
  light, 2–5
  ultraviolet, 2–5, 26–1
  x-rays, 2–5, 26–1
Electron, 2–4, 37–1, 37–4 ff
  charge on, 12–7
  radius of, classical, 32–4
Electron cloud, 6–11
Electron-ray tube, 12–9
Electron volt (unit), 34–4
Ellipse, 7–1
Energy, chemical, 4–2
  conservation of, 3–2, 4–1 ff
  elastic, 4–2, 4–6
  electrical, 4–2
  electromagnetic, 29–2
  gravitational, 4–2 ff
  heat, 4–2, 4–6, 10–7, 10–8
  kinetic, 1–7, 4–2, 4–5 f, 39–4
  mass, 4–2, 4–7
  nuclear, 4–2
  potential, 4–4, 13–1 ff, 14–1 ff
  radiant, 4–2
  relativistic, 16–1 ff
Energy levels, 38–7 f
Energy theorem, 50–7 f
Enthalpy, 45–5
Entropy, 44–10 ff, 46–7 ff
Eötvös, 7–11
Equilibrium, 1–6
Euclid, 5–6
Euclidean geometry, 12–3
Evaporation, 1–5 f
  of a liquid, 40–3 f, 42–1 ff
Expansion, adiabatic, 44–5
  isothermal, 44–5
Exponential atmosphere, 40–1 f
Eye, compound, 36–6 ff
  human, 35–1 f, 36–3 ff

Farad (unit), 25–7
Fermat, P., 26–3
Fermi (unit), 5–10
Fermi, E., 5–10
Fields, 2–2, 2–4, 2–5, 10–9, 12–7 ff, 13–8 f, 14–7 ff
  superposition of, 12–9

Focal length, 27–1 ff
Focus, 26–5
Force, centrifugal, 7–5, 12–11
  components of, 9–3
  conservative, 14–3 ff
  Coriolis, 19–8 f
  electrical, 2–3 ff
  gravitational, 2–3
  molecular, 1–3, 12–6 f
  moment of, 18–5
  nonconservative, 14–6 f
  nuclear, 12–12
  pseudo, 12–10 ff
Fourier, J., 50–2 f
Fourier analysis, 50–2 ff
Fourier transform, 25–4
Four-vectors, 15–8 f, 17–5 ff
Fovea, 35–1
Frank, I., 51–2
Frequency, angular, 21–3, 29–2
  of oscillation, 2–5
Fresnel's reflection formulas, 33–8
Friction, 10–5, 12–3 ff
  coefficient of, 12–4

Galileo, 5–1, 7–2, 9–1, 52–3
Galilean relativity, 10–3
Galilean transformation, 12–11
Gauss (unit), 34–4
Gell-Mann, M., 2–9
Geometrical optics, 26–1, 27–1 f
Gravitation, 2–3, 7–1 ff, 12–2
Gravitational acceleration, 9–4
Gravitational coefficient, 7–9
Gravitational energy, 4–2 ff
Gravitational field, 12–8 ff, 13–8 f
Gravity, 13–3 ff
  acceleration of, 9–4
Green's function, 25–4
Gyroscope, 20–5 ff

Harmonic motion, 21–4, 23–1 ff
Harmonic oscillator, 10–1, 21–1 ff
  forced, 21–5 f, 23–3 ff
Harmonics, 50–1 ff
Heat, 1–3, 13–3
Heat energy, 4–2, 4–6, 10–7, 10–8
Heat engines, 44–1 ff
Heisenberg, W., 6–10, 37–1, 37–9, 37–11, 37–12, 38–9
Helmholtz, H., 35–7
Henry (unit), 25–7
Hooke's law, 12–6
Huygens, C., 15–2, 26–2
Hypocycloid, 34–3

Ideal gas law, 39–10 ff
Impedance, 25–8 f
  complex, 23–7
Incidence, angle of, 26–3
Inclined plane, 4–4
Index, of refraction, 31–1 ff
Inductance, 23–6
Inductor, 23–6
Inertia, 2–3, 7–11
  moment of, 18–7, 19–5 ff
  principle of, 9–1
Infrared radiation, 23–8, 26–1
Integral, 8–7 f
Interference, 28–6, 29–1 ff
Interfering waves, 37–4
Interferometer, 15–5
Ion, 1–6
Ionic conductivity, 43–6 f
Ionization energy, 42–5
Isothermal atmosphere, 40–2
Isothermal compression, 44–5
Isothermal expansion, 44–5
Isotopes, 3–4 ff

Jeans, J., 40–9, 41–6 f
Johnson noise, 41–2, 41–8
Joule (unit), 13–3
Joule heating, 24–2

Kepler, J., 7–1
Kepler's laws, 7–1 f, 9–1, 18–6
Kerr cell, 33–5
Kinetic energy, 1–7, 4–2, 4–5 f, 39–4
  rotational, 19–7 ff
Kinetic theory, 42–1 ff
  of gases, 39–1 ff
Kirchhoff's laws, 25–9

