

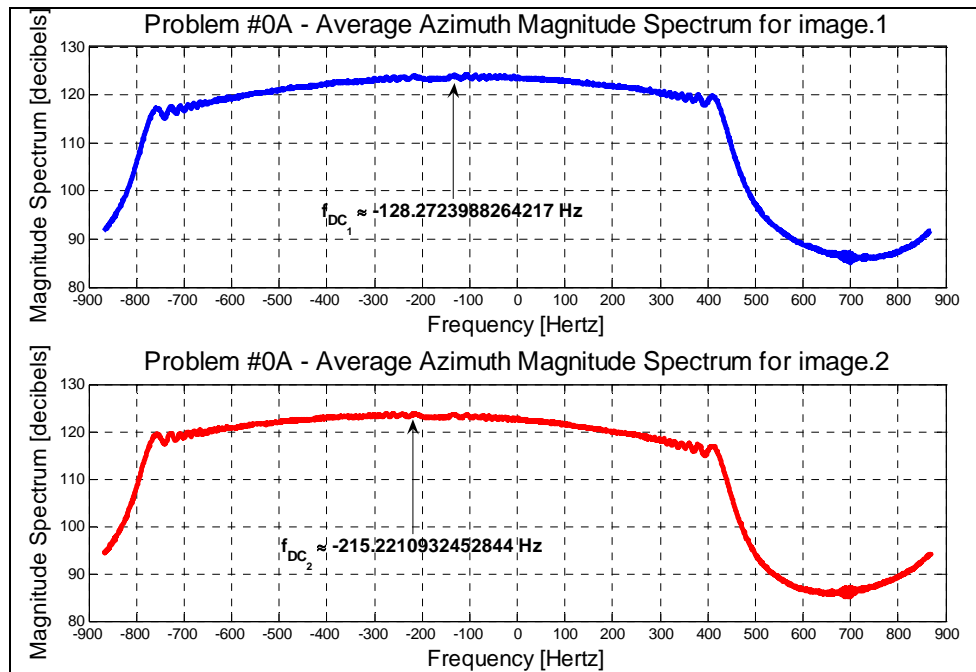
Problem Set VIII

Radar Interferometry and Correlation Images

Problem #0 – Determination of the Doppler Centroid

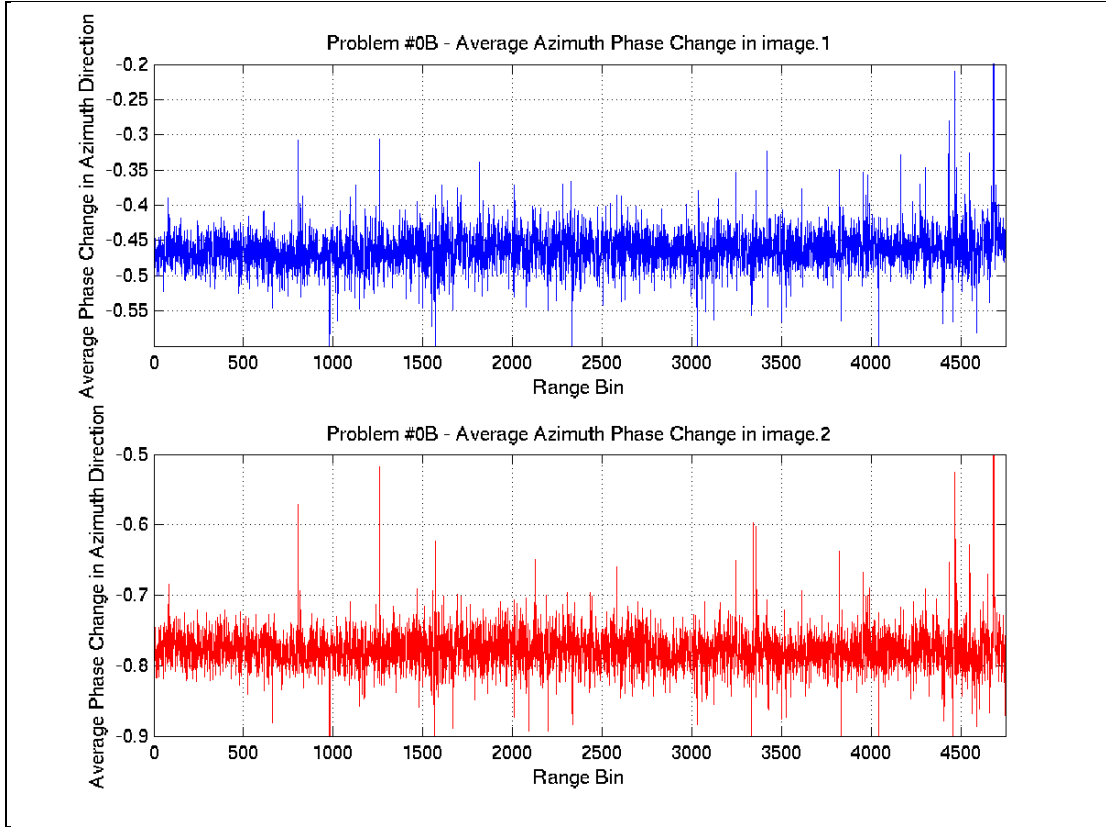
Before we begin forming the interferogram, we should examine the azimuth spectra along the vertical dimensions of our two images to ascertain their Doppler centroid frequencies. Because we will eventually sample the second image onto the grid defined by the first image, we must ensure that we also standardize the Doppler centroid frequency so that we sensibly resample from the same azimuth reference; such azimuthal alignment requires that we know the Doppler centroid of both images, so we can properly account for their difference.

