Christopher Tsai
February 5, 2007
EE 355 – Imaging Radar and Its Applications

Problem Set IV
Azimuth Processing and Unfocused SAR

Problem #1 – Azimuth Spectra and the Doppler Centroid

After removing the header (412 bytes) and combining adjacent pairs of the remaining 9806 raw data columns to form complex numbers, we obtain 4903 range bins:

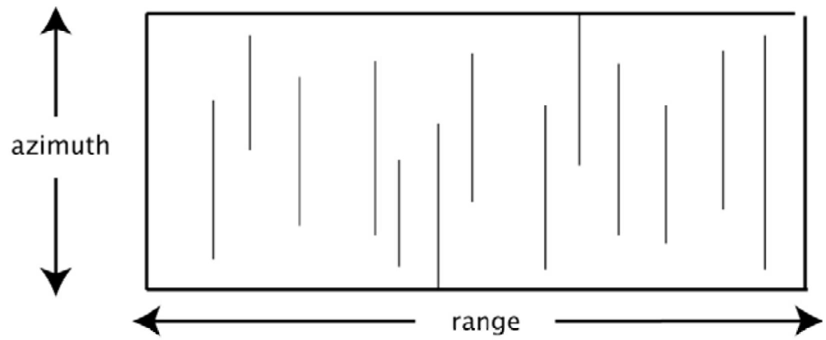

Even though we can extract azimuth spectra from the raw data, we first compress the range by correlating each range column with a matched reference chirp having the same slope (4.189166×10^{11} Hz/sec), pulse duration (37.12 μs), and center frequency (0 Hz) as the European Remote-Sensing Satellite (ERS) transmitted signal. Row by row, we move across the azimuth and compress along the range. Because the number of range bins (4903) is not a power of two, we zero-pad our range dimension to $8192 = 2^{13}$ so that we can apply the Fast Fourier Transform (FFT); after multiplying each range spectrum with our reference spectrum, we inverse-transform to return to the time domain. However, having implemented correlation in the discrete-time domain, we must eliminate extraneous values at the end of the inverse-transformed bins, so we truncate our range-compressed lines after 4903 values and, furthermore, remove $\lfloor f_s \tau \rfloor$ samples from the transform to preclude wrap-around effects from the length of our artificially generated reference chirp. Thus, we obtain 4200 complex-valued range bins, each of which supplies 10,100 samples in the azimuth direction (which we transform):

