Christopher Tsai
February 22, 2007
EE 355 – Imaging Radar & Its Applications

Problem Set VI
Range Migration

Problem #1 – The Sub-Aperture Shift Algorithm

In compressing both the range and the azimuth signals with a matched filter, we must estimate the slope or rate of the chirp response to both high accuracy and fine precision. First, we discovered in Problem Set III (Problem #3) that even a slight slope mismatch of 2% between the range compression matched filter and the transmitted range chirp can result in rapid deterioration of Peak Side Lobe Ratio (PSLR) as well as dramatic decreases in the Integrated Side Lobe Ratio (ISLR); if we overestimate or underestimate the range chirp slope by even 5%, our compression no longer boasts the well-defined bandwidth and sharp peak power that range compression can theoretically achieve.

However, despite its importance in range compression, accurate determination of the chirp rate matters most in azimuth compression, where the ultra-fine resolution that the Synthetic Aperture Radar (SAR) compression algorithm offers depends instrumentally and crucially on proper matched filtering of the received chirp; it is this chirp rate that controls the phase of our reference chirp, and ultimately the quality and authenticity of our image. At the most fundamental level, azimuth "slow" time and position are inextricably intertwined with the Doppler frequency spectrum, and hence the phase history of the chirp response that we must match in our reference function.

The sub-aperture shift algorithm gives us an iterative procedure through which we can refine our slope estimate until it matches the desired response. In this problem, we actually know the true chirp slope that we hope to match, but, in general, we can apply the method iteratively, comparing our estimate to the raw data chirp spectra and correcting our estimate upon correlation with the true and reference chirps.

When we correlate the true chirp response with our estimated reference chirp, the mismatch in chirp slope will broaden the main lobe of the filter output, but, if we divide the compressed spectrum into positive and negative frequencies – respectively the left and right halves of the fast Fourier transform (FFT) – then the two "sub-aperture" time responses will peak at different instances in time, displaced by an amount proportional to the chirp slope mismatch between scene and processor. The precise time locations of these peaks is irrelevant, since only time *difference* matters in the auto-correlation of stationary processes; to facilitate our determination of this time difference, we can cross-correlate the two sub-aperture images and locate its peak. Even though the two images arise from two mutually exclusive Doppler sub-bands that partition the true chirp spectrum, each image reaches its own peak at a location dependent on the degree of similarity between reference slope and real slope; nearly matched slopes will peak at similar locations and hence yield smaller peak separations, while disparate slopes will result in very different sub-aperture images. The resultant cross-correlation of sub-aperture images turns estimation error calculation into simple peak detection, as our algorithm output demonstrates for a variety of slope mismatches:

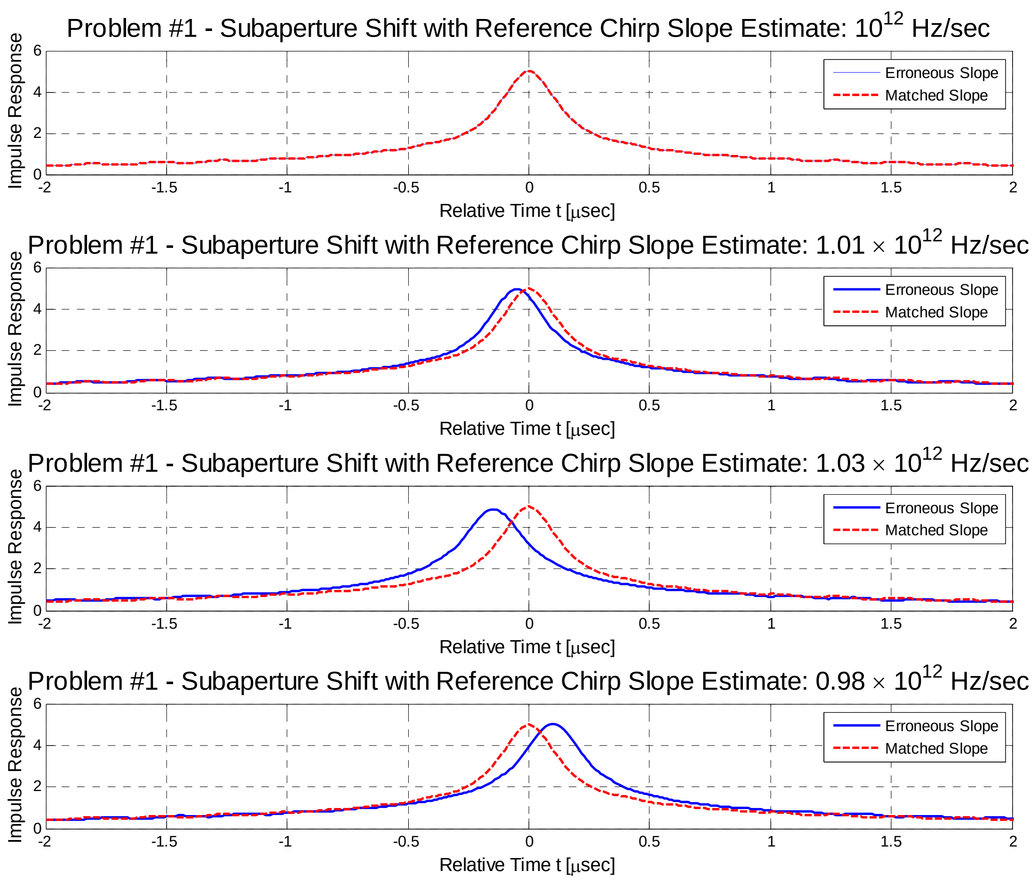

The mismatched correlations actually result in *complex* images because the mismatched slope results in mismatched chirp phases, leading to a matched filter whose complex conjugate no longer completely removes the phase of the true chirp; normally, the filter's conjugated phase would remove all complexity from the received signal, leaving a purely real image, but mismatching will introduce residual phase in the compression. However, since we cannot readily view both a real and an imaginary part, we instead graph the magnitude and locate the peak. Remark that a perfectly matched filter resulting in the dotted compression yields no sub-aperture shift, since both the positive and negative true spectra overlap completely with the reference spectra, independent of the sub-aperture; thus, upon comparison with this ideal, our mismatched chirp compressions all deviate from the perfect correlation in the sense that the sub-aperture cross-correlations lie away from the time origin, indicating a difference in the alignment between negative frequencies and positive frequencies when the chirp slope estimate deviates slightly from the actual value. We convert this time deviation into a slope mismatch estimate by assuming approximate stationarity in chirp phase, allowing us to relate the cross-correlation time difference to the chirp slope difference between actuality and estimate through $\Delta s = \frac{2\Delta t \cdot s^2}{\text{bandwidth}}$. In our problem, we know the exact value of true chirp slope *s*, but, even if we did not, we could approximate the slope error by using the slope estimate instead of *s*, allowing us to iteratively improve our estimate with each cross-correlation. When we do know the true chirp slope, then the sub-aperture shift algorithm allows us to correct our estimate in a single step, as the resultant error calculations reveal:

Reference Chirp Slope	**True Chirp Slope**	**Computed Pixel Offset**	**Computed Chirp Slope Error**
[Hertz/second]	**[Hertz/second]**	**[pixels]**	**[Hertz/second]**
1×10^{12}	1×10^{12}	0	0 (matched)
1.01×10^{12}	1×10^{12}	-5	-1×10^{10} (overestimate)
1.03×10^{12}	1×10^{12}	-15	-3×10^{10} (overestimate)
0.98×10^{12}	1×10^{12}	+10	$+2 \times 10^{10}$ (underestimate)

Notice how the *sign* of the offset and slope error even reveals whether we overestimated or underestimated the chirp slope; estimates that exceed the actual slope result in negative errors that differ in polarity from positive errors arising from slope underestimates, which possess greater deviation along the positive frequencies. In both cases, we can add the errors to our estimates to rectify the mismatch.

