Detecting and Geo-locating Moving Ground Targets in Airborne QuickSAR via Keystoning and Multiple Phase Center Interferometry

P. K. Sanyal, D. M. Zasada
The MITRE Corp., 26 Electronic Parkway,
Rome, NY 13224 (psanyal, dmzasada)@mitre.org

R. P. Perry
The MITRE Corp., 202 Burlington Road, Rte. 62,
Bedford, MA 01730-1420, rpp@mitre.org

***Abstract*— Without some form of motion compensation, SAR images experience significant range walk and can be quite blurred. In 1997, MITRE reported development of the Keystone Process [1, 2]. The first stage of Keystone Formatting simultaneously and automatically compensates for range walk due to the radial velocity component of each moving target, independent of the number of targets or the value of each target's radial velocity with respect to the ground.**

As is well known, target radial motion also causes moving targets to appear in synthetic aperture radar images at locations offset from their true instantaneous locations on the ground. In an appropriately configured multi-channel radar, the interferometric phase values associated with all non-moving points on the ground can be made to appear as a continuum of phase differences while the moving targets appear as interferometric phase discontinuities [3, 4, 5, 6]. By multiple threshold comparisons and grouping of pixels within the intensity and the phase images, we show that it is possible to reliably detect and accurately georegister moving targets within short duration SAR (QuickSAR) images.

***Index Terms*—Airborne SAR, Geolocation, Interferometry, Keystoning, Surface Moving Targets.**

I. INTRODUCTION

Without motion compensation, Synthetic Aperture Radar (SAR) images of the ground are generally blurred. In 1997, MITRE reported the development technique called the Keystone Process [1] for removing the range migration caused by the radial velocity component of each pixel's movement within the scene, whether moving or stationary with respect to the ground.

The Keystone Process removes the range migration effect of the constant radial velocity component of each pixel's motion relative to the platform from the processed image, whether of the ground itself or from moving targets embedded within the scene. This first pass process allows for automated detection of moving targets via phase thresholding, as have been reported previously in [2, 3].

II. KEYSTONE FORMATTING FOR RANGE WALK CORRECTION

MITRE has developed a technique called Keystone Formatting for motion compensation of targets, the advantage of which is that it can

compensate for several targets moving at different radial velocities simultaneously.

Keystone formatting can be derived by noting that the spectrum of a single received pulse is given by,

$$S_r(f) = P(f)\exp[-i\frac{4\pi}{c}(f + f_0)R(t)] \qquad (1)$$

where

$P(f) = \text{spectrum of transmitted pulse}$

$f = \text{baseband frequency}\,(\frac{-B}{2} \le f < \frac{B}{2}),$

$f_0 = \text{carrier frequency.}$

Expanding $R(t)$ in a Taylor series, we get:

$$R(t) = R(t_0) + \dot{R}(t_0)t + \frac{1}{2}\ddot{R}(t_0)t^2 + \cdots . \qquad (2)$$

Substituting (2) into (1) and dropping cubic and higher order terms,

$$S_r(f) = P(f)\exp[-i\frac{4\pi}{c}(f + f_0)R - i\frac{4\pi}{c}(f + f_0)\dot{R}t - i\frac{2\pi}{c}(f + f_0)\ddot{R}t^2] . \qquad (3)$$

The second term in the brackets containing the product $f\dot{R}t$ gives rise to range walk. This term becomes zero when we use the temporal transformation $t = (\frac{f_0}{f + f_0})t' .$

With the above substitution, (3) can be written as,

$$S_r(f) = P(f)\exp[-i\frac{4\pi}{c}(f + f_0)R - i\frac{4\pi}{c}f_0\dot{R}t' - i\frac{2\pi}{c}(f + f_0)\ddot{R}(\frac{f_0 t'}{f + f_0})^2]. \qquad (4)$$

Since Keystone formatting does not solve the quadratic (or higher order) motion problem, let us also drop the quadratic term in (4) simplify it to:

$$S_r(f) = P(f)*\exp[-i\frac{4\pi}{c}(f + f_0)R - i\frac{4\pi}{c}f_0\dot{R}t']. \qquad (5)$$

Notice that the substitution of t′ for t has removed the phase term that varied with both time and frequency and this removes the linear component of the range-walk. Thus no matter what velocity the target is moving at, it will remain in a given range cell determined by its position at the center (t=0) of the coherent processing interval. Figure 1 shows the 'keystone' nature of the transformation.

Figure 2 shows the effect of Keystoning on the range walk in a real SAR data collection. The left insets show the range-time-intensity (RTI) plots while the right insets show the locations of the peaks in the RTIs.

Standard motion compensation (see, for example, [7]) corrects the range walk for one target at a time. The Keystone process compensates for the different radial velocities of all moving targets simultaneously. Figure 3 shows the result of a simulation to illustrate how targets moving at two different velocities are simultaneously corrected for range-walk.