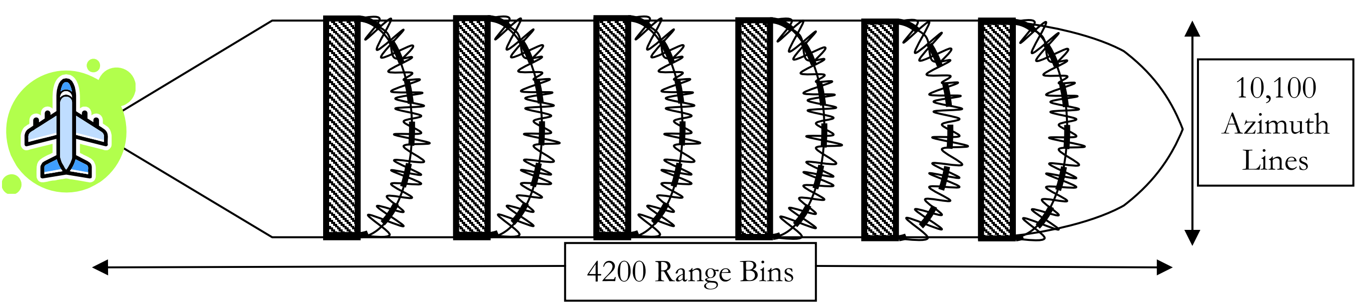

Because the azimuth samples in each range bin could arise from any location along the azimuth beam width, we cannot extract any positional information from a plot of range return and azimuth return in time; instead, we apply the FFT to samples in each range bin to obtain one azimuth signal *spectrum* from each range bin. Each range bin certainly contains enough time samples (10,100) to provide sufficiently fine frequency resolution, regardless of where the samples occur along the azimuth beam width, so the resultant spectrum should accurately portray the distribution of frequency content that we receive from signals in the azimuth beam width. The Doppler frequency content from varying positions along the azimuth directly determine which frequencies exist in this spectrum and how spectrally strong or dense they appear, so we can create our image along the azimuth not from the haphazardly drawn time samples but from the azimuth magnitude spectrum; the relative proportion of Doppler frequencies will reveal the relative azimuth locations and reflectivity of our scattering targets along the beam width, since the Doppler frequency directly relates to position: $f_D = -\frac{2v}{\lambda}\frac{x}{r} = -\frac{2v}{\lambda} sin\theta \cdot sin\varphi$, where x represents the azimuth coordinate.

In tapping the FFT, we must round the number of azimuth samples to the nearest power of two, either ceiling (10,100 → 16,384) and zero-padding, or flooring (10,100 → 8192) the number of samples we transform. I chose to floor (round down) to expedite computation; the choice hardly changes the resulting spectra since 8192 samples already constitute a sufficiently large pool.

For the purposes of this problem, we seek not an azimuth image but the offset squint angle (and therefore offset Doppler frequency) of our beam antenna. In order to facilitate processing and production of a properly oriented image, we will eventually need to correct for squinting to center our azimuth spectrum on zero; proper steering therefore requires knowledge of the exact offset Doppler "centroid" frequency, so we know how to cancel its effect in an image. To ascertain this Doppler centroid, we incoherently average the azimuth spectra from all 4200 range bins to obtain one single magnitude spectrum representative of spectral response from all targets in the beam width. As variations in reflectivity across the swath tend to cancel when averaged over large areas, the *average* magnitude spectrum remains relatively immune to scattered bright spots that might distort certain frequencies in any *individual* magnitude spectrum. Following `fftshift` in Matlab, the average magnitude spectrum appears offset by the Doppler centroid frequency, which arises from off-center antenna squinting:

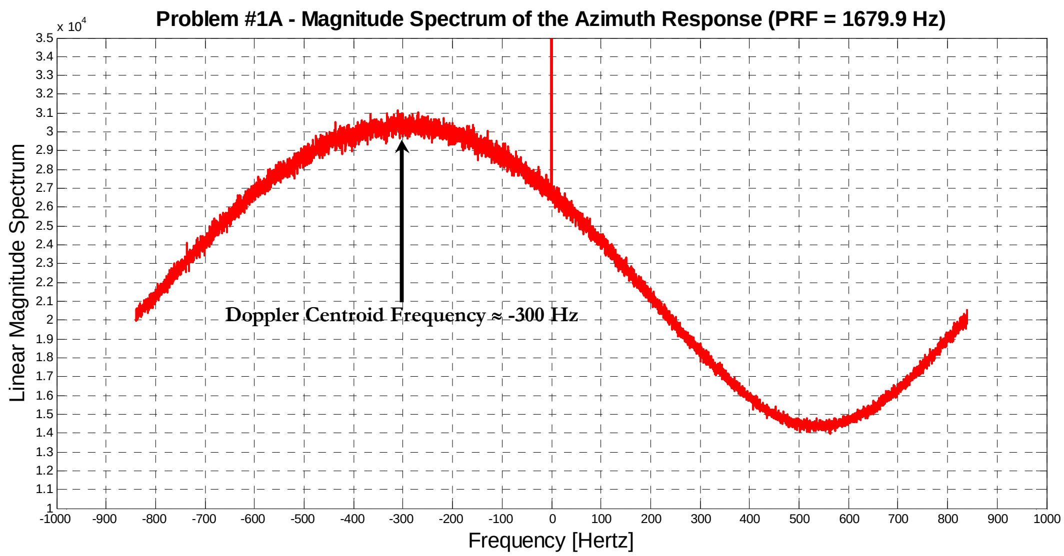

The azimuth spectrum offset frequency appears to be approximately -300 Hz, since we know that the antenna pattern attains its maximum at the center of the beam, where it sees the most backscatter. Thus, without any squint, or if we properly steer the spectrum by the Doppler centroid, the magnitude spectrum should peak at zero frequency:

