Computational Burden in Moving Target Detection in Multi-Channel SAR using Phase Interferometry

P. K. Sanyal, D. M. Zasada, R. P. Perry
The MITRE Corp., 26 Electronic Parkway, Rome, NY 13441
Ph: (01) 315-838-2663, Fax: (01) 315-838-2678; (psanyal, dmzasada, rpp)@mitre.org)

***Abstract*—The authors have shown in several earlier papers that large ground moving targets can be detected fairly reliably via phase-interferometry in multi-channel Synthetic Aperture Radar (SAR) images of the ground. In a interferometric phase image, all points on the ground nominally appear as a continuum of phase differences while the moving targets appear as discontinuities. By detecting these anomalies, it is possible to detect moving targets in the SAR.**

When the target radar cross section (RCS) is small, the target return is dominated by and often lost in the ground clutter. The authors have shown that the detectability of small-RCS moving targets in multi-channel SAR can be enhanced by a direct, channel-to-channel cancellation of the ground clutter.

In this paper we discuss the computational burden in the moving target detection process involving direct channel-to-channel clutter cancellation process.

Index Terms— Computational burden, Surface Moving Target Indication (SMTI), Keystone Formatting, Clutter Cancellation, SAR Phase Interferometry.

I. Introduction

In a SAR image of the ground, the target motion causes the moving targets to appear at locations different from their true instantaneous locations on the ground. If the target RCS is large enough, it simply appears as an observable bright spot, but usually there is no indication that it is displaced from its actual position in the image. However, in the corresponding interferometric phase image, the phase-difference values at all pixels corresponding to the points on the ground nominally appear as a continuum of phase differences while the phase difference values at the pixels containing the returns from moving targets appear as discontinuities. By detecting these phase discontinuities or anomalies, we can detect moving targets in the single frame of the multi-channel SAR [1-5]. An example is shown below using real radar data collected by Lincoln Lab (LiMIT data).

The top insert in Figure 1 shows the GPS-based tracks of seven moving targets on a runway. They start from the upper right end of the runway and proceed to the lower left end of the runway. In the middle, two, #s 5 and 6, of the targets turn left on a crossing road while the remaining

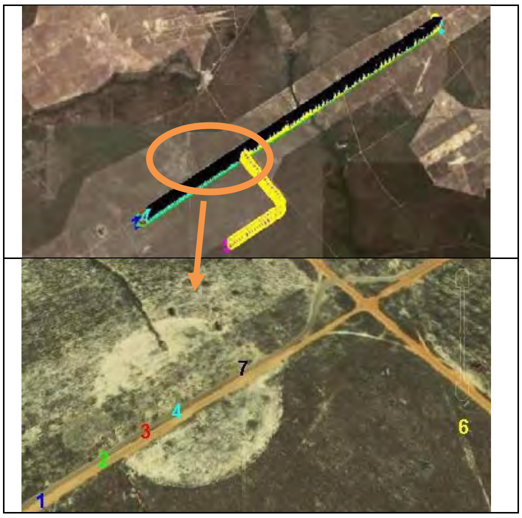

Figure 1. The GPS tracks of a convoy of moving targets on a complex of roads

five continue on the original path. Of these five, the #s 1-4 are military vehicles and #7 is a civilian pickup truck. The bottom insert shows a GPS-indicted snapshot of their positions at an instant after the split in the convoy.

The top insert in Figure 2 shows the SAR image of the area near the cross-roads at about this time, formed from one channel of the 8-channel LiMIT system. The cross-roads and the looping roads around the crossing are easily discerned (note: there is a smaller back-scatter from the relatively smooth surfaces of roads and they appear as low-RCS objects relative to the surrounding vegetation-covered and/or rough terrain).

Recall that the moving targets, if observed, would not have appeared on the roads they were travelling on but would appear horizontally displaced, i.e., displaced in Doppler or the cross-range direction. However, none of the moving targets seem to be observed in the body of this SAR image.

The middle insert in Figure 2 is the interferometric image of the same area, i.e., the scaled color plot of the pixel-by-pixel phase-differences between the images from two channels, (channel #1 and #8 in this case).

