Old Section 30 on oversampling retired 3.17.13

30. Oversampling and Decimation

Consider a standard D/A output system operating at sampling period T_1 and $\omega_1 = 2\pi/T_1$. Assume there are a few output latches before the D/A converter, and that the D/A is glitch-free. Assume there is some aperture time τ on the output of the D/A with $\tau/T_1 = D$, the aperture duty cycle, and that this is corrected by an analog post-filter with sin(x)/x correction. This has been discussed in detail in Section 26.

The output of the D/A converter is signal w(t) which is an amplitude modulated pulse train approximation to signal y(t). We know the spectrum of w(t) from (26.3),

$$W(\omega) = (1/T_1)\, X_{\mathbf{pulse}}(\omega,\tau)\, [\, Y(\omega\,) + \sum_{m \neq 0} Y(\omega - m\omega_1)]$$

where $X_{\mathbf{pulse}}(\omega,\tau) = \tau\, \mathrm{sinc}(\omega\tau/2)\, \exp(-i\omega\tau/2)$

Note that the image spectra are separated by distance ω_1, and $Y(\omega)$ is the Fourier spectrum of y(t).

Suppose now the digital clocking frequency going to the final latch before the D/A converter and the converter itself is raised by a factor of 4, and that the duty cycle of the D/A aperture is kept at D, but no other changes are made. What happens to the output signal w(t) and its spectrum $W(\omega)$?

Certainly w(t) now looks different. The old w(t) was a set of duty cycle D pulses separated by T_1 tracking the signal y(t), as shown in Figure 26.1. The new w(t) has closer pulses spaced by $T_1/4$, and they are not tracking y(t) any more because the new pulses come in groups of 4 which have the same amplitude. Only every fourth pulse correctly "tracks" y(t). This is because back in the input latch to the D/A converter we clock in the same identical data four times since we raised this latch's clock rate by a factor of 4.

The new pulse is the same as the original pulse except τ is 1/4 what it was before, so we get $X_{\mathbf{pulse}}(\omega,\tau/4)$.

What is the spectrum $W(\omega)$ of this complicated new signal? We compute it in our standard way, by superposing time-shifted pulses. We get:

$$W(\omega) = X_{\mathbf{pulse}}(\omega,\tau/4)\, [\ldots + y_0 + y_0\, e^{-i\omega(T_1/4)} + y_0\, e^{-2i\omega(T_1/4)} + y_0\, e^{-3i\omega(T_1/4)} + \ldots.\,]$$

Here we have shown the contribution of one group of four pulses of the same amplitude. The entire sum is groups of four like this. We can combine these four terms to get

$$[\, 1 + e^{-i\omega(T_1/4)} + e^{-2i\omega(T_1/4)} + e^{-3i\omega(T_1/4)}\,]\, y_0 = F(\omega)\, y_0\ .$$

Of course every group of four terms will have this same common factor $F(\omega)$, so we can factor it out of the entire sum. The sum now looks like this:

$$W(\omega) = X_{pulse}(\omega,\tau/4)\, F(\omega)\, [\ \ldots\ \ y_0 + y_1\, e^{-4i\omega(T_1/4)} + y_2\, e^{-8i\omega(T_1/4)}\ ]$$

But now the thing in brackets is exactly $Y'(\omega)/T_1$, our original digital Fourier spectrum of y(t). Go look at the definition in (25.5) . Thus we conclude that

$$W(\omega) = (1/T_1)\, X_{pulse}(\omega,\tau/4)\, F(\omega)\, [\ Y(\omega) + \sum_{m\neq 0} Y(\omega - m\omega_1)\ ] \qquad \text{D stayed fixed}$$

So this is the spectrum of the output of our new circuit, where we have increased the output latch and D/A clock rate by 4, and have kept the D/A duty cycle fixed, and have done nothing else. Look at the result. There are many observations that can be made about it:

(a) the $X_{pulse}(\omega,\tau/4)$ is fine, it is just the aperture factor that we know is present. This is the spectrum of a pulse that is 4X narrower than our original pulse, so the sinc function is now 4X broader than it was before. This is an improvement, because we have less aperture distortion on the main spectrum.

(b) Of course if we made the above change and held the aperture time τ fixed (assuming it was less than 25% of T_1 to start with), then we would still have our original $X_{pulse}(\omega,\tau)$. In this case, the above result would be

$$W(\omega) = (1/T_1)\ X_{pulse}(\omega,\tau)\, F(\omega)\, [\ Y(\omega\ ) + \sum_{m\neq 0} Y(\omega - m\omega_1)] \qquad \tau \text{ stayed fixed}$$

(c) The image spectra are still at spacing ω_1. They are *not* spaced at $4\omega_1$.

(d) We have a new messy factor $F(\omega)$ sitting in the result! At the new higher clock rate, you can think of this as the action of some mysterious filter $F''(z) = 1 + z^{-1} + z^{-2} + z^{-3}$ that appeared out of nowhere. As we saw above, you can interpret this thing as the cause of the little pulses repeating four times. This new factor $F(\omega)$ is bad news, because it is *distorting* our output signal. After all, the whole purpose of our circuit including the analog post-filter is to try and recover $W(\omega) = Y(\omega)$. In the case that we held aperture time τ fixed, the factor $F(\omega)$ is the only "new" thing introduced into our result by the 4X clock speedup. Thus, by doing this speedup *with a constant aperture time* τ, we have definitely made things *worse*.

(e) Here is an interesting limit. Suppose the aperture were 100% before and after our change. This would be the case in the simplest circuit you could design. In this case, the signal w(t) *does not change at all* when we make our 4X speedup. It continues to be a stepwise fit to y(t). Thus $W(\omega)$ must not change. Since the first formula above must apply, and since $Y(\omega)$ certainly cannot have changed, it must be true in this case that:

$$X_{pulse}(\omega,\tau/4)\, F(\omega) = X_{pulse}(\omega,\tau) \qquad \text{for } \tau = T.$$

That this is true is certainly not obvious , but when you write it all out, you find that it is indeed true. If you let $z = \exp(-i\omega T_1/4)$, replace sinc(x) with sin(x)/x, and replace sin(x) with exponentials, the thing boils down to this:

$$(1 - z)\,(1 + z + z^2 + z^3) = (1 - z^4)$$

which we know is true. In this case of maintaining 100% aperture time, we have not made things any worse by our 4X speedup. Things stay exactly the same.

Looking back at the origin of F(ω) in the above development, it is clear that you could get rid of the F(ω) distortion and make F(ω) = 1 if you were to kill off the last 3 pulses in every group of 4 pulses. It is easy to imagine a hardware circuit that would do this as part of the 4X speedup business. This process is called **decimation**, and the resulting spectrum for W(ω) would be this:

$$W(\omega) = (1/T_1)\, X_{\mathtt{pulse}}(\omega,\tau/4)\ \Big[\, Y(\omega) + \sum_{m \neq 0} Y(\omega - m\omega_1)\Big] \qquad \text{D stayed fixed+ decimation}$$

So if we do the 4X speedup with decimation, and keep the duty cycle of the D/A fixed, we get the above result that has no F(ω) distortion, and also has reduced aperture distortion on the main spectrum because the aperture is 4X less than it was before.

If we now think of aperture = 100% again, then w(t) is the exact shape of the digital signal *before* the D/A converter. That is, w(t) is the output of a latch clocked at 4X . Between clockings, the signal holds a constant level.

So now forget the D/A converter and just think about this digital domain signal. If you speed up the clock 4X *and* decimate, the spectrum of the digital signal at the output of the decimating latch is given by

$$W(\omega) = (1/T_1)\, X_{\mathtt{pulse}}(\omega,\tau/4)\ \Big[\, Y(\omega) + \sum_{m \neq 0} Y(\omega - m\omega_1)\Big]$$

The factor $X_{\mathtt{pulse}}(\omega,\tau/4)$ just arises from the fact that the signal is constant between sample points, so we are in effect superposing a bunch of abutting square wave pulses (3 out of 4 have zero amplitude).

The digital signal at this point is in good shape. There is no F(ω) distortion because we decimated. The image spectra are still ω_1 apart, they are not $4\omega_1$ apart.

What happens now if you send this digital signal into a digital low pass filter designed to do roughly a brick wall just above $\omega_1/2$. We assume that y(t) is Nyquist-OK and has no components beyond this point.

We presume of course that this digital filter is going to run at the 4X rate. As we know from earlier analysis, the image spectra of this filter *are* spaced at $4\omega_1$. So the first thing that happens is that the two image spectra of W(ω) which lie between the main spectrum and the m=4 spectrum are completely knocked out. So we get

W(ω) DigitalFilter(ω) = Output(ω) has spectra spaced by $4\omega_1$.

If we now take Output(ω) into a D/A converter running at 4X, here is what we get:

$$\text{Output}(\omega) \approx (1/T_1)\, X_{\mathbf{pulse}}(\omega,\tau)\, [\, Y(\omega) + \sum_{m \neq 0} Y(\omega - m4\omega_1)] \quad 4X \text{ with decimation}$$

Here we use ≈ because the digital filter is probably not a perfect brick wall. Compare the above result to where we started this section, ie, a system without 4X and decimation:

$$\text{Output}(\omega) = (1/T_1)\, X_{\mathbf{pulse}}(\omega,\tau)\, [\, Y(\omega) + \sum_{m \neq 0} Y(\omega - m\omega_1)] \qquad 1X$$

The only difference is that factor of 4 in the image spectra separation.

The method being discussed here is called "4 times oversampling". If you do N times oversampling, including decimation and a digital filter, the image spectra are pushed out by a factor of N. Then you can run the D/A output through an analog filter that is optimized for constant phase, like a 3rd order Bessel filter that has 3 components including only 1 inductor Such filters tend to have rather amorphous magnitude dropoff, but that is fine here because there is lots of room between the main spectrum and the first image spectrum which is now at $\omega = N\omega_{\cdot}$. The net result of doing all this is the following:

$$\text{FinalOutput}(\omega) \approx Y(\omega)$$

The digital convolution filter is hopefully a cheap digital integrated circuit, and the analog filter is also cheap. This is why 4X oversampling CD players cost $100.

There are several ways of dealing with the potential aperture distortion. If you have a narrow-pulse ground-dumping sample and hold circuit after the D/A, you can make the aperture τ very small, and pretty much get rid of the aperture distortion represented by $X_{\mathbf{pulse}}(\omega,\tau)$. Or you can pre-correct for it in the digital filter by designing in a sinx/x boost. In the 1X situation, you might even set the aperture at full 100% since this puts the first zero of the sinc function at the first image spectrum, helping the analog post-filter to crush it out.

Without the oversampling, this analog filter must try to meet both the brick wall and the constant phase requirements. Basically, this is an impossible task and one must compromise. The steep edge of the brick wall causes ringing , and ringing is distortion in the form of group delay dispersion, which is just a way of saying the filter has phase problems.

Moreover, the filter is expensive because a brick wall requires high order. Looking in the charts, in order to have magnitude down below 100 dB at the first image spectrum, a Chebyshev filter with 0.1 dB passband ripple would have to be 11th order, which means 11 reactive components in a passive implementation.