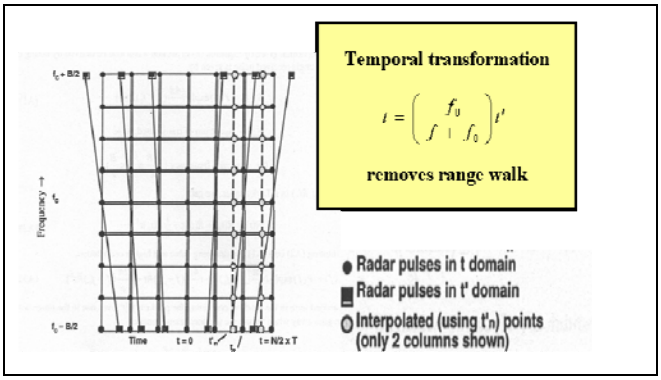

Figure 1. Keystone Formatting Performs Motion Compensation for Targets Moving at Different Velocities

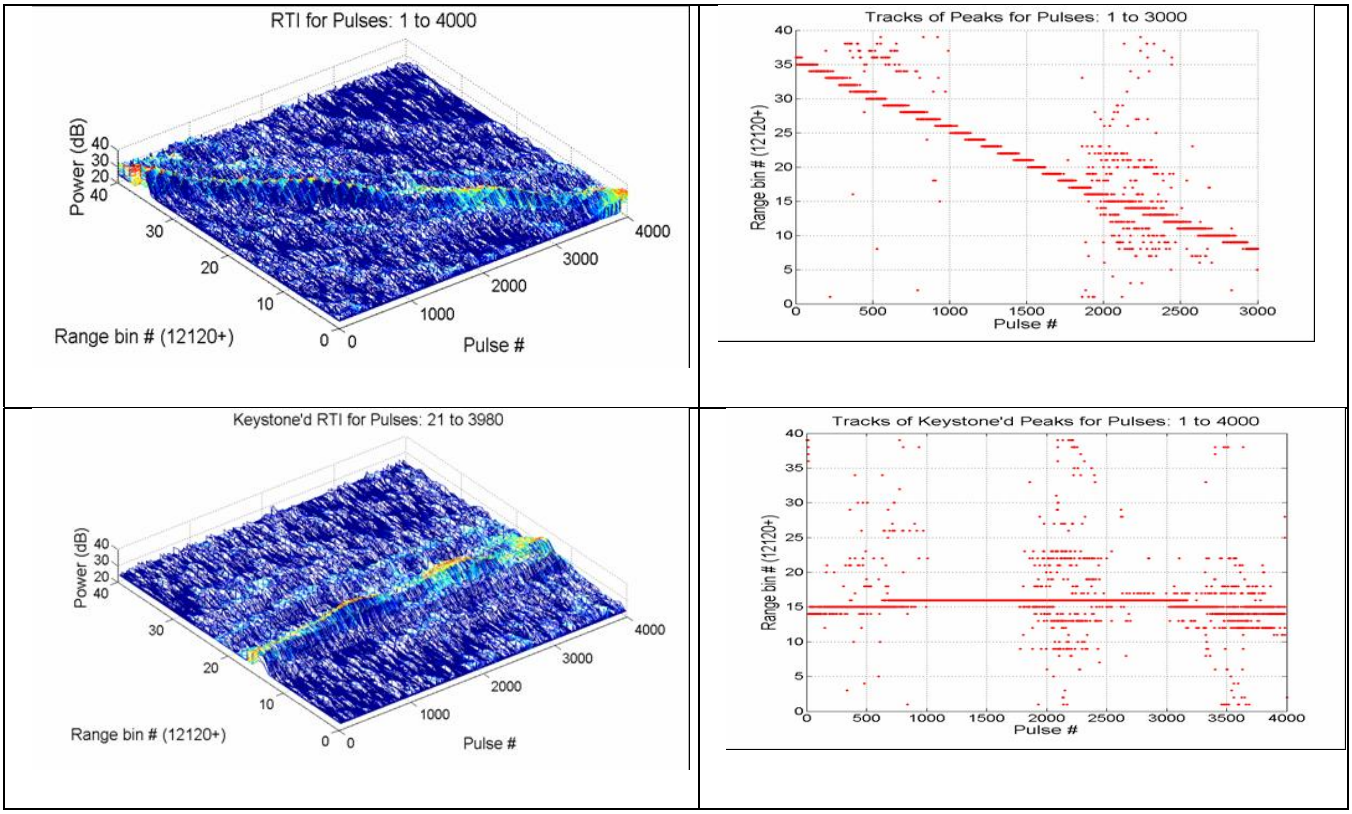

Figure 2. Keystoning removes range walk

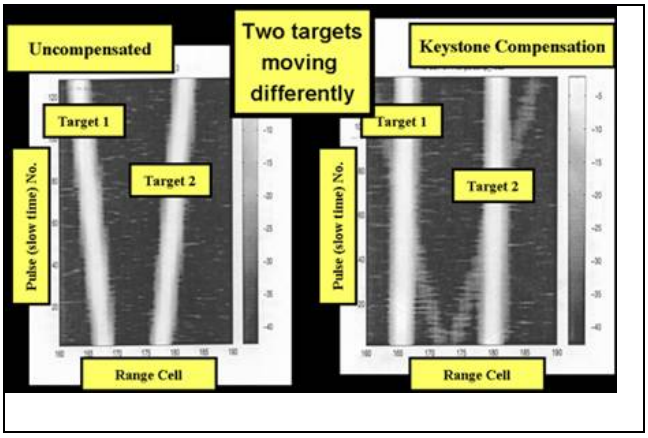

Figure 3. Keystone formatting performs motion compensation for targets moving at different velocities

III. Acceleration Correction

If only Keystone Formatting were applied, the SAR image would only be partially focused due to the quadratic and higher terms in equation 2 above that we have dropped. Therefore, beyond range walk correction, the SAR data also has to be acceleration-compensated to produce reasonably well-focused images. Even when the radar platform is moving at a constant velocity, i.e., in a straight line and at a fixed speed, a point on the ground experiences a significant pseudo-acceleration with respect to the phase center of the radar. Consider the radar-target geometry shown in figure 4.