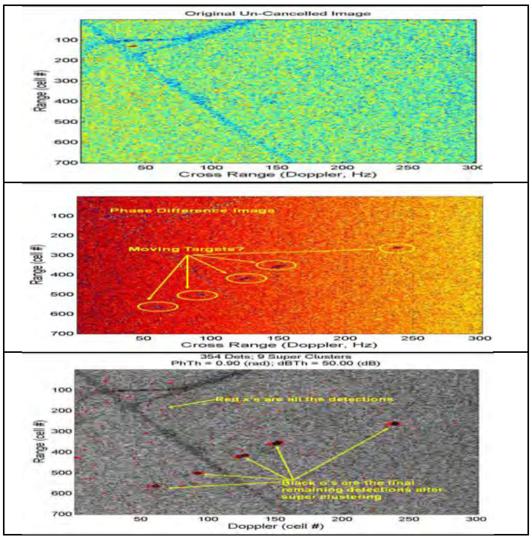

Figure 2. Detecting the convoy of moving targets via phase interferometry

Dark spots corresponding to the five moving targets are clearly visible. All the spots are displaced from the road on which they are travelling. The lower four spots appear to be parallel to the road and separated by about the same distance, indicating that they are travelling at about the same speed. These must be the targets #1-4. The top right spot therefore must be corresponding to target #7. Its range separation from #4 is larger and it its horizontal displacement is clearly more than the others.

Since the purpose of the SAR, etc., is to obtain situational awareness, a bit of speculation about what we are observing in Figure 2 is in order. Without even referring to the detailed GPS records, we can assume that the convoy starts off with seven targets ordered according to their numbers and with about the same separation between the elements and each moving with the same speed and thus maintaining their separations. After the targets #5 and #6 leave the main group, a larger space opens up between the targets #4 and #7 and we can speculate that #7 speeds up temporarily to narrow down the gap. This situation is reflected in the interferometric image where the last target has a larger range separation and, because of its temporary higher speed, has a larger Doppler displacement. As to why the targets #5 and #6 do not appear in the phase-difference image, there are two possible answers: a) they were outside the area being painted by the radar or b) their response is displaced farther to the right and is outside the area displayed here.

While the moving objects are clearly observable in the phase interferometry image by a human analyst, it is important to have the machine recognize the phase anomalies and thus detect moving targets automatically. The last insert in Figure 2 shows the result of automatic detection. The process of automatic detection is explained in the next section.

Notice also that the roads are discernible in the phase-difference image. This is because we have arbitrarily set the phase differences of the pixels with power levels below a chosen threshold to the lowest phase value, i.e., $-\pi$. This action not only reduces speckle in the image, it also, fortuitously, retains some of the useful features of the SAR image, e.g., the roads and other low-RCS objects and thus keeps the image recognizable.

II. COMPUTATIONAL BURDEN

In the example shown above, the moving targets were observable without clutter cancellation because apparently the clutter in the region where the displaced targets appear is not very strong. However, clutter cancellation may be required when this is not the case. The functional diagram of the total process of detecting moving ground targets via phase interferometry is shown in Figure 3.

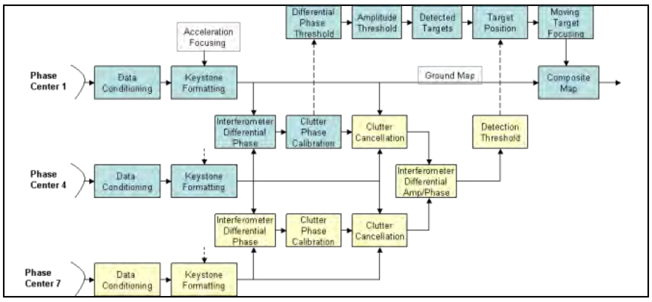

Figure 3. Block Diagram of Moving Target Detection and Tracking Using Multiple Phase Centers

The very first step is data conditioning, which is standard in almost any data processing endeavor and is included mainly for completeness.

The next step is the Keystone Formatting, which removes range walk, as shown in Figure 4. The Keystone formatting involves re-sampling the wide-band I&Q data in the Fourier domain according to the instantaneous frequency, i.e., closer sampling intervals at the lower end

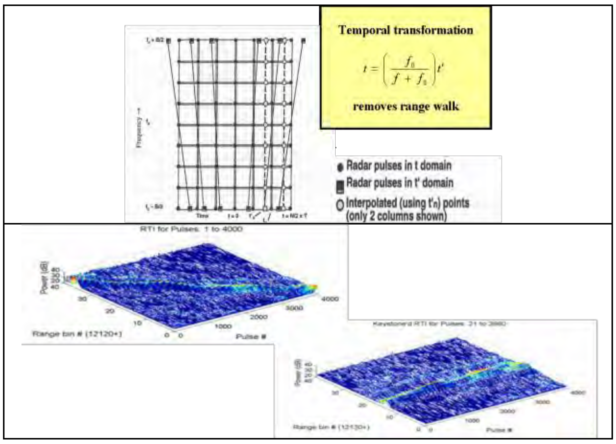

Figure 4. Keystone Formatting removes range walk

of the spectrum and larger sampling intervals at the upper end of the spectrum, as shown in the top insert in Figure 4. The bottom insert shows how the range walk implicit in the slanting ridge line has been changed to a horizontal ridge-line after Keystoning.

Sometimes the IQ data available for processing is not really 'raw' in that it has already been preprocessed to remove gross range walk, as in 'scene-center-stabilized' data. In this case, the Keystone formatting is not necessary.

After the range walk has been removed from the data, the IQ data is converted to a SAR image by Doppler processing across the pulses via an FFT. This step is not explicitly shown in the processing diagram.

The next step in the processing chain is the computation of the interferometric differential phases. Note that the SAR image is the magnitude part of the complex data that results from the SAR processing; the phase part is usually completely ignored. The interferometric differential phase is simply the pixel-to-pixel phase difference. The middle insert in Figure 2 is a plot of this differential phase and we refer to as the phase difference image.

The basis for the differential phase is shown in the left insert in Figure 5.

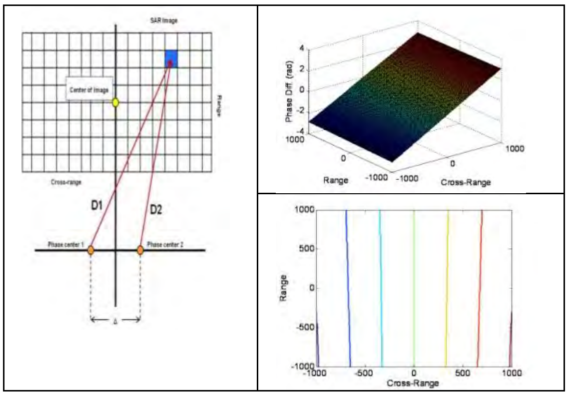

Figure 5. Computing the interferometric differential phase

The top right insert plots, as a surface, the differential phase over a 1km x 1km region at a distance of 15 km, with the phase centers separated by one wavelength at X-band; the lower right insert plots the contour of this surface. The maximum of the magnitude of the phase difference is about 3.07 radians in this case.

A close examination of the contour plot reveals that the contour lines are not exactly parallel to each other; they tend to flare out towards the top, indicating that the surface is not exactly a plane. We have purposely chosen the range of 15 km to show that the differential phase surface is not strictly a plane, but at longer ranges, e.g., ~35 km, as in the LiMIT data, the surface, for all practical purposes, is a perfect plane.

If the radar data was completely noise-free, the computed differential phase would indeed constitute a plane. But the radar data is noisy and after the SAR processing, each pixel ends up with a different amount of noise. The edge-view of the differential phase surface corresponding to the image in Figure 2 is shown in the top insert of Figure 6.

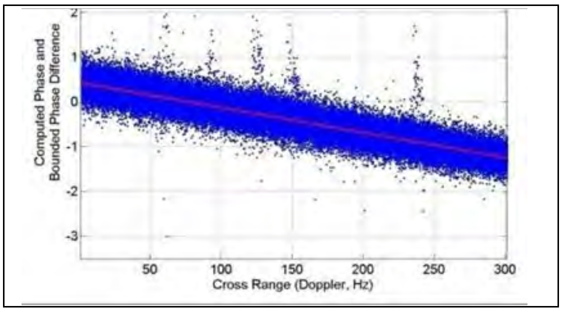

Figure 6. Deviations from the interferometric differential phase plane

Clearly, the differential phases do not all lie on a plane; they all deviate from this ideal plane. But the differential phases at the pixels containing the displaced returns from the five moving targets display clearly observable larger deviations.

For the machine to recognize deviations from the ideal plane, it has to know the ideal plane. This is done by fitting a plane in a least-square-error (LSE) sense, through the computed differential phase. This plane is derived as follows. Think of the phase difference image as an mxn matrix of phase difference numbers, where m is the number of range cells and n is the number of Doppler cells, i.e.,

$$\Delta\phi(i, j),$$
$$range\,cell\;\; i = 1, \ldots, m;$$
$$Doppler\,cell\;\; j = 1, \ldots, n$$

We fit a plane to this observed surface:

$$\Delta\phi(i, j) = c_0 + c_{range} i + c_{Doppler} j$$

where the c's are the regression coefficients. Having obtained the coefficients, we can generate the ideal LSE plane. The edge view of this plane is shown by the red line in Figure 6. The phase deviations from this ideal plane can now be computed. Pixels whose magnitude exceed a chosen power (dB) threshold and whose phase deviation exceeds a chosen phase threshold are declared as detections.

The bottom insert in Figure 2 displays such detections as red dots. It is seen that there are a large number of red dots, indicating a fairly large number of false alarms. As in all threshold-based detection processes, the number of false alarms could be reduced by raising the threshold. Here we have employed another approach - that of clustering. It stands to logic that most true 'false alarms' appear as isolated pixels or at the most a few contiguous pixels while returns from most extended targets appear as a large group of pixels. Thus performing a clustering on the detections and setting a threshold on the minimum number of detected pixels in a cluster leads us to final result shown in the bottom insert in Figure 2.

Referring to Figure 2, we have noted earlier that the RCS of the moving targets must have been large enough not to be dominated by the clutter return in the pixels where they appear and are therefore discernible in the interferometric differential phase image without the need for clutter cancellation. This is somewhat fortuitous but, based on the many data sets looked at, does not appear to be rare occurrences either. However, the detection performance can clearly be improved if the clutter level is reduced. Even when the moving targets are detected without clutter cancellation, the clutter cancellation improves the geo-location process, which is the last step in the processing chain in Figure 3.

The direct channel-to-channel clutter cancellation is based on the Displaced Phase Center Antenna (DPCA) principle. In the absence of any signal noise, the complex SAR images (as noted earlier, SAR images display only the magnitude) from two channels of a multi-channel system differ only in phase, with the phase varying across the pixels and represented by the plane displayed in Figure 5. Therefore, in an ideal case, if the complex image from one channel, appropriately phase-modified, is subtracted from the complex image from another channel, the difference will be zero, i.e., the image, which is also the 'clutter', will be completely cancelled. Referring the two images from, say, channels 1 and 8, the ideal differential image

$$I_{k,l}^{1-8} = I_{k,l}^{1} - I_{k,l}^{8} exp\left(-j\Delta\emptyset_{k,l}^{1-8}\right) = 0, \text{ for all k, l}$$

will be zero. But while the ground returns or clutter would cancel out completely, the pixels which contain the moving target returns will not cancel completely because the differential phase for those pixels do not lie on the phase plane which is the value we are using for the phase weighting one of the channels.

As we have already seen, because of the noise in the data, the computed differential phases deviate from the plane. However, we cannot simply use the computed differential phases for the weighting and achieve the near perfect clutter cancellation because this would also completely cancel the moving targets. So we need to estimate the ideal phase plane. But we have done that already in the previous step and this step does not require any extra computation for estimation. In this step, we only weight the image from one channel and subtract from the image from the other channel, as indicated in the above equation where we use the estimated phase differential.

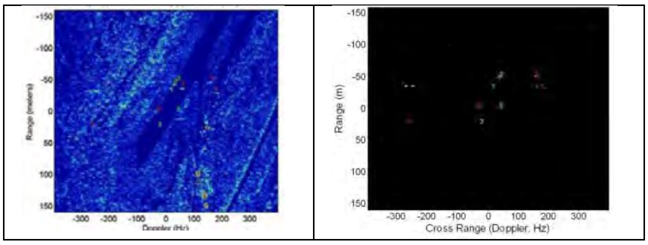

Figure 7. Pre- and post-clutter-cancellation SAR images

Figure 7 shows the results of applying the direct channel-to-channel clutter cancellation to SAR image of the end of a runway (General Dynamics DCS data) where there were several moving targets. The right insert shows that most of the clutter, i.e., the returns from the ground, has been eliminated, leaving the slightly attenuated returns from the moving targets.

When estimating the phase plane, the computed differential phases should be weighted inversely by the power in the pixel because low-power pixels have more noise. Alternatively, as we have done here, one can simply ignore pixels whose power level is below a chosen threshold.

III. SUMMARY

In this paper we have described the computational burden in the phase interferometric technique for detecting moving targets in synthetic aperture images (SAR). We have included results from processing real data.

REFERENCES

[1] Sanyal, P. K., Perry, R. P., Zasada, D. M., ‘Detecting Moving Targets in SAR Via Keystoning and Phase Interferometry, IRSI-2005, Bangalore, India, December 2005

[2] Sanyal, P. K., Perry, R. P., Zasada, D. M., ‘Detecting Moving Targets in SAR via Keystoning and Multiple Phase Center Interferometry’, IEEE-2006 Radar Symposium, Verona, NY, April 2006

[3] Stockburger, E. F., Held, D. N., Interferometric Moving Target Imaging, IEEE International Radar Conference, 1995

[4] Zasada, Perry and Sanyal, ‘Detecting Moving Targets in Clutter in Airborne SAR via Keystoning and Multiple Phase Center Interferometry’, SPIE Defense and Security Symposium 17-21 April 2006 in Orlando, Florida.

[5] Sanyal, P. K., Zasada, D. M., Perry, R. P., ‘Tracking Moving Ground Targets from Airborne SAR via Keystoning and Multiple Phase Center Interferometry’, IEEE-2007 Radar Symposium, Boston, MA, April 2007