Numerical Simulation of Pink Noise

D.McClain, Sr. Scientist
Raytheon Missile Systems Co., Tucson, AZ

May 1, 2001

Abstract

We present considerations for accurate simulation of $1/f$ shaped noise records used for system simulation. This is a difficult subject because the inverse square root function has both an "infrared catastrophe" at zero frequency, and because it decays so slowly that it is not strictly integrable over the infinite frequency domain.

1 Numerical Simulation of Pink Noise

There are two published methods for generating $1/f$ shaped noise[1]. One method utilizes a hierarchical assembly of maximum length sequence (MLS) generators, and the other utilizes cascaded 1-pole filters to approximate the -3 dB/octave rolloff. This paper summarizes another technique based on the forced spectral shaping of noise. We will attempt to describe the limitations of this method and give expressions for error bounds in the resulting approximations.

The generation of $1/f$ noise is fraught with difficulty. Researchers have often mistakenly produced $1/f^2$ noise instead of their intended results. The generation of $1/f$ noise involves amplitude scaling in the frequency domain by $1/\sqrt{f}$, and not by $1/f$ itself. The production of $1/f$ noise implies that its power spectrum has this frequency roll-off, and hence, that the amplitude itself rolls off at $1/\sqrt{f}$.

The function $1/\sqrt{f}$ has an "infrared catastrophe" at zero frequency, making it difficult to simulate numerically. It is not strictly integrable over the infinite frequency domain because it rolls off so slowly, yet it does possess a well defined Fourier Transform.

I will coin a new term "iso-transmorphism" to describe functions that retain their character under a Fourier Transform. Scaling and translation alter the precise forms in the conjugate domains, but the overall character of such functions is preserved. It is well known that Gaussian envelopes behave in this

[1] An excellent summary of these two methods can be found on the Web site http://www.firstpr.com.au/dsp/pink-noise. An excellent bibliography of available papers on the subject of $1/f$ noise in general, can be found at Web site http://linkage.rockefeller.edu/wli/1fnoise.

manner. The function $1/\sqrt{f}$ also behaves this way, since its transform is $1/\sqrt{t}$, with t being the conjugate variable to frequency f. There may be others as well, but these two are the most well known, with the square-root function being less commonly understood.

In numerical simulations we depart from strict Fourier Transforms, to implied periodic functions which have discreet spectra at integer multiples of their inverse period. The Fourier Transform is commonly replaced by Fast Fourier Transforms (FFT). The convolution theorem continues to hold in this simulation domain, but it describes "circular convolutions" instead of the ofttimes implied "linear convolutions". One result of this is that "edge effects" can contaminate the results unless care is taken to minimize this possibility.

Edge effects are usually mitigated by means of "windowing" applied to the input record prior to transformation, and by zero-padding the input record so that the edge effects of convolution can be absorbed in the added null regions.

Windowing a function ahead of transformation produces a convolved spectrum between the spectra of the input data and that of the windowing function. Commonly used windowing functions are the Hanning Window, which consists of a biased 1-cycle cosine function, and variants of it that produce varying degrees of smoothness in low-order derivatives. Sharp transitions in various orders of derivatives of the input data cause ringing phenomena in their spectra that "bleed through" the implied narrow-bandpass filters of the FFT.

Zero padding one's data produces interpolated spectra with more fine-grained detail in the spectra, but with no more information than was originally present.

In all cases, unless the input data are sufficiently band-limited, there will be aliasing of high frequency information to the lowest bandpass of the FFT. Windowing and zero padding can help reduce the effects of aliasing, but aliasing occurs ahead of these mitigating processes and is always present unless the input data is sufficiently band limited.

1.1 Forced Spectral Shaping

We can generate noise records of a variety of statistical character, e.g., uniform and Gaussian distributions. But generally, these noise records have "white" spectra, implying that the noise retains its statistical character at all frequencies. Forced spectral shaping is performed by multiplying the spectrum of a noise record by the desired frequency domain amplitude function. The results are then inverse transformed back to the time domain for use in simulations.

Multiplication in the frequency domain is equivalent to (circular) convolution in the time domain by the inverse transform of the amplitude shaping function. When this inverse transform has limited extent in the time domain we can extend its form with zero padding to absorb the end effects resulting from the circular convolution. One can also force a function to have limited extent in the time domain by application of windowing. But realize that in both cases, the frequency domain amplitude function will depart from the originally desired shape. Windowing the time domain "filter" function smears out the fine details of the amplitude frequency response. Zero padding produces a band-limited interpolation of the amplitude frequency response.

