Circular SAR GMTI

Douglas Page, Gregory Owirka, Howard Nichols[a],
Steven Scarborough[b]
[a]BAE Systems Technology Solutions, 6 New England Executive Park, Burlington, MA 01803
[b]AFRL/RYA, 2241 Avionics Circle, WPAFB, OH 45433

ABSTRACT

We describe techniques for improving ground moving target indication (GMTI) performance in multi-channel synthetic aperture radar (SAR) systems. Our approach employs a combination of moving reference processing (MRP) to compensate for defocus of moving target SAR responses and space-time adaptive processing (STAP) to mitigate the effects of strong clutter interference. Using simulated moving target and clutter returns, we demonstrate focusing of the target return using MRP, and discuss the effect of MRP on the clutter response. We also describe formation of adaptive degrees of freedom (DOFs) for STAP filtering of MRP processed data. For the simulated moving target in clutter example, we demonstrate improvement in the signal to interference plus noise (SINR) loss compared to more standard algorithm configurations. In addition to MRP and STAP, the use of tracker feedback, false alarm mitigation, and parameter estimation techniques are also described. A change detection approach for reducing false alarms from clutter discretes is outlined, and processing of a measured data coherent processing interval (CPI) from a continuously orbiting platform is described. The results demonstrate detection and geolocation of a high value target under track. The endo-clutter target is not clearly visible in single-channel SAR chips centered on the GMTI track prediction. Detections are compared to truth data before and after geolocation using measured angle of arrival (AOA).

Keywords: GMTI, SAR, STAP, MRP, Gotcha Radar, Change Detection

1. INTRODUCTION

Performing GMTI in SAR radar systems may be challenged by several factors. Moving targets, especially if they are accelerating, can produce defocused responses in SAR images[1]. Correspondingly, the peak signal to noise ratio (SNR) of a defocused target response will be reduced compared to that of the same target when it is stationary. If significant clutter is also present in the pixel(s) of the image containing the target response, additional degradation of the SINR will occur, causing detection and tracking performance to suffer.

In this paper, we describe application of a set of specialized techniques for compensating for such effects in Gotcha radar systems. The Gotcha concept[2] employs a high revisit rate and fine radar resolution, along with a wide antenna beam to perform continuous radar surveillance. In Gotcha radar systems, significant clutter competing with moving targets of interest will likely be present over a wide area spanned by the antenna footprint. As a consequence, clutter cancellation techniques are needed, especially in urban environments. As an example of prior work, Deming[3] studied along-track interferometry (ATI) and displaced phase-center antenna (DPCA) for Gotcha radar systems.

Our approach employs a combination of MRP to compensate for moving target defocus and STAP to mitigate the effects of clutter interference. MRP[1] in its original form reprocesses phase history data in order to produce a focused response from a moving target having known motion parameters. Here, we utilize a more efficient form of MRP that operates on localized SAR chips. STAP[4,5] performs adaptive filtering based on local clutter statistics and an assumed target "steering" vector in order to maximize SINR for a target of interest. We define a set of adaptive DOFs for STAP tailored to the SAR GMTI problem. In addition to STAP and MRP, we use data from previous platform orbits and employ a noncoherent form of change detection in order to reduce false alarms from clutter discretes. Finally, we make use of tracker feedback from a high value target track. The track state is used to localize the processing and separate the target of interest from other movers and false alarms.

An outline of the paper is as follows. We first describe the SAR MRP processing in Section 2 and show the results of processing simulated target and clutter responses. Section 3 discusses selection of adaptive DOFs for STAP, and

describes detection and parameter estimation of moving targets. The results of applying MRP and STAP to the simulated target in clutter example are shown. SINR loss curves using our approach are compared to results using more standard algorithm configurations. Section 4 shows the results of processing a measured CPI from collected Gotcha radar data. The change detection technique used to reduce clutter false alarms is described and shown to be effective on the processed CPI. Detection and geolocation of a truthed endo-clutter target under track are demonstrated.

2. DESCRIPTION OF SAR MRP

Our SAR MRP model assumes that a polar format procedure has been applied to a set of complex phase histories for multiple antenna phase centers located on a moving radar platform. Each phase history is a function of frequency (or wavelength) and radar pulse (or slow-time). The polar format algorithm produces a set of functions $P(n,\mathbf{k}_x,\mathbf{k}_y)$ defined on a two-dimensional rectangular k-space grid, where *n* labels the antenna phase center. Standard image formation would be performed by applying a two-dimensional discrete Fourier transform over the variables $\mathbf{k}_x,\mathbf{k}_y$, producing a SAR image for each of the antenna phase centers.