Laplace, P., 47–7
Laser, 32–6, 42–10
Least time, principle of, 26–3 ff, 26–8
Leibnitz, G. W., 8–4
Lens formula, 27–6
Leverrier, U., 7–5
Light, momentum of, 34–10 f
  polarized, 32–9
  scattering of, 32–5 ff
  speed of, 15–1
Light waves, 48–1
Linear momentum, conservation of, 4–7, 10–1 ff
Linear systems, 25–1 ff
Logarithms, 22–4
Lorentz, H. A., 15–3
Lorentz contraction, 15–7
Lorentz transformation, 15–3, 17–1, 34–8, 52–2

Magnetic field, 12–9 f
Magnetic induction, 12–10
Magnetism, 2–4
Magnification, 27–5
Maser, 42–10
Mass, 9–1, 15–1
  center of, 18–1 f, 19–1 ff
  relativistic, 16–6 ff
Mass energy, 4–2, 4–7
Mass-energy equivalence, 15-10 f
Maxwell, J. C., 6–1, 6–9, 28–1, 40–8, 41–7, 46–5
Maxwell's equations, 15–2, 25–3, 47–7
Mayer, J. R., 3–2
Mean free path, 43–3 f
Mean square distance, 6–5, 41–9
Mendeléev, 2–9
Metastable atom, 42–10
Meter (unit), 5–10
Mev (unit), 2–9
Michelson-Morley experiment, 15–3 ff
Miller, W. C., 35–2
Minkowski, 17–8
Modes, 49–1 ff
Mössbauer, R., 23–9
Mole (unit), 39–10
Molecular attraction, 1–3, 12–6 f
Molecular diffusion, 43–7 ff
Molecular motion, 41–1
Molecule, 1–3
Moment, dipole, 12–6
  of force, 18–5
  of inertia, 18–7, 19–5 ff
Momentum, 9–1 f, 38–2 ff
  angular, 7–7, 18–5 ff, 20–1, 20–5
  of light, 34–10 f
  linear, 4–7, 10–1 ff
  relativistic, 10–8 f, 16–1 ff
Monatomic gas, 39–5
Motion, 5–1, 8–1 ff
  circular, 21–4
  constrained, 14–3
  harmonic, 21–4, 23–1 ff
  parabolic, 8–10
  planetary, 7–1 ff, 9–6 f, 13–5
Music, 50–1

Nernst heat theorem, 44–11
Neutrons, 2–4
Newton, I., 8–4, 15–1, 37–1
Newton·meter (unit), 13–3
Newton's laws, 2–6, 7–3 ff, 7–11, 9–1 ff, 10–1 ff, 11–7 f, 12–1, 39–2, 41–1, 46–1
Nishijima, 2–9
Nodes, 49–2
Noise, 50–1
Nuclear cross section, 5–9
Nuclear energy, 4–2
Nuclear forces, 12–12
Nucleus, 2–4, 2–8 ff
Numerical analysis, 9–6
Nutation, 20–7

Ohm (unit), 25–7
Ohm's law, 25–7, 43–7
Optic axis, 33–3
Optic nerve, 35–2
Optics, 26–1 ff
  geometrical, 26–1, 27–1 ff
Oscillation, amplitude of, 21–3
  damped, 24–3 f
  frequency of, 2–5
  period of, 21–3
  periodic, 9–4
  phase of, 21–3
Oscillator, 5–2
  harmonic, 10–1, 21–1, 21–5 f, 23–3 ff

Pappus, theorem of, 19–4
Parabolic antenna, 30–6 f
Parabolic motion, 8–10
Parallel-axis theorem, 19–6
Parallel-plate capacitor, 14–9
Paraxial rays, 27–2
Partial derivative, 14–9
Pascal's triangle, 6–4
Pendulum, 49–6 f
Pendulum clock, 5–2
Period, of oscillation, 21–3
Periodic time, 5–1 f
Perpetual motion, 46–2
Phase, of oscillation, 21–3
Phase shift, 21–3
Phase velocity, 48–6
Photon, 2–7, 26–1, 37–8
Physiochemistry, of color vision, 35–9 f
Planck, M., 41–6, 42–8, 42–9
Planck's constant, 5–10, 6–10, 17–8, 37–11
Planetary motion, 7–1 ff, 9–6 f, 13–5

Poincaré, H., 15–3, 15–5, 16–1
Polarization, 33–1 ff
Polarized light, 32–9
Potential energy, 4–4, 13–1 ff, 14–1 ff
Power, 13–2
Pressure, 1–3
Probability, 6–1 ff
Probability density, 6–8 f
Probability distribution, 6–7 ff
Proton, 2–4
Pseudo force, 12–10 ff
Ptolemy, 26–2
Purkinje effect, 35–2
Pythagoras, 50–1

Quantum electrodynamics, 2–7, 28–3
Quantum mechanics, 2–2, 2–6 ff, 6–10, 10–9, 37–1 ff, 38–1 ff

