CFSystems

Displaced Frequency FFTs

CFS-224

July 8, 2003

David Dunthorn
www.c-f-systems.com

Note

The Precise Signal Component method as described in USP 6,751,564 is as powerful a method as that patent implies, but I did find it to have one area of weakness in precision (the only such weakness of which I am aware). This document describes a method for completely eliminating that weakness. I believe this method would have been patentable in its own right, but I offer it here as a public domain addition to the Precise Signal Component method. I intend this to serve as an example of the "obvious" improvements that will arise as Precise Signal Component becomes more familiar and normal science and engineering methods are applied to it. This method depends upon quite venerable technology and certainly will be "obvious" to practitioners in the field once it has been explained. As to whether it would have been obvious given just a statement of the weakness to be addressed, I feel that it should have been, but my own experience was that it was not. It required exploration of a number of blind alleys spread over two years before the concepts aligned in my mind in a way that made it obvious to use an extremely low frequency carrier wave.

This document is viewable only. I believe this will be adequate for people who do not intend to study it. Please contact me through our web site if you need a printable version. I am aware that the no-print can be defeated, but again I ask that you contact me instead. I really need to know if and how people are finding these documents useful, and this seems one of the few ways I have to encourage feedback.

Working with the Precise Frequency Component method, single, precise frequency components which do not coincide with the bin numbers of an FFT will produce "leakage," with the single frequency showing up significantly in the frequency "bin numbers" adjacent to the bin number nearest the single precise frequency. The leakage has the following pattern when the precise frequency is 0.4 of the way between two frequency bins. The graph below shows the amplitudes reported for a single precise frequency component versus k, the relative FFT frequency bin number. The graph shows only points at each discrete frequency bin number rather than a continuous, peaked curve because such a curve does not exist $-$ values exist only at the discrete bin numbers. As previously shown, with the signal spreading out like that, the Precise Frequency Component method can easily find the precise frequency as well as its precise amplitude and phase using a least squares fit to the known pattern over the range of frequency bin numbers.

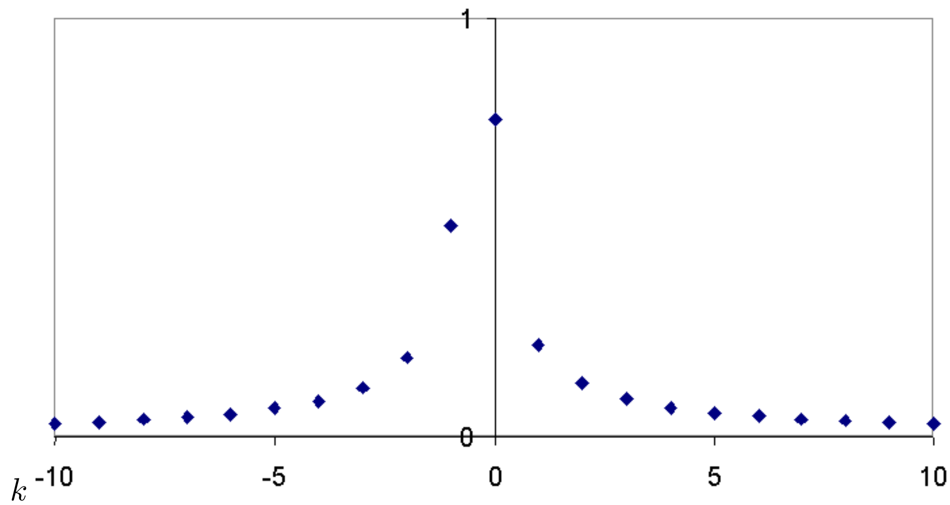

$\epsilon = 0.4$

When a precise frequency signal nearly coincides with an FFT frequency bin number, the peak is much sharper with very little leakage, as shown in the figure below. The least squares approach owes much of its high accuracy to finding a pattern, usually partly hidden by noise, that is consistent over a range of frequency bin numbers. Because there are few or no significant contributions from the signal to adjacent bin numbers, the least squares approach becomes very dependent on the single largest value, losing much of its power and losing accuracy of the amplitude and phase in particular.

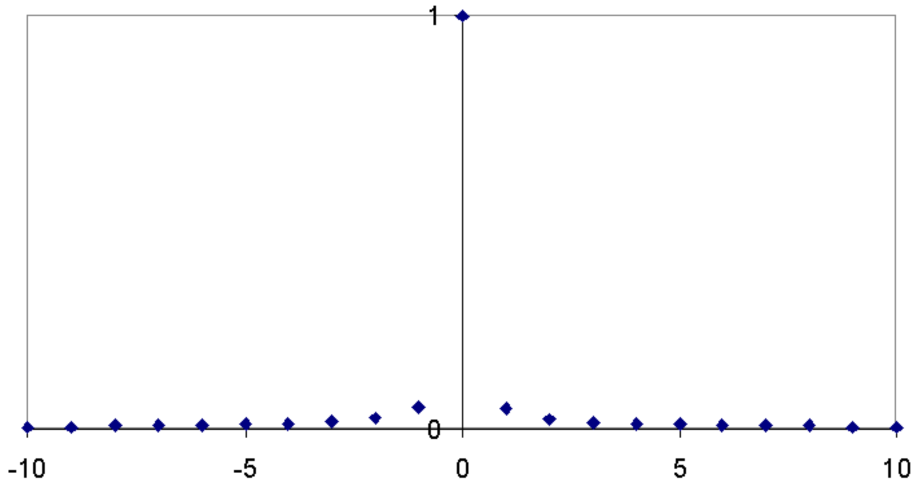

k

$\epsilon = 0.05$

As a different view of this, suppose an FFT is taken over a sample of length one second (sample interval). The FFT bin numbers will be spaced at intervals of 1/(sample interval), or 1 Hz in this case. When trying to find a frequency of 245.674 Hz, it falls very roughly halfway between the bins at 245 Hz and 246 Hz. The leakage will be large, much like the first graph above, and the precise frequency component can easily be found. When trying to find a precise frequency component of 245.08 Hz, almost all of that will fall in the 245 Hz bin, similar to the second graph. With very small effects in adjacent bins, the least squares method has nothing to work with and loses accuracy.

Since the bin numbers are set by the sample interval, the obvious approach to this problem is to use a different sample interval. While possible, this would require either using sample point counts other than a power of 2 (necessary for an efficient FFT), actual resampling of the same signal, or interpolating samples. Those choices compromise accuracy, speed and/or practicality, and the end result must still be uneven, with some frequencies inevitably coming close to bin numbers in both the original FFT and the resampled FFT.

It would be ideal to be able to just move all the frequency bin numbers for a second FFT. That is, for the example above, have an alternative FFT in which instead of the bin numbers being 1, 2, 3, 4, ..., the effective bin numbers would be 1.5, 2.5, 3.5, 4.5, ... Then both FFTs could be analyzed and every precise frequency will be at least 0.25 away from the nearest bin number in one of the two FFTs. The results could be weighted in favor of the FFT in which the frequency is farthest from the nearest bin number.

With the bin numbers being set by the sample interval, it would seem impossible to do anything that would affect the bin frequency numbers that even-handedly. Surprisingly, it turns out that it is possible and even relatively easy to do. The Fourier Transform of $g(t)$, determining the frequency spectrum f is:

$$\mathfrak{F}(g(t), f) = \int_{-\infty}^{\infty} g(t) e^{-2\pi i f t} dt \tag{1}$$

We wish to have the frequency shifted by μ, which will be 1/2 in the case of primary interest:

$$\mathfrak{F}(g(t), f - \mu) = \int_{-\infty}^{\infty} g(t) e^{-2\pi i (f-\mu) t} dt \tag{2}$$

$$\mathfrak{F}(g(t), f - \mu) = \int_{-\infty}^{\infty} g(t) e^{2\pi i \mu t} e^{-2\pi i f t} dt \tag{3}$$

But note that the integral in Equation 3 is the Fourier Transform of $g(t)e^{2\pi i \mu t}$

$$\int_{-\infty}^{\infty} g(t) e^{2\pi i \mu t} e^{-2\pi i f t} dt = \mathfrak{F}(g(t) e^{2\pi i \mu t}, f) \tag{4}$$

and thus:

$$\mathfrak{F}(g(t), f - \mu) = \mathfrak{F}(g(t) e^{2\pi i \mu t}, f) \tag{5}$$

There were two versions of this report, the second done with the finite version of the Fourier transform. When a signal $g(t)$ is subjected to a Discrete Fourier Transform (DFT), a frequency spectrum of values distributed in discrete frequency bin numbers f is obtained. The DFT is derived from the integral representing the Fourier transform evaluated over only a finite sample time τ:

$$\mathfrak{F}(g(t), n) = \frac{1}{\tau} \int_{0}^{\tau} g(t) e^{-\frac{2\pi i n t}{\tau}} d\tau \tag{1}$$

Here the frequencies used in the DFT are represented as $f = n/\tau$.

$$\mathfrak{F}(g(t), f) = \frac{1}{\tau} \int_{0}^{\tau} g(t) e^{-2\pi i f t} d\tau \tag{2}$$

We wish to have the frequency shifted by μ, which will be 1/2 in the case of primary interest:

$$\mathfrak{F}(g(t), f - \mu) = \frac{1}{\tau} \int_{0}^{\tau} g(t) e^{-2\pi i (f-\mu) t} d\tau \tag{3}$$

$$\mathfrak{F}(g(t), f - \mu) = \frac{1}{\tau} \int_{0}^{\tau} g(t) e^{2\pi i \mu t} e^{-2\pi i f t} d\tau \tag{4}$$

But note that the integral in Equation 4 is the Fourier Transform of $g(t)e^{2\pi i \mu t}$

$$\frac{1}{\tau}\int_0^\tau g(t)e^{2\pi i\mu t}e^{-2\pi ift}d\tau = \mathfrak{F}(g(t)e^{2\pi i\mu t}, f) \tag{5}$$

and thus:

$$\mathfrak{F}(g(t), f-\mu) = \mathfrak{F}(g(t)e^{2\pi i\mu t}, f) \tag{6}$$

Which says that the Fourier Transform of $g(t)e^{2\pi i\mu t}$ is the same as the Fourier Transform of $g(t)$ with the frequencies displaced by μ. In the case above, if μ is taken as 0.5, the frequency bin numbers are all shifted by 0.5, which is exactly what is wanted. In operation, first an FFT is taken of a sample $g(t)$, then the sample is multiplied by $e^{2\pi i\mu t}$ and a second FFT taken. Of course it also is possible to apply different vales of μ. For example, if it were decided that 0.25 bin number is still too close for best accuracy of the least squares, $\mu = 0.25$ and $\mu = 0.75$ could be used in addition to $\mu = 0.5$.

This approach has been tested and does obtain higher accuracy for frequencies near the frequency bin numbers of the original FFT while maintaining all the desirable characteristics of the original Precise Frequency Component method.

It is worth pointing out that what we have done here is apply the venerable Modulation Theorem. Normally, this theorem is used to show that when a signal is used to amplify (i. e. multiply) a carrier wave, the modulated signal retains the relative identity of all its frequency components, but all the frequencies are shifted by an amount equal to the carrier wave frequency. That is, the differences in frequency between the various frequency components represented in the frequency spectrum remain the same but all frequencies are raised the same value; the value of the carrier frequency. Such carrier frequencies are normally much higher frequencies than any in the modulating signal, but here the "carrier" is low frequency; 0.5 Hz in the example.