Multi-Resolution Processing To Enhance Knowledge-Aided STAP

G.A. Showman and W.L. Melvin
Georgia Tech Research Institute
Sensors & Electromagnetic Applications Laboratory
7220 Richardson Rd., Smyrna, GA 30080
(Voice) 770.528.7719, (Email) greg.showman@gtri.gatech.edu

Abstract – The need to detect slowly moving ground targets in strong clutter has motivated extensive research into space-time adaptive processing (STAP) algorithms and architectures. While these techniques perform very well in simulated clutter that is stationary, homogeneous, Gaussian, and target-free, many suffer tremendous performance losses on recorded data. Degradation occurs because clutter measured in flight tends to be heterogeneous and non-Gaussian, contains returns from strong discretes and multiple moving targets, and can suffer non-stationarity due to non-ideal collection geometries. Robust STAP algorithms, and STAP architectures incorporating *a priori* knowledge, are being developed to contend with these challenges. In this paper we examine another approach for attacking clutter heterogeneity, one exploiting long coherent dwells. Multi-resolution processing (MRP) uses motion compensation and image formation techniques developed by the synthetic aperture radar (SAR) community to eliminate the range, angle, and Doppler migration that clutter experiences over extended CPIs, permitting traditional STAP techniques to be applied to long-dwell data. We demonstrate that SAR imaging has the additional benefit of arresting the cross-track motion of movers. MRP employs post-Doppler STAP, which provides a convenient bridge between fully-adaptive short-dwell STAP and adaptive image processing of multi-channel SAR. We show that MRP improves the detectability of slow movers in heterogeneous clutter through coherent gain, suppression of discretes, and high fidelity calibration of the array.

The material in this paper was presented at the 2003 KASSPER Workshop, 14-16 April 2003, Las Vegas, NV.

I. Introduction

Timely detection of ground moving targets is a critical requirement for the modern war fighter. However, slow ground moving target indication (SGMTI) remains a challenging task for modern airborne radar due to the high power and large Doppler spread of clutter returns. A possible solution to this problem involves optimal filtering of spatial and temporal degrees-of-freedom (DOFs) to suppress clutter returns and maximize the target signal-to-interference-plus-noise ratio (SINR). This method, known as space-time adaptive processing (STAP), shows great promise in theory but has had mixed success in practice, largely due to the need for perfect knowledge of the interference characteristics, principally in the form of an ideal covariance matrix. Because clutter statistics are not generally known *a priori*, they must be estimated from secondary data. Unfortunately, clutter heterogeneity, discretes, internal clutter motion (ICM), and targets in the secondary data (TSD) conspire to corrupt the covariance estimate, thereby degrading STAP performance.

One approach to tackling heterogeneity involves the use of *a priori* knowledge. Under the Knowledge Aided Sensor Signal Processing with Expert Reasoning (KASSPER) Program, maps, land use data, transportation data bases, synthetic aperture radar (SAR) imagery, and ownship INS/GPS measurements are being used to guide STAP training, null discretes, regularize estimated covariance matrices, and pre-filter measured data [1]. An alternative scheme uses wide bandwidths and long dwells to sub-resolve discretes, cultural discontinuities, TSD, and other sources of heterogeneity. The fine resolution of these waveforms has the added benefit of greatly reducing the clutter energy in a target cell. In this paper we propose a multi-resolution architecture that merges wideband, long dwell, GMTI processing with SAR imaging and knowledge-aided STAP. Multi-resolution processing (MRP) is tailored to detect slowly moving ground targets in severely heterogeneous environments.

The remainder of this paper is organized as follows: Section II describes the multi-resolution processing concept of operation. Section III examines clutter migration effects over long dwells, and introduces compensatory processing from a SAR motion compensation and image formation perspective. Three benefits of the MRP mode are examined in detail in Section IV. Finally, Section V discusses the impact of target motion on MRP performance and the manner in which SAR imaging arrests cross-track target migration.

II. MULTI-RESOLUTION PROCESSING CONCEPT OF OPERATION

A multi-resolution timeline is shown in Figure 1. The radar begins in a baseline wide area surveillance (WAS) mode, represented by the gray boxes at left. The WAS mode consists of a series of short dwells over the ground-referenced coverage area (GRCA), each dwell consisting of nominally three coherent processing intervals (CPIs). Processing times appear below the "Short Dwell GMTI" boxes and indicate 75 ms allocated to each dwell. Waveform diversity, in the form of CPI-to-CPI RF and/or PRF changes, are used to resolve range and Doppler ambiguities and mitigate target RCS nulls. The values at the bottom of Figure 1 are the number of pulses processed, in this case three CPIs of 32 to 48 pulses per CPI for a PRF of 2,000 Hz, parameters consistent with the notional KASSPER system [2]. The baseline mode employs narrowband waveforms only, 5 to 20 MHz of bandwidth, so that convolutional matched filtering can be used for pulse compression, thereby allowing wide ground swaths to be processed. The small bandwidths and integration times employed under WAS limit the measured knowledge of the environment, so we expect that *a priori* knowledge will be exploited extensively to combat clutter heterogeneity. Knowledge-aided training, knowledge-constrained optimization, and knowledge-derived pre-whitneing of the data all figure prominently here [3].

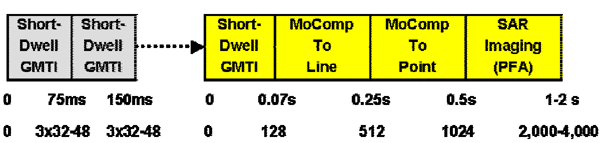

Figure 1. Multi-Resolution Processing (MRP) timeline

Occasionally the baseline mode will transition to the multi-resolution processing (MRP) mode. Reasons for this transition include:

- *A priori* knowledge suggests the current dwell area will be challenging for the adaptive processor, due to severe clutter heterogeneity, large number of moving targets, etc;
- A desire to detect very slow moving ground targets;
- A requirement to continuously track targets that are moving into the clutter notch, one consequence of a "cradle-to-grave" tracking philosophy.