The discretized bin index of the cross-correlation peak with respect to the origin yields the pixel offset, which would appear in our image as a shift in the compressed impulse response in any direction along which our compression matched filter fails to match the true response in that direction; the pixel offset would appear in azimuth if we mismatched the azimuth response chirp rate, and in range if we misrepresented the range chirp slope.

Problem #2 – Range Migration and the Azimuth Ambiguity

The simlband.dat file demonstrates the effect of range migration not only in range curvature and bin-to-bin energy propagation, but also along the azimuth direction, where the azimuth ambiguity in pulse repetition frequency (PRF) exacerbates the image distortion and degraded resolution that range migration correction might otherwise easily remove.

Before we even compress the raw data along range, we can detect migration in the image values as the slight slanting of the target brightness; without even resolving the scattering object, we can tell that its energy does not align properly along vertical bins:

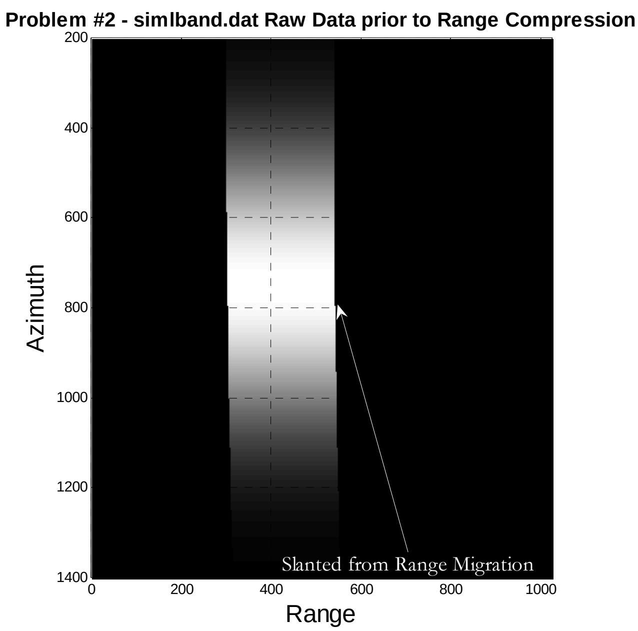

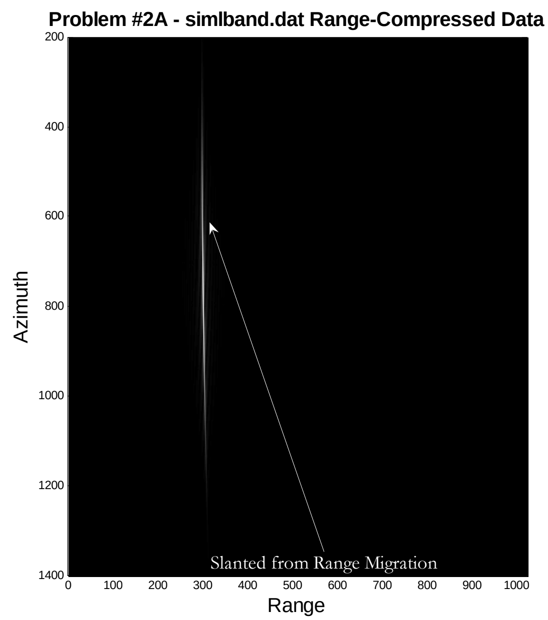

The grid lines reveal a slight "migration" or "walk" of bright backscatter into bins that should otherwise display pure darkness. At the edges of the image, we expect range bins with no data values to be completely devoid of brightness, but some of the energy from the target migrates into the supposedly dark bins due to the relative motion of the target relative to our imaging radar. This motion manifests itself even more clearly in the range-compressed image; we can easily perceive the slanting of the thin line of range-compressed data from the lower-right to the upper-left; instead of showing backscatter from a single point target in a single range bin, the image exhibits brightness in several adjacent columns (or range bins), with slight shadows that appear even more prominently after transforming the data into azimuth:

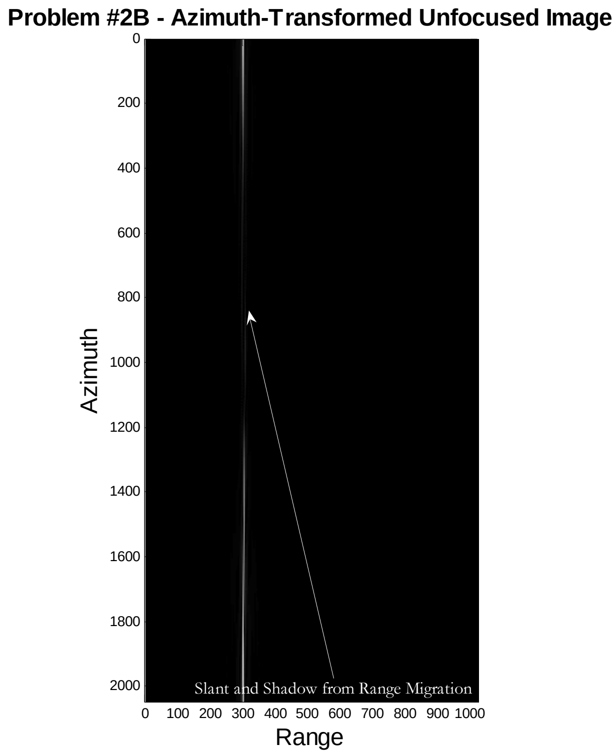

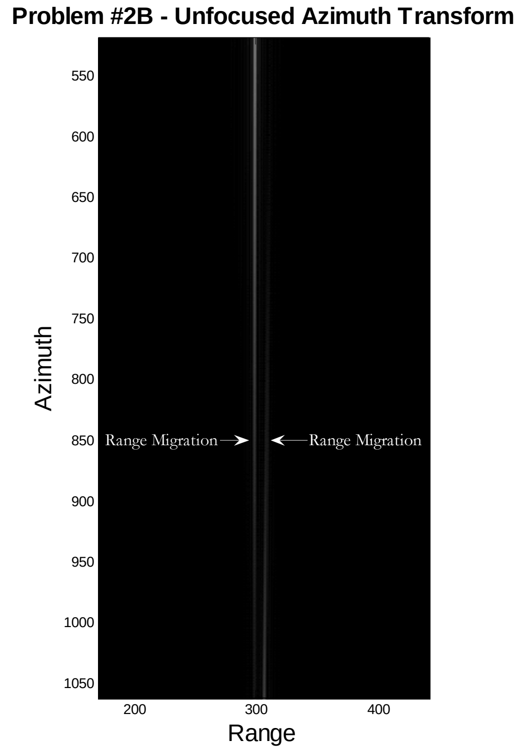

Notice that the thin line in our range compression now splits into two – or perhaps even more – thin lines representing two (or more) range bins in our image. Thus, because the backscatter has propagated or "range-walked" into a second bin, the azimuth transform erroneously interprets it as a second scatterer, and the unfocused image now shows two distinct returns at two distinct range bins, in addition to the ever-present bin slanting. These shadows represent only the beginning of the manifestations of range migration along the azimuth; these visible imaging problems, which we see here as mere brightness duplication or range curvature, worsen as the resolution improves. For example, when we upgrade our unfocused SAR processing – the simple Fourier transform of our azimuth samples – to focused SAR processing, theoretically pressing our currently stretched azimuth resolution to its limit, range migration rears its head as even greater adversity.

In order to improve resolution along the azimuth direction, we must determine the approximate Doppler centroid of our imaging geometry so that we can compress the azimuth image by correlating azimuth samples with a matched filter with the same quadratic phase as our return. Assuming that our azimuth beamwidth is short enough so that the Doppler center frequency at the center of any range bin can apply to the entire range bin – f_{DC} depends only weakly on azimuth slow time or position – we can locate an approximate center Doppler frequency associated with the squint angle. The average of these center frequencies, which also happen to be the carrier frequencies for our azimuth chirps, follows from either the average phase change algorithm or simple graphical analysis of the average magnitude spectrum:

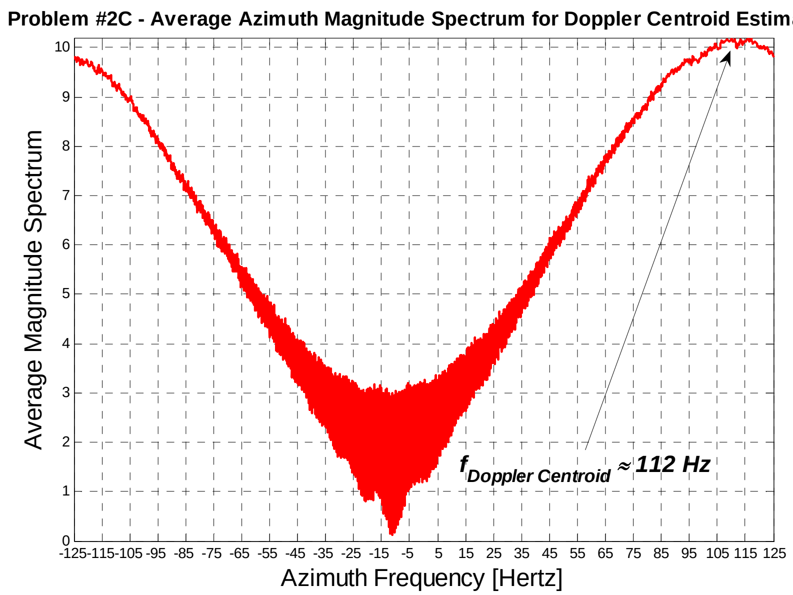

We observe an azimuth Doppler centroid frequency of approximately 112 Hertz, but because we discretely transform (FFT) our azimuth samples collected at the PRF, the resultant spectrum is periodic with period PRF, meaning that our determination of the Doppler centroid might not represent the true centroid since we view the spectrum over only one period. Our value for f_{DC} is correct only by an integral number of PRFs, so the true Doppler centroid could assume any one of a countably infinite set of values, spaced by the PRF of our azimuth sampling (250 Hz):

$$\vdots$$
$$f_{DC} \approx 112.4058\text{ Hz} - 3(\text{PRF}) \approx 112.4058\text{ Hz} - 3(250\text{ Hz}) \approx -637.5942\text{ Hz}$$
$$f_{DC} \approx 112.4058\text{ Hz} - 2(\text{PRF}) \approx 112.4058\text{ Hz} - 2(250\text{ Hz}) \approx -387.5942\text{ Hz}$$
$$f_{DC} \approx 112.4058\text{ Hz} - 1(\text{PRF}) \approx 112.4058\text{ Hz} - 1(250\text{ Hz}) \approx -137.5942\text{ Hz}$$
$$f_{DC} \approx 112.4058\text{ Hz} + 0(\text{PRF}) \approx 112.4058\text{ Hz} + 0(250\text{ Hz}) \approx 112.4058\text{ Hz}$$
$$f_{DC} \approx 112.4058\text{ Hz} + 1(\text{PRF}) \approx 112.4058\text{ Hz} + 1(250\text{ Hz}) \approx 362.4058\text{ Hz}$$
$$f_{DC} \approx 112.4058\text{ Hz} + 2(\text{PRF}) \approx 112.4058\text{ Hz} + 2(250\text{ Hz}) \approx 612.4058\text{ Hz}$$
$$f_{DC} \approx 112.4058\text{ Hz} + 3(\text{PRF}) \approx 112.4058\text{ Hz} + 3(250\text{ Hz}) \approx 862.4058\text{ Hz}$$
$$\vdots$$

