Chapter 3 Overview: Sampled Signals PhL 3.15.03

Section 20: Image Spectra

In this section, I show that if you multiply a smooth signal y(t) by a function which is an infinite sum of delta functions spaced T_1 apart, the resulting product function w(t) is really a PAM signal sampled by delta functions and the spectrum of the product function w(t) has a main spectrum and an infinite number of image spectra. The spectrum of the smooth function y(t) is finite, but each delta function has an infinitely wide constant spectrum, just as we find with Gaussian white noise. This infinite spectral width then propagates into the product function w(t), and that is why those image spectra are there. If they don't overlap, you can isolate the main spectrum and regenerate your initial function y(t), as shown in Section 34.

Section 21: Digital Filters and their Image Spectra

Here, I show that if you try to approximate a continuous time convolution relation a = b*c with a finite digital filter, you find a tangled relation between the true (non-primed) Fourier spectra. Instead of getting $A(\omega) = B(\omega)C(\omega)$, you get $A(\omega) = B'(\omega)C(\omega)$, where $B'(\omega)$ is $B(\omega)$ plus all the image spectra of B. The filter appears to have all kinds of high-frequency pass bands! These allows high frequency aliased versions of an intended baseband signal to pass right through the filter. The filter cannot do anything about it, because it is blind to variations of the input signal between the sample points. The blindness of the filter is the reason that these image spectra are there in the ω-space relation.

Section 22: The Digital Fourier Transform

The results of this section appear in a box on page 10 (of Section 23). The idea is to define *new* expansion and projection formulas that work well with *sampled* functions of time, as opposed to the Fourier Transform which works well with *continuous* functions of time. The transform looks very much like the Fourier transform, but for the following key differences: (1) the projection has $\int dt$ replaced by $\Sigma\ \Delta t$, still an infinite sum; (2) the projection only detects values of x(t) at the sample times; (3) the DFT projection $X'(\omega)$ has all those image spectra; (4) the expansion formula integration range is $-\omega_1/2$ to $+\omega_1/2$, instead of being from $-\infty$ to $+\infty$, and it only reconstructs x(t) at the sample times. I denote this new spectral projection with a prime. A sampled-time convolution relation a = b*c, even with an infinite sum, results in a simple diagonalized form $A'(\omega) = B'(\omega)\ C'(\omega)$. As $T_1 \rightarrow \infty$, the DFT transform becomes identical with the Fourier Transform, and those image spectra move infinitely far away and vanish.

Section 23: How $X'(\omega)$ of the DFT relates to $X(\omega)$ of the FT.

It is no great surprise that $X'(\omega) = X(\omega)$ + the images, as shown in 23.1. If we analyze a finite sampled convolution relation a = b*c, we can write the diagonalized form as either (23.3) which we already saw in Section 21, or as (23.4) where everything is primed.

Section 24: The Z Transform

I show here that the Z transform is exactly the same as the DFT just written with a change in variables in the frequency domain. Instead of having a finite real-line-segment ω-space integral over dω, we wrap that into an integral around a unit circle of a new complex variable $z \equiv \exp(i\omega\Delta t)$. Although the Z transform is really the same as the DFT, it looks very different! In the projection (24.3) our former $\exp(-i\omega t_n)$ has now morphed into z^{-n}, and in the expansion our $d\omega \exp(+i\omega t_n)$ has become $dz\, z^{n-1}$ with a unit circle contour integral. We also change the scale of the projection X"(z) as shown in 24.2 to get the standard result. This scale change results in an extra factor in the diagonalized convolution, as shown page 12.

There is no question that the Z transform is easier to write, and as with the Laplace Transform and its variable s, we have the possibility of doing analytic continuation in the **complex z-plane**, and we can then do the contour by picking up pole residues, and all that good complex variable stuff, so perhaps the Z transform should be given a little more respect by me.

Secondly, in addition to this complex-plane aspect of variable z, we now really have the idea of an N-tap no-feedback filter being a **polynomial** in z^{-1} . We knew in the DFT or Fourier Transform world that a delay in time caused an exponential phase shift in ω, and here that is expressed as simply as possible, because here a delay of one sample creates a "phase factor" of z^{-1}. It is the same phase shift, but easier to write. Now when we have a no-feedback FIR digital filter $b(t_n)$ or B"(z), we can think of B"(z) as a polynomial in z^{-1} with weights. If the filter has feedback, then the equations show that B"(z) is an inverse polynomial, and then we have poles in the z-plane and we have an IIR.

We can model a derivative as shown on the top of page 13 -- which would become a real derivative in the limit. Then we can think about digital *versions* of analog filters such as an RC filter by using this thing for the derivative. I think my Lam book does a lot of this work.