Under MRP the radar collects data continuously from an area of interest. (Thus, the MRP mode is not consistent with WAS operation.) The transmit frequency and PRF are fixed to allow long-dwell coherent processing of the scene. MRP is, in essence, the GMTI analog to the SAR "spotlight" mode. At any given point in time, all of the data recorded back to the beginning of the dwell are available for processing. As more data are collected higher resolutions are possible and more samples are available for training. As the dwell time increases measured knowledge of the target and clutter properties of the area of interest also increase, and we expect *a priori* knowledge will be progressively de-emphasized.

The yellow boxes on the right side of Figure 1 show the MRP timeline. The mode begins with short-dwell GMTI processing very much like that used in the WAS mode, but without waveform diversity. As the dwell time increases target and clutter returns begin to experience significant range, Doppler, and angle migration. As Figure 1 suggests, this can happen after as few as 128 pulses, around 70 ms. Migration motivates the use of motion compensation ("MoComp") and focusing techniques used in SAR imaging. These will be described in the next section.

A key component of the multi-resolution philosophy is reduced dimension adaptive processing in the form of post-Doppler STAP, specifically, the extended factored algorithm (EFA) [4]. EFA exploits full spatial DOFs and a subset of Doppler DOFs about the Doppler bin under test. EFA is a well-known and understood reduced dimension STAP technique. EFA using only a few Doppler DOFs realizes SINR performance comparable to the fully adaptive joint-domain-optimum (JDO) limit under most GMTI conditions; indeed, EFA and JDO performance is equivalent when all Doppler DOFs are used. However, when only a few Doppler DOFs are used the resulting covariance matrices are much smaller and, therefore, lend themselves to localized training and real-time inversion.

EFA has another advantage in that it provides a convenient model to bridge the gap between STAP and multi-channel SAR image processing for GMTI. For example, moving targets have been detected in SAR imagery by performing adaptive spatial processing on each crossrange bin [5]. Because crossrange in SAR is simply Doppler (preceded by some compensation, on the part of the SAR image former, for clutter range and Doppler migration), this approach is equivalent to applying EFA with a single Doppler DOF. In STAP parlance this technique is also known as Factored-Time-Space (FTS) [6] and is not a true STAP technique. Nevertheless, FTS may suffice on relatively benign long-dwell data. Over shorter dwells, or on clutter with severe internal clutter motion (ICM), FTS performance will suffer and the more general EFA is preferred. Even in high-resolution multi-channel SAR, the multiple Doppler DOFs of EFA are expected to improve performance in the presence of ICM and discretes.

Finally, two bandwidth modes are envisioned for MRP. The narrowband mode uses convolutional matched filtering of linear FM (LFM) waveforms. Though limited to bandwidths below 20 MHz, convolutional filtering permits moderate resolution measurements over swaths many kilometers deep. The wideband mode, in contrast, requires "stretch" processing of LFM. Stretch bandwidths can be as high as 1 GHz so that very fine range resolutions are achievable, but swath coverage is on the order of only one to two kilometers. The utility of large bandwidths in GMTI remains to be determined. Some advantages of wideband operation are

- Finer range bins, leading to a potential increase in sample support for a fixed training interval;
- Frequency diversity is available, so that fluctuating target signatures can be countered;

while disadvantages include

- A limited surveillance area, as stated earlier;
- Exacerbated target and clutter migration;
- Targets are sub-resolved, which complicates detection processing.

III. LONG DWELL PROCESSING

Ordinarily, STAP models are developed using far-field geometries and planar wavefronts, properties consistent with the short CPIs typical of GMTI operation. These assumptions are violated over long dwells however, causing migration of returns in range, Doppler, and angle. Traditional STAP techniques can yet be applied to long dwells, but only after the data has been compensated for migration effects. In this section we discuss three corrections for migration within the context of multi-resolution processing. These are motion compensation to a line, motion compensation to a point, and polar formatting [7].

A. MoComp to a Line (MCL)

Early in a multi-resolution dwell data would be processed as in the radar's baseline WAS mode. Given short records, it suffices to simply Fourier Transform the temporal data to form Doppler filters, followed by EFA processing.

The first migration effect is in range and occurs when a side-looking array is scanned away from broadside. When the mainbeam is scanned to angle θ away from broadside, ground returns drift along the line-of-sight (LOS) at the aircraft ground-speed times a component $\sin\theta$. Let L be the baseline length the aircraft flies over the CPI. If we wish to limit migration to one-half a downrange bin width Δ_{DR}, then

$$L \leq \frac{\Delta_{DR}}{2\sin\theta}.$$

For example, if the range bins are 15 meters wide (10 MHz bandwidth), and the squint angle is 60°, then the baseline length L is limited to 8.6 meters. For an aircraft speed of 125 m/s, the CPI is limited to 0.07 s, around 128 pulses at a PRF of 2,000 Hz. These limits appear in Figure 1 below the right corner of the MRP short-dwell GMTI box.

Range migration causes temporal amplitude modulation of returns within a given range bin. For targets, the result is an SINR loss due to a steering vector mismatch. For clutter, amplitude modulation is manifested chiefly as a temporal covariance matrix taper (CMT) [8]. To illustrate this effect, a simulation was used to generate the ideal covariance matrix over a 16-pulse dwell without range migration, and with migration equal to one range bin. Comparing the two matrices yielded the CMT shown in Figure 2. Note that the upper-right CMT entry is near zero, indicating the first and last pulses are almost completely uncorrelated. This is consistent with migration equal to one-range-bin over the dwell: a range bin will contain a completely new set of scatterers at the end of the CPI. CMTs are well-known to give rise to subsapce leakage [8], as shown by the eigen-spectrum of Figure 3. In short, range migration has an effect very similar to ICM.

