Christopher Tsai
January 31, 2007
EE 355 – Imaging Radar and Its Applications

Problem Set III
System Design and Sidelobe Reduction

Problem #1 – In-Phase/Quadrature vs. Offset Video

For the In-Phase/Quadrature (I/Q) system, we process the complex-valued chirp signal centered on the carrier frequency f_c = 10 MHz. We specify this center frequency when generating the chirp signal from our Matlab function, and the resultant chirp has both real and imaginary components; we plot the real part:

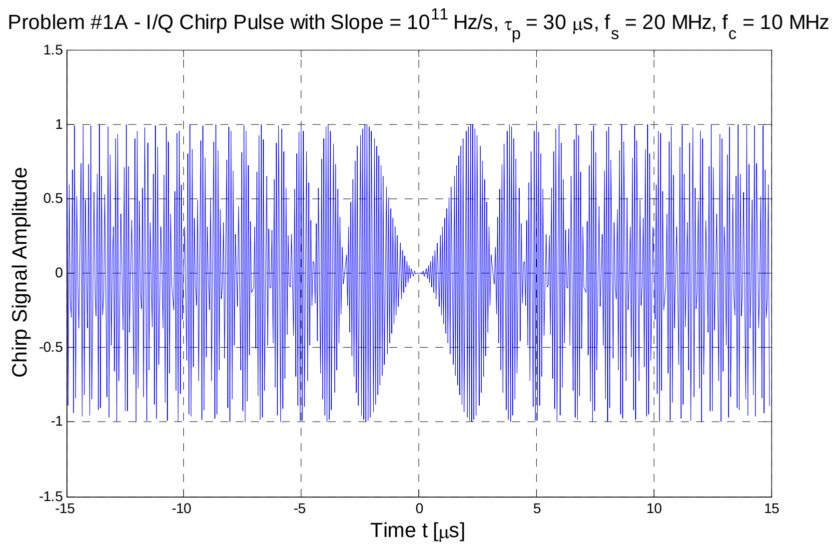

Applying the Fast Fourier Transform (FFT) to the complex signal, we observe a magnitude spectrum centered at 10 MHz:

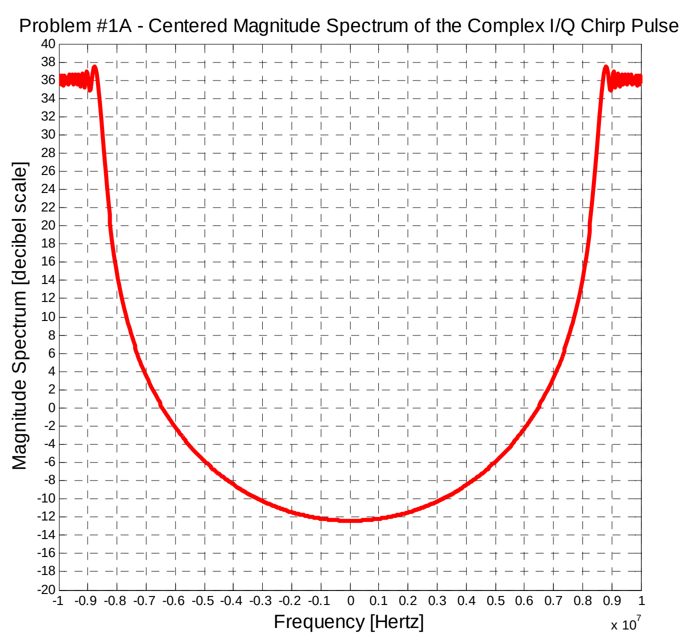

In our plot, we also see spectral content around -10 MHz, but this merely reflects the cyclical nature of the FFT rather than actual frequency content in the negative spectrum. Because our signal is complex (I/Q), its magnitude spectrum must necessarily contain only one side. By correlating the complex chirp with a matched filter through multiplication in the frequency domain, we obtain a compressed pulse.