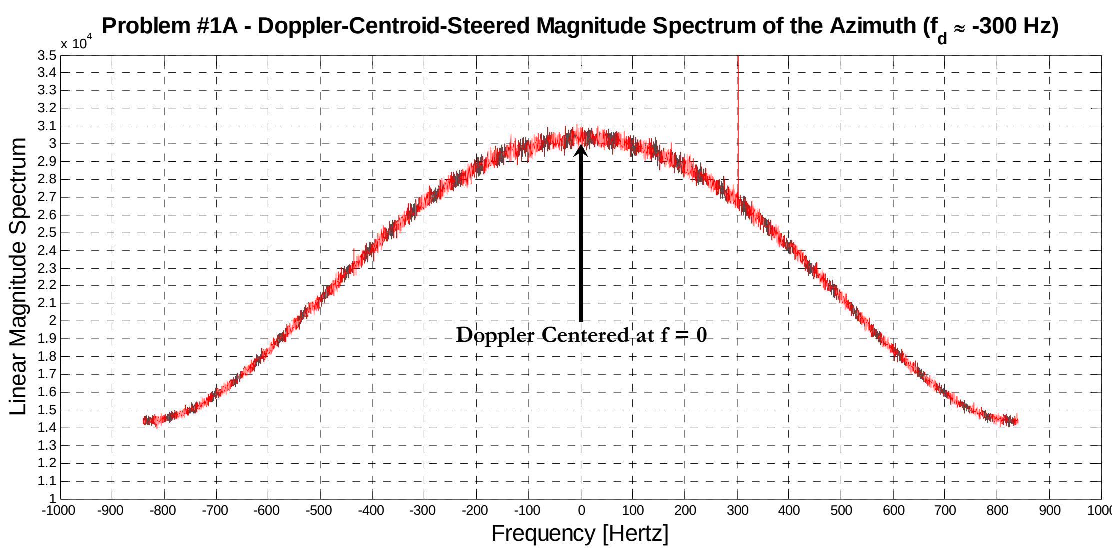

Indeed, we obtain this spectrum if we add a frequency shift of f_D = -300 Hz, ascertaining that our antenna squint angle did, in fact, produce a Doppler centroid around -300 Hz, which we have now successfully offset.

As further verification of this centroid frequency, we can employ the average phase shift algorithm; by computing the net complex difference between samples in each range bin and averaging its phase across all range bins, we can determine the offset squint angle (and hence offset frequency). First, within each range bin, we compute the total phase difference between samples by calculating the complex difference between all adjacent samples and obtaining the phase of the sum of differences:

$$\Delta\varphi_{\text{range bin i}} = \tan^{-1}\frac{\text{Im}\{\sum_{\text{azimuth bin j}} r(i,j)\cdot r^*(i,j-1)\}}{\text{Re}\{\sum_{\text{azimuth bin j}} r(i,j)\cdot r^*(i,j-1)\}}$$