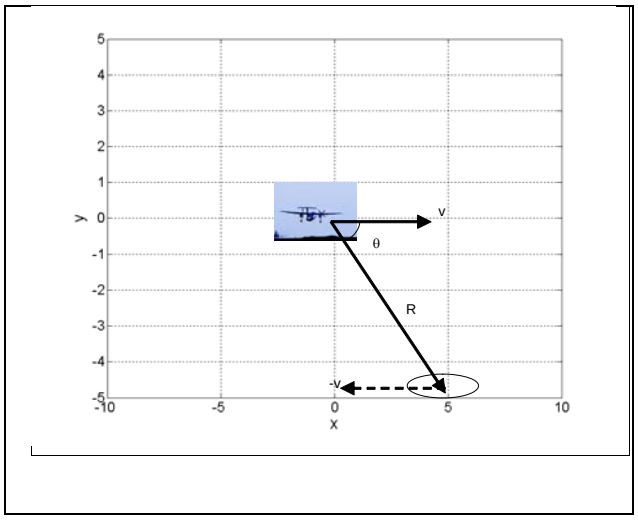

Figure 4. A typical radar-target geometry in SAR

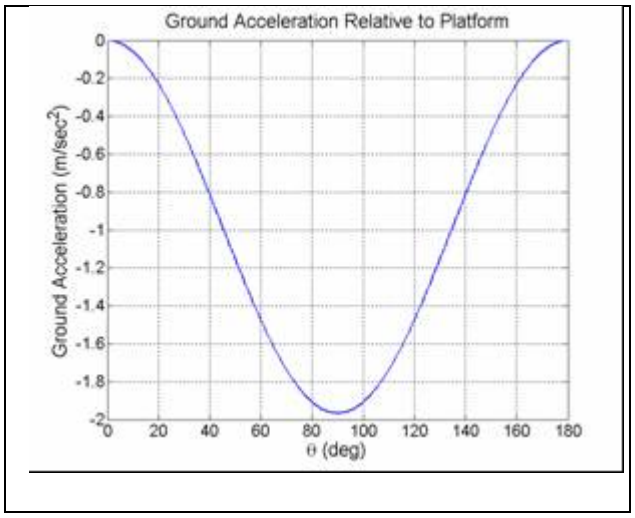

Figure 5. Acceleration of ground points relative to platform

The range, R, of a point located at (x,y) is

$$R = (x^2 + y^2)^{\frac{1}{2}}$$

The radial velocity is then

$$\dot{R} = \frac{x\dot{x} + y\dot{y}}{\sqrt{(x^2 + y^2)}} = \frac{x\dot{x} + y\dot{y}}{R}$$

For the case where $\dot{y} = 0$ and $\dot{x} = -v$,

$$\dot{R} = -v\frac{x}{R} = -v\cos(\theta)$$

and the radial acceleration is then

$$\ddot{R} = -v\left[\frac{\dot{x}}{\sqrt{(x^2 + y^2)}} - x\frac{x\dot{x}}{\left(\sqrt{(x^2 + y^2)}\right)^3}\right]$$
$$= \frac{-v^2}{R} + \frac{v^2x^2}{R^3} = \frac{v^2}{R}\left(\cos^2(\theta) - 1\right)$$

For the LiMIT data [6] discussed here, R = ~22 km, V = ~208 m/s and the ground acceleration as a function of the look angle is shown in Figure 5.

This acceleration has considerable effect on the focusing of the SAR image. Figure 6 shows a range-Doppler-intensity image, i.e., a QuickSAR image formed from 4000 range cells over 4000 pulses (2 seconds) from a single channel from an 8-channel, 180-MHz bandwidth SAR data collected by Lincoln Laboratory on its LiMIT data collection campaign over Ft. Huachuca, AZ. Due to the relatively short duration of the integration time to form the QuickSAR image, the QuickSAR pixels are highly elliptical (approximately 4.5:1 ellipticity in this image).

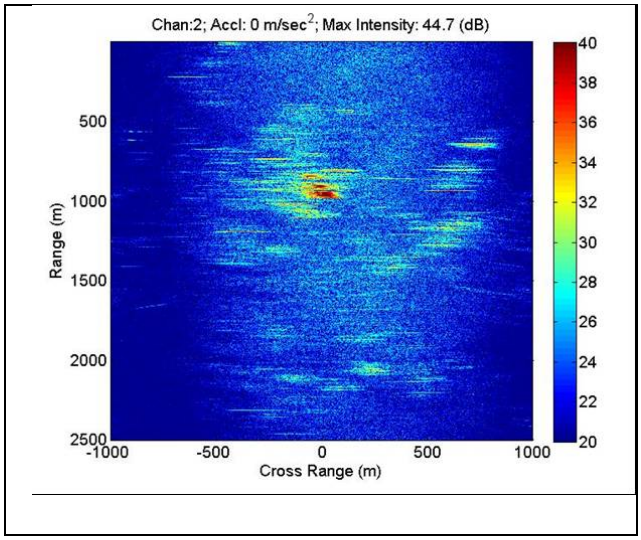

Figure 6. A LiMIT image with keystoning but without an acceleration correction

In this image, the Keystone formatting has been applied but no acceleration correction has been performed.

As figure 5 illustrates, if the radar-target geometry is known, the acceleration correction to be applied can be computed. For our purpose, we adopted a recursive search method for determining the optimum acceleration correction. Figure 7 shows the variation of the maximum of RDI image as a function of the acceleration correction applied.

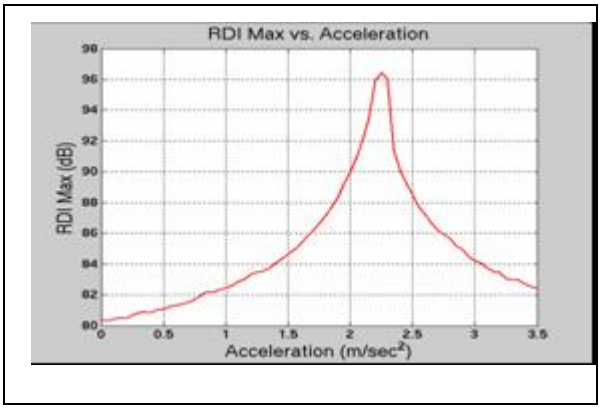

Figure 7. RDI maximum vs. acceleration correction applied to image

It is seen that the RDI maximum increases by almost 16 dB from no correction to the optimum acceleration correction of -2.225 m/s2. Figure 8 shows the image with the optimum acceleration correction along with the optical image for the corresponding area as available from Google.

The high degree of correspondence between the QuickSAR image and the optical image is excellent and major features of the QuickSAR image are readily matched to the optical image.

Since each moving target in the scene may have a different acceleration, the moving targets can be individually and automatically focused after detection using the same procedures discussed above.

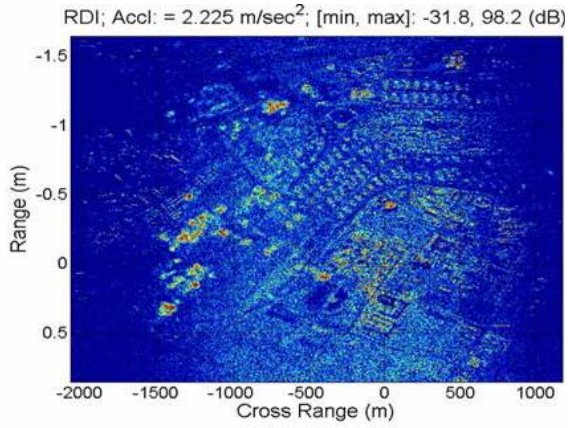

Figure 8. The acceleration-corrected QuickSAR and the corresponding image from Google

IV. Moving Target Detection with Phase Interferometry in QuickSAR

The target motion causes the moving targets to appear at locations different from their true instantaneous locations on the ground in the SAR image. This is due to the coupling of the cross-range position to the target radial velocity and the fact that the moving target and the ground under it have different radial velocities relative to the platform. The result is the well known 'train-off-the-track' or 'boat-off-the-wake' phenomenon (Figure 9).

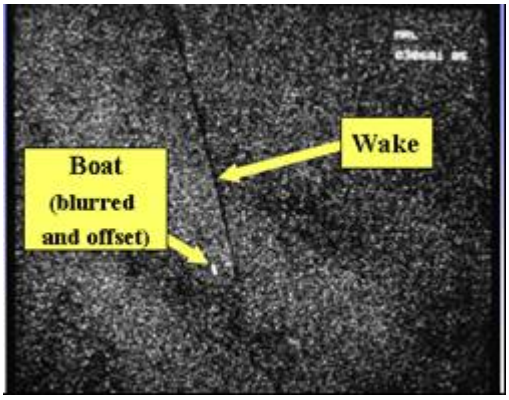

Figure 9. Moving Boat Appears Displaced from its Actual Position in a SAR Image (Source: www.sandia.gov)

Shift in cross-range due to Doppler processing of the CPI

The Doppler, f_D, of a ground point at θ radians from the normal to the radar velocity vector is, for a small angle, given by

$$f_D = \frac{2V\theta}{\lambda}$$

If a moving target has the Doppler f_{targ}, then it will appear shifted in cross-range by an angle θ such that

$$f_{t\arg} = \frac{2V\theta}{\lambda},\ or, \frac{2v_{t\arg}}{\lambda} = \frac{2V\theta}{\lambda},\ \ or, \theta = \frac{v_{t\arg}}{V}$$

At a range R, this amounts to a linear shift of

$$cross-range-shift = R\theta = \frac{Rv_{t\arg}}{V}$$

For R = 22 km, platform V = 208 m/s and v_{targ} (radial) = 30 m/s,

$$cross-range-shift = \frac{22*10^3*30}{208} = 3.173\,\text{km}$$

which is almost the size of the cross-range dimension of the image. As such, the moving target could potentially wrap around the window once.