Originally, MRP was applied directly to phase history data prior to image formation[1] in order to focus moving target responses. In the approach discussed here, a more efficient implementation of MRP is used that applies a filtering operation in k-space after polar format, but prior to the Fourier transform:

$$G(n,\{f\},x_{im},y_{im}) = \sum_{\mathbf{k}_x=\mathbf{k}_{x,\min}}^{\mathbf{k}_x=\mathbf{k}_{x,\max}} \sum_{\mathbf{k}_y=\mathbf{k}_{y,\min}}^{\mathbf{k}_y=\mathbf{k}_{y,\max}} \exp(-j\cdot[\mathbf{k}_x\cdot x_{im}+\mathbf{k}_y\cdot y_{im}])\cdot F(\{f\},\mathbf{k}_x,\mathbf{k}_y)\cdot P(n,\mathbf{k}_x,\mathbf{k}_y) \quad (1)$$

Here, $\{f\}$ represents a set of derived target MRP motion parameters which specify the MRP filtering function *F*. Note that a filtering function that is constant across k-space will produce standard SAR images for each antenna phase center. By applying the MRP filtering in the k-space domain rather than the raw phase history domain, reduced processing time can be achieved when tracking high value targets. This is because the k-space grid that is summed over in Eq. (1) can be chosen as a subset of the full k-space domain, corresponding to the discrete Fourier transform domain of a small image chip containing the target under track.

The number of elements in $\{f\}$ and the form of the filter function *F* used in Eq. (1) depends on the target motion model assumed. For a specified radar imaging geometry, an assumed motion model leads to a specific k-space dependence of the function *P* appearing in Eq. (1). Qualitatively, for stationary targets close to the radar aimpoint the phase of *P* will have a linear dependence on the k-space coordinates. This in turn leads to a focused scene around the aimpoint, without any additional filtering. For moving targets, the phase of *P* will in general have a significant nonlinear dependence on the k-space coordinates. This nonlinearity leads to defocused SAR responses if no compensating filtering is performed. The purpose of the MRP filter function *F* is to compensate for the nonlinearity of the phase of *P* and produce a focused moving target SAR image response.

We illustrate how MRP compensates for defocused SAR responses using a simulation example. Figure 1 shows the effect of applying MRP to the simulated SAR responses of a moving target and a stationary point scatterer representing clutter that is present in the target pixel. The target is also modeled as a point scatterer, but is given a constant velocity and acceleration. The MRP filter is calculated based on the known target model and SAR imaging parameters. The defocus of the target response prior to applying MRP is evident in the upper left corner of Figure 1. The upper right corner of Figure 1 shows that, after MRP, the target is well focused in a single image pixel. The bottom half of Figure 1 shows the effect of applying the MRP filter (again based on the moving target motion model) to the simulated response of the stationary clutter discrete. The lower left plot in Figure 1 shows that the clutter scatterer response is well focused before MRP, while the lower right plot shows that the clutter response is defocused after applying MRP. This behavior is expected, since the SAR imaging filter (i.e. SAR without MRP) is an approximate matched filter for a stationary scatterer. The MRP filter, which was based on the motion model of an accelerating target, is a highly mismatched filter for the stationary clutter scatterer, and therefore produces defocus and a drop in peak amplitude.

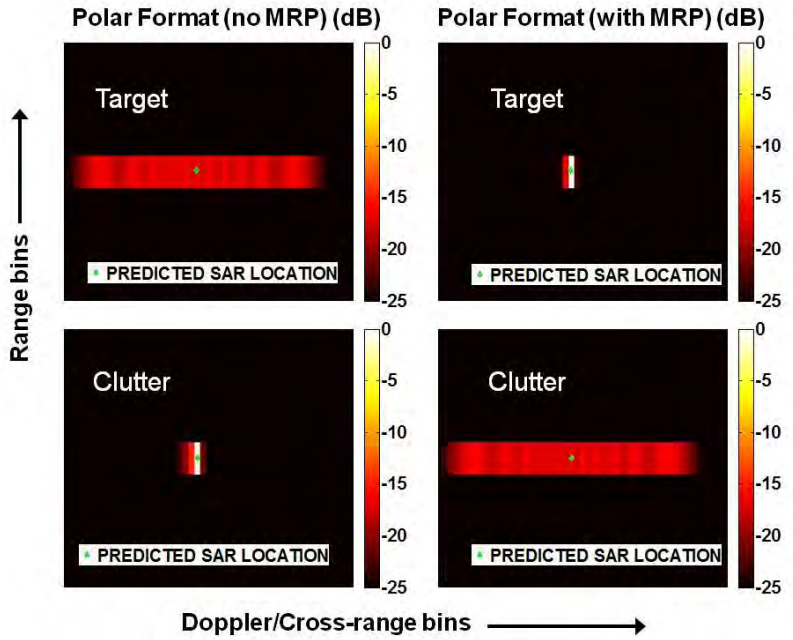

Figure 1 Simulated SAR image responses of a slowly moving accelerating target (upper subplots) and single stationary point clutter scatterer (lower subplots). Response is shown before (left subplots) and after (right subplots) MRP is applied: i) upper left subplot shows target response before MRP ii) upper right subplot shows target response after MRP iii) lower left subplot shows clutter response before MRP iv) lower right subplot shows clutter response after MRP. Predicted SAR locations based on the known scatterer parameters are indicated by the green diamonds in each subplot.