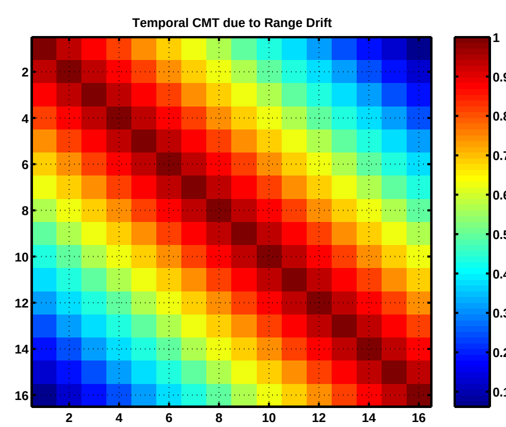

Figure 2. Example Temporal Covariance Matrix Taper (CMT) due to range migration

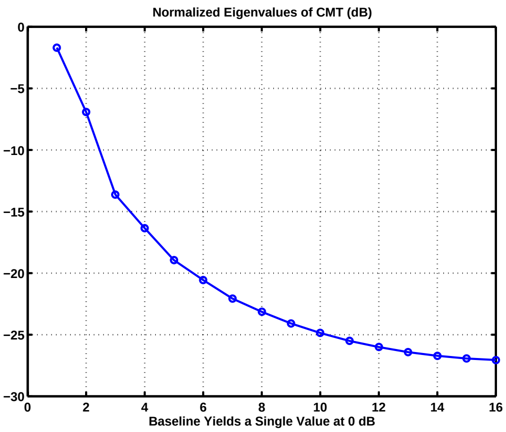

Figure 3. Subspace leakage in temporal eigen spectrum due to CMT

A simple solution to this kind of range migration is to apply fast-time (range) delays to the multi-channel data records. These delays progressively increase as a function of slow-time (pulse number) so as to exactly cancel the linear component of range migration of scatterers at the center of the mainbeam. In SAR parlance, this is "MoComp to line" (MCL), in this case, a line that is normal to the mainbeam LOS. The new virtual baseline intersects the original baseline with an angle equal to the mainbeam scan angle. Because the delays generally contain some fraction of the range bin time, they would be most easily realized by modulating (fast-time) frequency domain data by linear phase functions.

B. MoComp to a Point (MCP)

When the CPI baseline (either true, or virtual after MCL) becomes significant relative to the range to scene center, the LOS angle to a scatter will drift over the CPI. This results in some angle migration of scatterers over the dwell. In addition, the effective Doppler of a scatterer is related to LOS angle, so there will be Doppler migration as well. The combined angle-Doppler migration is coupled so that scatterers appear to drift along the clutter ridge in the angle-Doppler domain over a long dwell. Finally, the range to the scene increases towards the beginning and end of the dwell, so there is also some range migration

The range migration is not significant unless fine range resolutions are generated over large scene sizes. Angle migration is manifested in spatial steering vectors only with very large arrays. Doppler migration then is the phenomenon of most concern. For radar systems and flight geometries consistent with the KASSPER program, Doppler migration can become an issue at about 0.25 seconds, or 512 pulses, as shown in Figure 1. The simulated SINR-loss example in Figure 4 shows a 4-dB loss due to Doppler migration (magenta), a loss due almost entirely to the mismatch between the temporal components of the target signal and steering vector (red).

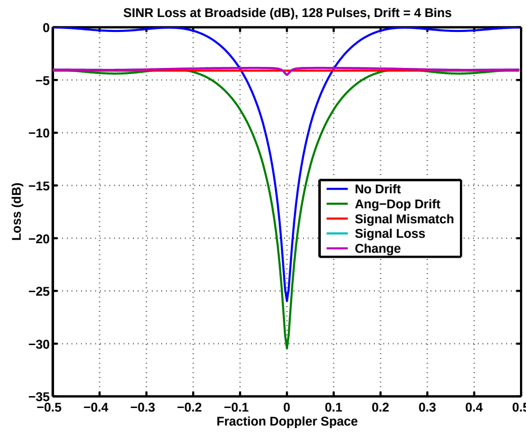

Figure 4. Example SINR-loss due to Doppler migration

Doppler migration can be removed from the data by applying fast-time delays so that the phase of the return from scene center is unchanged pulse-to-pulse. In SAR parlance, this is "MoComp to point" (MCP), in this case, the point being the scene center. Implementation is similar to MCL, except that the delay history is not simply linear. (It is, in fact, hyperbolic in pulse number.)

C. Polar Formatting

Still longer dwells cause clutter motion through resolution cells (MTRC) that simple motion compensation procedures cannot account for. MTRC includes both range and Doppler migration and becomes more severe as resolutions become finer and scene size increases. Figure 1 suggests that MTRC becomes significant at about half a second, or after one thousand pulses, although this calculation assumed the use of a wideband waveform providing range resolutions on the order of 1.5 meters. MTRC causes smearing of returns over multiple range and Doppler bins. SINR-loss will occur because clutter range migration results in a temporal CMT (as discussed in Section A above), and because range and Doppler migration yield a mismatched target steering vector (as discussed in Section B above).

MTRC is well understood in the SAR community. SAR image formation techniques such as the polar formatting (PFA) algorithm serve to remove these higher-order effects. In essence, SAR processing acts to "focus" clutter returns, thereby removing range and Doppler migration. Under the MRP concept, PFA would be applied to the data record for each spatial channel should MTRC become a concern.

Finally, a combination of very long dwells (over several seconds), wideband operation, and large scene sizes can result in migration effects so severe that even the PFA cannot completely compensate for them. Under these circumstances it would be wise to bypass MCP and PFA and go straight to an errorless inversion like the range migration algorithm (RMA) [9]. However, we anticipate this will not be necessary for everyday SGMTI requirements. Indeed, MCL and MCP alone will usually suffice if narrowband operation is employed.

On the other hand, the MCL, MCP, and PFA procedures serve to perfectly remove migration effects at scene center *only*. Correction will be adequate over the scene area as defined by the mainbeam, but will be in error for sidelobe and backlobe returns. Indeed, these algorithms can actually exacerbate the migration of scatterers in the sidelobe and backlobe. One implicit assumption then on the part of MRP is that returns outside the mainbeam have negligible impact on GMTI performance as compared to mainbeam returns.

IV. BENEFITS

MRP offers a number of advantages over conventional short-dwell STAP for detecting slowly moving ground targets. In this section we present three benefits in detail, these being coherent integration gain, suppression of clutter discretes, and array calibration.

A. Coherent Integration Gain

An elemental benefit of MRP derives directly from processing long coherent dwells. This has less to do with any narrowing of the SINR clutter notch, and is simply a benefit of the coherent integration gain on a target. Figure 5 illustrates the former point, showing SINR-loss as a function of Doppler for a hypothetical GMTI radar for different numbers of pulses. Once the number of pulses exceeds eight, there is little change in the SINR-loss curves, and little improvement in the detectability of slow movers near the clutter notch. Indeed, our analysis suggests that the array size must increase jointly with the number of pulses in order to realize increasingly narrower clutter notches.

The curves in Figure 5 have been normalized to the 0-dB level, the thermal noise level for each pulse number condition. The losses, then, are against the *relative* SINR for each case. However, changing the number of pulses changes the thermal noise level and the target output power so that increasing pulses should increase overall SINR. Figure 6 shows SINR-loss curves plotted against *absolute* SINR. The thermal noise level for 256 pulses is at 0-dB, while that for 16 pulses is at –12 dB. The dashed line at –17 dB represents a –5 dB SINR loss for 16 pulses. Looking at where this line intersects the 16-pulse curve (red) defines the minimum detectable velocity (MDV), we see a notch width of about 7 m/s exists. When this black dashed line is used to define an absolute SINR requirement, its intersection with the other curves suggest a clutter notch that decreases with increasing pulse number. For example, with 256 pulses the width reduces to 4 m/s.

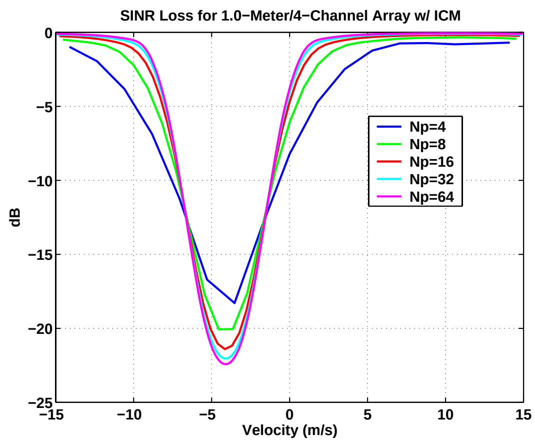

Figure 5. Relative SINR-loss as a function for pulse number of notional GMTI radar

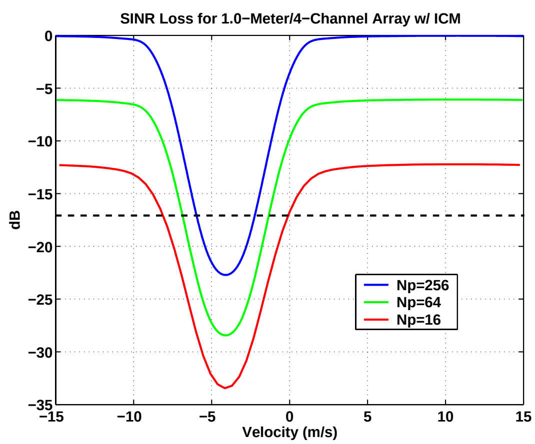

Figure 6. Absolute SINR-loss as a function of pulse number (16, 64, and 256) for notional GMTI radar

Another interpretation is that increasing the number of pulses narrows the Doppler filter width and results in finer crossrange bins. Smaller resolution cells contain less clutter power, but the target power from a coherent target will not change. (Until the cells are so small that the target is sub-resolved.) Long dwells provide the ability to “burn through” clutter to some extent. The benefit is only 3 dB for every doubling of the pulse number, so this gain comes at substantial cost.

B. Suppression of Discretes

While strong clutter discretes are widely acknowledged to be deleterious to STAP performance, there is some disagreement as to whether their impact is greater on short or long dwells. Figure 7 attempts to quantify the issue with a simple STAP simulation example. Absolute SINR is shown as a function of fractional Doppler frequency. The lower blue curve shows the SINR-loss against homogeneous clutter for a four-pulse waveform, while the red curve nearby shows the same when training on homogeneous data but testing on a cell with heterogeneous characteristics, specifically, a single strong discrete in homogeneous clutter. The difference is shown by the broad cyan curve at the top of the figure. The presence of the discrete in the test cell caused significant SINR-loss over a good fraction of the Doppler space. However, let us examine the impact on the clutter notch as defined by the –5 dB SINR-loss level for the four-pulse waveform, corresponding to the black dashed line at –17 dB absolute SINR. This line intersects both the red and blue curve near ± 0.14, so that both the homogeneous and the heterogeneous case have about the same notch width.

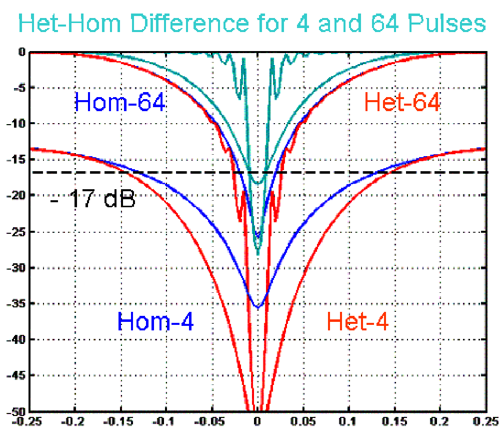

Figure 7. Homogeneous (random clutter only) and Heterogeneous (CUT contains a discrete) SINR-loss for 4 and 64 pulses

The red and blue lines towards the top of the Figure 7 are similar curves but for a 64-pulse waveform. The difference is again in cyan at the top, but is sharper and deeper than the cyan difference for the four-pulse case. Examining where the dashed line interests the 64-pulse SINR-loss curves we see again that the presence of the discrete has little impact on MDV; the clutter notch width is about ± 0.025 in either case. However, any difference between the homogeneous and heterogeneous notches is amplified now because the notch itself is so very small. Indeed, while the change in usable fractional Doppler space (UFDS) is small, the fractional increase in the clutter notch width is substantial, close to 40%.

We concede that results are sensitive to the SINR level chosen for MDV measurements. In general discretes in the cell under test (CUT) cause loss over a wide portion of the UFDS but have little impact on MDV for short dwells. For long dwells SINR loss is concentrated near the clutter notch, and while effect on the UFDS might be negligible, the fractional MDV can increase substantially.

Discretes in the CUT can generate false alarms in the target detector that follows the STAP stage. The detection threshold is typically biased upwards to provide a degree of robustness to the detector and reduce false alarms rate due to discretes. Though not addressed in this study, the resulting desensitization of the detector due to this "threshold loss" can be considered an additional SINR loss.

Several options are available to counter discretes. One can simply include the CUT in the training data; this has the advantage of working at all resolutions, but is known to aggravate losses due to TSD. An alternative that is available only at moderate and finer resolutions consists of detecting and removing the discretes from the data after Doppler processing and prior to STAP. Figures 8 and 9 show the benefit of applying CLEAN [10], an iterative deconvolution algortihm, to remove discretes. Figure 8 shows a range-Doppler image collected by the Raytheon Adaptive Processing Technology Investigation (APTI) radar test-bed [11]. Resolution is 1.5 meters downrange (vertical) by 6.75 meters in crossrange (horizontal), and strong discrete returns are apparent along fence lines (upper half of image) and due to retro-directive scattering off a hanger structure (at lower right). Figure 9 shows the same data after CLEAN; the discretes are clearly absent, so this data is better conditioned for EFA processing.

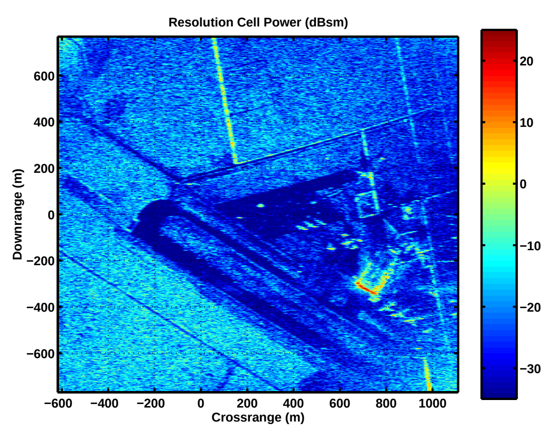

Figure 8. Baseline 1.5-by-6.75-meter resolution APTI range-Doppler image of Mojave Airport

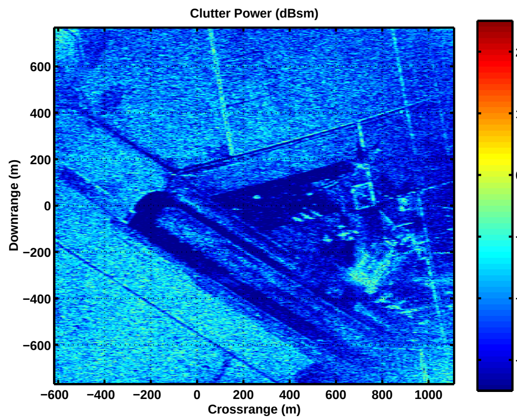

Figure 9. APTI range-Doppler data after discrete removal by CLEAN

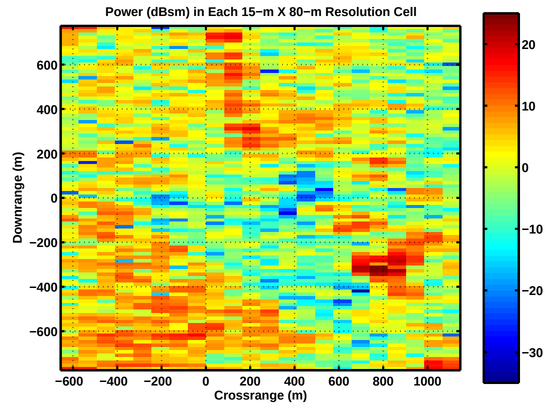

Figure 11. APTI image spoiled to 15-by-80-meter resolution

CLEAN possesses a number of advantages for discrete removal purposes. It is simple to implement and, given judicious choice of parameters, is quite stable. It removes the strongest scatterers first, and need be run only as long as strong discretes are present, or until real-time processing demands become a constraint. It does not operate well on coarse resolution imagery, however. To illustrate this point, Figures 10 and 11 shows the same APTI data spoiled to downrange-by-crossrange resolutions (in meters) of 15-by-15 and 15-by-80, respectively. Figure 11 is consistent with a range-Doppler map generated from a short GMTI dwell. As resolution grows coarser discretes tend to fade into any continuous background clutter present, making discrete detection and parameter estimation problematic.

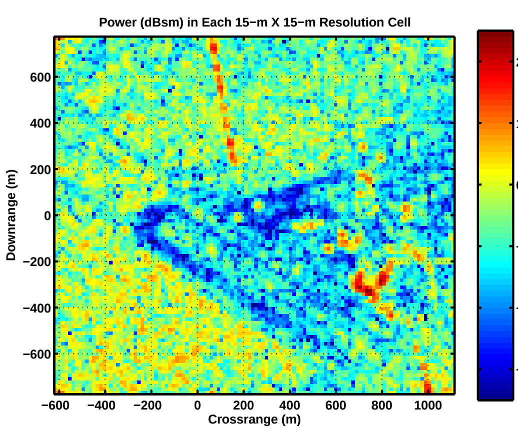

Figure 10. APTI image spoiled to 15-by-15-meter resolution

C. *Array Calibration*

A number of STAP techniques have been developed that exploit interference covariance matrices generated from *a priori* knowledge to combat clutter heterogeneity, including colored loading [2] and data pre-filtering [12]. A widely-acknowledged limitation to these approaches hinges on errors in the knowledge of the radar array response, chiefly the scalar amplitude and phase gains of the spatial channels. Even very small channel errors can greatly diminish the benefit of colored loading and pre-filtering when strong clutter is present. Therefore, these techniques require very accurate array calibration information.

Relevant work on array calibration has been performed under the auspices of the SGMTI Active Electronically Scanned Array (AESA) Demonstration (SAD) program. SAD is funded by the Air Force Research Laboratory, with Georgia Tech Research Institute (GTRI) as primary contractor and Raytheon the sub-contractor. This program is primarily concerned with geolocating slowly moving ground targets with airborne fire-control radar. Detailed information on SAD's technical approach can be found in [13].

A key component in the SAD processing chain is the use of wideband long-dwell data to simultaneously achieve fine direction-of-arrival (DOA) estimates on targets, and to estimate and compensate for the array response. For DOA estimation the GMTI outputs of multiple sub-bands and sub-CPIs are combined to maximize SINR and, therefore minimize the error in the DOA estimate. The procedure is reminiscent of MRP as presented here and in [14], though there are significant differences, as SAD attempts to maximize SINR for DOA estimation purposes, while MRP aims to maximize detectability in clutter heterogeneity. In any event, here we are more interested in the SAD approach to array calibration.

SAD calibrates an array using a "cal-on-clutter" procedure. First, high-resolution multi-channel SAR imagery is formed and accurately geolocated in an absolute coordinate system. Then the spatial response of the geolocated pixels is processed to generate an estimate of the array response as a function of absolute location. In this section we illustrate the procedure using multi-channel APTI data. Figure 12 contains eight range-Doppler images, one for each of the eight APTI spatial channels. Figure 13 shows pixel phase angles relative to the pixel phase in the first channel. Linear phase progressions are apparent in both crossrange and downrange, a consequence of the horizontal and vertical separation of the channels. (APTI is a two-dimensional array.)

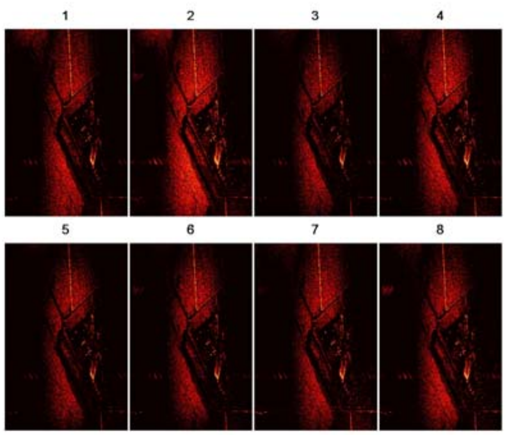

Figure 12. APTI range-Doppler images (power in dB) for all eight channels

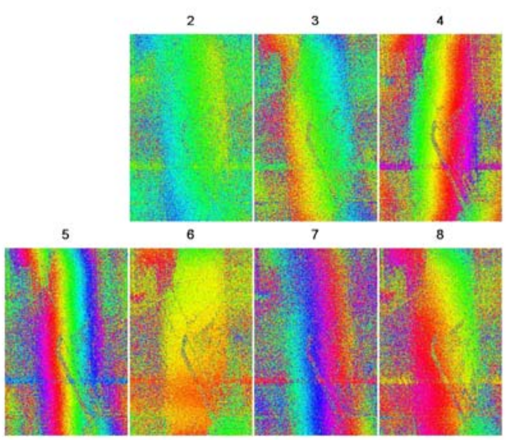

Figure 13. Range-Doppler phase angle differences for each APTI channel relative to first channel

A first-order estimator was developed to determine the phase gain and linear phase slopes in each channel. The estimator used a weighted-least-squares (WLS) solution, with weighting by pixel power. The estimated phases are shown in Figure 14, and compare favorably to Figure 13. The first-order WLS estimator is superior to a moving average estimator as it explicitly accounts for the presence of the linear phase components, thereby increasing the accuracy of the zeroeth-order term (the overall phase bias). In addition, the linear phase estimate can be used to account for errors in the knowledge of the subarray locations. Finally, a moving average estimator yields poor phase estimates in low power clutter regions (such as the airport runways and tarmacs in the APTI image), but the WLS processes all the available pixels at once, and low power regions have little influence on the resulting estimates.

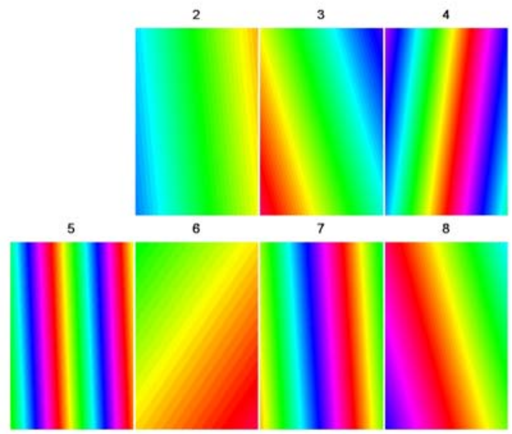

Figure 14. First-order estimator phase angle differences for each APTI channel

The array calibration parameters generated by long-dwell records would be applied to the measured data so as to minimize any statistical mismatch between the knowledge-derived and estimated covariance matrices.

V. ADDRESSING MOVING TARGETS

Long-dwell processing is the heart of the multi-resolution architecture and has a number of inherent drawbacks. The most significant of these is caused by the motion of moving targets over the course of the CPI. Over several seconds, even very slowly moving targets can walk through many downrange and crossrange/Doppler resolution cells, smearing the target response and thereby diminishing detection performance. SINR loss is due directly to the mismatch between the target steering vector and the target range-angle-Doppler history.

The effects of cross-track (motion "downrange," to or away from the radar flight path) were examined using simulated high-resolution imagery. The simulation was performed at X-band and with bandwidth and collection geometry consistent with 3-meter resolution. The image in Figure 15 shows an array of stationary targets centered at the origin, while Figure 16 shows an array that moved toward the radar 5.25 meters during the CPI. One consequence of motion was to shift the apparent position of the targets in crossrange (vertical axis) by nearly 1000 meters. The more significant effect was to blur the scatterers in both downrange and crossrange. Such blurring would result in substantial SINR loss for GMTI.

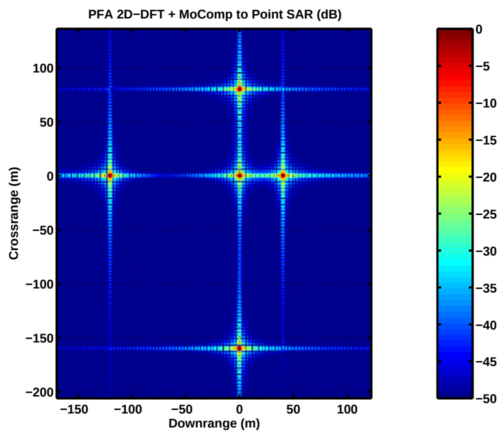

Figure 15. X-band SAR image for stationary array of scatterers located at origin

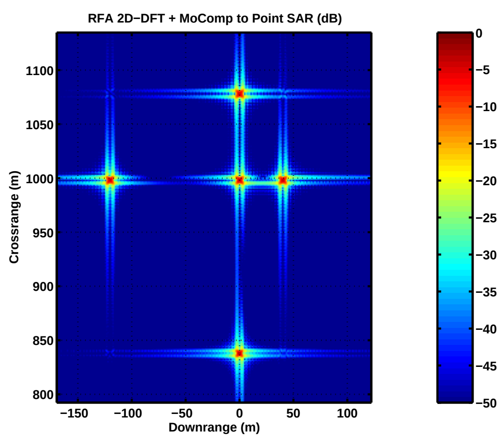

Figure 16. Range-Doppler image for array of scatterers located at origin moving 5.25 meters cross-track during the CPI

Figure 16 is a simple range-Doppler map, generated using only MCP followed by Fourier Transforms over frequency and slow-time. Using the same measurements, Figure 17 shows the image produced by a SAR image former, in this case PFA. While SAR image formation has done nothing to correct the crossrange offset, it has effectively removed the downrange walk of the targets, yielding returns focused in both downrange and crossrange. Indeed, the quality of the compensation is such that Figure 17 is indistinguishable from the SAR image of a stationary array placed at 1000 meters in crossrange.

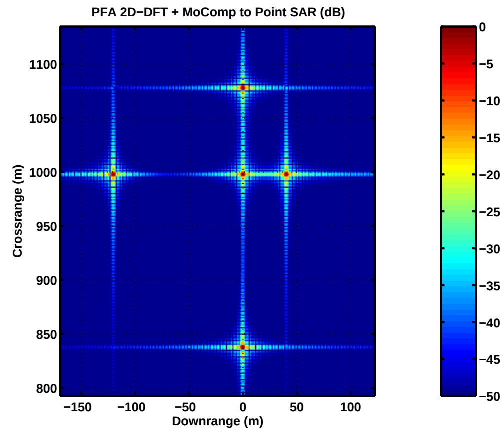

Figure 17. PFA image for array of scatterers located at origin moving 5.25 meters cross-track during the CPI

The ability of SAR image formers to remove most cross-track target motion was maintained even when the target effectively migrated through many resolution cells (fine resolutions and/or high-speed targets). This property held across all sophisticated image formation algorithms, including PFA, RMA, back-projection [15], range-Doppler techniques [16], and keystone formatting [17]. The phenomenon is not surprising when we consider that SAR image formers tie Doppler to range migration, usually implicitly but sometimes explicitly [16]. When Doppler is consistent with range migration, as it is with unaliased returns from stationary clutter and moving targets, range migration is properly compensated. When range migration is not consistent with Doppler, as with a return aliased in Doppler, SAR imagers do not completely remove range migration and, under unfavorable conditions, can actually exacerbate it.

Unfortunately, SAR image formers do not compensate for any along-track motion of targets. Figure 18 is a PFA image of the array at the origin with a 5.25-meter movement along track during the CPI. When compared to Figure 15, we see that the along-track motion has caused one-dimensional blurring of the targets in crossrange. No SAR image former will account for this motion. On the other hand, the blurring is one-dimensional, as compared to the two-dimensional blurring caused by cross-track motion. In addition, this blurring can be shown to consist primarily of a quadratic phase error (QPE) over the CPI. One-dimensional QPE can be easily removed after EFA processing by performing one-dimensional low-order FIR filtering over Doppler [18], or by Fourier Transforming to slow-time, applying a conjugate QPE, and inverse Fourier Transforming.

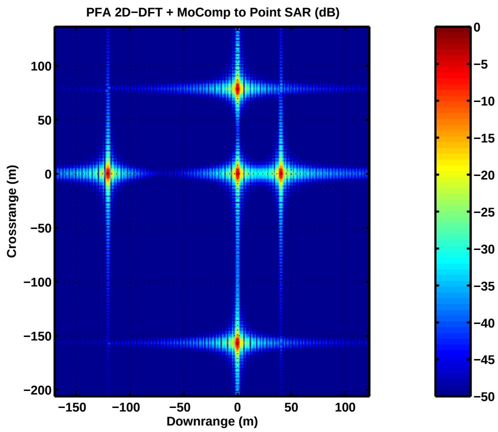

Figure 18. PFA image for array of scatterers located at origin moving 5.25 meters along-track during the CPI

Putting it all together, an MRP procedure that accommodates target motion consists of the following steps:

1. Multi-channel SAR imaging;
2. EFA processing of the SAR imagery with complex output, i.e., phase is retained and generation of a magnitude or power detection statistic deferred;
3. Application of candidate QPE compensations as described above, designed to cover the expected range of target along-track motion;
4. Magnitude or power detection statistics are finally generated.

A detection statistic is generated for every candidate along-track speed. In short, step 3 can be thought of as a filter bank for along-track motion.

VI. SUMMARY

MRP is a specialized mode, part of a larger architecture for SGMTI in heterogeneous environments that includes the use of *a priori* knowledge. Baseline operation consists of a WAS mode using conventional short-dwell STAP, waveform diversity, and knowledge, while the multi-resolution mode is activated on a limited area of interest. Early in the dwell, MRP is similar to baseline operation. As the dwell time increases SAR-like processing becomes significant, and heterogeneity is attacked less by *a priori* knowledge and more by the fine resolution and large signal integration afforded by long data records.

Future work will concentrate on quantifying the detection benefits of MRP, and examining target motion effects in more detail. We will also explore the tradeoffs of wideband operation; some of these were listed in Section II. In addition, issues regarding area coverage must be resolved. Initially, multi-resolution processing will be reserved for GMTI in small areas of high interest. Alternatively, digital array radar (DAR) architectures might permit long dwells over large areas, but the technical hurdles to implementing this approach on airborne systems are significant. Finally, the potential of setting dwell lengths in real-time is intriguing. In this concept, target detections and environmental metrics are generated continuously during the CPI. The CPI is ended only after the processor has concluded that the dwell length is sufficient for the level of heterogeneity in the clutter as determined by the recorded data.

ACKNOWLEDGMENT

This work was supported by DARPA and the US Air Force under Contract Number F30602-02-C-0011.

The authors gratefully acknowledge Trung Nguyen and Jeffrey Hoffner of Raytheon, El Segundo, for describing the signal processing details behind the SAD program, and for providing APTI multi-channel data sets.

REFERENCES

[1] J. Guerci, “ KASSPER: A framework for next generation adaptive sensors,” this proceedings (*Proceedings of the 2003 KASSPER Workshop*).
[2] J. Bergin, “High-fidelity data generation and algorithm development for knowledge-aided sensor signal processing,” *Proceedings of the 2002 KASSPER Program Review*, Lincoln Laboratory, September 23, 2003.
[3] W. L. Melvin, G. A. Showman, “A knowledge-aided STAP architecture,” this proceedings (*Proceedings of the 2003 KASSPER Workshop*).
[4] R. C. DiPietro, “Extended factored space-time processing for airborne radar,” *Proceedings of the 26th Asilomar Conference*, Pacific Grove, California, October 1992, pp. 425-430.
[5] J. H. G. Ender, “Space-time processing for multichannel synthetic aperture radar,” *Electronics & Communication Engineering Journal*, February 1999, pp. 29-38.
[6] J. Ward, *Space-Time Adaptive Processing for Airborne Radar*, Lincoln Laboratory Technical Report, ESC-TR-94-109, December 1994.
[7] W. G. Carrera, R. S. Goodman, and R. M. Majewski, *Spotlight Synthetic Aperture Radar*, Artech House, Boston-London, 1995.
[8] J. R. Guerci, “Theory and application of covariance matrix tapers for robust adaptive beamforming,” *IEEE Transaction on Signal Processing*, vol. 47, no. 4, April 1999, pp. 977-985.
[9] C. Cafforio, C Prati, E. Rocca, “SAR data focusing using seismic migration techniques,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 27, no. 2, March 1991, pp. 194-206.
[10] G. M. Skofronick-Jackson and A. J. Gasiewski, “Nonlinear multispectral statistical CLEAN-based precipitation parameter retrieval algorithm,” *International Geoscience and Remote Sensing Symposium (IGARSS)*, Vol. 3, pp. 1674-1676, 1998.
[11] T. Nguyen, J. Hoffner, APTI flight-test data, flight test #6, electronic data, 1999.
[12] P. Techau, “Reduced degree-of-freedom STAP with knowledge-aided data pre-whitening,” this proceedings (*Proceedings of the 2003 KASSPER Workshop*).
[13] T. Nguyen, J. Hoffner, E. Day, R. Latter, E Baik, “ERFT/SAD Program: Geolocation of slow moving ground targets,” *Proceedings*

of the Association of Old Crows Radar/EW Conference, U.S. Army Research Laboratory, March 19-20 2003, http://www.crows.org/FTP/030319radar/030319radar.htm.
[14] J. Bergin, "Multi-resolution signal processing techniques for airborne radar," this proceedings (*Proceedings of the 2003 KASSPER Workshop*).
[15] R. Rau, J. H. McClellan, "Analytic models and postprocessing techniques for UWB SAR," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 36, no. 4, pp. 1058-1074, October 2000.
[16] Richard Bamler, "A comparison of range-Doppler and wavenumber domain SAR focusing algorithms," *IEEE Transactions on Geoscience and Remote Sensing*, vol 30, no. 4, July 1992, pp. 706-713.
[17] R. P. Perry, R. C. DiPietro, R. L. Fante, "SAR imaging of moving targets," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 35, no. 1, January 1999, pp. 188-199.
[18] N. Pulsone, "Improving ground moving target indication performance," this proceedings (*Proceedings of the 2003 KASSPER Workshop*).