As usual, the mean azimuth magnitude spectrum, averaged across all range bins, attains its maximal value at the Doppler centroid. Because the azimuth has already been matched-filter-compressed in both of these pre-processed image files, the antenna weighting protrudes distinctly from the expected rect spectrum that accompanies any matched-filtered focused image; locating the frequency at which this pattern peaks will immediately reveal the Doppler centroid. Upon juxtaposing the average magnitude spectra from the two images, we perceive a slight difference in the antenna weighting and conclude that the two antennae must point in slightly different directions, with slightly different viewing angles:



Applying the average phase change algorithm, we can verify these graphical estimates; 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 Doppler 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:

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



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 with 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

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

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

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 along azimuth – both of them derive from the same set of azimuth samples – our centroid frequency calculations should match, allowing us to further refine our graphical estimates to the values:

$$\boxed{f_{D1} \approx -128.2723988264217 \text{ Hz}} \text{ and } \boxed{f_{D2} \approx -215.2210932452844 \text{ Hz}}$$


First, we notice that the Doppler centroid supplied in the problem statement lies approximately halfway between the two actual Doppler centroids observed in the image but fails to describe either true centroid accurately; in other words, the average value would prove useless if we wanted to correct for the Doppler squint in either individual image, so our investigation here will prove vital to proper resampling. More importantly, we note that the *difference* between the two Doppler centroids is

$$\boxed{f_{D2} - f_{D1} \approx -86.94869441886274 \text{ Hz.}}$$

This difference is the factor by which we must steer the second image before we resample it onto the grid of the first image.

Problem #1 – Formation of the Raw Interferogram

After loading 6144 azimuth lines of 5120 range bins and combining complex pairs accordingly, we multiply the first image by the complex conjugate of the second image to yield not only the amplitude product but, more importantly, also the phase difference. To refine the accuracy of our phase estimation at the expense of spatial resolution, we also average groups of four range bins and four azimuth lines in a 4×4 multi-look algorithm, thereby reducing our multiplied image matrix fourfold as well. Upon downsizing the product image with this average, we display the phase instead of the amplitude:




We perceive range differences as phase differences in the cycle of colors, but, to our dismay, the resolution leaves much to be desired, especially in the lower-right corner of the interferogram, where the lines of constant phase begin to blend together. The lines contrast sharply in the left side of the image, but their exactitude and separability degenerate as we move rightward, past the central circular feature that represents our mountainous region.

Problem #2 – Offset Determination

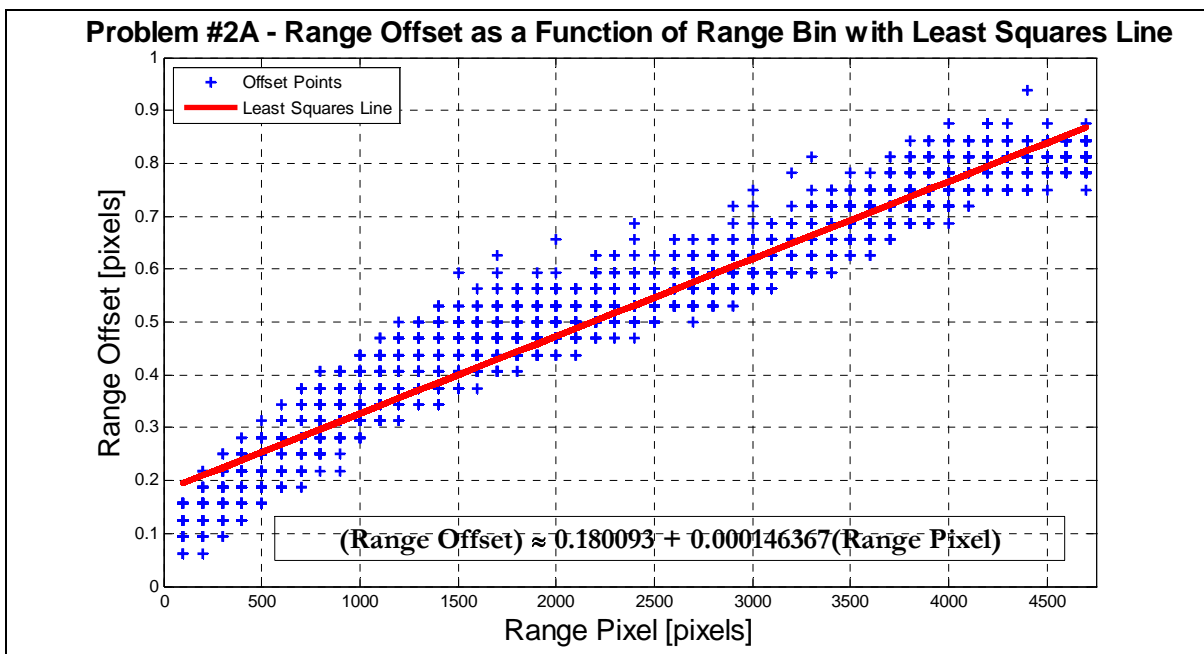
In preparation for our resampling of the second interferometric image, we compute the offset between regions in the first image and their corresponding areas in the second image. By cross-correlating a small region in the first image with a slightly larger region centered around the same bins in the second image, we produce a two-dimensional response that peaks where the two corresponding regions are most similar; however, because the region in the second image is slightly offset from its counterpart in the first image, the peak location will not occur at the center of the region, and its displacement yields the offset we need to correct for the slight deviations. Repeating this cross-correlation across several regions in the two images, we obtain a set of offsets that we can then fit with a least squares line to estimate the offset of any pixel. First, we run the pre-compiled offset code with the following instructions:

<pre>image.1 5120 image.2 offset10.out 500 4500 10 500 5500 10 0 0 2. 0 0</pre>	<pre>image.1 5120 image.2 offset20.out 500 4500 20 500 5500 20 0 0 2. 0 0</pre>	<pre>image.1 5120 image.2 offset.out 100 4700 47 100 6000 60 0 0 2. 0 0</pre>
---	---	---

In the first two inputs, we specify the number of sample region pairs we would like to study, and the program performs cross-correlation and peak detection, outputting the resultant offsets. In the third input file, I decided to extend the regions of interest closer to the edges of the image while also increasing the number of boxed regions to sample; while this lengthened the duration of our offset determination runtime, I obtained more samples through which I could fit a better, more representative line of least squares. We begin by fitting lines of least squares to the ten-box and twenty-box sample sets, both across range and along azimuth:



Even though the acquisition of more region offsets consumes more time during cross-correlation, we see that the additional samples produce a better representation of the actual offset behavior over a more refined grid in our image; as the correlation coefficient of the linear fit increases as we add samples, we are better able to predict the offset at any location in the image when we sample more frequently. Following this improvement further, we increase the number of boxes as well as extending our regions into the fringes of the image, taking special care that we do not cross the boundary into invalid data (near the 4750th range bin):




This regression line for the range is $(Range\ Offset) \approx 0.180093 + 0.000146367 \cdot (Range\ Bin\ Index)$.

This linear model improves our offset determination's resistance to outliers while also providing a much more continuous range of data values, so we adopt this range offset linear relationship for baseline calculations.

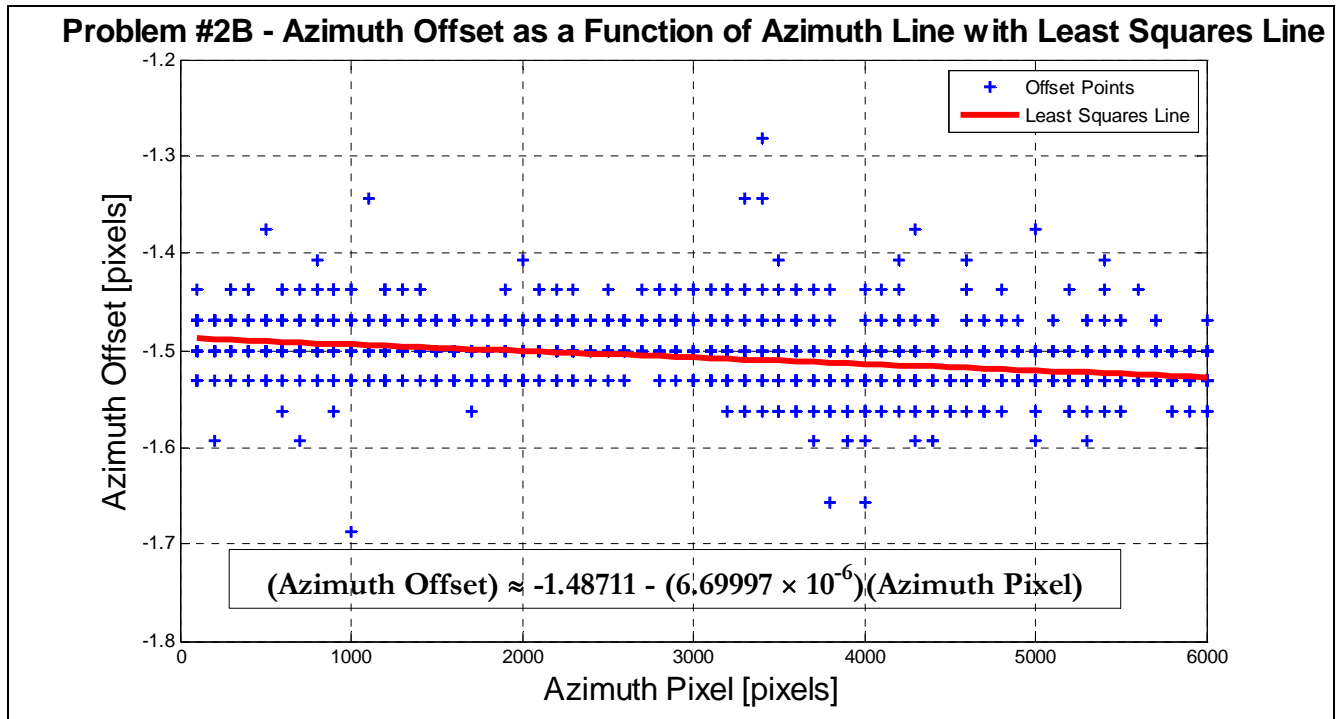
Either linear model reveals why the raw interferogram blurred more noticeably as we moved from the left side of the image to the right; the range offset between the two images increases with range bin, with the most deviation in the rightmost bins, where the range offsets reach nearly pixel-size differences. As a result, the phase resolution of our interferogram naturally suffers, as phase varies quite sensitively with range; the slight range offset proves sufficient to derail the precise phase measurement that we expect to see in an interferogram.

Similarly, for the azimuth, we begin with the ten-box and twenty-box regression models for expedience:



Because the offsets are now negative values, we notice that, in general, the azimuth offsets occur in the opposite direction as the range offsets. Nevertheless, similar to but not nearly as severe as the range, the azimuth offsets also worsen slightly with bin number. As we witnessed in our raw interferogram, the phase fringes appeared sharpest in the **upper**-left corner, where the azimuth deviations were least negative; however, as we moved toward the **lower**-right corner, the image degraded, with the worst resolution occurring at the image bottom (highest azimuth line index) and at the image right (highest range bin index). Because the slope of the azimuth line is not nearly as steep as the range line, the image offset is much more sensitive to range bin index, explaining why the degradation varied mainly from left to right, with only slight degeneration from top to bottom.

As with the range, we obtain a refined linear model when we acquire more samples from the offset program:



Observe that the number of points reduces the effect of outlier data and refines our calculation of the slope and intercept of the best-fit line. Along the azimuth, the relationship between offset and line number is

$$(Azimuth\ Offset) \approx -1.48711 - 6.69997 \times 10^{-6} \cdot (Azimuth\ Line\ Number)$$

The mean value of this azimuth offset is approximately -1.50754 pixels. Even though this average offset seems generally greater than the range offset (which hardly reaches a full pixel), it is the slope or derivative of the offset that produces the most noticeable difference in our fringe pattern; while the constant offsets still undoubtedly affect the accuracy of our multiplied interferogram, the *changes* in offset precipitate the most blatant smearing, as the error between our image points varies from bin to bin and line to line, deteriorating not only the accuracy but also the *precision* of our interferogram values and leading to the blurred lines and dulled edges that we witness prominently on the right side of our image.

Problem #3 – Baseline Component Calculations

Before we proceed with baseline calculations, we must define references in our image by which we will measure all of our range offsets. If we place the reference in the central range bin (index 2560), then the reference point of our least squares line has a reference range of

$$r_{reference} = r_0 + (i_{reference} - 1) \cdot \delta_r$$

$$r_{reference} = r_0 + (i_{reference} - 1) \cdot \frac{c}{2f_s}$$

$$r_{reference} = (309,735 \text{ m}) + (2560 - 1) \cdot \frac{(299,795,637.7 \frac{\text{m}}{\text{sec}})}{2(44.997 \times 10^6 \text{ Hz})}$$

$$r_{reference} \approx 318,259.7576155848 \text{ m.}$$

Likewise, we must establish a reference look angle at the center of our range span. Applying the Law of Cosines,

$$\theta_{reference} = \cos^{-1} \left(\frac{r_{reference}^2 + (R_{Earth} + z)^2 - R_{Earth}^2}{2 \cdot r_{reference} \cdot (R_{Earth} + z)} \right) \approx 46.875856^\circ.$$

Having designated the central pixel as our reference point along with its range and look angle, we compute the baseline components from a combination of our offset regression line parameters and geometry:

$$B_{\parallel} = [(offset \ intercept) + (offset \ slope) \cdot i_{reference}] \cdot \frac{c}{2f_s} \approx 1.8481735 \text{ m}$$

$$B_{\perp} = (offset \ slope) \cdot r_{reference} \cdot \tan(\theta_{reference}) \approx 49.737373 \text{ m}$$

$$B = \sqrt{|B_{\parallel}|^2 + |B_{\perp}|^2} \approx 49.771699 \text{ m}$$

$$\alpha = \tan^{-1} \left(\frac{B_{\perp}}{B_{\parallel}} \right) + \theta_{reference} - 90^\circ \approx 44.747801^\circ$$


Problem #4 – Interferogram Resampling

In order to resample the second image onto the same grid as the first image, we must first place both azimuth spectra on the same Doppler centroid frequency. Thus, we steer the second image by a complex exponential chirp along the azimuth: $e^{-j \cdot 2\pi \frac{f_{D2} - f_{D1}}{PRF} \cdot x}$, where x represents the azimuth slow-time coordinate that we must steer to center frequency. Recall that $f_{D2} - f_{D1} \approx -86.94869441886274 \text{ Hz}$ from our initial graphical and average phase change estimates.

After removing this Doppler centroid from each azimuth line, we resample the second image pixel-by-pixel with the offsets given from our linear regression models along the range and azimuth. Meanwhile, because the offsets specified in the linear model are not always integers, we employ an eight-point two-dimensional sinc interpolation algorithm to reconcile non-integral pixel numbers with a sinc weighting of neighboring integer nodes. We compute the range coefficients by sinc-interpolating the offset values across eight adjacent range bins centered on each bin. Likewise, we determine the azimuth coefficients by sinc-interpolating the azimuth offset values across eight successive azimuth lines surrounding each processed line, although we can assume from the offset sample outputs shown in Problem #2 that the azimuth offset remains relatively constant at the mean value of approximately -1.50754 pixels; recall that the slope was meager: -6.69997×10^{-6} pixels of offset per line, allowing us to re-shift every azimuth line by the same constant amount with no regard to offset drift.

Finally, upon resampling the second image, we reintroduce the difference in Doppler frequency that we removed prior to sampling, since we need to preserve the precise frequency content of the second image in our interferogram; even though we had to resample our image from the same azimuth domain as the first image, we do not want to alter the phase differences inherent in our two images, so we undo the Doppler standardization through multiplication with the conjugate of our original steering exponential: $e^{+j \cdot 2\pi \frac{f_{D2} - f_{D1}}{PRF} \cdot x}$.

Dot-multiplying the untouched first image with the resampled second image element-by-element and taking four looks in range and four looks in azimuth as before, we see a marked improvement in the phase resolution of our resultant interferogram:



Improvement manifests itself most prominently on the right side of the image, where the nuances in phase now appear clearly, with relatively infrequent blurring when compared to the amalgamation of colors seen in the lower-right corner of the raw interferogram. By correcting for the range and azimuth offsets and interpolating the results, we have made the pixels in the first image correspond nearly one-to-one with the pixels in the second image, with nearly non-existent offset migration, allowing our product of images to display phase continuously and crisply from

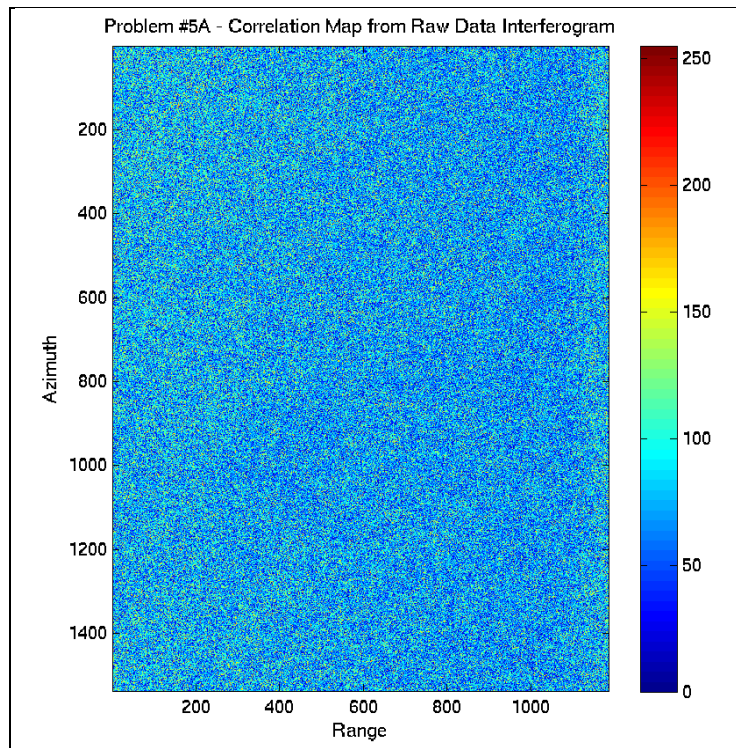
bin to bin and line to line. Furthermore, the image also contains much less noise, as we can more clearly perceive the detailed edges of each fringe and the mountainous feature outline in the center. In essence, our interferogram measures the differences in phase – and hence also range – with much greater precision and accuracy following resampling.

Problem #5 – Correlation Images

As a measurement of our interferogram quality, we can also compute correlation images to display how closely the phase in our second image tracks the phase of the first. We compute the correlation point-by-point:

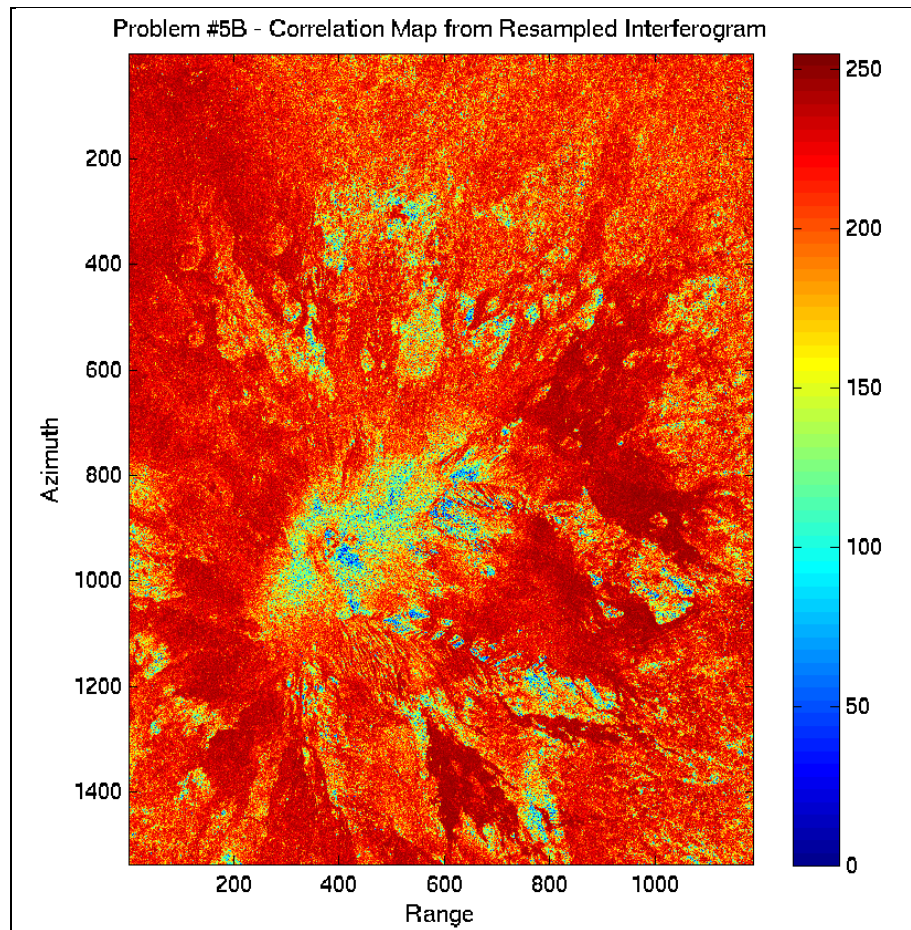
$$\rho = \frac{\sum_{\text{azimuth line } j} \sum_{\text{range bin } i} \langle \text{Image}_1(i, j) \cdot \text{Image}_2^*(i, j) \rangle}{\sqrt{\sum_{\text{azimuth bin } j} \sum_{\text{range bin } i} |\text{Image}_1(i, j)|^2} \cdot \sqrt{\sum_{\text{azimuth bin } j} \sum_{\text{range bin } i} |\text{Image}_2(i, j)|^2}}$$

Furthermore, by stretching the normalized correlation coefficient to a scale of [0 255], we can even display the scaled correlation coefficients as an image; the quality and resolution of the final correlation map indicate not only how well our interferogram records phase but also how brightness varies across the imaged surface, since reflective areas returning more signal will naturally exhibit higher correlation and hence greater image brightness. Before resampling, we draw the following correlation map from the raw data:



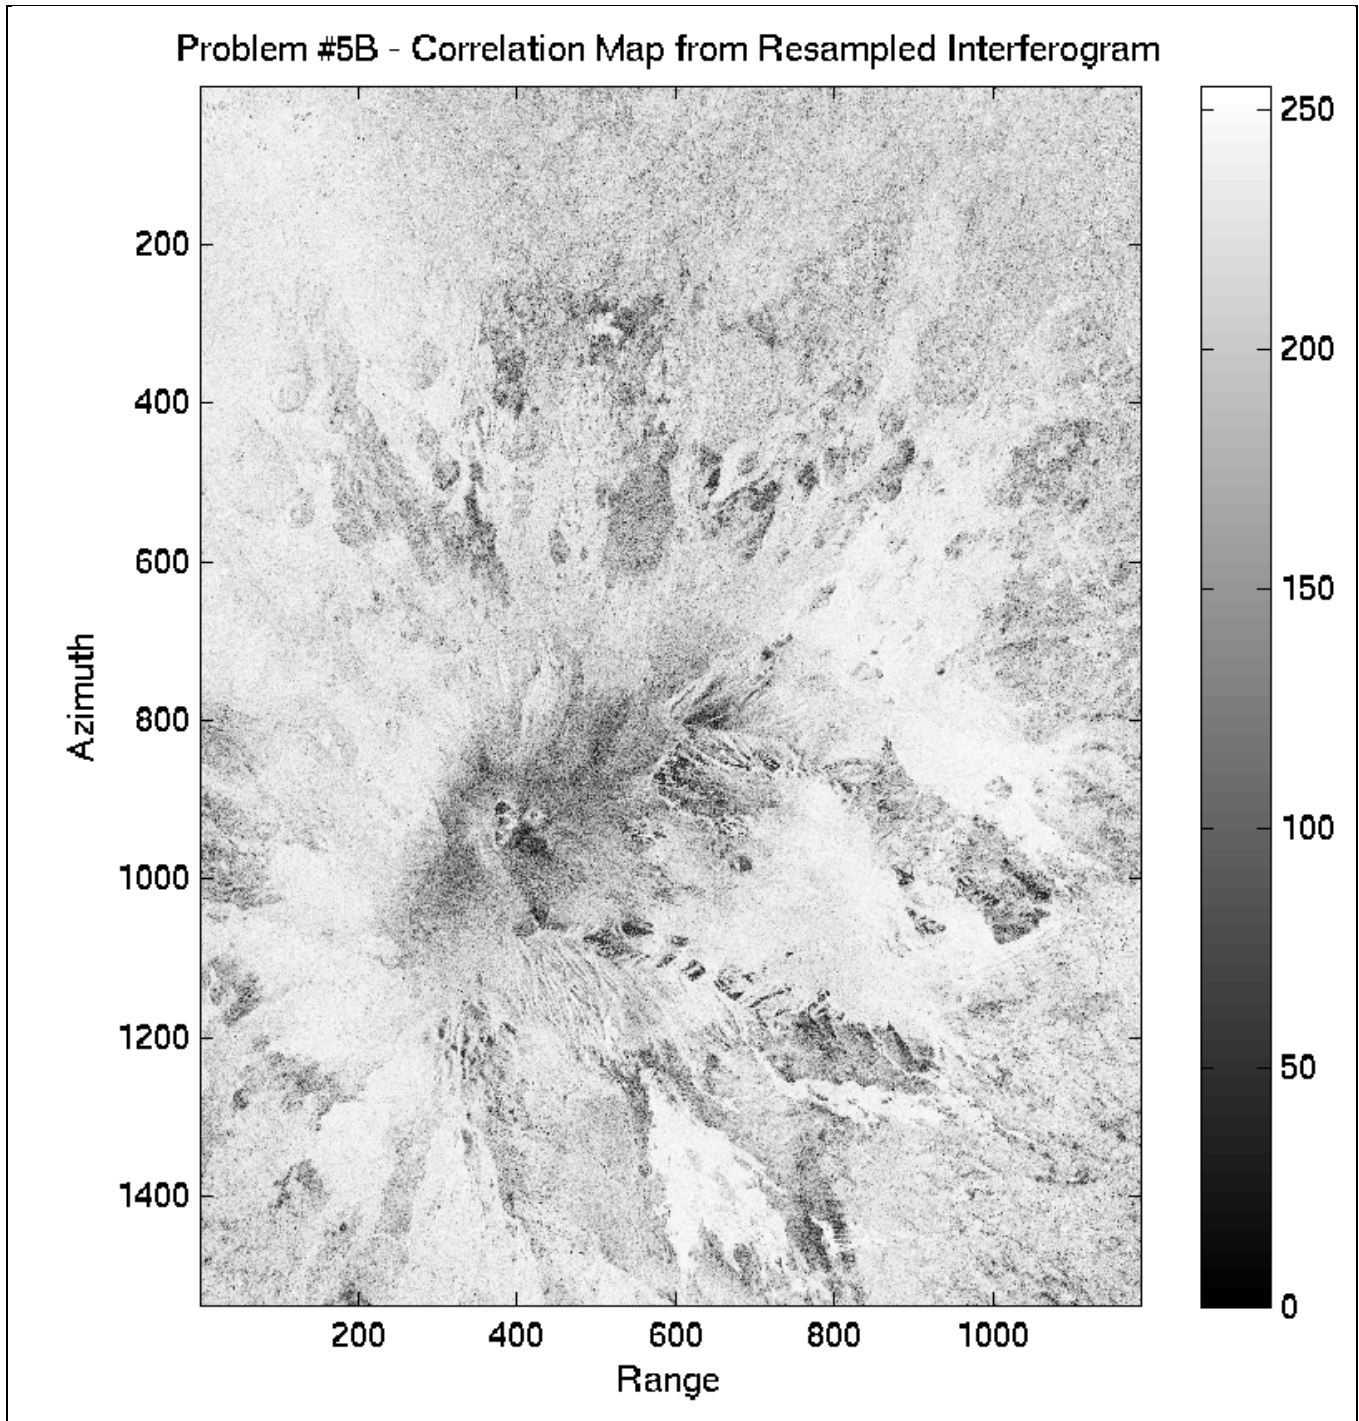
The resolution is poor and the surface features indistinguishable because corresponding pairs of points (sharing the same indices in the two images) are generally uncorrelated ($\rho < 0.5$), resulting in the predominantly blue (100-level) speckle seen above. Moreover, in addition to relative uncorrelation, the offset drift along range effectively randomizes the correlation so that the correlation image contains no visual rhyme or reason; the offset and hence correlation assume a different value at every single point, leading to the generally disparate speckle that seen above.

However, resampling improves not only the interferogram but also the correlation image:



The first aspect that leaps from the resampled interferogram is the intense sea of crimson indicating the general prevalence of high correlation; in other words, the interferogram tracks phase well throughout most of the image, with the exception of the central region featuring the sudden altitude change.

Furthermore, in the resampled correlation image, the imaged surface begins to appear. We begin to see the mountainous region as the area in the correlation map with the lowest correlation. Because two images acquired at slightly different viewing angles respond quite differently to sloped surfaces or rapid changes in altitude, a mountainous region such as that imaged centrally would naturally lead to low correlation (high decorrelation) between images from differently oriented antennae. Indeed, for a relatively flat surface (relative to the imaging wavelength), all of the signals would add coherently, resulting in a large echo, strong specular return, and therefore generally high correlation, as seen as red in our resampled map, but a sloped surface will inevitably cause the phase from two imaging geometries to differ wildly, appearing in our map as incoherent addition, or low correlation.



Because we have removed the range and azimuth offsets from our interferogram in resampling the second image, we have simultaneously improved the resolution of our correlation map, providing us an alternative view of the imaged surface in detail. Because phase difference measurements integral to computing correlation generally have much higher resolution than distance difference measurements, these correlation images can reveal more than just the significant accuracy improvement of our resampled interferogram...