3. STAP FILTERING OF SAR MRP DATA

3.1 Formation of adaptive degrees of freedom

STAP[4,5] applies filtering across a set of adaptive DOFs in order to improve target SINR in multi-channel radar data. The selection of DOFs for STAP involves a tradeoff between several factors. More DOFs generally will produce better SINR, if the interference covariance matrix can be estimated accurately. However, the amount of training data required to produce accurate covariance estimates scales with the number of DOFs. Moreover, the computational processing load increases rapidly as the number of DOFs increases. Reduced-DOF approaches are therefore of interest. Factored post-Doppler STAP[5] defines the adaptive DOFs as multiple antenna phase centers within a range-Doppler bin. This choice of spatial-only adaptive DOFs can produce good performance, but improved SINR can result by adding adaptive temporal DOFs. Some typical choices[5] are the "PRI-staggered" DOF set constructed from different time-staggered sub-CPIs, or the "multi-bin" approach which stacks data from multiple Doppler bins.

We define here a DOF set tailored to the SAR-GMTI problem as follows:

$$\vec{\mathbf{g}}(x_{im}, y_{im}) = \begin{bmatrix} G(1, \{f\}_0, x_{im}, y_{im}) \\ G(2, \{f\}_0, x_{im}, y_{im}) \\ \vdots \\ G(N_n, \{f\}_0, x_{im}, y_{im}) \\ G(1, \{f\}_1, x_{im}, y_{im}) \\ G(2, \{f\}_1, x_{im}, y_{im}) \\ \vdots \\ G(N_n, \{f\}_1, x_{im}, y_{im}) \\ \vdots \\ G(N_n, \{f\}_{N_f-1}, x_{im}, y_{im}) \end{bmatrix} \tag{2}$$

In each SAR pixel, complex data samples for different antenna channels (n) and different MRP parameter sets ($\{f\}_k$, $k=0,1,...N_f-1\}$) are stacked to form the adaptive DOFs. Qualitatively, the goal of using the different MRP parameter sets appearing in Eq. (2) is to reduce the overlap of normalized target and clutter response vectors for a given target motion state. More precisely, we desire to reduce the projection of the clutter response onto the one-dimensional subspace defined by a normalized target response vector. The normalized target response vector, in turn, is a function of the motion model and motion state estimate of the target under track.

3.2 Adaptive matched filter (AMF) calculation

The adaptive matched filter (AMF)[6] can be calculated for a given adaptive DOF set in order to maximize SINR and perform threshold detection. The equation defining the AMF utilizes an estimated covariance matrix; the equation we use defining this covariance estimate is:

$$\hat{\mathbf{R}}(x,y) = \frac{1}{N_{y'}} \sum_{y' \neq y} \vec{\mathbf{g}}(x,y') \vec{\mathbf{g}}(x,y')^H \tag{3}$$

The "H" superscript indicates that the conjugate transpose is taken. The sum is over range pixels (y') in the same cross-range column (x) as the test cell, but excluding the test range cell (y). This training technique is a standard post-Doppler approach that excludes the test cell to reduce the effect of target contamination of the clutter covariance estimate.

The target to be detected has an assumed response across the adaptive DOF set defined by a steering vector $\vec{\mathbf{s}}(\alpha)$, where α refers to a set of target parameters. For our processing, a bank of steering vectors is defined corresponding to different angles of arrival (AOAs) and MRP motion states (i.e., α is a vector parameter set containing AOA and the elements of the MRP parameter set f). A search over α will be performed in order to obtain maximum likelihood estimates, as will be shown below.

With these definitions, the parameter-dependent AMF in a pixel (x,y) is given by

$$AMF(\alpha,x,y) = \frac{|\vec{\mathbf{s}}(\alpha)^H \hat{\mathbf{R}}(x,y)^{-1} \vec{\mathbf{g}}(x,y)|^2}{\vec{\mathbf{s}}(\alpha)^H \hat{\mathbf{R}}(x,y)^{-1} \vec{\mathbf{s}}(\alpha)} \tag{4}$$

Potential target detections are identified by applying a threshold to the AMF after maximizing over the target parameters:

$$\max_\alpha AMF(\alpha,x,y) >= T \tag{5}$$

The AMF is known to have a CFAR property[6], and also acts as a ML parameter estimator[7]. For our processing, estimates of the AOA and MRP motion state are obtained as the corresponding parameter values that maximize the AMF:

$$\hat{\alpha}(x,y) = \arg\max_\alpha AMF(\alpha,x,y) \tag{6}$$

3.3 SINR Loss in a simulated example

The so-called SINR loss factor is often used to characterize the relative STAP performance using different adaptive DOFs. We can calculate the SINR loss for the simulation example of the last section, using the known target parameters to define the steering vector and data vector appearing in Eq. (4). Additionally, the estimated covariance matrix in Eq. (4) is replaced by a calculated ideal covariance matrix for the simulated clutter scatterer plus thermal noise, and the data vector $\vec{\mathbf{g}}(x,y)$ is replaced by a simulated target data vector across the adaptive DOFs. With these substitutions, the value of the AMF in Eq. (4) gives the simulated SINR of the target after MRP and STAP. The SINR loss factor is then obtained by dividing the result by a scale factor given by the "full-DOF" SINR for thermal noise interference (i.e. no clutter) using all degrees of freedom (i.e. all channels, pulses, and wavelength samples). Since the ideal covariance matrix is used, the resultant SINR loss will not include the effects of training losses due to finite training samples and inhomogeneous clutter. However, it can give insight into the relative STAP performance that can be expected using different adaptive DOF sets.

Figure 2 shows the effect of employing different DOFs on the output SINR loss for the simulated example shown in Figure 1, as a function of the clutter to noise ratio (CNR) of the clutter scatterer. Two of the curves shown, namely the

black and red curves, correspond to performing only STAP or only MRP respectively. Before discussing the combined STAP-MRP approaches, it is useful to consider the STAP-only and MRP-only results.

The black curve shows the performance of factored STAP, without any MRP applied, while the red curve shows the performance of MRP without any STAP (but with conventional beamforming applied across the antenna channels). Note that the gap between the black and red curves varies from > 10 dB at low CNR to nearly 0 dB at high CNR. The gap between the two curves at low CNR corresponds to the improvement in SNR that is achieved by performing MRP. By focusing the target response, SNR is greatly improved using MRP. Clutter is not an issue at low CNR, so that SINR is approximately equal to SNR. As the CNR is increased, note that the red curve falls off more steeply than the black curve. This is because the red curve was produced without any clutter cancellation, while the black curve was produced using STAP. STAP performance is seen to be less sensitive to the CNR, due to adaptive nulling of the clutter. However for these simulation parameters and CNR extent, better overall performance is seen using MRP-only than STAP-only. For higher values of CNR than shown in Figure 2, the red curve would eventually cross below the black curve.

The magenta curve in Figure 2 corresponds to a non-adaptive application of MRP, followed by factored STAP across the antenna channels. Note that the magenta curve is well above the black curve at low CNR, again showing the benefits to SNR of applying MRP. The magenta curve is also seen to be less sensitive to CNR than the red curve, which shows the benefits of adaptive clutter nulling. MRP followed by factored STAP is therefore seen to perform better than either the STAP-only or MRP-only algorithms.

The best performance however is seen in the blue curve, which was produced using Eq. (2). Utilizing different MRP states as well as different antenna channels in forming the adaptive DOFs is seen to produce more effective clutter nulling, so that essentially no drop in SINR is seen as the CNR increases. STAP with adaptive MRP DOFs is therefore observed to produce the best performance in this simulated example. While this example is a highly simplified case, it does contain essential aspects of real-world situations. Often in Gotcha radar systems, moving targets are well focused using a constant acceleration model, and in urban environments there are bright specular returns that appear point-like.

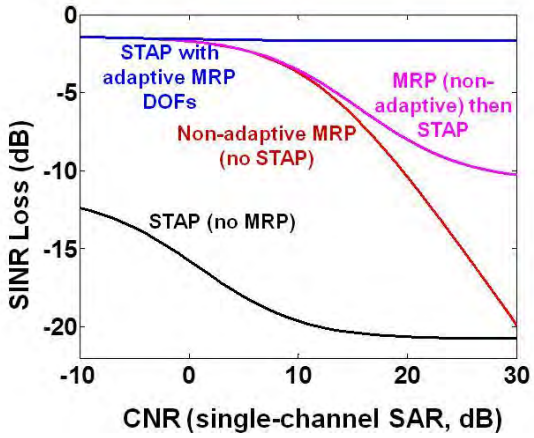

Figure 2 SINR Loss vs. CNR for simulated constant acceleration MRP example

4. DETECTION PROCESSING OF A MEASURED DATA CPI

Figure 3 shows the extracted SAR image chip for a single antenna channel produced by processing measured Gotcha data. The imaging geometry showing the velocity vectors of the radar platform and estimated motion state of a target under track are shown in the left half of Figure 3, while the actual chip is shown on the right half of Figure 3. GPS truth was available for the tracked target, and the truth location of the target (including Doppler shift due to target velocity) is shown as the green diamond just above the center of the chip. The actual center of the chip was determined by the predicted track state, so due to track errors as well as GPS truth errors the green diamond does not lie exactly in the center of the chip. The horizontal lines on the right half of Figure 3 show the search area over which detection processing was performed. Generally, the range extent of the extracted chip is larger than the detection search area, in order to give more training data for both image registration (to a reference orbit) and covariance estimation (for STAP).

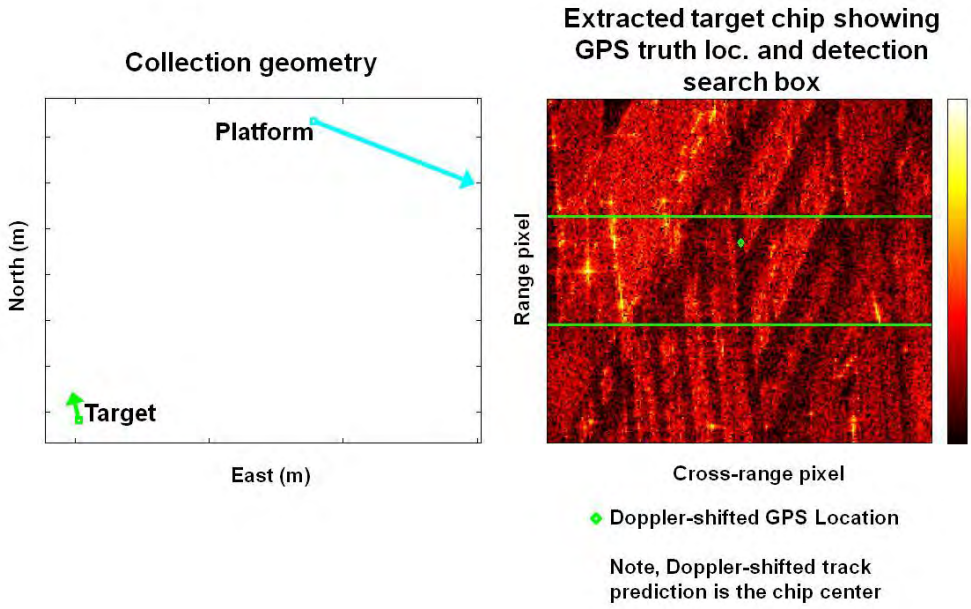

Figure 3 Collection geometry (left) and single-channel SAR chip (right) for measured data CPI example

The extent of the detection search area shown on the right half of Figure 3 is determined by the covariance matrix of the track state. The covariance measures the uncertainty of the target track state, and the processing uses that covariance to determine the area over which to search for GMTI detections from the target under track. The search area determined by the track covariance is enlarged to allow for systematic errors, as well as the maximum expected SAR defocus of the target response. Compared to a wide area surveillance mode, the use of tracker feedback to localize the detection search area allows lower detection thresholds to be used, for a given false alarm rate. This is because for a given threshold setting the probability of a false alarm over a localized chip is smaller than for a large image.

In addition to forming multi-channel "mission" pass SAR images containing the vehicle of interest, "reference" pass SAR images from a previous platform orbit are also formed on each CPI. Reference SAR chips are extracted and aligned with the mission chips on each CPI. Figure 4 compares the mission and reference orbit SAR images that were formed for the measured data example. The images are shown zoomed over the detection search box contained within the green lines in the plot on the right half of Figure 3. Note that even though the target is only expected in the mission orbit image, the Doppler-shifted GPS truth location is shown on both images to make comparisons in the vicinity of the expected target response easier to see.

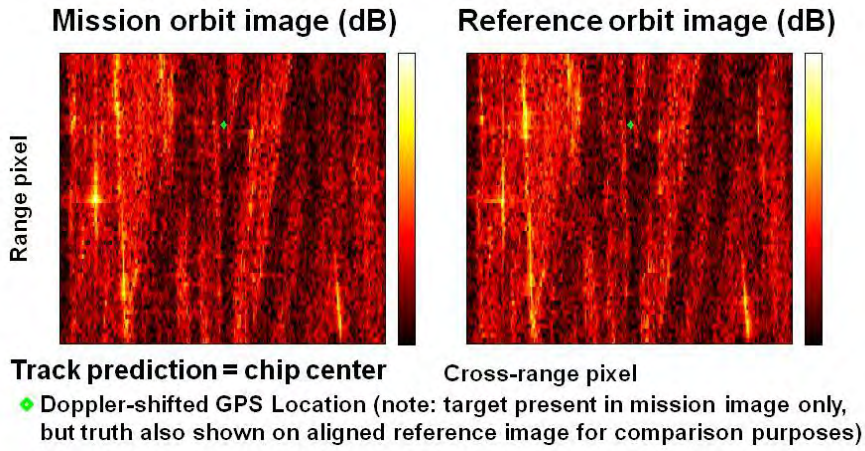

Figure 4 Single-channel SAR images for mission (left) and reference (right) orbits for the measured data example shown in the previous figure (images are shown zoomed over detection search box contained within green lines on previous figure)

The reference pass is separated in time from the mission pass, so that the moving target of interest will not be present in the reference pass images. A bright target with a focused SAR response well above the clutter background would appear in the mission orbit SAR image, but not the reference orbit SAR image. No such clear target response is seen in the left half plot in Figure 4. This is an indication that the target may be defocused with weak amplitude, and additional processing is needed to raise the target above the level of the background clutter.

As described in the last section, our approach employs a combination of MRP and STAP to increase the SINR of the target returns so that they can be detected above the clutter background. After MRP and STAP are performed on complex mission-orbit SAR data, thresholding of the AMF as shown in Eq. (5) produces an initial set of potential target detections. This initial set can contain false alarms from stationary clutter, as the estimated covariance estimate (Eq. (3)) will always contain errors, which may lead to undernulling of strong discretes. Observe that the image chips in Figures 3 and 4 show strong spiky returns typical of urban radar environments, as well as weaker returns in areas containing roads and shadows. In environments containing such strongly inhomogeneous clutter returns, it is difficult to obtain a very accurate covariance matrix in any given test pixel. As a consequence, undernulling occurs, which will lead to false alarms if additional processing to mitigate false alarms is not performed.

In order to reduce false alarms from clutter discretes, we perform a noncoherent form of change detection. The same MRP and STAP processing is performed on both the mission orbit SAR data and reference orbit SAR data, resulting in two corresponding AMF images. Figure 5 shows mission and reference orbit AMF images, after maximizing over target steering vector AOA and MRP parameters, for the measured CPI example. The extent of the two plots shown in Figure 5 is the same as in Figure 4 (i.e. the detection search area bounded by the horizontal green lines in Figure 3).

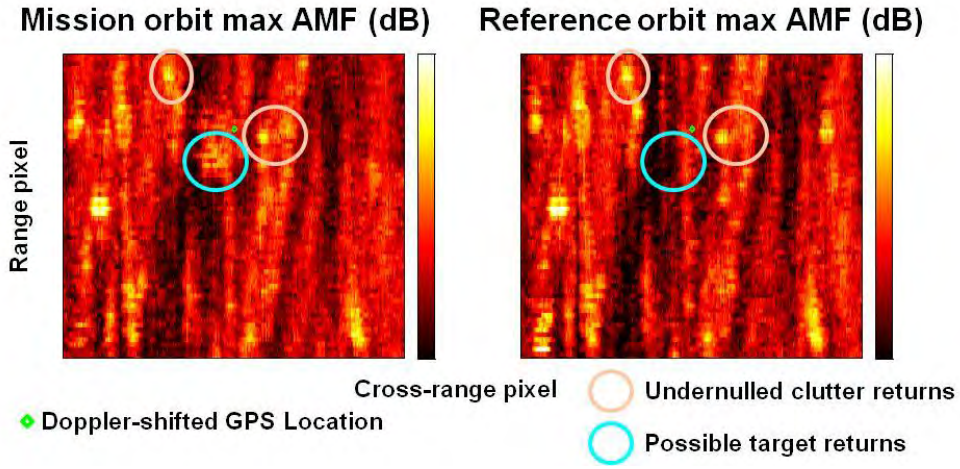

Figure 5 Mission-orbit (left) and reference orbit (right) AMF images (maximized over target parameters) for the measured data example over the same image extent shown in the previous figure. Examples of undernulled clutter areas are shown enclosed by peach ovals, while a potential target response is shown using the cyan ovals.

Note that some of the bright areas shown in the mission-orbit AMF images in Figure 5 are correlated with the locations of strong bright clutter SAR returns seen in Figure 4. However, the clutter AMF responses in Figure 5 are defocused relative to the clutter SAR responses in Figure 4, due to the defocus of clutter by MRP filtering (as was seen in Figure 1). As discussed in the previous paragraph, these clutter AMF responses are due to undernulling of bright clutter returns by the STAP filter, and would produce false detections if only mission orbit data was processed. However, undernulled AMF responses due to bright stationary objects will appear in both the mission and reference AMF images. Since the target of interest is moving and therefore not expected to be present in the reference orbit image, any detections crossing the AMF threshold for both mission and reference orbits can be removed from consideration. This allows a significant reduction in clutter false alarms to be achieved. The peach ovals in Figure 5 show examples of two such areas of undernulled clutter that can be removed from consideration.

The bright mission orbit AMF response contained within the cyan oval on the left half plot in Figure 5 shows a potential target detection, as the corresponding area on the reference AMF image on the right half of Figure 5 does not contain a bright AMF response. Moreover, the potential target response is also located fairly close to the center of the chip, at least on the scale of the image shown, which as mentioned before corresponds to the track covariance matrix. This indicates a

motion state statistically "close" to that of the mean track state prediction. The response is also close to the GPS truth location (green diamond). Note that the target AMF response is somewhat defocused. Due to the presence of MRP cross-range sidelobes, a finite search extent over MRP motion state, and the fact that the AMF image is maximized over MRP state in each pixel, a defocused maximum AMF target response is expected.

In order to perform geolocation as well as perform further mitigation of false alarms, we measure the "clutter ridge." The clutter ridge characterizes the response of clutter across the antenna phase centers as a function of the clutter cross-range coordinate. It is used as a reference in order to translate measured target channel-to-channel phase shift, which is a function of target AOA, to geolocated target position. The channel-to-channel phase shift is converted to an AOA index, and the clutter ridge is shown on an AOA index versus cross-range plot.

For the measured data example, the left half of Figure 6 shows the detections remaining after thresholding and false alarm mitigation on the mission orbit maximum AMF image. The right half of Figure 6 shows the measured clutter ridge, along with the measured AOA index vs. cross-range of the detections remaining after thresholding and false alarm mitigation.

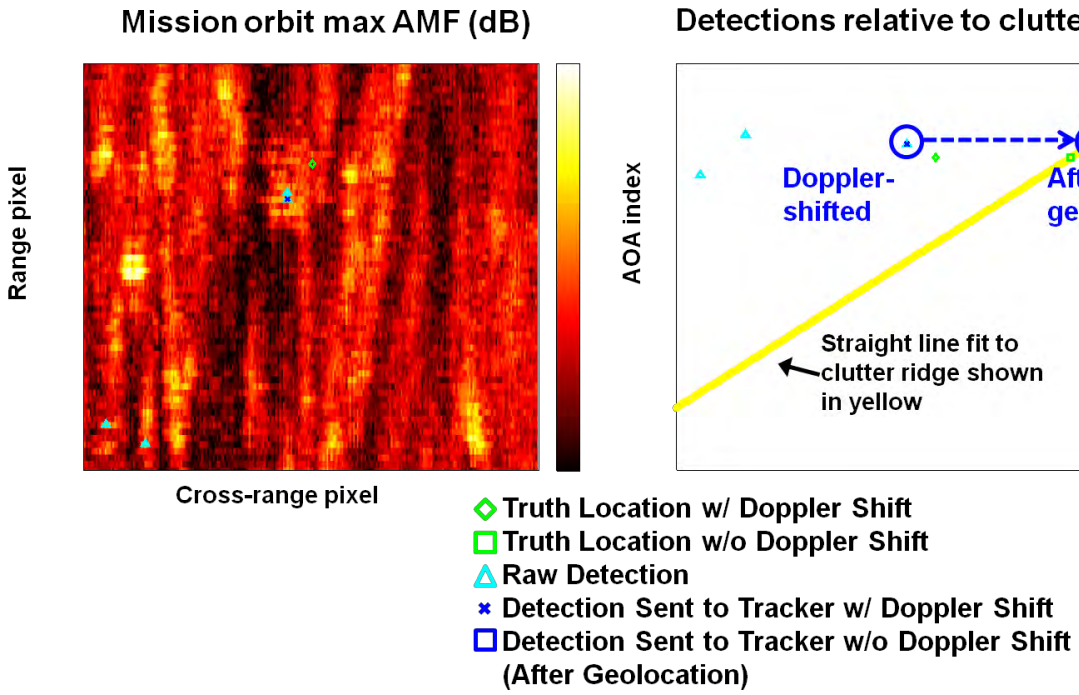

Figure 6 Detections after thresholding and false alarm mitigation (cyan triangles) shown on mission-orbit max AMF Image (left) and on AOA index versus cross-range plot (right). Measured clutter ridge shown as yellow line on right plot. Detection sent to the tracker is shown in blue both before (blue x) and after (blue square) geolocation.

Observe that detections are seen in three well separated areas. Two of these are in the lower left corner of the chip, but the detections just below and to the left of the GPS truth are much closer to the center of the chip. Qualitatively, since the predicted track state corresponds to the center of the chip, the latter set of detections appear to be much more likely from the target under track than the two detections in the lower left corner of the chip. As a final processing step, a quantitative likelihood score is computed in order to select a single detection on each CPI correlating best with the track state of the target under track. The selected detection for the measured data CPI example is shown in blue in Figure 6.

The blue square on the right half of Figure 6 shows the estimated cross-range location of the selected detection after performing geolocation. This "true" cross-range location is determined as the cross-range pixel for which the clutter ridge AOA index matches the measured AOA index of the detection. A horizontal line is drawn from the Doppler-shifted detection location to determine where it intersects the clutter ridge. This is shown by the blue arrow in the plot on

right half of Figure 6. Note that the cross-range determined in this manner is equivalent to the estimated location of the detected target if it were stationary (or did not have any Doppler shift).

The geolocated detection location (blue square) can be compared with the corresponding GPS prediction (green square) on the right half of Figure 6. The difference between the two is due to both errors in the measured AOA as well as GPS truth errors. The track prediction is not shown on the right half of Figure 6, but it basically lies on top of the GPS truth prediction. The likelihood score used to select a single detection is a function of how well each detection correlates to the track prediction in the measured Doppler-shifted x,y location (x,y, detection location in left half of Figure 6 relative to the center of the chip), as well as the measured AOA (y detection coordinate in the right half of Figure 6 relative to the track prediction). It is clear that the two detections on the left half of Figure 6 that are in the lower left corner of the image will not correlate as well with the track prediction as the selected detection shown in blue. The selected detection is seen to be relatively close to the track prediction in all three measured quantities (Doppler shifted x,y location and AOA). Use of the likelihood score to select a single detection on each CPI has the effect of reducing the chances of updating the tracker with a detection from the wrong target or from a false alarm. This in turn improves time in track on high value targets.

5. SUMMARY

We have described GMTI techniques for compensating for the effects of moving target defocus and clutter interference in multi-channel SAR systems. Using simulated data for a moving target competing with clutter interference, we demonstrated application of MRP filtering to focus the moving target SAR response. We also discussed formation of a set of adaptive DOFs for performing STAP on MRP filtered data. For the simulated target in clutter example, SINR loss was shown as a function of clutter CNR and compared to the performance of more standard algorithm configurations. The results showed a substantial benefit to using our adaptive DOF set.

Tracker feedback, parameter estimation, and false alarm mitigation techniques for improving GMTI tracking of high value targets performance were also described. Processing of a measured data CPI from a Gotcha radar system showed that applying MRP and STAP was able to detect a truthed endo-clutter target that was not visible in the single-channel SAR images. Applying false alarm mitigation based on a noncoherent change detection technique was shown to greatly reduce the effect of false alarms due to strong clutter discretes that were present in the data. Comparison of the estimated high value target location to truth data on the processed CPI was also presented.

REFERENCES

[1] S. Scarborough, C. Lemanski, H. Nichols, G. Owirka, M. Minardi, T. Hale, “SAR Change Detection MTI,” Algorithms for Synthetic Aperture Radar Imagery XIII, Edmund G. Zelnio, editor, Proceedings of SPIE (2006).

[2] LeRoy Gorham, Steven Scarborough, Capt. Matthew Judge, “Gotcha Radar Program Processing and Data Products”, ISR Systems and Technology Workshop, Lexington, Mass (November 17-19, 2009).

[3] R. Deming, S. MacIntosh, M. Best, "Three-channel processing for improved geo-location performance in SAR-based GMTI interferometry”, Proc. SPIE. 8394, 83940F (2012).

[4] I. S. Reed, J. D. Mallett, and L. E. Brennan, “Rapid convergence rate in adaptive arrays,” IEEE Trans. Aerospace Elec. Sys., Vol. 10 No. 6, pp. 853-863 (Nov. 1974).

[5] J. Ward, “Space-Time Adaptive Processing for Airborne Radar,” Tech. Rep. F19628-95-C-0002, MIT Lincoln Laboratory (December 1994).

[6] Frank C. Robey, Daniel R. Fuhrmann, Edward J. Kelly, Ramon Nitzberg, “A CFAR Adaptive Matched Filter Detector,” IEEE Transactions on Aerospace and Electronic Systems, Vol. 28, pp. 208-216, (1992).

[7] James Ward, “Angle and Doppler Estimation with STAP Radar,” 4th Annual Adaptive Sensor Array Processing (ASAP) Workshop (14 March 1996).