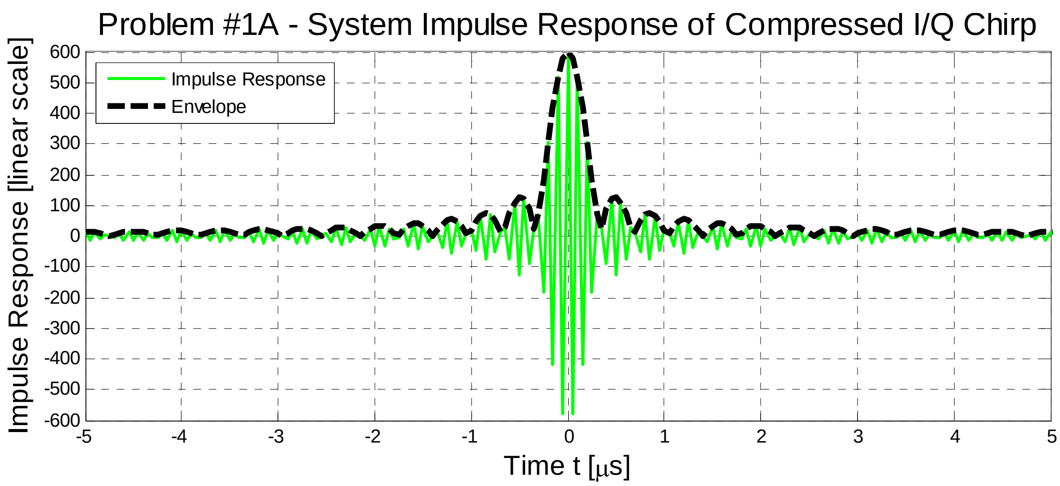

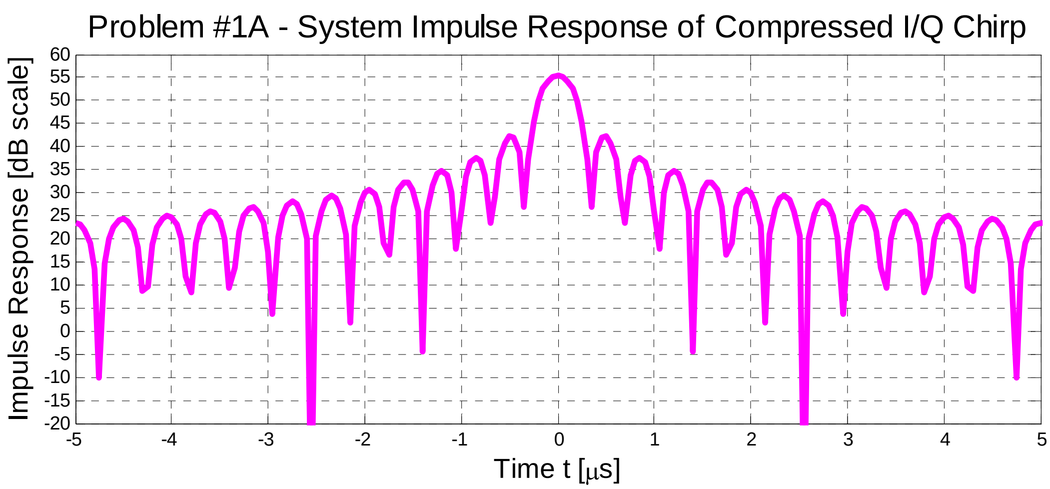

The resultant signal is completely real since matched filtering removes the chirp's complex phase. The pulse envelope resembles a sinc function since the spectrum is a rectangular band centered on a carrier frequency. Note the actual impulse response oscillates with a high frequency since the signal rests on a 10 MHz carrier.

Instead of processing complex data in I/Q format, suppose our system samples purely the real part. We can still reconstruct the complete signal and compress it identically, but we must first sample at twice the bandwidth to accommodate the negative frequency replica that results from a real signal; because our sampled chirp signal is now real rather than complex, its spectrum is hermitian, so we must extend the bandwidth to contain the hermitian conjugate side so that the positive and negative sides do not alias. By extracting only the real part, we see the two-sided spectrum:

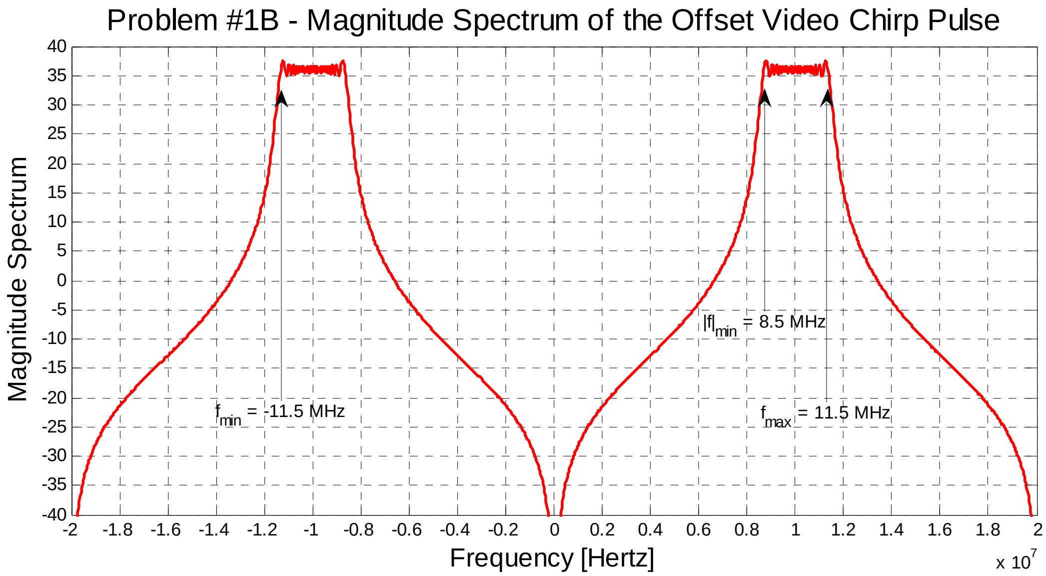

The maximum frequency in the spectrum is approximately 11.5 MHz, while the negative side extends to approximately -11.5 MHz. Since each side spectrum's bandwidth is still the product of chirp frequency slope and pulse width [$s\tau = (10^{11}$ Hz)(30 μs) = 3 MHz], the total effective bandwidth is approximately 6 MHz, forcing us to sample at twice the frequency to accommodate the additional spectral content. First, notice that the two side spectra interfere very little because our carrier frequency centers each portion well above half a bandwidth from zero frequency, with DC phase cancelling nearly perfectly; meanwhile, our higher sampling frequency precludes FFT aliasing. While the magnitude spectrum appears symmetric, the phases of the two sides are actually opposites due to the hermitian symmetry of a real signal's spectrum, so we must take care to extract the right half-spectrum if we plan to use a complex reference chirp for which the phase matters. Upon halving the spectrum (through truncation in Matlab), we discard the negative half-spectrum and multiply the basebanded right half-spectrum (shifted about zero frequency) with our reference spectrum, obtaining the same pulse compression seen in the I/Q complex chirp case:

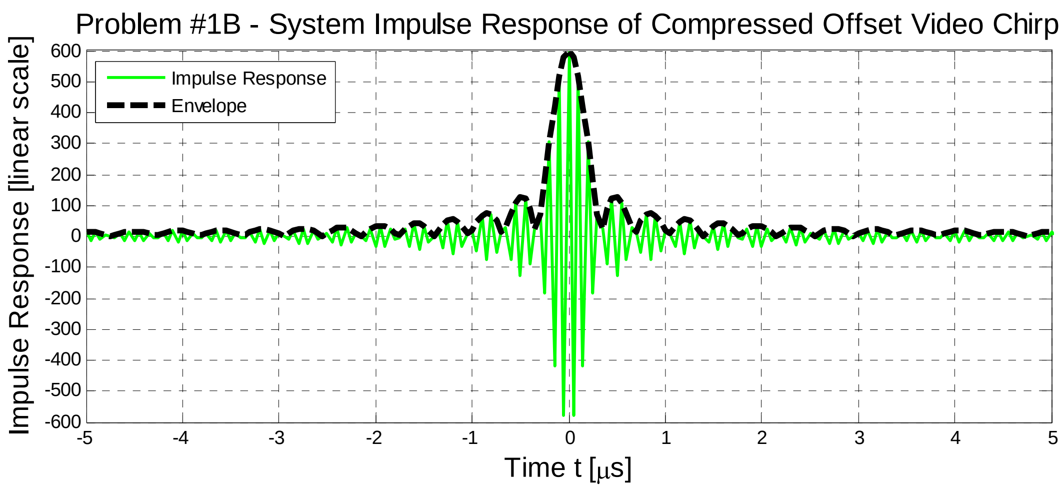

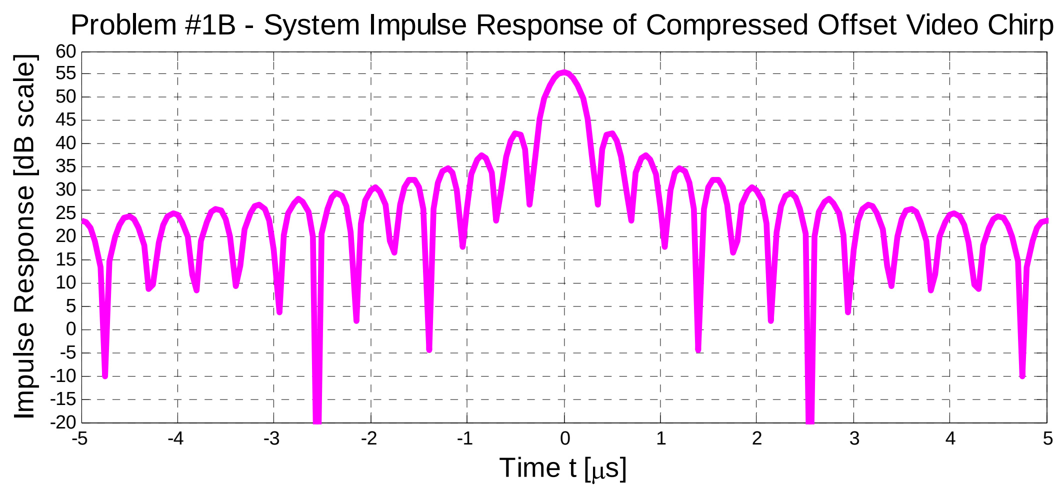

Thus, as expected, the two processing methods – I/Q and Offset Video – produce exactly the same impulse response compressed pulse, in both magnitude and phase. Our choice largely depends on the type of signal our receiver or processor must accept, and other system design specifications. Each one possesses its own advantages. For example, the I/Q system requires a much lower sampling frequency equivalent to the signal bandwidth, but sampling both real and imaginary components will demand two samplers offset exactly 90° in phase, thus placing significant timing and synchronization restraints on our sampling hardware. The Offset Video system, on the other hand, requires twice as much bandwidth and sampling frequency while necessitating more complicated coding (from basebanding the right half-spectrum) as well as care (due to the asymmetry in phase of the two half-spectra), but loosens the system constraints, since we need only one analog A/D converter, which also tends to operate at higher speeds without the need to synchronize with any other sample.

Problem #2 – Peak and Integrated Sidelobe Ratios

To compute the Peak Sidelobe Ratio (PSLR) and Integrated Sidelobe Ratio (ISLR), we could employ the decibel scale plot or simply approximate our impulse response as a sinc function that we can numerically integrate. For the PSLR, we inspect the plot and obtain a PSLR of -13.36 dB ratio when defining

$$PSLR \equiv \frac{\textit{first sidelobe peak}}{\textit{main lobe peak}} \equiv \left[\textit{first sidelobe peak}\right]_{dB} - \left[\textit{main lobe peak}\right]_{dB}.$$

Notice that this definition uses the *first* sidelobe peak rather than the *maximum* sidelobe peak. To compute the ISLR, we either integrate a model sinc function (done analytically in hand calculations), or employ Matlab's numerical integration function `trapz` on the compressed pulse absolute envelope with the definition:

$$ISLR \equiv \frac{\int_{\textit{first null}}^{\infty} g(t)\,dt}{\int_{0}^{\textit{first null}} g(t)\,dt} \equiv \frac{\int_{\textit{first null}}^{\infty} (\textit{compressed pulse})}{\int_{0}^{\textit{first null}} (\textit{compressed pulse})} \equiv \frac{\int \textit{sidelobes to right of main lobe}}{\int \textit{right half of main lobe}}.$$

This definition yields an ISLR of approximately 0.10756 ≈ -9.68338 dB for the compression in Problem #1.

If we first multiply our reference spectrum by some weighting or window function, then we can decrease the PSLR and ISLR, allowing us to obtain higher time resolution and lower noise floors in our compressed pulse. We apply a variety of windows assuming the form $w + (1-w)\cos\left(\frac{2\pi f}{BW}\right)$ and compute the resultant PSLR and ISLR. The optimal weighting factor *w* appears to be 0.57 for PSLR and 0.51 for ISLR:

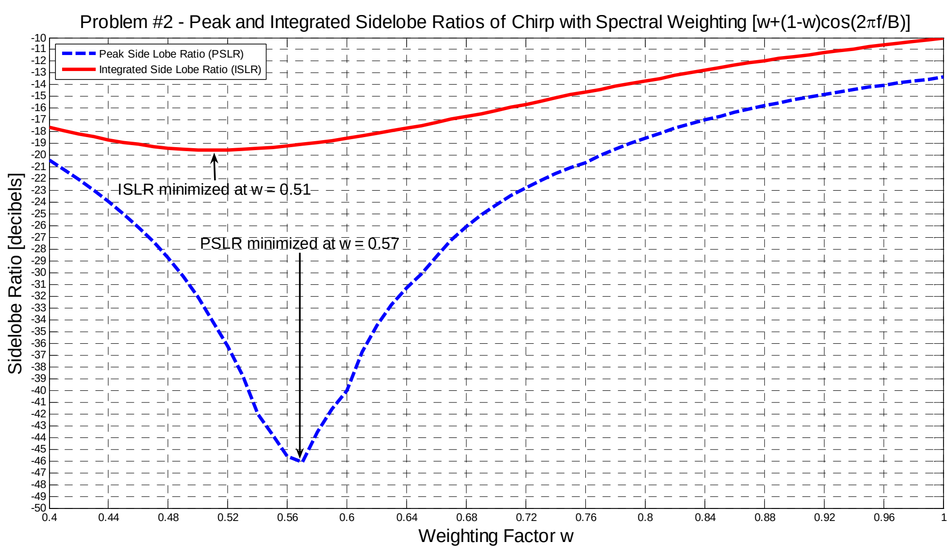

The Hamming window $0.54 + 0.46\cos\left(\frac{2\pi f}{BW}\right)$ uses the average of these two optimal values: w = 0.54, thereby compromising between optimal PSLR (for sharp peak detection) and optimal ISLR (for minimal sidelobe power). Interestingly, weightings around these optima yield sidelobes that increasing in magnitude farther from the main lobe, contrary to the typical decreasing trend. Take 0.54, for example:

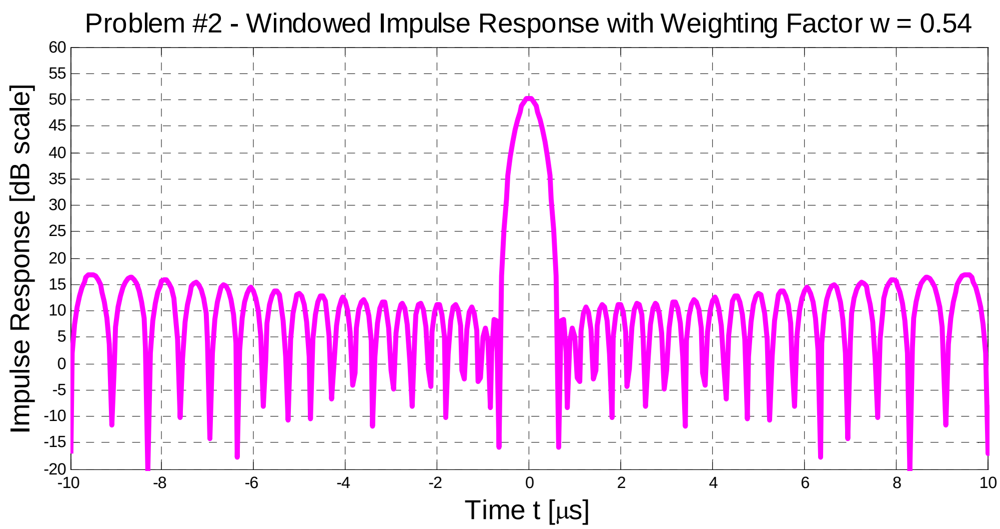

Remark the rise in sidelobe power due to the heavier weighting near the main lobe; cosine windows with intermediate weighting factors no longer average sidelobes arithmetically. As a result, the optimal PSLR differs if we instead define it with the *maximum* sidelobe peak rather than the *first* sidelobe peak:

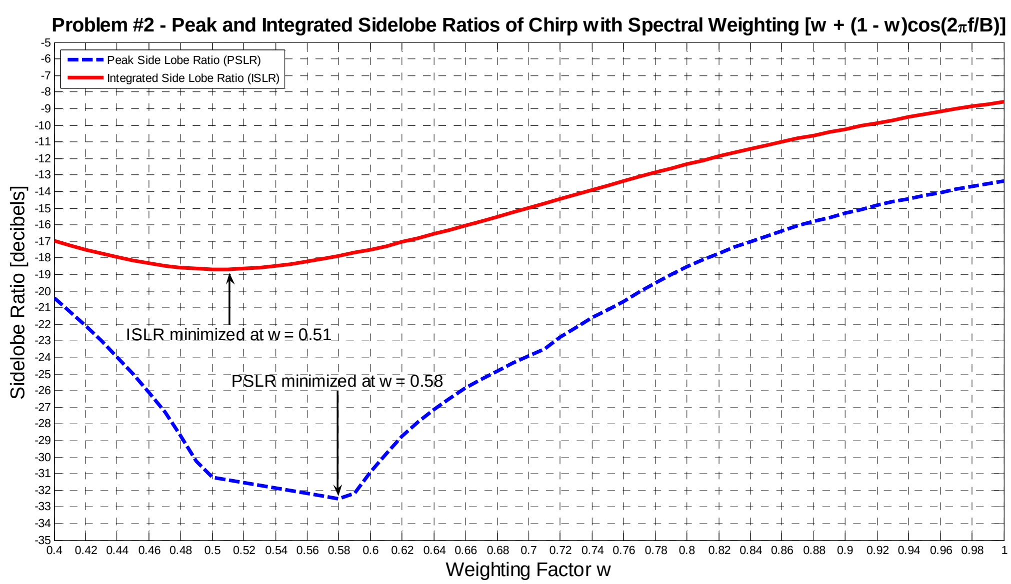

The ISLR attains minimum at the same location, but the PSLR now reaches its optimal value for w = 0.58.

