Time and Frequency Domain Windowing of LFM Pulses

Mark A. Richards

September 29, 2006

1 Frequency Domain Windowing of LFM Waveforms in *Fundamentals of Radar Signal Processing*

Section 4.7.1 of [1] discusses the reduction of time (range) sidelobes in the output of the linear FM (LFM) waveform matched filter using weighting in the frequency domain. The approach is to form a modified filter frequency response

$$H'(F) = w(F) X^*(F), \tag{1}$$

where $w(F)$ is a window function (Hamming, Kaiser, Taylor, *etc.*), instead of the usual matched filter choice $H(F) = X^*(F)$.

By performing the matched filtering in the frequency domain, we can combine the matched filter operation with the windowing to reduce the required computation. A block diagram of one way to arrange the flow of operations is shown in Figure 1. This diagram shows the window applied to the filter frequency response before the actual filtering step. This has the advantage of allowing the windowing to be done off-line, *i.e.* not in real time. We could just as well apply the window to the product of the signal and matched filter frequency responses. The end result at the output would be identical, but the windowing would have to be repeated every time there was a new signal to filter, and in real time.
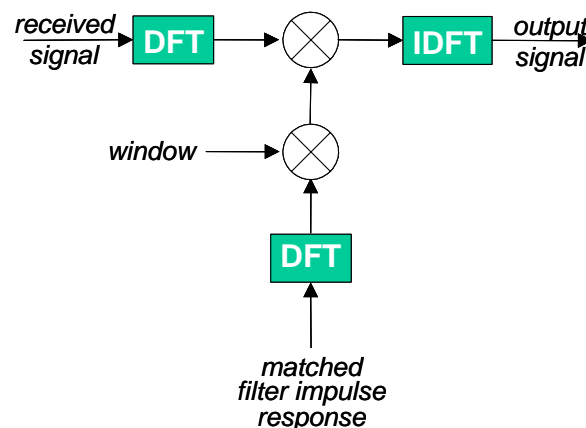

Figure 1. Flow diagram for LFM sidelobe suppression by frequency domain windowing.

Figure 2, which also appears as Figure 4.32 in [1], illustrates the application of a Hamming window function to the LFM chirp oversampled by 1.2x. The DFT of the chirp has been rotated to place the zero frequency point in the middle of the plot. Note that the window function has been aligned so that its center is aligned with the center of the LFM spectrum. Furthermore, the width of the window corresponds to the nominal width of the LFM spectrum, namely $\pm\beta/2$ Hz.
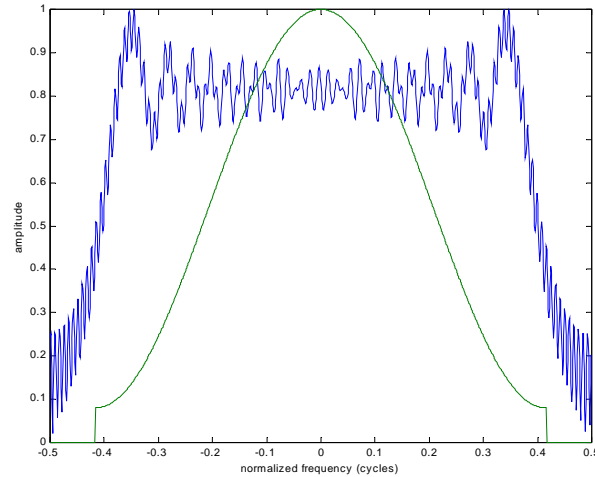
Figure 2. Hamming window function aligned with respect to chirp spectrum and cut off at $\pm\beta/2$ Hz.

Figure 3 shows the output obtained by applying the operations shown in Figure 1 to the echo from a single point scatterer, that is, to a replica of the transmitted waveform. An oversampling rate of 10x was used to get good definition of the response details. The results are shown with (blue curve) and without (green curve) Hamming weighting.

As discussed in [1], the windowed filter response suffers a loss in the absolute response of the peak. This is inevitable because the window modifies the matched filter frequency response, so that it is no longer exactly matched to the transmitted waveform. Consequently, there must be a loss in peak response. In Figure 3 the peak is reduced from 60 dB to 54.65 dB, a loss of 5.35 dB. The predicted loss is given by [1]

$$LPG = \frac{1}{K^2}\left|\sum_{k=0}^{K-1} w[k]\right|^2$$

A 16,384 point FFT was used for the DFTs in Figure 3 to get good detail in the spectrum. Because the data is sampled at 10x the signal bandwidth, the LFM spectrum is largely confined to only $1/10^{th}$ of the full DFT output, so that the size of the Hamming window is

(16,384/10), which rounds to 1,638 samples. Evaluating the equation for *LPG* for a 1638-point Hamming window gives a predicted loss of 5.36 dB, in excellent agreement with the measured value of 5.35 dB.
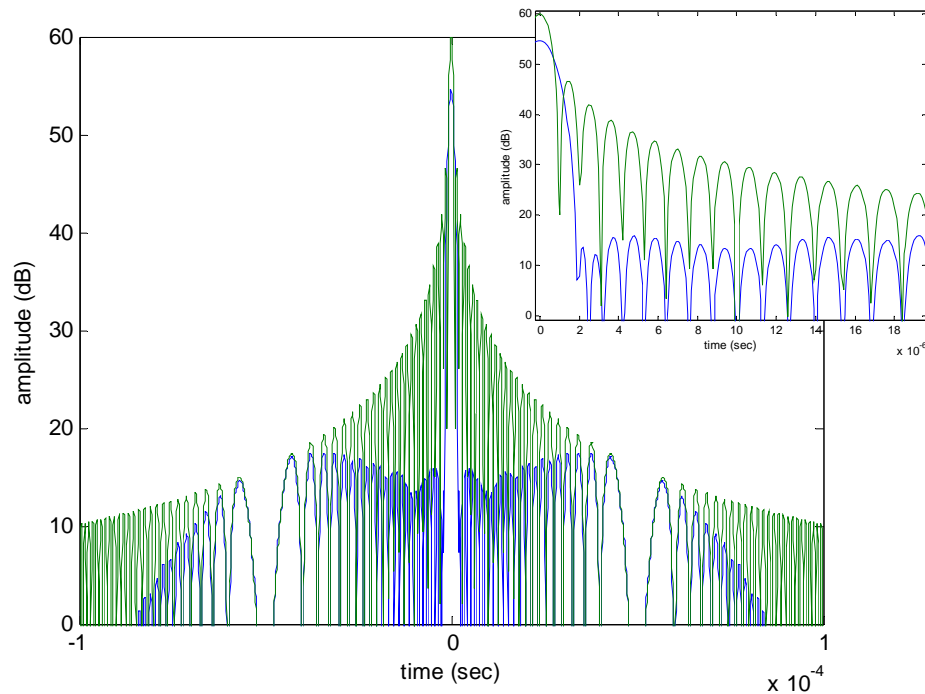

Figure 3. Output of frequency domain matched filter, with and without Hamming weighting.

The primary purpose of windowing is reduction of sidelobes. The first sidelobe is 13.5 dB below the peak for the unwindowed case, and 41 dB below the peak for the windowed case. The peak sidelobe (not the same as the first sidelobe in the weighted case) is 37.2 dB down with weighting. Thus, use of the Hamming weighting has improved peak sidelobe suppression by 23.7 dB (37.2 – 13.5).

The final effect of interest caused by the windowing is the broadening of the matched filter mainlobe, which represents a loss of range (time) resolution. Inspection of the location of the first zero of the response for each case shows it to occur at about 1 μs in the unwindowed case, and 1.93 μs in the windowed case.[1] This latter value is a little less than the 2 μs anticipated due to the 2x broadening expected of a Hamming window.

There is a legitimate question as to whether the Hamming window should be chosen to cut off at $\pm\beta/2$ Hz as shown in Figure 2. This choice is obviously motivated by the swept instantaneous frequency range of the chirp, but because of the modest BT

---

[1] The oversampling rate was increased to 30x to get better definition of the first zero in the windowed case.

product of 100, the spectrum does not cut off sharply at $\pm\beta/2$. Some perhaps non-trivial energy outside of $\pm\beta/2$ is zeroed by the window in this case. Figure 4 repeats the experiment of Figure 3, but with the Hamming window expanded by 10% in frequency to cover more of the LFM spectrum tails.
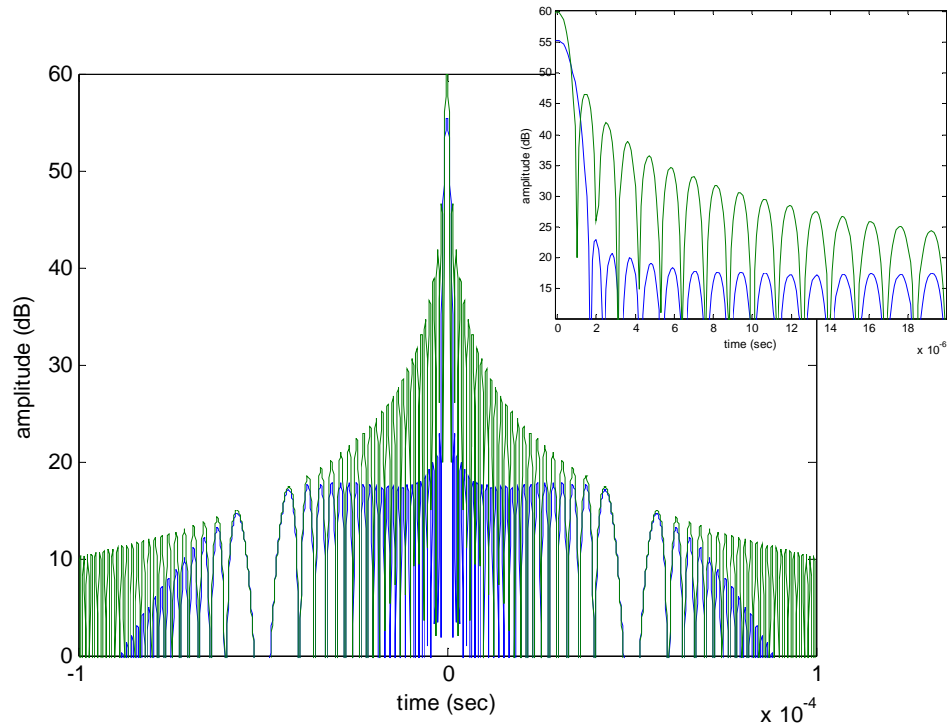

Figure 4. Same as Figure 3, but with 10% expanded bandwidth Hamming window.

Close inspection of Figure 4 shows that the peak is now reduced only to about 55.4 dB. The corresponding loss is 4.6 dB instead of the previous 5.35 dB. Thus, the loss has been reduced (improved) by 1.34 dB by not discarding the energy at the tails of the spectrum. On the other hand, the peak sidelobe (which is now also the first sidelobe) is about 32.5 dB down from the peak, not quite as good as the 37.2 dB for the case where the Hamming window cutoff was at $\pm\beta/2$. As the BT product gets larger, the cutoff of the signal spectrum becomes sharper, so that for large BT products, one should most likely cutoff the window at $\pm\beta/2$ Hz.

2   Time Domain Windowing of LFM Waveforms

Next, time-domain weighting of the receiver impulse response is considered. Section 4.6.2 of [1] showed, using the principle of stationary phase (PSP), that an LFM pulse with a time-domain amplitude $A(t)$ would have a spectrum whose magnitude followed the same shape as $A(t)$, but spread over the frequency range $\pm\beta/2$ Hz.

Specifically, from Eqn. (4.91), an LFM waveform with amplitude $A(t)$ in the time domain will have an approximate spectral magnitude given by

$$\left|X\left(\Omega\right)\right| = \left|A\left(\frac{\Omega}{2\alpha}\right)\right|, \quad \alpha \equiv \pi\frac{\beta}{\tau}$$

If $A(t)$ has finite support on $-\tau/2 \le t \le \tau/2$, it follows that $X(\Omega)$ will have finite support on $-\beta/2 \le F \le \beta/2$ and that, in that interval, $|X(\Omega)|$ has the same shape as the window $A(t)$. Thus, a Hamming-shaped (for example) spectrum can be obtained by applying a Hamming window to the impulse response $h(t)$ instead of the frequency response $H(F)$. Note that this result is specific to the use of linear FM.

The output of the resulting filter, shown in Figure 5, has the same general character as the frequency-domain weighting result of Figure 3, but with some differences in details of the sidelobe structure. The peak is reduced from 60 dB to 54.64 dB with weighting, a reduction of 5.36 dB that agrees with the predicted value. The peak sidelobe of the weighted response (which is the first sidelobe in this case) is 40.7 dB below the mainlobe peak, 3.5 dB better than the first frequency-domain case. The Rayleigh width of the unwindowed case remains 1 μs, while the windowed case Rayleigh width is 1.97 μs; closer to the 2 μs expected for the Rayleigh window than the 1.93 μs observed in the frequency domain weighted case.
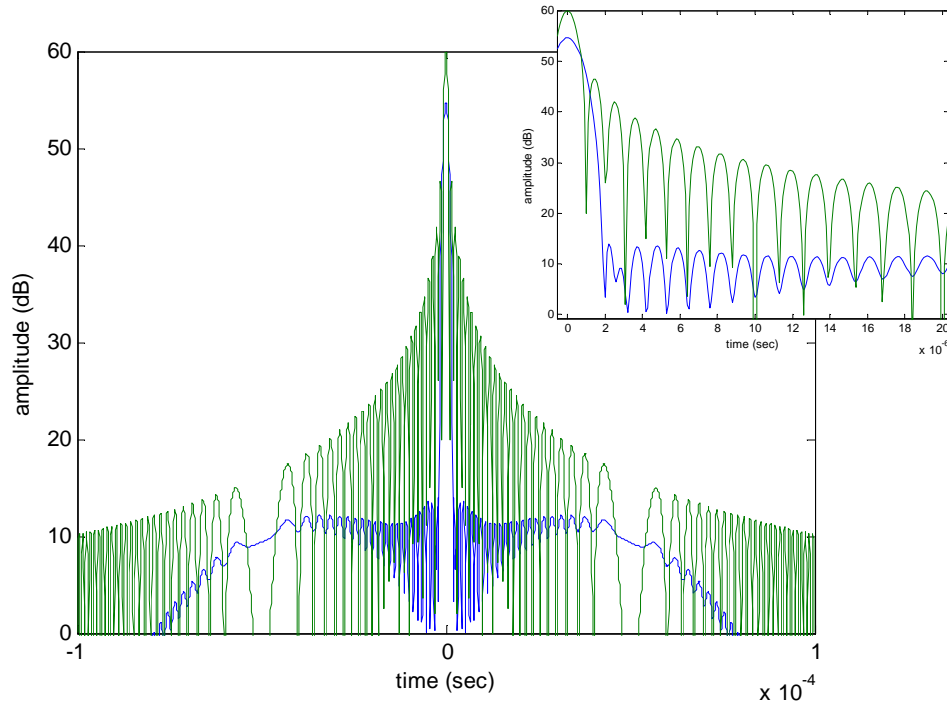

Figure 5. Matched filter output with and without time-domain weighting of the filter impulse response. Compare to Figure 3.

For convenience, Figure 6 plots the frequency- (blue curve) and time-domain (green curve) weighted responses on the same plot. The main part of the figure plots lines connecting the sidelobe peaks to enable an easy comparison of the sidelobe levels in the two cases. The inset is the full response. The difference in sidelobe levels ranges from about 2.5 to as much as about 5 dB.
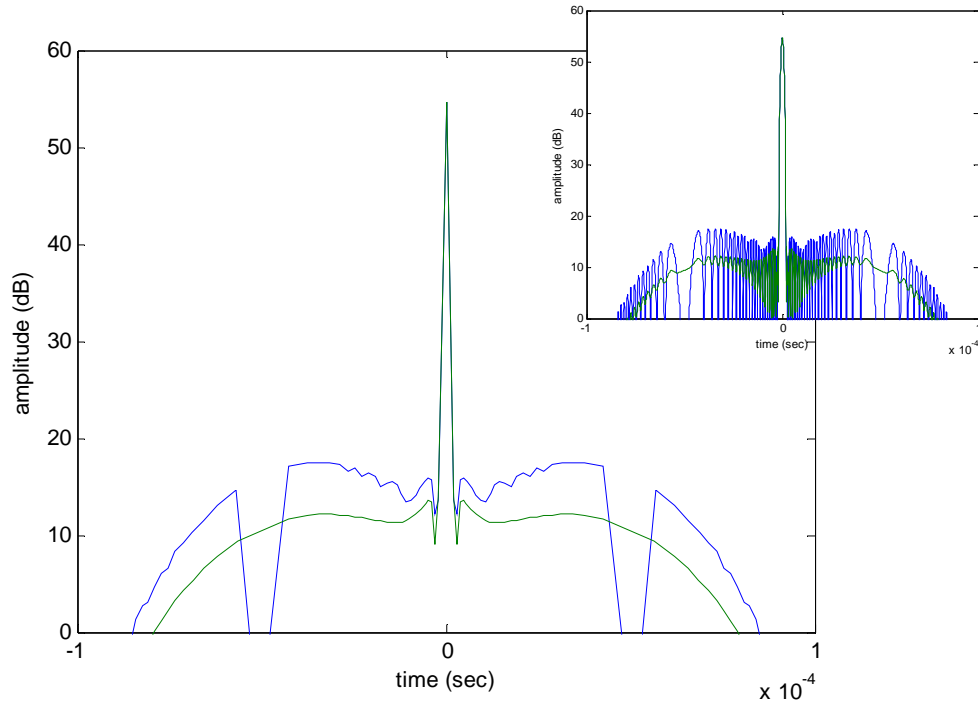

Figure 6. Comparison of sidelobe levels of frequency-domain weighted (blue) and time-domain weighted (green) matched filter outputs.

3   Additional Options

It is worth noting a third option. Both a Hamming-weighted spectrum at the output of the matched filter, and a truly matched condition, can be obtained by allowing amplitude modulation of the transmitted pulse. The transmitted LFM pulse is amplitude modulated to have a shape corresponding to the square root of the desired weighting function, $\sqrt{A(t)}$. A true matched filter is then used, which therefore has the same amplitude shaping. According to the PSP approximation, the spectrum of the waveform and the frequency response of the filter both also have magnitudes proportional to the square root of the desired weighting function, namely $\sqrt{A(\Omega/2\alpha)}$. When they are multiplied to form the filter output spectrum, it will have the desired shape $A(\Omega/2\alpha)$, producing the desired reduced-sidelobe response. However, this approach requires a transmitter with amplitude modulation capability. Furthermore, amplitude modulation

implies that less than the maximum energy is transmitted in a given pulse, since the transmitter will not be at full power throughout the pulse duration.

Finally, note that receiver filter output spectrum shaping can also be achieved while using both a true matched filter and avoiding amplitude modulation by resorting to nonlinear FM modulation, as discussed in Section 4.7.2 of [1]. However, as noted there, NLFM waveforms tend to have poor Doppler tolerance. Additional details on NLFM waveforms are available in [2].

4   References

[1]   M. A. Richards, *Fundamentals of Radar Signal Processing*. McGraw-Hill, New York, 2005.

[2]   N. Levanon and E. Mozeson, *Radar Signals*. Wiley, New Yori, 2004.