Radiant energy, 4–2
Radiation, infrared, 23–8, 26–1
  relativistic effects, 35–1 ff
  synchrotron, 34–3 ff, 34–6
  ultraviolet, 26–1
Radiation damping, 32–3 f
Radiation resistance, 32–1 ff
Radioactive clock, 5–3 ff
Radius, of electron, 32–4
Ramsey, N., 5–5
Random walk, 6–5 ff, 41–8 ff
Ratchet and pawl machine, 46–1 ff
Rayleigh's criterion, 30–6
Rayleigh's law, 41–6
Reciprocity principle, 30–7
Rectification, 50–9
Reflection, 26–2 f
  angle of, 26–3
Refraction, 26–2 f
  anomalous, 33–9 f
  index of, 31–1 ff
Relativistic dynamics, 15–9 f
Relativistic energy, 16–1 ff
Relativistic mass, 16–1 ff
Relativistic momentum, 10–8 f, 16–1 ff
Relativity, special theory of, 15–1 ff
  Galilean, 10–3
  theory of, 7–11, 17–1
Resistance, 23–5
Resistor, 23–5
Resolving power, 27–7 f, 30–5 f
Resonance, 23–1 ff
  electrical, 23–5 ff
  in nature, 23–7 ff
Resonance interaction, 2–9
Retarded time, 28–2
Retina, 35–1
Rigid body, 18–1
  angular momentum of, 20–8
  rotation of, 18–2 ff
Ritz combination principle, 38–8
Rods, 35–1, 36–6
Roemer, O., 7–5
Root-mean-square distance, 6–6
Rotation, of axes, 11–3 f
  plane, 18–1
  of a rigid body, 18–2 ff
  in space, 20–1 ff
  in two dimensions, 18–1 ff
Rushton, 35–9
Rydberg (unit), 38–6

Scalar, 11–5
Scattering, of light, 32–5 ff
Schrödinger, E., 35–6, 37–1, 38–9
Scientific method, 2–1 f
Screw jack, 4–5
Second (unit), 5–5
Seismograph, 51–5
Shannon, C., 44–2
Shear wave, 51–4
Side bands, 48–4 f
Simultaneity, 15–7 f
Sinusoidal waves, 29–2 f
Smoluchowski, 41–8
Smooth muscle, 14–2
Snell, W., 26–3
Snell's law, 26–3, 31–2
Sound, 2–3, 47–1 ff, 50–1
  speed of, 47–7 f
Space, 8–2
Space-time, 2–6, 17–1 ff
Special theory of relativity, 15–1 ff
Specific heat, 40–7 f, 45–2
Speed, 8–2 ff, 9–2
  of light, 15–1
  of sound, 47–7 f
Spontaneous emission, 42–9
Standard deviation, 6–9
Statistical fluctuations, 6–3 ff
Statistical mechanics, 3–1, 40–1 ff
Stevinus, S., 4–5
"Strangeness" number, 2–9
Striated muscle, 14–2
Superposition, of fields, 12–9
  principle of, 25–2 ff
Symmetry, 1–4, 11–1 ff
  of physical laws, 16–3, 52–1 ff
Synchrotron, 2–5, 15–9, 34–3 ff, 34–6

Tamm, I., 51–2
Temperature, 39–6 ff
Thermal conductivity, of a gas, 43–9 f
Thermal equilibrium, 41–3 ff
Thermal ionization, 42–5 ff
Thermodynamics, 39–2, 45–1 ff
  laws of, 44–1 ff
Thompson scattering cross section, 32–8
Three-body problem, 10–1
Tides, 7–4 f
Time, 2–3, 5–1 ff, 8–1, 8–2
  retarded, 28–2
  standard of, 5–5
  transformation of, 15–5 ff
Torque, 18–4, 20–1 ff
Transformation, Fourier, 25–4
  Galilean, 12–11
  linear, 11–6
  Lorentz, 15–3, 17–1, 34–8, 52–2
  of time, 15–5 ff
  of velocity, 16–4 ff
Transient, 24–1 ff
  electrical, 24–5 f
Transient response, 21–6
Translation, of axes, 11–1 ff
Twin paradox, 16–3 f
Tycho Brahe, 7–1

Ultraviolet radiation, 26–1
Uncertainty principle, 2–6, 6–10 f, 37–9, 37–11, 38–8 f
Unit cell, 38–5
Unit vector, 11–10

Vector, 11–5 ff
Vector algebra, 11–6 f
Vector analysis, 11–5, 52–2
Vector product, 20–4
Velocity, 8–3, 9–2 f
  components of, 9–3
  transformation of 16–4 ff
Vinci, Leonardo da, 36–2
Virtual work, principle of, 4–5
Vision, 36–1 ff
  binocular, 36–4
  color, 35–1 ff
Visual cortex, 36–4
Visual purple, 35–9

Wapstra, 52–10
Watt (unit), 13–3
Wave, shear, 51–4
  sinusoidal, 29–2 f
Wave equation, 47–1 ff
Wavefront, 47–3
Wavelength, 19–3, 26–1
Wave number, 29–2
Waves, 51–1 ff
  light, 48–1
Weyl, H., 11–1
Work, 13–1 ff, 14–1 ff

X-rays, 2–5, 26–1

Young, 35–7
Yukawa, H., 2–8
Yustova, 35–8

Zeno, 8–3
Zero, absolute, 1–5
Zero mass, 2–10