In the case of $1/\sqrt{f}$ this function has neither a limited extent time domain conjugate function (its trans-isomorphism is $1/\sqrt{t}$), and there is a gross discontinuity in the zeroth derivative (the function itself) near the origin in both domains. Furthermore, this function is not strictly defined for either negative times, nor for negative frequencies. So we face the need to form an analytic continuation for negative values of its domain variable.

Real valued signals in the time domain have spectra with Hermitian symmetry in the frequency domain. The real part of the spectrum is an even function of frequency, while the imaginary part of the spectrum is an odd function. We must preserve this symmetry condition when applying forced spectral shaping, in order that we may continue to operate with real-valued time domain signals.

There are two simple ways of forming an analytic continuation of $1/\sqrt{f}$: make it a real, even, function of frequency, or make it a purely imaginary, and odd, function. In both cases, subsequent multiplication against a signal spectrum preserves the originally Hermitian symmetry of the signal spectrum. When the shaper is a real, even, function of frequency, the character of the signal spectrum is unchanged. Its real part remains real and even, while its imaginary component remains imaginary and odd. When the analytic continuation is an imaginary, odd, function of frequency, subsequent multiplication against the signal spectrum interchanges the real and imaginary components of that signal spectrum, while force its originally real and even character to become imaginary and odd. Its imaginary and odd component becomes real and even, possibly with negation. The effect of an imaginary shaping function for the analytic continuation of $1/\sqrt{f}$ merely produces a phase shift in the original signal in addition to the spectral shaping desired. This may or may not be of any significance, depending on the use made of the shaped noise record in subsequent simulation activities.

What should be done at the zero frequency? We can't operate with infinities, even though they integrate just fine. The spectral shaping is accomplished by means of pointwise multiplication (complex multiplication) between the signal spectrum and corresponding elements of the shaping spectrum. If we arbitrarily set the zero frequency value to some value then we will introduce an error in the overall approximation. Yet at the same time, since we cannot operate with infinities, we have no choice but to force an approximation.

What value should be used for the zero frequency value of the shaping function? Zero? One? Ten? If you perform an FFT on a signal record formed as the even extension of $1/\sqrt{t}$ then you get a family of curves for its spectrum depending on the value chosen for its zero time point. The same is true of an inverse FFT applied to a similarly extended $1/\sqrt{f}$ function. In either case, it seems desirable to generate a conjugate function as free as possible from gross artifacts. Figure 1 shows an even extension of $1/\sqrt{t}$ for a total of 4096 samples. Figure 2 shows what happens when its value at zero is forced to zero. Notice the grotesque artifact produced at around 0.1 cycles/period. The blue line drawn in that figure illustrates the desired $1/\sqrt{f}$ behavior.

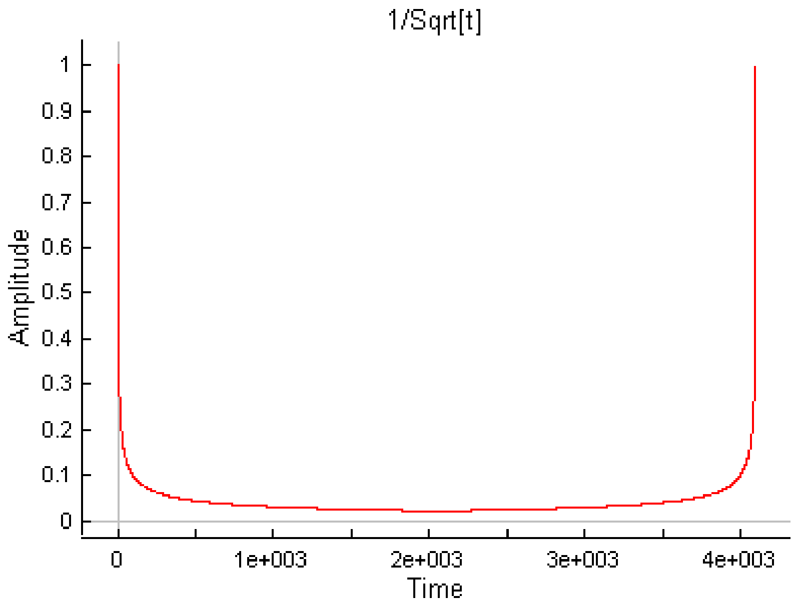

Figure 1: Graph of $1/\sqrt{t}$ analytically extended as an even function of time, and forced to become periodic over 4096 samples. Values of time above 2048 correspond also to negative time values.

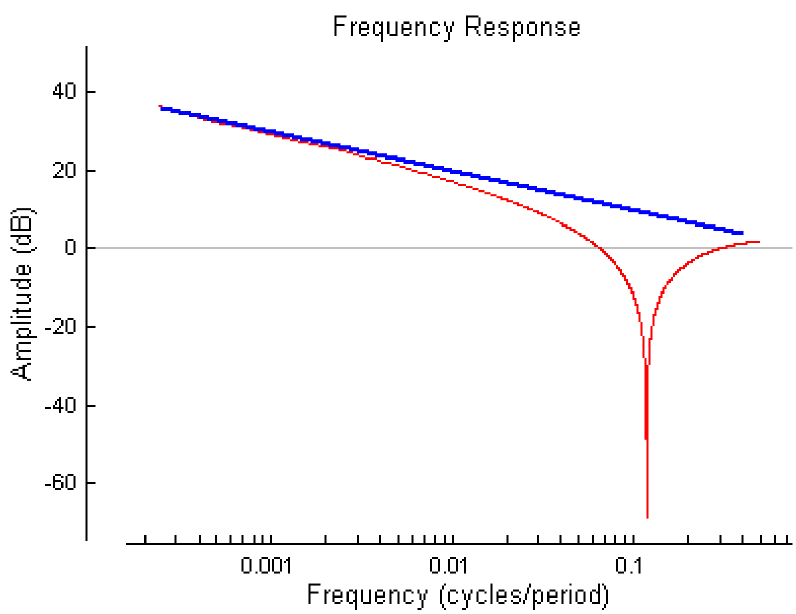

Figure 2: Discreet (FFT) spectrum of the function $1/\sqrt{t}$ showing a gross artifact in its spectrum resulting from forcing the zero time value to be zero. The blue line indicates the expected $1/\sqrt{f}$ behavior.

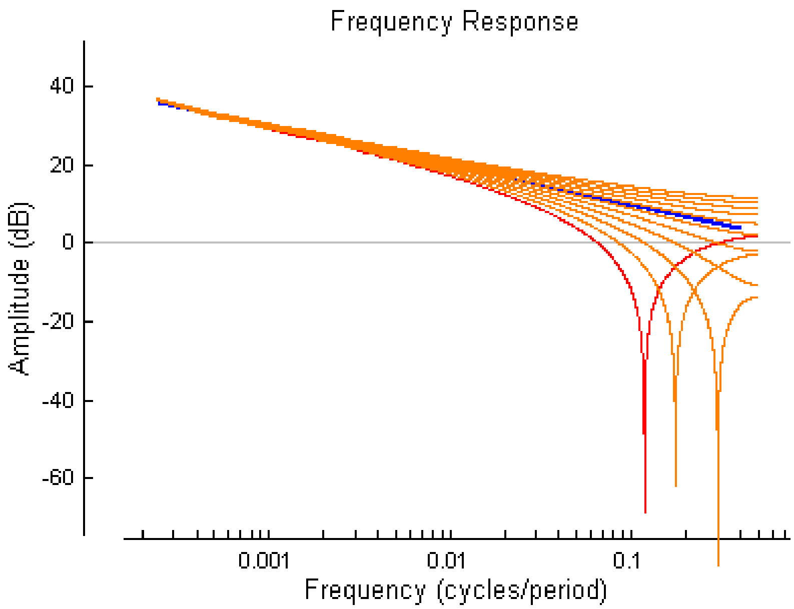

Figure 3: A family of discreet (FFT) spectra of the function $1/\sqrt{t}$ for various values of its zero time amplitude. The leftmost gross artifact corresponds to the value zero, and all others are for ascending values of the zero point progressing in 0.5 steps up to 5.

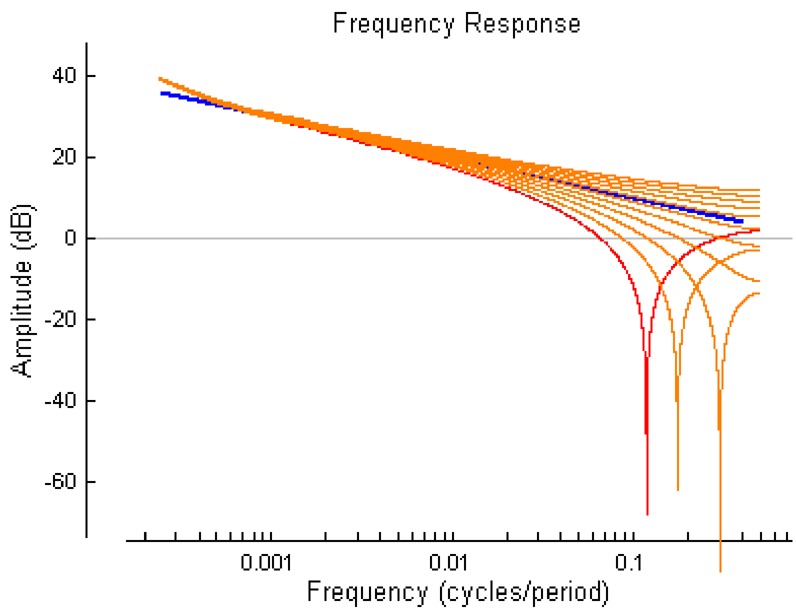

Figure 4: A graph of the same family of curves as shown in Figure 3, but with a Hanning window applied ahead of the FFT's.

Figure 3 shows the family of curves obtained as the zero time amplitude is adjusted from a value of zero up to a value of 5. Notice that there is some value near the middle of this series that has a least departure from the desired (blue) behavior at higher frequencies. All of these approximations become pretty good at low frequencies. None of these functions were windowed ahead of the FFT. When a Hanning window is applied the results shown in Figure 4 indicate a slightly worse behavior at low frequencies than without windowing.

Since the departure from ideal behavior at low frequency is nearly identical for all curves in the family when windowed by the Hanning window, let's just examine the nature of the departure at high frequencies for the unwindowed functions.

By plotting the departure at high frequencies as a function of the zero time amplitude, we can find one curve that demonstrates nearly "minimax" error behavior. That is to say, the maximum absolute error is minimized and is roughly equal for departures above and below the ideal behavior. The curve that most closely demonstrates this behavior has a zero point amplitude of 2.73. This is shown in Figure 5. Its error behavior is shown in Figure 6.

So what if there is an artifact in the spectrum when transforming $1/\sqrt{t}$? We could simply multiply the signal spectrum by $1/\sqrt{f}$ since that is the correct shaping (ignoring the "infrared catastrophe" for the moment). Well, just as the spectrum resulting from the forward FFT shows a gross artifact when we use zero

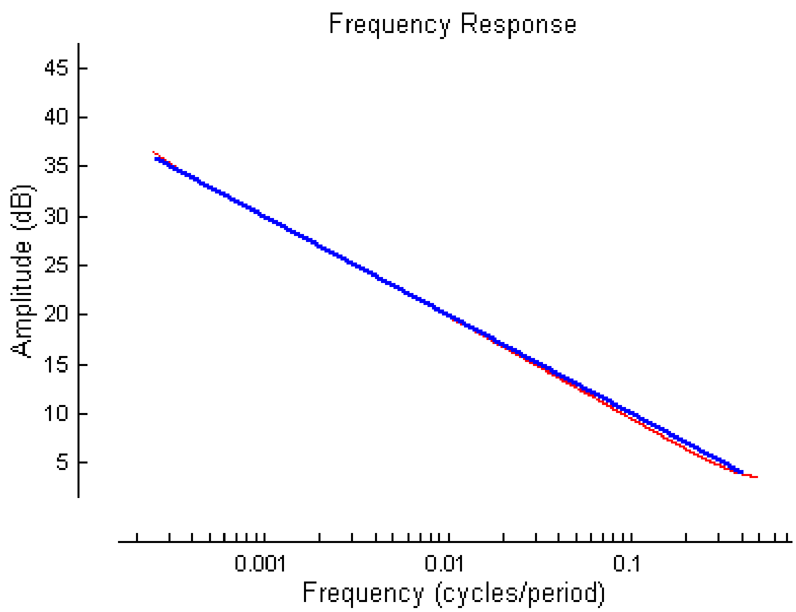

Figure 5: A graph of the even analytic extension of $1/\sqrt{t}$ which has its zero time value specified as 2.73.

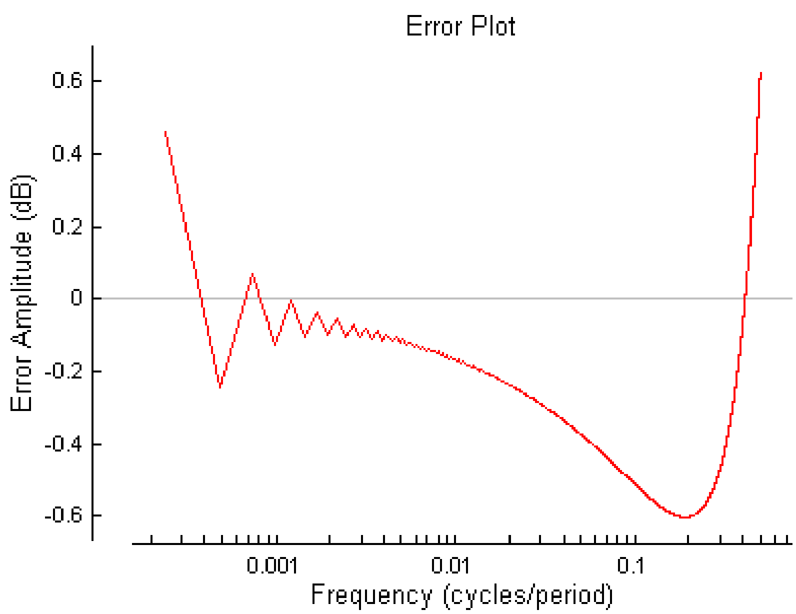

Figure 6: Graph of the error between the transform of the even analytic extension of $1/\sqrt{t}$ and the desired $1/\sqrt{f}$ behavior when the zero time amplitude is specified as 2.73. Notice that the error extremes are roughly equal, corresponding to error values of -0.60 dB and +0.63 dB.

for the value at zero time, so too will the time record resulting from convolution with such a function. We would prefer to have more uniform temporal behavior in the statistics of the output time record. It is not immediately clear what the nature of a temporal artifact would be. But why settle for anything more than minimal departure from ideal behavior?

1.2 Dealing with Non-Integrability

We have already shown that the application of windowing in an effort to mitigate the effects of infinite temporal response produces an error at low frequencies that is nearly independent of the value chosen for mitigating the "infrared catastrophe". This error at low frequencies must be traded against the error due to edge effects from circular convolution. Since the temporal behavior of $1/\sqrt{t}$ is not limited in time there will be edge effect induced artifacts in any circular convolution with this function.

Windowing generally imposes a reduction to small amplitudes at the edges of one period of the signal. We could arbitrarily use a shortened window that goes gracefully to zero at its endpoints and remains zero up to the edges of the overall signal record. This is equivalent to zero padding fore and aft of the windowed values. This will reduce edge effects to inconsequential values, but then the approximation in the conjugate domain must suffer. Escape from edge effects is complete when the convolution kernel is restricted to the central half of the record. What will be the effect in the conjugate domain of such an approximation? Surely, the error will be lessened when we utilize a sufficiently long kernel so that the values of $1/\sqrt{t}$ have themselves diminished to negligible values. How long a record should we use for the convolution kernel?

It should be clear that using a zero padded time record for a convolution kernel will induce departures from ideal behavior at the lowest frequencies. The longer the overall time record for the kernel, the lower these departure frequencies will be. So when asking how long a time record should be used, we are also asking above what frequencies is the approximation within some tolerable error level?

Figure 7 shows the effects of successively shorter windowed convolution kernels. Each convolution kernel was confined to the central half of their time record, with zero padding on either side sufficient to absorb the end effects resulting from circular convolution. As expected, as the kernel shrinks in length the minimum usable frequency increases. At a record length of 32K the departure becomes noticeable at around $f = 0.000125$ cycles/period, while for records as short as 128 elements, it has increased to a minimum usable frequency of $f = 0.03$ cycles/period. In fact, with the logarithmic frequency axis shown in Figure 7 it is apparent that doubling the length of the kernel record shifts the minimum usable frequency by roughly 0.5. In other words, the minimum usable frequency appears to be related to zero-padded kernel length as

$$f_{\min} \gtrsim 4/N \text{ cycles/period} \tag{1}$$

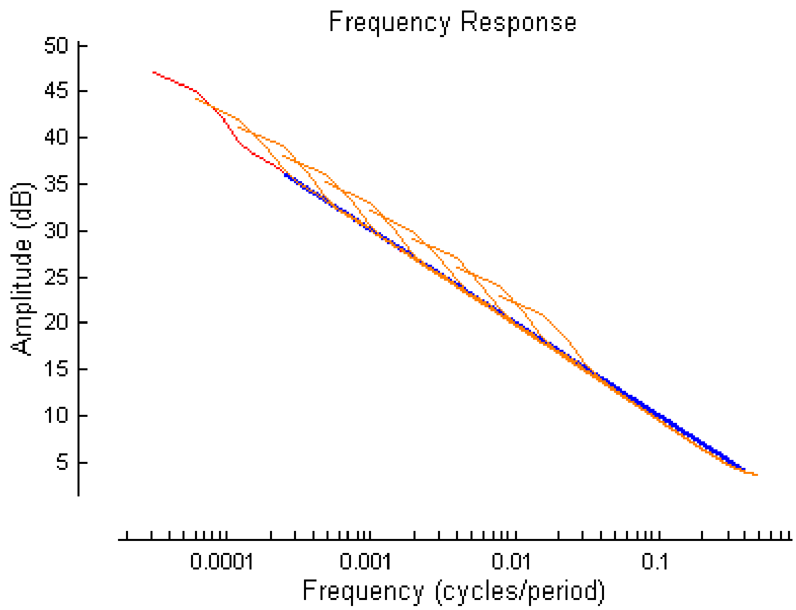

Figure 7: Family of graphs showing the effect of changing the length of the convolution kernel. The leftmost (red) departure shows what happens for a record length of 32768 elements with the kernel confined to the central 16384 samples. Error curves to the right are for successively shorter kernels of 16384, 8192, 4096, 2048, 1024, 512, 256, and 128 elements.

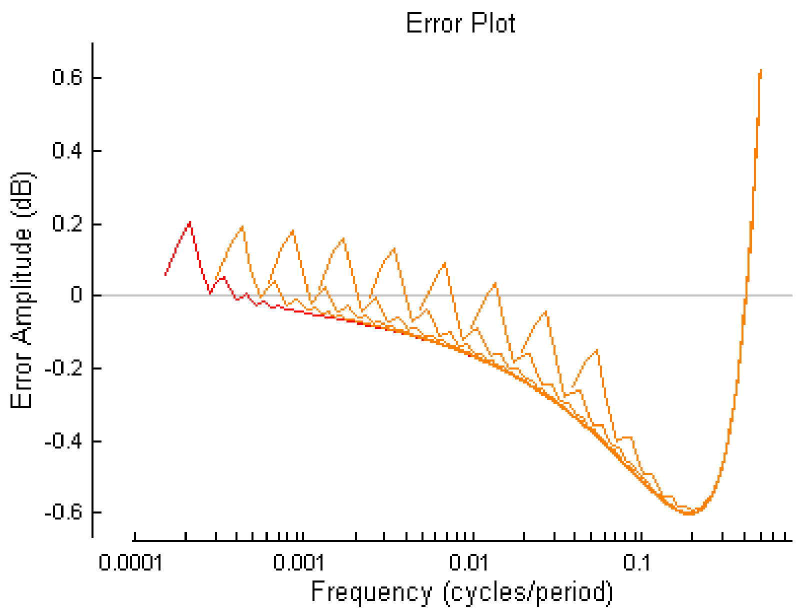

Figure 8: Plot of errors between the spectral approximation of zero-padded, interior windowed, analytic continuations of $1/\sqrt{t}$ as a function of kernel size, for frequencies above $4/N$.

where N is a power of 2 length for a zero padded kernel with non-zero values restricted to the central half of its time record.

Figure 8 shows the error committed by the windowed approximations for frequency values above the minimum stated in Relation 1 above, when the zero time amplitude in all cases is set to 2.73. In every case the absolute error for frequencies above $f_{\min}$ is less than the earlier bounds established for the unwindowed case shown in Figure 6.