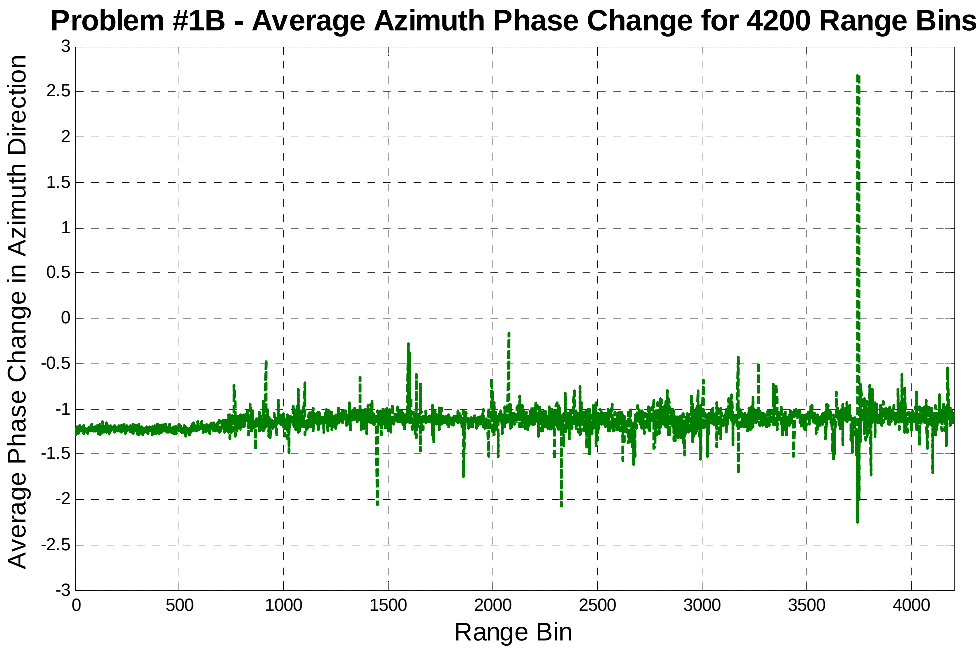

Once we have a phase difference for each range bin, we again average that value across all range bins to suppress the effect of data outliers, such as extremely bright spots (high reflectivity). Because a phase shift of 2π corresponds to a net frequency shift by the pulse repetition frequency (PRF), we can relate the *average phase difference* in our azimuth samples to the *average Doppler frequency shift* that our squinted antenna imparts to the azimuth spectrum using

$$\frac{\sum_{\text{range bin i}} \Delta\varphi_i}{N_{\text{range bins}}} = \frac{2\pi f_D}{\text{PRF}}$$

$$f_D = \frac{\text{PRF}}{2\pi}\cdot\left(\frac{\sum_{\text{range bin i}} \Delta\varphi_i}{N_{\text{range bins}}}\right).$$

Since the average Doppler frequency shift in the azimuth spectra arising from this phase difference is the same Doppler frequency shift that we observed in the visual average in azimuth – both of them derive from the same set of azimuth samples – our centroid frequency calculations should match. Indeed, they do; with the average phase change algorithm, we obtain $\boxed{f_D \approx -302.452026\text{ Hz} \approx -300\text{ Hz}}$, just as we observed in the averaged magnitude spectrum.

Problem #2 – An Unfocused SAR Processor

We begin by evaluating the parameters of the unfocused processor for our ERS radar:

$$Chirp\ rate: s = 4.189166\ \times 10^{11}\ Hz/s$$

$$Pulse\ duration: \tau_p = 37.12\ \mu s$$

$$Sampling\ frequency: f_s = 18.96\ \times 10^{6}\ Hz$$

$$Pulse\ repetition\ frequency: PRF = 1679.9\ Hz$$

$$Minimum\ range: r_0 = 830{,}000\ \mathrm{m}$$

$$Azimuth\ velocity: v_x = 7550\ m/s$$

$$Wavelength: \lambda = 0.0566\ m$$

$$Antenna\ length: l = 10\ m$$

$$Azimuth\ resolution: \delta_{Az} = \sqrt{\lambda r_0} \approx 216.744\ m$$

$$Pulse\ spacing = \frac{v_x}{PRF} \approx 4.4943\frac{m}{pulse}$$

$$Minimum\ burst\ pulses = \frac{\delta_{Az}}{Pulse\ spacing} \approx 48.226\frac{pulses}{burst} \rightarrow 64\frac{pulses}{burst}$$

$$Azimuth\ beamwidth = \frac{\lambda r_0}{l} \approx 4697.8\ m$$

$$Pulse\ repetition\ interval = \frac{Azimuth\ beamwidth}{Azimuth\ velocity} \approx 0.6222\ s$$

$$Frequency\ resolution: \delta_f = \frac{PRF}{Burst\ pulses} \approx 26.2484\ Hz$$

$$Pixel\ spacing: \Delta x = \frac{\lambda r_0 \delta_f}{2v_x} \approx 81.6622\frac{m}{pixel}$$

$$Burst\ length\ [m] = (Burst\ pulses) \times (Pulse\ spacing) \approx 287.6362\ \frac{m}{burst}$$

$$Patch\ spacing = \frac{Burst\ length\ [m]}{\Delta x} \approx 3.52227\ pixels$$

$$Number\ of\ patches = \left\lfloor \frac{Number\ of\ azimuth\ lines}{Minimum\ burst\ pulses} \right\rfloor = 157\ patches$$

$$Azimuth\ lines\ in\ image = Burst\ pulses + round[(Number\ of\ patches - 1) \cdot (Patch\ spacing)\ ] = 613$$

Our final unfocused Synthetic Aperture Radar (SAR) image will comprise the superposition of 157 patches along the azimuth swath, with each patch shifted from the others in azimuth by 3 or 4 (3.52227) pixels. Thus, we devise a procedure that will operate identically on each patch with the exception of its placement in the final image matrix.