The 30 m/s target radial velocity produces a Doppler of 1944 Hz, which is close to the PRF of 2000 Hz.

Phase Image

The upper inset of Figure 10 shows a portion of the range-Doppler-intensity (RDI) image, i.e., the SAR image, of the urban area of Ft. Huachuca, Arizona.

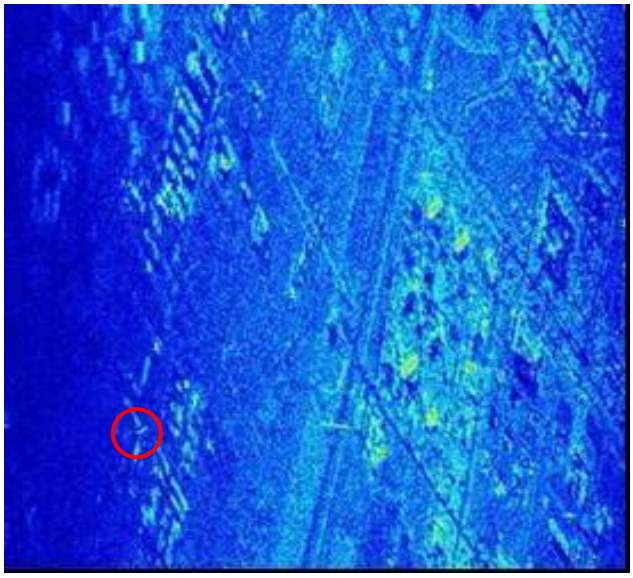

(a) RDI Image from channel # 1

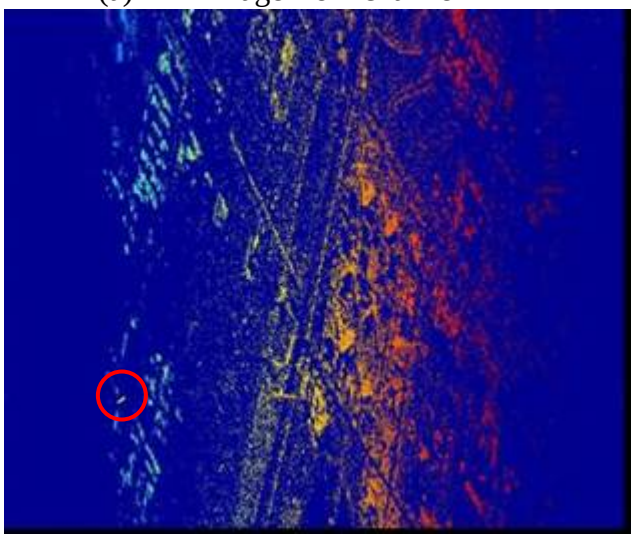

(b) Phase-Difference Image, Channels 1 & 8

Figure 10. RDI and Phase-Difference Images of Ft. Huachuca, Arizona

In this 2-second integration QuickSAR image one can easily discern the roads, buildings and other features of reasonably large size.

To make use of the above-mentioned cross-range displacement phenomenon to detect moving targets in SAR images, we also form a phase interferometry image as a complement to the SAR image.

The lower inset of figure 10 shows the phase-difference image created from the complex images from channels one and eight. It is the color-coded plot of the pixel-to-pixel phase difference between channels one and eight. The phase difference values vary between $-\pi$ and $+\pi$, with red for $+\pi$, gradually changing to blue for $-\pi$.

In an interferometric phase image, all points on the ground should nominally appear as a continuum of phase differences from left to right, i.e., vertical stripes changing in color from blue to red to blue and so on. In the phase image above, we see, as expected, the vertical stripes, changing in color from left to right.

Notice that we also see the roads and buildings and the other features we saw in the SAR image. This is so because we have arbitrarily set the phase difference values for pixels in the SAR which are below a selected dB value (here it is 10 dB) to $-\pi$ with the purpose of reducing speckles. Pixels of low intensity usually consist of system noise rather than ground-returns and hence the channel-to-channel phase difference becomes un-related to the indicated location on the ground and can randomly vary between $-\pi$ and $+\pi$ and the phase-difference image becomes filled with speckles. Setting the above limit reduces the speckles but also helps to delineate areas of low reflectivity and hence the same features we see in the SAR are also visible in the phase-difference image.

Now, while all the points on the ground appear as a continuum of phase differences from left to right of the SAR image, the moving targets appear as discontinuities. This is so because, as discussed above, a pixel that represents a moving object appears displaced from the ground point on which it is instantly situated. This is because the rate of change of the phase to the moving target is different from the rate of change of the phase to the ground point. And yet, the difference in the phases of the moving-target pixels in the two channels is the same as the difference in the phases of the ground point. This causes a phase discontinuity. In the color-coded phase difference image, it shows up as a color anomaly. By searching for color, i.e., phase, discontinuities, one can easily identify several moving targets in the interferometric image. One such target is circled in red. The color anomaly may not be obvious in this highly shrunk image but is quite clear in a higher resolution version and, in particular, easily detected by a phase discontinuity detection algorithm. In this case, we did not have the truth data from instrumented moving targets to verify if the detections correspond to known moving targets, but this correspondence has been verified in other data sets with truth data.

Automated Phase Discontinuity Detection

A trained operator would be able to detect moving targets in a color-coded phase difference image as described above. To avoid dependence on human operators, we have implemented the following technique for automated target detection.

1. Apply an appropriate acceleration correction to focus the ground plane
2. Fit a reference plane, in a least-square error sense, to the phase difference image.
3. Look for pixels whose phase differences from the reference plane exceed a given threshold.
4. Of the above, select those that also exceed an amplitude threshold.

We illustrate this technique with results from another data set (General Dynamics DCS 8-Channel, 160-MHz bandwidth data from Eglin, FL.). In this scenario, there were seven targets moving over the runway at an airfield, and the GPS data from the vehicles as well as the motion data from the radar platform are available.

The top inset in figure 11 shows a 0.83 sec QuickSAR of the portion of the runway where the

targets were moving around. This QuickSAR is generated from the channel 2 data.

The lower inset in figure 11 is the phase-difference image using channels 2 and 5. Note that the runway is clearly discernable. The numbers in green indicate the actual positions of the moving targets at the time the data was collected while the displaced red numbers indicate where they should appear in the range-Doppler SAR image. It is interesting to note that target #5, which is a large target, shows up more prominently in the phase-difference image (red circle) than in the QuickSAR image (though the color anomaly is not quite as obvious in the small scale used here; it is much clearer in a larger display of the image). This visual observability is very useful to a human operator because he/she can easily place the detections in the 'context' of the background terrain, i.e., is it near a building, a road, etc.

In the following, we describe how a machine detects the moving targets. The blue dots in the upper left inset in figure 12 are the row-by-row plots of the phase differences shown in figure 11. The red dots (which appear as a line because of the compactness) are the phases from the plane fitted to the actual phase differences.

Ideally, all points on the ground would have phase differences that lie on a plane and only moving targets will have phase differences that deviate from the plane. But in reality, because of the system noise, variations in terrain elevation, etc., the actual phase differences, of course, do not all lie on the plane. Smaller the pixel power, more likely it is to be affected by system noise and thus more likely it is to deviate from the plane. The upper right inset of figure 12 shows that a very large number of pixels exceed the selected 1-radian phase deviation threshold from but a significant number of them are below 15 dB.

On the other hand, the pixels corresponding to the moving targets we expect to detect, should be of sufficient strength as well as deviate significantly from the phase plane. Thus the moving targets have to meet the dual thresholds: one in power and the other in phase deviation. If the power threshold is set at 15 dB in this case, the number of pixels that deviate from the plane by more than the 1-radian phase deviation threshold drops to 85, as shown in the lower left inset of figure 12.

The lower right inset in figure 12 shows these pixels as the detected targets. These are indicated with red 'x's. In this particular image, there were as many as 85 detections. While some of these are likely to be false alarms, most of these are possibly multiple detections from 'extended' targets. In fact, an observer can readily group many of these detections into possible clusters that belong to the same targets.

Clustering of the detections

We have used an automatic clustering algorithm from Matlab, which results in eight clusters. These clusters are depicted with red circles. Note that the circles only indicate the center of the clustered detections; they do not indicate which detections belong to the cluster or how many detections were included in the cluster.

Notice that one cluster coincides with the expected location of target #2, two appear to coincide with target #1, four clusters appear to coincide with target #5 (which is known to be a large target) and the eighth cluster appear to comprise one or more false alarms.

A human observer examining the four clusters near target #5 or the two near target #1 would have readily combined them into single clusters each. The Matlab fuzzy-logic clustering algorithm used could possibly have been tweaked to yield similar results but we did not attempt to do so at this time.

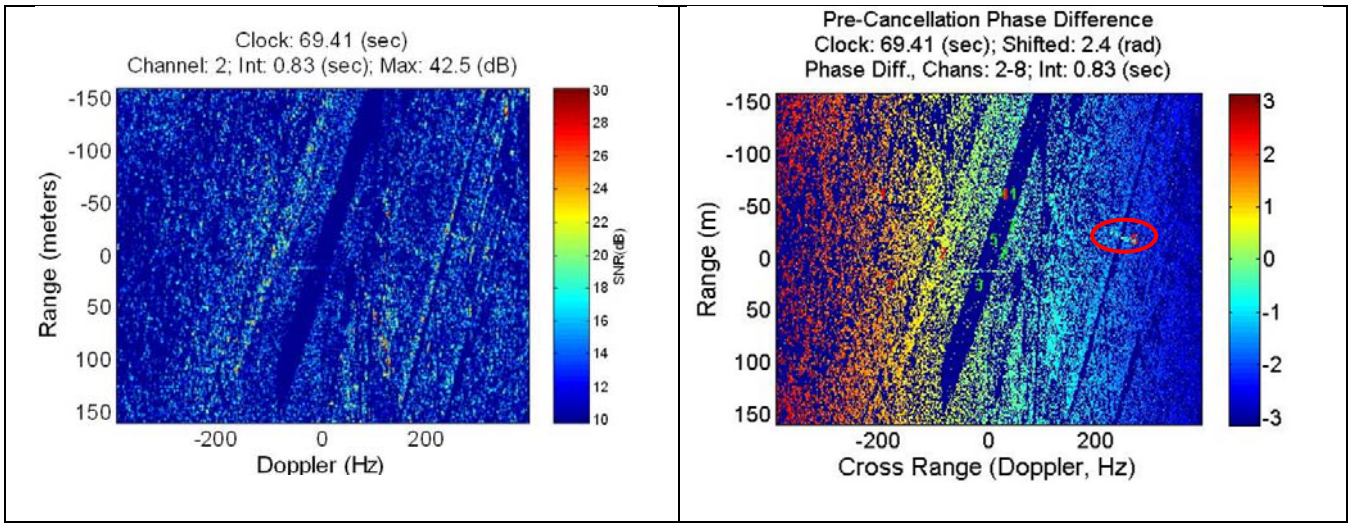

Figure 11. QuickSAR and Phase-Difference image of the Eglin runway with vehicular traffic.

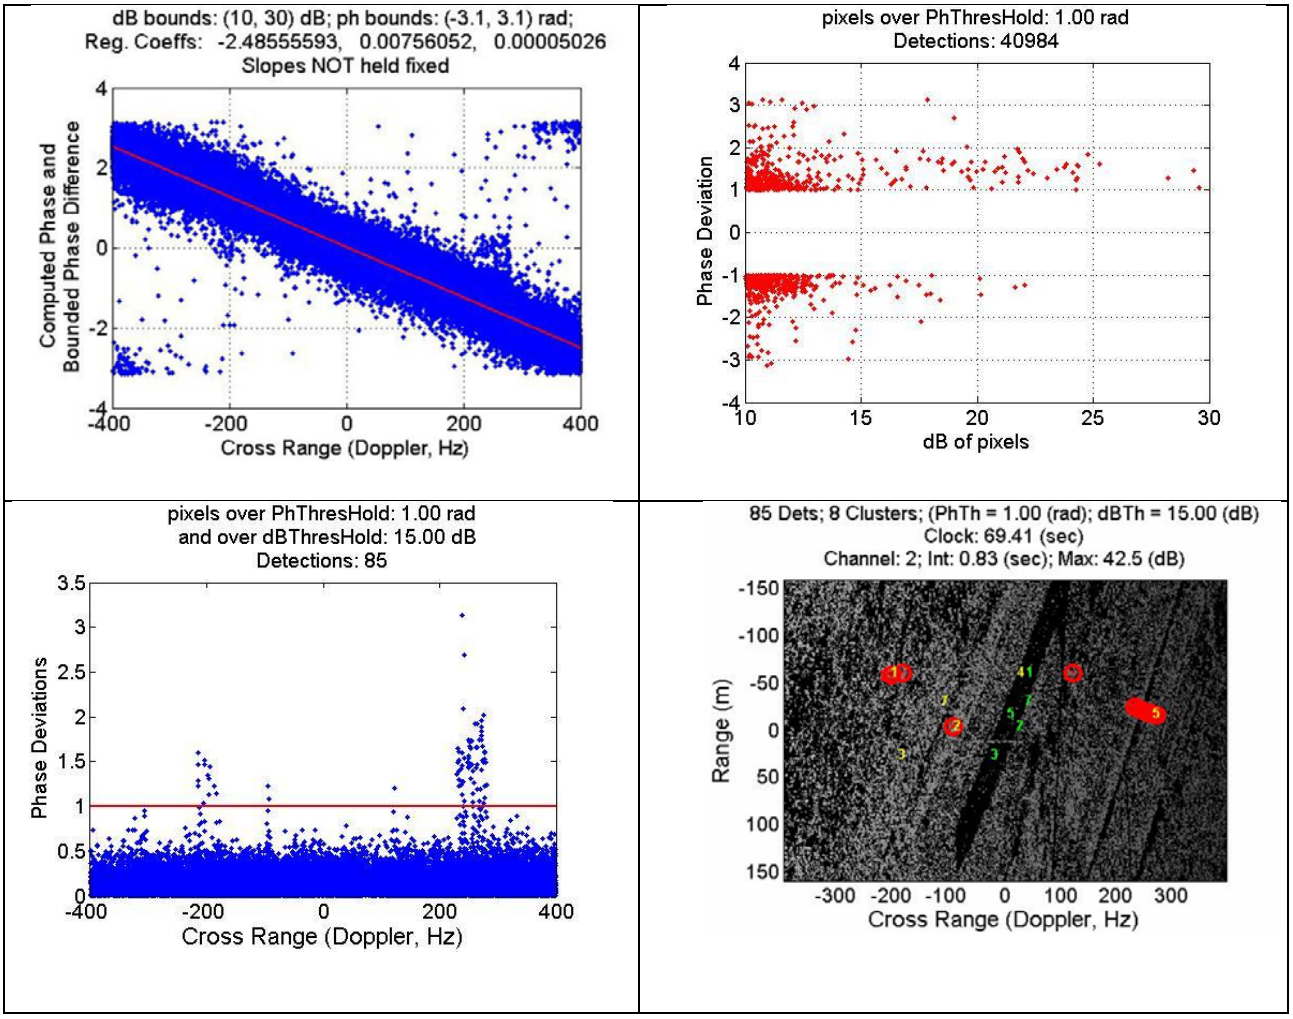

Figure 12 Machine detection of moving targets in a QuickSAR

V. Geolocation of Detected Moving Targets in QuickSAR

The 'apparent' locations of the moving targets detected in the SAR image are not their actual locations on the ground. In fact, recall that the detection of the moving targets in the SAR image depends on the fact that they are displaced, in cross-range only, from their actual locations. Therefore, after detection, they need to be 'geo-located', i.e., they need to be placed correctly at their actual locations.

Fortunately, the geo-location is almost a by-product of the detection process. The left inset in figure 13 shows the phase difference of the detected pixels (blue dots) and the plane fitted to the phase differences of all the ground points (red dots, which have merged to appear as a line).

Since the moving targets are also on the ground (that is the assumption), the actual locations of the detected pixels are found by displacing them in cross-range until they lie on the red line, as indicated in the left inset of figure 13.

This displacement may be done on a detection-by-detection basis or a cluster basis. In a detection-by-detection-basis geolocation, each detection is moved independently. As may be obvious, because of the various amounts of noise in the individual pixels, the pixels that appear to be in a cluster before geolocation can become scattered after the above-described displacement.

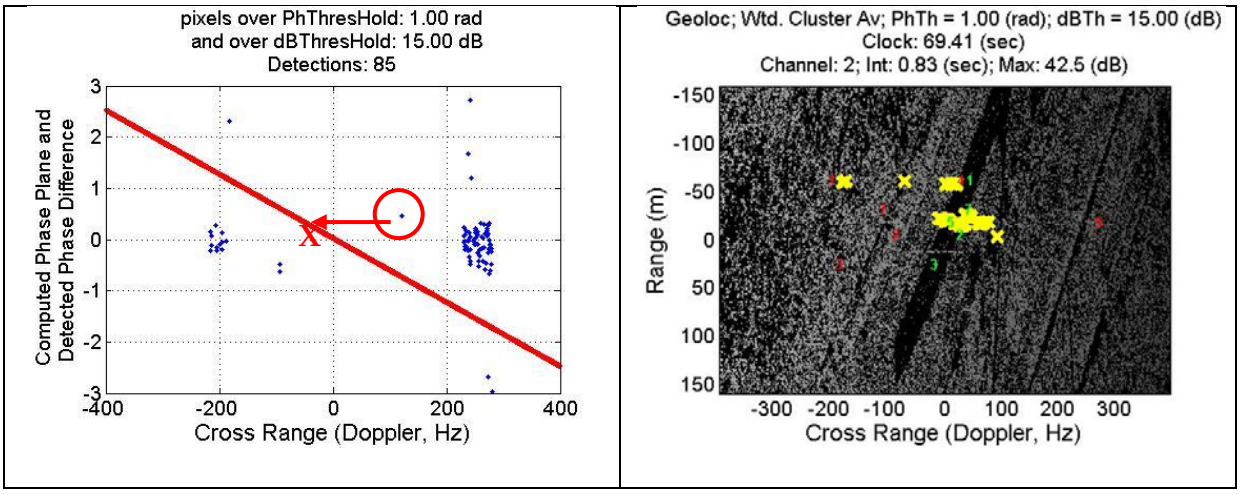

Figure 13 Geolocation of moving targets in a QuickSAR

Cluster move

Since a clustering has already been performed, it is easy to move all the detections in a cluster together to appear as the same cluster after the geo-location. For this, one computes the average phase difference for the cluster and uses this value to relocate the whole cluster on the red line.

Weighted cluster move

A slight twist on the above technique is to find a weighted average of the phase difference using the pixel power to weight the phase difference values. This plausibly reduces the effect of low-power pixels which are likely to be noisier.

The right inset in figure 13 shows the geo-located detections using the weighted average method. It is seen that the majority of the detections are now appear on the runway where they actually belong.

Here we have shown only one image 0.83-second frame of the data set from this SAR data collection scenario which lasts for about 250 seconds. We have applied this technique to the 300 or so frames form this data set and have obtained consistent detection and geolocation

results.

VI. Conclusion

In this paper, we have described an interferometric scheme for detecting and geo-locating surface moving targets in multi-channel SAR. The usual ground moving target Indication (GMTI) schemes can detect moving targets when they have sufficient velocity to appear clear of the ground clutter. The interferometric scheme described here is able to detect moving targets well within the ground clutter.

Since the detections take place within a SAR image, and further, they are properly geo-located, this has the additional benefit of 'contextual' detection, i.e., the detections are already 'in context', that is to say, one can see if they are on roads, runways, near buildings, etc.

All the results included here are from real multi-channel radar data collected by two different organizations.

Acknowledgment

We would like to thank AFRL Rome Research Site in Rome, NY for providing us with the various radar data sets used in our work.

References

[1] Richard P. Perry, Robert C. DiPietro and Ronald L. Fante, The MITRE Corporation, 1999. 'SAR Imaging of Moving Targets'; IEEE Transactions on Aerospace and Electronic Systems, Vol. 35, No. 1, pp. 118-199

[2] Richard P. Perry, Robert C. DiPietro and Ronald L. Fante, 'Coherent Integration With Range Migration Using Keystone Formatting' IEEE Radar conference, April 2007, Waltham, MA

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

[4] 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

[5] 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

[6] 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 in Orlando, Florida.

[7] Preiss, Gray, and Stacy, 'The Effect of Polar Format Resampling on Uncompressed Motion Phase Errors and the Phase Gradient Autofocus Algorithm' IEEE-2001 0-7803-7031 July 2001.