Question: Is there some simple power rule which says that the sum of the energy in all the coherences is constant, so suppressing ones you don't want is helpful?

I guess you could say the energy in the final system is tr(ρH') where H' is the Ham. In our state basis and in secular approximation, H is diagonal. Thus, we have $tr(\rho H') = \rho_{ij}H'_{ji} = p_{ii}H'_{ii}$ which says only the populations affect <E>, which I of course know is right. But, the FID signal is going to be proportional to both population difference between the photon states and the transition rate. Recall this result

$$d\rho_{mn}(t)/dt = -i\, H'_{mn}\, \{\, \rho(t)_{nn} - \rho(t)_{mm}\, \}$$

which says that your coherence builds in proportion to H' which here means the RF pulse. The larger you build up the ρ_- coherence during your last RF pulse, the more FID you will have. But that last pulse only adds more ρ_{mn} to what was already there. Well, above is differential, a pulse is not so.

Are there any conservation laws that apply to each stage of the pulse sequence? Ignoring relaxation, you would think nothing happens in the gaps, things just precess with no friction lets say.

Well, I suppose the RF pulse does cause some double and single photon absorption. Classically, suppose you pulse sequence generates some $I_{1+}I_{2+}$ output at FID start. That means there is some ρ_{++} in the FID-start density matrix. So for example you have $\langle\alpha\alpha|\, \rho\, |\beta\beta\rangle \neq 0$.

Suppose ββ is the upper-most state and you end up with a large $\rho_{\beta\beta}$ because your pulse sequence caused it to build up. There will then be spontaneous emission between this state and the αα state and that is going to represent an energy *loss* from the spin system so maybe that would be something to avoid. But perhaps the system absorbed more energy in the first place.

The goal I would say is to maximize the amplitude of ρ_- at FID start, and you think of it as just a carrier of frequency information that will make peaks. I meant that in this sense. Suppose you have

$$\rho = I_{1x}(S_{\theta 1-\phi} + S_{\theta 1+\phi})/2 \;\; -I_{1z}I_{2y}\,(C_{\theta 1-\phi} - C_{\theta 1+\phi}) \qquad + \;\; (1\leftrightarrow 2)$$

as in COSY. The I_{1x} acts as the "carrier" while $(S_{\theta 1-\phi} + S_{\theta 1+\phi})/2$ acts as the "signal". A pulse causes a rotation which does not change overall amplitude, and the same seems true for a gap.

So I don't really have an answer here, but I think the question is a good one. In a general theory of pulse sequence design, you would include maximizing output power per pulse as one of your goals.

However, as we go through pulses and gaps, there is not much you can do to enhance the coefficient multiply a matrix-rep term in the final ρ. You might arrange for multiple terms to contribute.