Accepting range-compressed data from the previous problem, we process one patch – one burst of pulses or samples – at a time. Since each burst along azimuth comprises 64 chirp pulses and hence 64 azimuth samples, we extract 64 lines of data to process at once. Before we apply the FFT along each range bin, we multiply each complex data value with a complex exponential of unit magnitude and phase that is the conjugate of the Doppler centroid, therefore *steering* the frequency domain in each data point away from the Doppler centroid frequency. As a result, the magnitude spectra that we obtain upon transforming the samples along the azimuth lines will be centered about zero, allowing us to construct the image much more intuitively. Finally, since we have applied the FFT algorithm in code, we need to move the right half of the resultant magnitude spectrum to the beginning of the array in order to keep zero frequency (baseband) at the center of our azimuth arrays; essentially, we must `fftshift` each azimuth spectrum. We transform each range bin exactly the same way to obtain one azimuth spectrum – one line in our image – from each range column. The array comprising these range columns along one dimension and their transformed azimuth lines along the other dimension constitute one patch of our unfocused SAR-processed image:

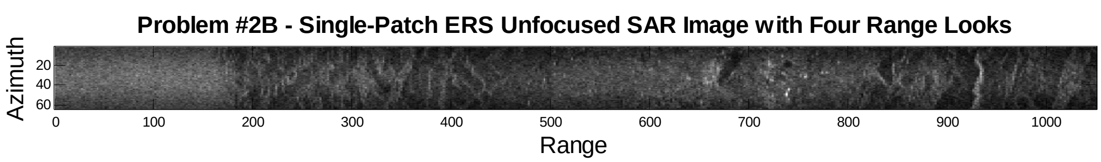

After iterating our azimuth transform across all range bins, we extract the next 64 lines in azimuth and repeat the steps outlined above on the second patch. As we move along the azimuth, processing 64 lines at a time, we must slightly offset each patch by either 3 or 4 pixels in the large image matrix; instead of applying the same offset to each patch, we round the incrementing multiple of patching spacing (3.52227 pixels), since a constant offset of 3 or 4 pixels would grow more and more erroneous as the 0.52227 differences accumulate. Thus, even though we cannot space our patches by exact 3.52227 pixels, we never deviate far from this ideal spacing since we round the total spacing accumulated so far; for example, we offset the first patch by 4 pixels, the second patch by 7 pixels, the third patch by

11 pixels, and so on. Shifting our azimuth indices by this rounded accumulation, we overlay the patches in the image array (declared to hold 613 total azimuth lines) to obtain a composite image representative of the ground imaged by 157 consecutive bursts during the spacecraft's motion.

As a final, purely aesthetic step, we move along the range bins in our image array and average groups of four adjacent range bins, replacing each quartet of values with their average. As a result, we slightly blur the range, but our final image display will appears squarer, since our averaging has effectively decreased the pixel length in range to correspond more closely to the pixel width in azimuth. We choose to average *four* pixels along the range since it offers the aspect ratio closest to unity; the resulting pixels in the image appear most square if we combine four range bins. The full unfocused SAR processed image recognizably displays the San Francisco area, with mirrored orientation:

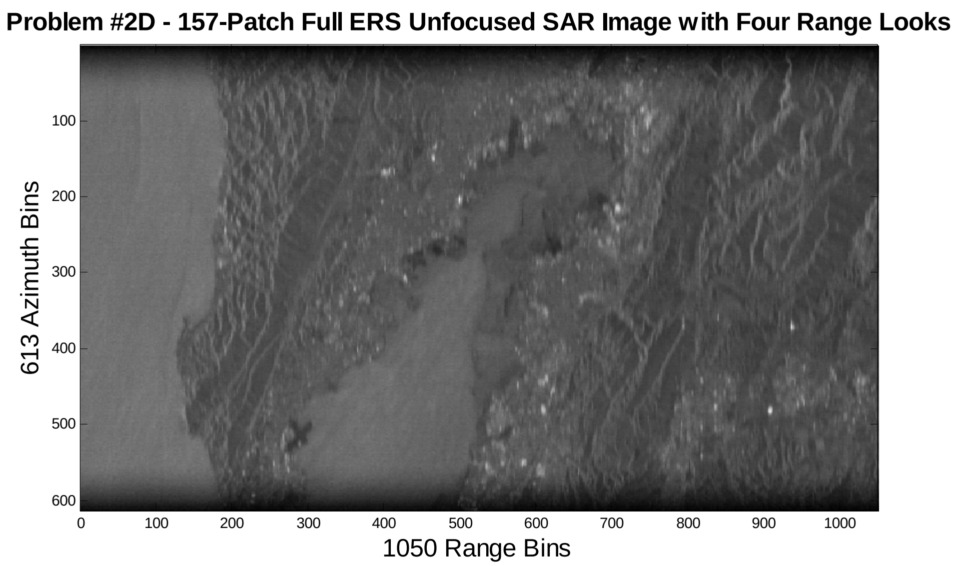

With resolution in both the range and the azimuth dimensions, we can now begin to identify parts of the image; we have sufficient resolution to distinguish bright points in the populated areas, the water of the bay and Pacific Ocean, and even the X-mark that appears along the lower shoreline. However, the resolution still can improve, as we will likely strive to do in *focused* SAR processing. Furthermore, remark that the image reverses the azimuth (vertical) orientation we are accustomed to viewing on maps. Even though the range bins align, our processing appears to have inverted the azimuth bins. Compare the following images:

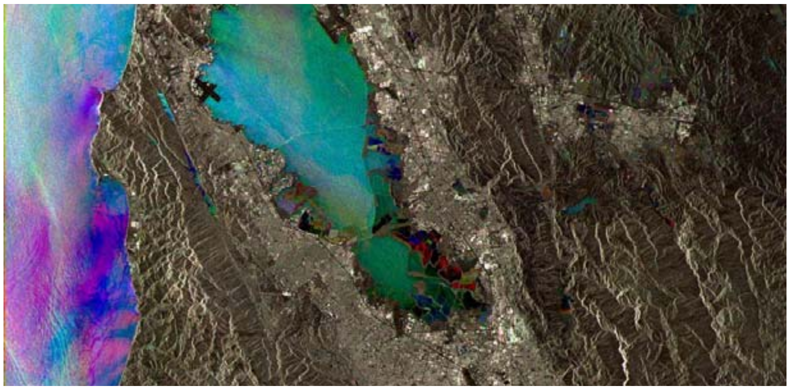

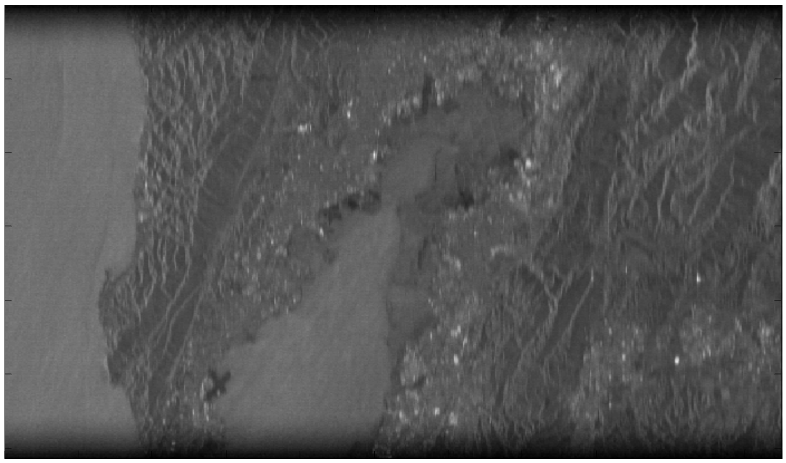

Notice how the vertical direction in our image is the mirror image of the vertical direction in the usual image. Of course, we can easily correct for the inversion in another step of processing, now that we recognize that our image should portray the San Francisco bay: