(070886)

Required Accuracy of the Roots of Unity for the FX Correlator

B. G. Clark
July 1986

The accuracy of the roots of unity required for the FX correlator is set by the desired spectral purity. That is, the response of an arbitrary channel to a strong signal in another channel. For a rough estimate of the effect it suffices to have a look at the response of the system to DC. We make a further simplifying assumtion that all data and multiplications are continuous, except for the roots of unity, which are quantized with least significant bit b.

Interestingly, in this case, the transform has enough symmetry (at least in binary radix) that odd numbered channels are exactly zero.

If the numbers being added to the actual trig functions were random rather than systematic truncation errors, one could easily calculate the expected response in an arbitrary odd channel:

(N-1)*b*b/12

N is the number of stages of a binary transform. N-1 appears in the formula because in the first stage of the transform, all coefficients are unity, and therefore have no truncation error. Since for radices which are a power of two, N-1 is still the number of trig functions going into an arbitrary channel, I maintain that it is still the appropriate number for radix four as well as for radix two.

We do not have a formal specification of spectral dynamic range. However, the correct number must be of order 5000 (RMS, giving about 1500 peak, which is more interesting for actual spectra). This suggests that b should be about 1/64; that is, that six bits (including sign) are required for the trig functions. To tolerate the spectral dynamic range of only about 400 implied by five bits would seem rather shortsighted to me.

Since the number of bits required is still fairly modest, I would urge their actual physical implementation, rather than relying on "dither" or some such trick to fake them.

Having come to a fairly definite conclusion, I must now include the hedges against the case in which it is definitely proved wrong.

There may be cases in which the truncation errors all add together rather than roughly randomly, as assumed above. I have looked for such cases in the obvious places and not found them; this does not mean they do not exist. Nor would the fact that they do not exist for the radix two transform (in which I was looking) mean that they do not exist in a radix four transform.

Including six bits does not guarantee high dynamic range, but merely permits it. I have not calculated the expected loss of dynamic range due to the five bit quantization of the data. There must be some.

In theory, the effect above is deterministic and reproduceable (although a function of fringe phase) and could be removed by a complex deconvolution of the baseline spectra before further processing. In practice, to do so would be so computationally onerous that I would regard it as unacceptable.