Problem #3 – Peak and Integrated Sidelobe Ratios of Mismatched Chirps

If we overestimate our transmitted chirp signal's slope by even 2% or 5% and employ a matched filter that does not exactly match our chirp, then the resulting pulse compression degrades from the ideal compression; we can see the degeneration of quality in both an increased PSLR and ISLR, indicating that our mismatched filter output has not only higher sidelobes (relative to the main lobe) but also a greater fraction of power in the sidelobes.

Reference Chirp Slope	**1.02×10^{11} (2%)**	**1.02×10^{11} (2%)**	**1.05×10^{11} (2%)**	**1.05×10^{11} (5%)**
Hamming Windowed?	☒	☑	☒	☑
Peak Sidelobe Ratio	-9.92894 dB	-40.0027 dB	-0.8441 dB	-37.6393 dB
Integrated Sidelobe Ratio	-7.27684 dB	-19.3752 dB	2.1595 dB	-20.3705 dB

Mismatched pulse compression decreases the main lobe's relative height above the sidelobes while also placing more power and width in the sidelobes. For example, if we overestimate the signal slope by 2%, we obtain the following compressed pulse:

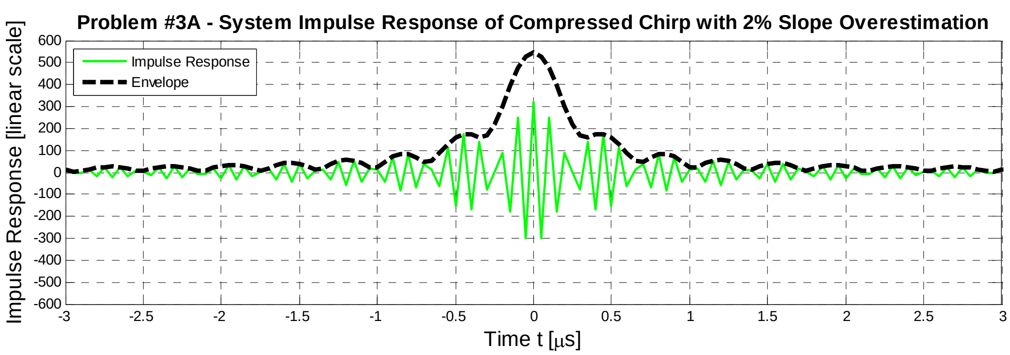

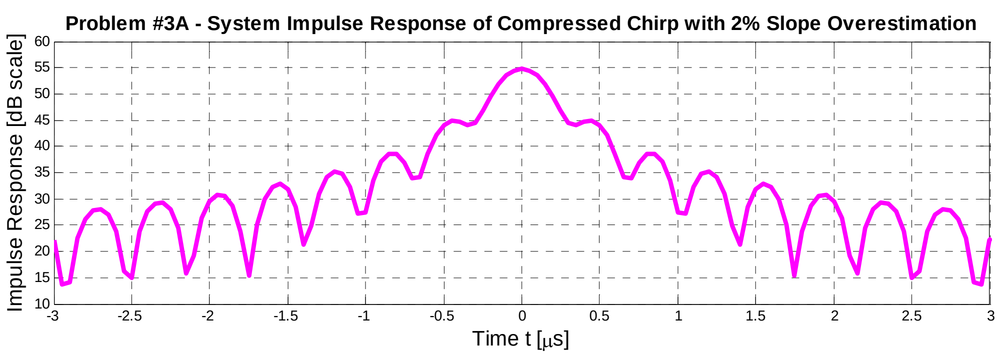

The first sidelobe is much closer to the main lobe, and the main lobe no longer appears narrow because much of the power resides in the sidelobes, which are wide and thick. At 5% slope overestimation, the pulse compression is virtually no longer a "compression":