In brief, the true Doppler centroid frequency could be any value satisfying $f_{DC} \approx 112.4058\text{ Hz} + k\cdot(\text{PRF})$, where $k \in \mathbb{Z}$ is an integer. All such values are consistent with the average spectrum that we observe in the data because the spectrum of a discrete-time (periodic) signal is periodic, therefore reaching a peak frequency in every spectral period of pulse repetition frequency (PRF = 250 Hz):

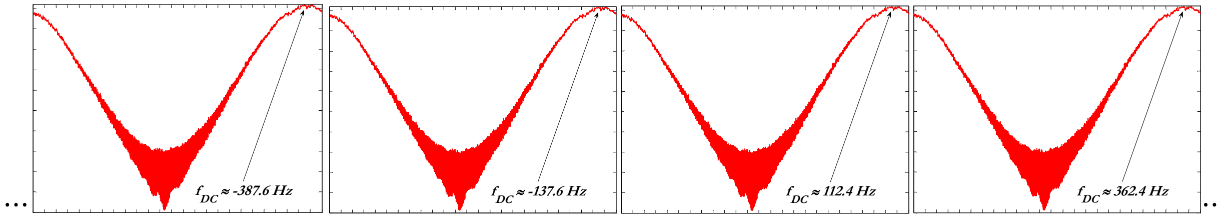

Assuming this center frequency holds across the entire azimuth beamwidth so that we need not re-compute its value in each bin, we construct a matched filter chirp with a center frequency at the Doppler centroid for every range bin; within each bin, the range to the aircraft differs slightly, so we must adjust the chirp rate of our reference chirp to correspond to the differing ranges, which lead to differing phase histories, and ultimately to different chirp rates in our received azimuth signal sample. The effective pulse length and Doppler centroid actually differ from range bin to range bin as well, but we assume that the Doppler centroid frequency does not deviate noticeably far from the average centroid previously computed and therefore adjust only the chirp rate and effective duration from range bin to range bin. Indeed, adding a line to compute a new carrier frequency in each range bin increases computation time with comparably little noticeable resolution gain in the final image. Compressing 80% of the azimuth beamwidth by correlating the azimuth spectrum with these range-varying azimuth reference chirps, we obtain the following migration-blurred impulse response:

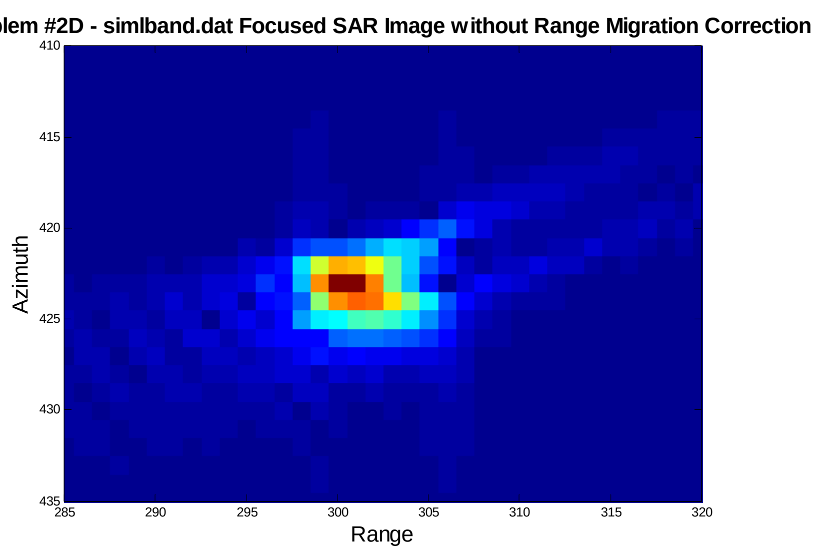

We first note that, although we have compressed the previously unfocused line down to a spot resembling our target reflection, significant blurring still occurs, not only about the center of the spot but also in several nearby range bins, which appear with a lighter tint of blue in reflection of the migration of scatterer energy across several adjacent range bins. Furthermore, despite the improvement in azimuth resolution, we have exacerbated the amount of migration (bin-wise) from the unfocused SAR image, since the spot now spreads over approximately 30 pixels across range and azimuth; range migration grows much more evident in the focused SAR impulse response because two parameters that we used to compress the azimuth – chirp rate f_R and Doppler centroid f_{DC} - depend strongly on range. Unlike the propagation or "range walk" values that insert themselves into one or two entries in the unfocused SAR Fourier-transformed azimuth spectra, range migration now creates a fundamental mismatch in chirp rate and center frequency, henceforth affecting the way we compress *every* point in the azimuth direction. Much like we discovered during our exploration of PSLR in range chirp mismatch (and providing the incentive for sub-aperture shift), slight discrepancies in reference chirp slope result in matched filter outputs with drastically and noticeably reduced PSLRs and ISLRs, explaining the degenerate impulse response resolution (blurring) that we witness here; the sharpness of the main lobe has dulled, and the comparably high side lobes have spread the backscatter energy across several range and azimuth bins. Thus, whereas range curvature appears merely as slants and shadows in unfocused SAR images, the strong dependence of chirp rate and center frequency on range position makes slight range migration errors much more noticeable in focused SAR images; even a slight migration of a few range bins could introduce enough deviation in our chirp rate calculation that the improperly matched filter correlation spreads our migrated energy much more widely in azimuth.

Henceforth, before we irrecoverably compress range-migrated points in the azimuth with improperly designed reference chirps, we must correct range migration. We employ the cut-and-paste algorithm for the `simlband.dat` image, but other, more sophisticated interpolation schemes also exist. However, as we begin to rectify range migration, we quickly realize that the amount of migration and hence our correction scheme depend heavily on the precise true Doppler centroid. Whereas the azimuth ambiguity (PRF periodicity) played no visible role in azimuth compression since we were already working with periodic spectra along the azimuth, our choice of f_{DC} (among all the possible values) numerically affects how we correct range migration; selecting the incorrect PRF ambiguity will result in potentially disastrous range migration rectification, possibly producing an even further distorted image as our cut-and-paste proves meaningless for an incorrect value of f_{DC}. Thus, only *one* value of f_{DC} will yield proper range rectification, and we must find it.

Radar veterans such as Barnett (1969), Chang, and Curlander (1992) have developed a myriad of algorithms that can precisely determine the Doppler centroid among the countable infinitude of possible PRF ambiguities; relying on the Chinese remainder theorem, some of these solutions chain backwards through Euclid's algorithm to solve a congruence relating spectra resulting from multiple PRFs. Along the same lines, we proceed empirically, observing four range-migration-corrected and azimuth-compressed impulse responses, processed at four different values of f_{DC} that we space by integer multiples of the PRF:

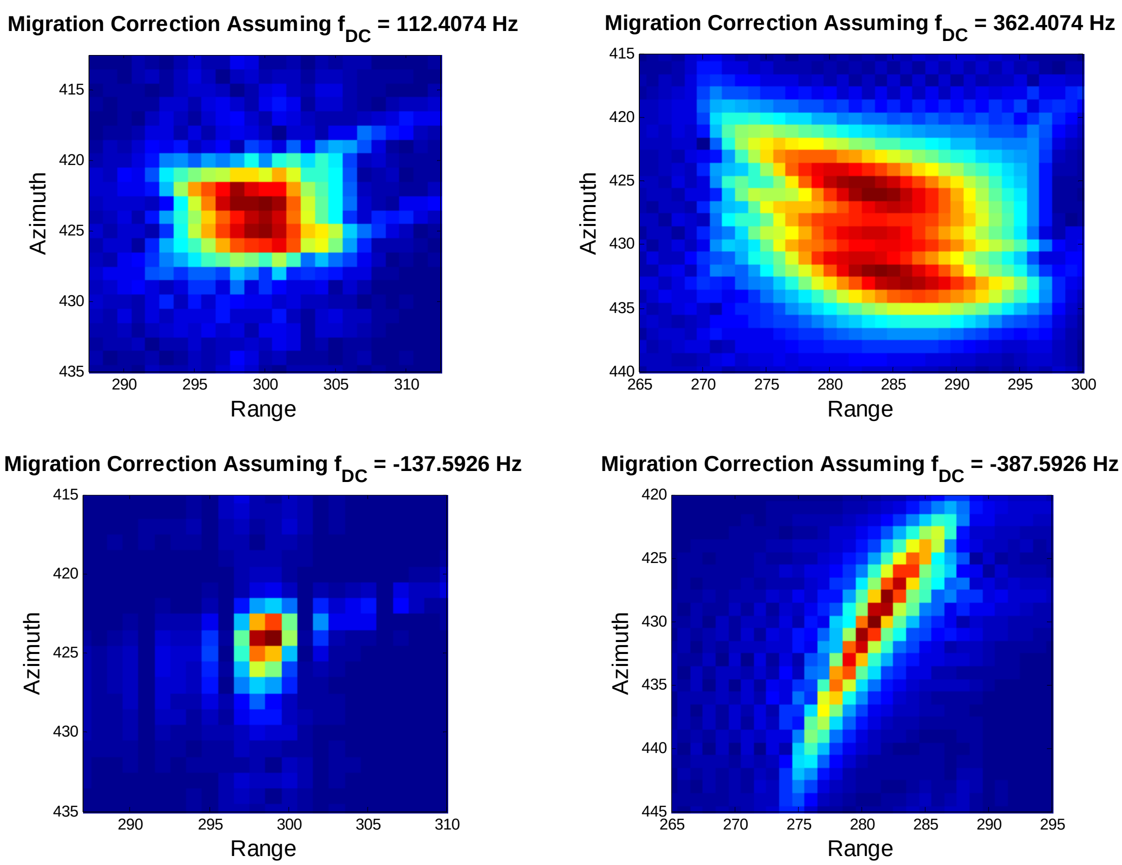

From the spread or sharpness of the bright spot in the impulse response, we can immediately discern that the true Doppler centroid frequency is $\boxed{\boldsymbol{f_{DC} \approx -137.5926\ \mathrm{Hz}}}$, or one PRF below the base centroid we computed previously from the average azimuth spectrum and average phase change algorithm. Beginning with the baseband centroid between $-\frac{\mathrm{PRF}}{2}$ and $\frac{\mathrm{PRF}}{2}$, we see a relatively distorted and even blurrier point than before migration correction. As we add integral multiples of the PRF to our Doppler centroid estimate, we notice that the bright spot splits further, occupying more azimuth bins and hence indicating that our azimuth compression must suffer from a chirp rate and/or centroid frequency mismatch. Moving below the baseband Doppler centroid estimate to negative frequencies, we observe that negative Doppler frequency estimates such as -387.5926 Hz and lower cause the spot energy to spread further along the azimuth, stretching the central bright reflectivity as it curves and migrates into other azimuth lines and range bins; we clearly have not rectified range migration with these cut-and-paste efforts, as we have exacerbated the mismatch and distorted the impulse response further.

The range migration correction made under the estimate of $f_{DC} \approx$ -137.5926 Hz clearly yields the best compression, since the impulse response reduces to the smallest point with the weakest side lobes and least amount of blurring of all our attempts. Because we precisely sought to compress the azimuth to as fine a resolution as possible without spreading energy into other range bins, this impulse response represents our most successful removal of range migration. Note, however, that the point still has imperfect resolution, likely because we still have slight mismatch in our chirp rate, and also since we assumed one Doppler centroid frequency for all range bins and azimuth lines within the beamwidth, whereas the true azimuth chirp carrier varies from range bin to range bin. Nevertheless, by performing sub-aperture shift on the chirp rate within each range bin, invoking more complicated sinc interpolation algorithms, and calculating a new Doppler centroid in each range bin, we can improve the sharpness of our impulse response and approach the theoretical limit on SAR resolution; we omit these additional refinement steps here since they involve greater computational complexity and intensity than Matlab can adroitly handle.

Please see hand-written section for ***Problem #3: Effective Velocity***.