25. PAM Pulse Trains

I don't think I realized it when writing this section, but a PAM pulse train as in (25.4) is really just a sampled-time convolution of the type considered in Section 21. We found there that the diagonalized convolution equation had the "tangled" result $A(\omega) = B'(\omega)C(\omega)$, where B'(ω) we later learned was the DFT transform, while A(ω) and C(ω) are regular Fourier Transforms. So this at once gives the boxed result which says that the spectrum of such a pulse train is the spectrum of the pulse times the multi-image version of the y(t) spectrum.

26. Aperture Correction

Here I consider a D/A converter (don't think of these much in Sklar!) with a rectangular partial hold window as shown page 18. The diagonalized spectrum is 26.3 which says that our image spectra series is multiplied by the sinc function. The idea is that if you sample really fast, like 8X oversampling, with zeros in between the rectangles as I drew it, then you make that sinc function nearly flat on the central main spectrum, and thus you minimize the spectral distortion. I keep forgetting that this is WHY we want to have spiky samples separated by zero space. If you still feel the need to do some correction, the filter type you use is called a "sine x over x filter".

27. The Discrete Fourier Transform

The Digital Fourier Transform described above in Section 22 was designed for use with sampled-time signals, as opposed to continuous-time signals. Recall how the Fourier Integral morphed into the Fourier Series in the case that the continuous-time signal was periodic. Here, we are going to have our Digital Fourier Transform (= Z Transform) morph into the *Discrete* Fourier transform for sampled-time signals which are periodic. The names Digital and Discrete are just labels, you have to remember what the things really are. This is the whole subject of power signals versus energy signals. For periodic signals, we need a transform that does not have an infinite projection series that diverges.

For T_1 I have suddenly changed notation unfortunately here. In earlier sections of this Chapter, T_1 was the spacing between samples. Now, however, T_1 is instead the period of our periodic function, and the spacing between samples is now $\Delta t = T_1/N$. I really should repair this somehow.

The Discrete FT is summarized in the box on page 23. I have persisted in my use of an infinite-extent Gaussian-like X_{pulse} function in writing the projection for the coefficient c_m', that is why you see an infinite sum for this projection. The corresponding expansion formula is a finite sum of these coefficients times a phasor. I show that the coefficients c_m' have a symmetry property.

This Discrete FT is really a lot more to the point if we choose the N values in an interval as our "pulse", then we have a finite sum in each direction. If you only have N points that you sample in a period, then you are only going to have N coefficients in your expansion sum. You can in fact think of this transformation as an NxN matrix transformation.

28 Alternative form of the Discrete FT.

I think DFT for most people really means the Discrete FT, not the Digital FT as I have used it. Most books rescale the DFT as I show here, it is a question of where to put factors of N.

29. Symmetry relation on the DFT coefficients.

I look again at the symmetry rule for the coefficients. For real x(t) sampled, not all N coefficients are different! In fact, only half of them are unique, give or take one. You of course are supposed to think of these coefficients c_m" as the "spectrum" of your sampled x(t). Thus, when you plot such a spectrum, there is no point in putting more than half the c_m on the x-axis, due to the mirror-image property of this spectrum due to this symmetry relation. Note that we do NOT have a similar situation in the Fourier Series! This is because Fourier Series has a continuous time variable. You cannot just shift ω by some amount and have the phasor repeat a previous value, see box on page 6 of Chapter 2.

30. Parseval's Theorem for the DFT.

This says that the sum of the squares of the sampled x(t) values in a period equals the sum of the abs-squared values of the coefficients. We have not been updating this theorem for our many transforms in this chapter.

31. Sample DFT.

I make a little sine wave with k periods fit into our time T_1. This obviously makes a continuous periodic signal. We find that in the lower half range of the c_m , only c_k is non-vanishing. In the full range, one other

coefficient is non-zero due to the symmetry relation. I ran a Maple program and made it plot the DFT spectrum for this case, and this showed the computational error on a log plot.

32. The FFT

Just a reminder that the FFT is just the Discrete Fourier Transform which you can compute very quickly when N is a power of 2.

33. Lee & Messerschmitt Notation

These guys are a little confusing, because they use the function X(..) to indicate both a Laplace transform and a Z transform, and you have to determine which one by the form of the argument. I don't like this much, but am grateful to L&M for ferreting out some errors I had in this chapter which are now fixed.

34. The Nyquist Sampling Theorm

Here I derive the little formula that lets you reconstruct signal x(t) from the samples $x(t_n)$ by weighting each one with a sinc function in time. This is the reconstruction formula that is the final statement of this theorem. Of course it only works if the image spectra do not overlap, all parts of the theorem.