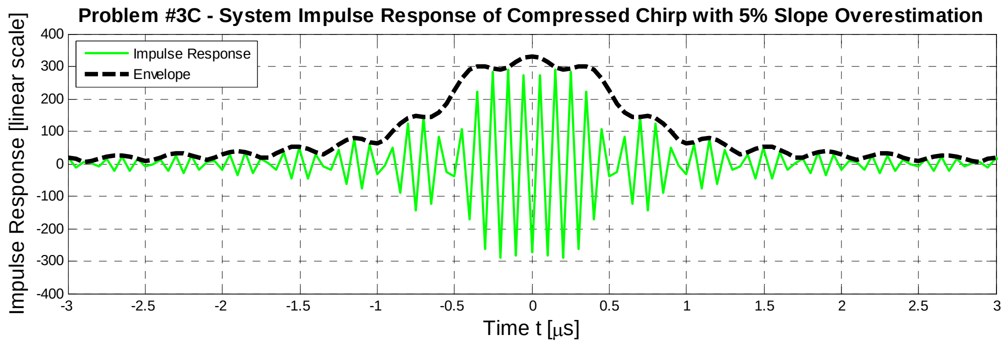

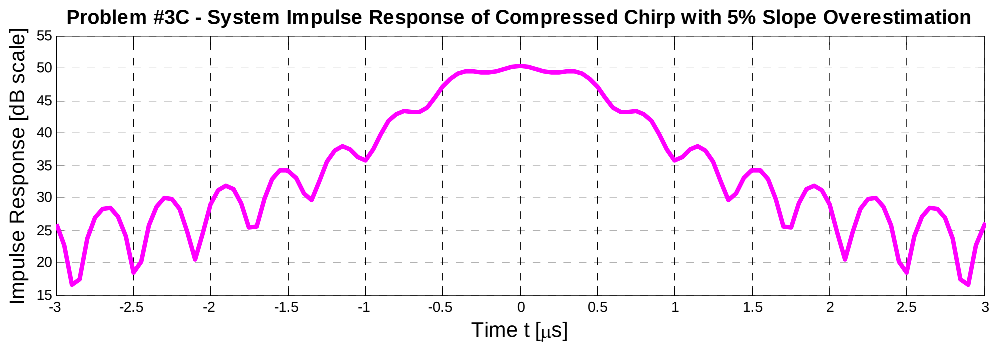

The main lobe does not visibly protrude, as the sidelobes rise to the point of obscuring it, and the effective pulse width now must include sidelobes since the heights are comparable, meaning that the pulse no longer compresses as much as it potentially can. Notice that PSLR is only -0.844094 dB in the 5% slope overestimation case; the sidelobes rise to virtually the same height as the main lobe! Meanwhile, the ISLR is a positive quantity, so the sidelobes contain even more power than the main lobe; this kind of sidelobe dominance would disrupt proper radar imaging with the pulse, as strong and powerful sidelobes would produce intense replica images, precipitating potentially irreversible interference if extremely bright targets appear adjacent to very dark targets.

Even though we can never attain the minimal PSLR and ISLR values achieved with the perfectly matched filter, we can nevertheless improve pulse compression by applying Hamming weighting to the reference spectra before convolution; this weighting still suppresses sidelobes, allowing the main lobes to protrude more prominently. Because we do not know the true chirp frequency slope, we cannot apply a window with the exact equivalent bandwidth *BW* in $0.54 + 0.46\cos\left(\frac{2\pi f}{BW}\right)$, but the resultant error makes virtually no noticeable difference. The improvement is noticeable even in 2% slope overestimation:

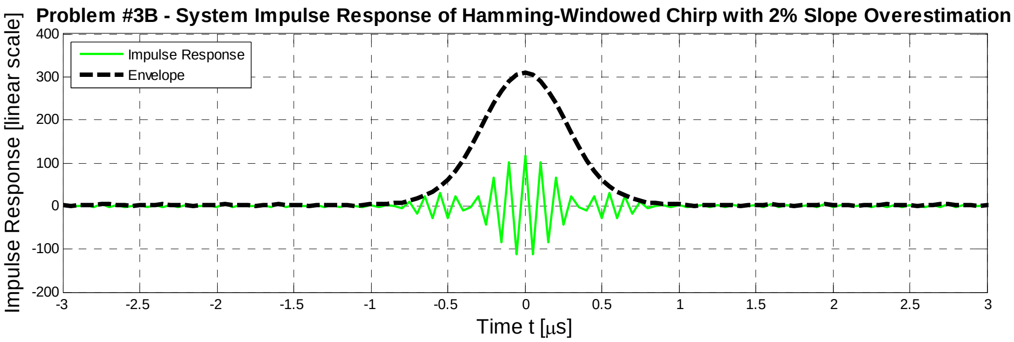

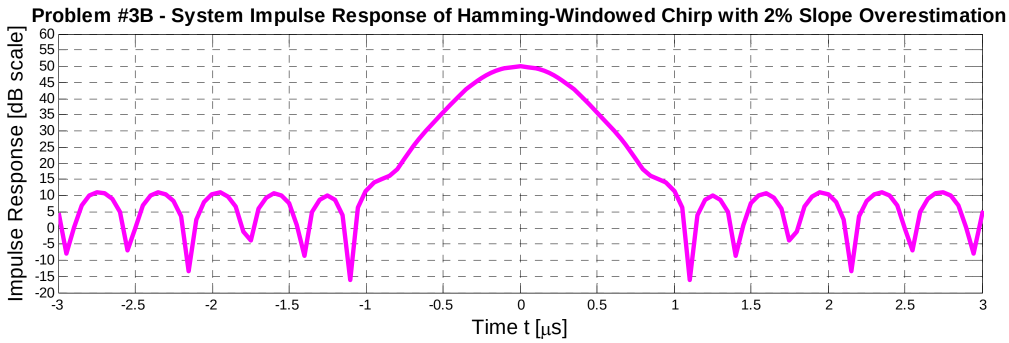

Because the Hamming window comprises a cosine spectrum, its time equivalent comprises a pair of delta functions surrounding the central impulse; as a result, convolution with the Hamming window in time widens the main lobe of the compressed pulse, since the two delta functions, spaced (1/Bandwidth) apart, produce slightly shifted replicas of the main sinc pulse, which add to form the smeared spectrum seen above. In fact, if we examine the main lobe carefully enough, we notice that the first sidelobe appears to have merged into the main lobe; thus, for PSLR and ISLR computations, we ignore this merging and assume that the first sidelobe appears farther out, where the crest is distinctly separate from the widened main lobe body. However, we *cannot* ignore the fact that the increase in PSLR and main lobe protrusion comes at the cost of resolution, since the window widens the main lobe, damaging our ability to resolve backscatter of point targets. The 5% overestimation filter output further accentuates this widening:

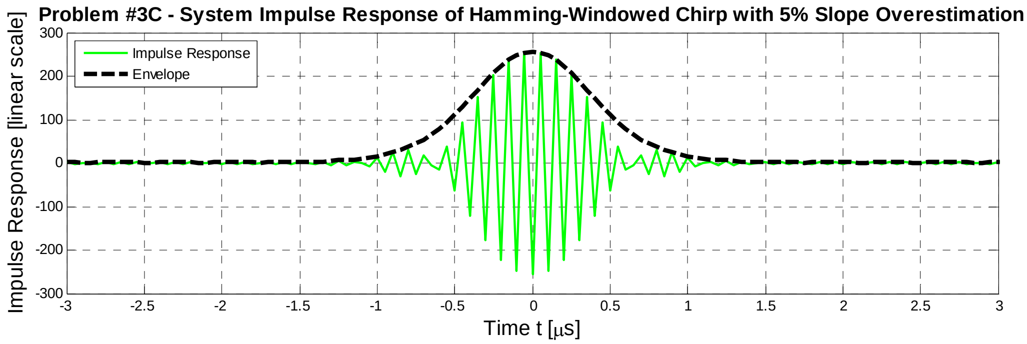

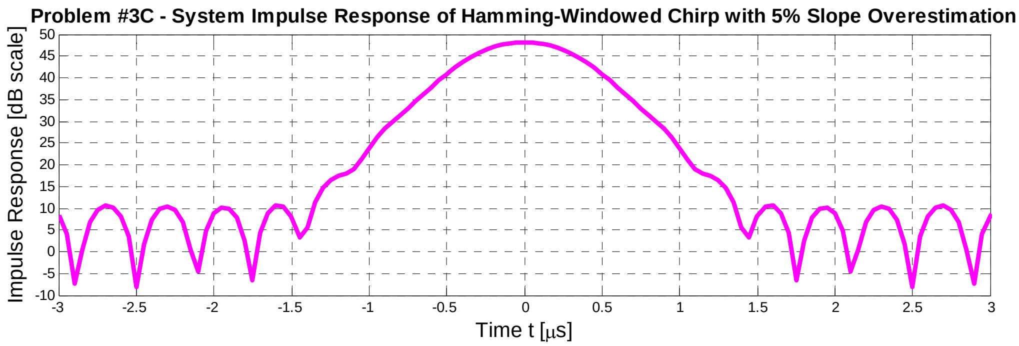

Here, the effective pulse width has increased from approximately 1 μs in the ideal matched reference chirp to nearly 2 μs in the windowed mismatched reference chirp. Most of this coarsening in the resolution arises from the smearing of sinc convolution with the window's time impulses, but we also see merged sidelobes around +/- 1.25 μs as well. However, we do not regret the loss of resolution; in return for lost resolution, the window has granted our main lobe with considerably more prominent protrusion from the sidelobes, which, despite their power, now sit noticeably below the main lobe; recall that these sidelobes rose comparably to the main lobe height before we applied windowing, leaving virtually no choice in this 5% case. Comparing the windowed cases of our two different mismatched reference signals, we see that the windowing also reduces the PSLR and ISLR sensitivity to error; in other words, the PSLR and ISLR remain comparable whether we overestimate by 2% or by 5%, so the resultant pulse changes less with successive error, making performance much more predictable for small deviations from the actual chirp slope.

All in all, we conclude that windowing becomes an absolute necessity in the event of even slight mismatches in reference chirps, for even 2% slope overestimation error results in much higher PSLR and ISLR in the compressed pulse, and hence potential sidelobe interference during imaging. By windowing the reference spectra before pulse compression, we can salvage the height of our main lobe at the cost of approximately doubling the effective duration (and hence resolution) of our pulse.