

Detection and Tracking of Prominent Scatterers in SAR Data

Benjamin Shapo^{*a}, Mark Stuff^b, Christopher Kreucher^a, Ron Majewski^a

^aIntegrity Applications, Inc., 900 Victors Way #220, Ann Arbor, MI USA 48108

^bMichigan Tech Research Institute, 3600 Green Ct., #100, Ann Arbor, MI USA 48105

ABSTRACT

Tracking prominent scatterers provides a mechanism for scene-derived motion compensation of Synthetic Aperture Radar (SAR) data. Such a process is useful in environments where GPS is unavailable and a lack of precise sensor position data makes standard motion compensation difficult. Our approach to sensor positioning estimates range histories of multiple isolated scatterers with high accuracy, then performs a geometric inversion to locate the scatterers in three dimensions and estimate the platform's motion.

For high-accuracy scatterer range tracking, we first detect prominent scatterers using a CFAR criterion automatic algorithm and then track them with a two-input Kalman Filter (KF) operation. These two steps provide accurate range estimates of multiple scatterers over a sequence of SAR pulses. The KF state space is range and range-rate. We derive data inputs to the KF algorithm from multiple SAR pulses, divided into Coherent Processing Intervals (CPI). Within each CPI, individual scatterer peak amplitudes and phases are available to the algorithm.

Our approach to scene-derived motion compensation combines the high accuracy range history estimates with a novel three-dimensional geometric inversion. This geometric inversion uses the range histories to estimate both 3D scatterer location and 3D relative motions of the radar. We illustrate our KF-based approach to high-accuracy tracking and demonstrate its application to estimating scene scatterer locations on synthetic and real collected SAR data.

Keywords: Synthetic Aperture Radar, SAR, Detection, Tracking, Kalman Filter, Nonlinear Optimization

1. INTRODUCTION

Conventional high bandwidth radar systems possess fine range resolution, but poor azimuthal (cross-range) resolution at useful stand-off distances. Synthetic Aperture Radar (SAR) is a radar imaging mode that achieves fine azimuthal resolution by formulating a large effective coherent aperture in the cross-range direction (a "synthetic aperture") [1]. This effective aperture results from the radar platform's motion.

One critical component in joining the radar returns at different synthetic aperture points is compensating for the phase changes associated with platform motion. This motion compensation (mocomp) procedure allows fully coherent combination of returning signals to create a high-resolution image.

Effective mocomp requires highly accurate knowledge of the platform's position. Even small errors in estimated position can lead to severely degraded images. For example, at the typical X-band frequency of 10GHz, the radar signal's wavelength is about 3cm. Limiting positional deviations to $\lambda/10$ for good coherent signal combination requires that the platform's motion be known to within 3mm.

The work described here attempts to estimate the platform's motion very accurately by using measured data from the scene. In particular, it uses **only** the range histories of multiple scatterers to evaluate both platform motion and scatterer locations in three dimensions. It combines the approach of one of this paper's authors [2] with a new method of estimating scatterer range histories based on optimal filtering.

This approach is also similar to the Prominent Point Processing (PPP) method [3]. The PPP approach utilizes returns from one, two, or three strong scatterers. The processing uses these scatterers to stabilize scene center, estimate and remove non-uniform scene rotation, and estimate the actual rotation rate of the scene.

Section 2 describes the SAR domain in which we operate. Then Section 3 shows our technical approach, including a novel, dual-input Kalman filter tracker to estimate scatterer ranges over time and our geometric inversion procedure that

estimates both platform trajectory and prominent scatterer locations. Section 4 provides scatterer location estimates from a real data set, with ground truth available. Finally, Section 5 concludes the paper.

2. TECHNICAL BACKGROUND

Tracking has been a key element of radar Command and Control systems for a long time [4]. Most early tracking algorithms rely on associating target detections (target amplitude exceeds a threshold) at the current scan with the correct target track. Once the system has determined the (noisy) measurement that goes with an ongoing track, it attempts to estimate the target's true state (usually position and velocity) based on all data received to date. Tracking methods in academic, commercial, and defense systems fall into both linear and nonlinear approaches [5].

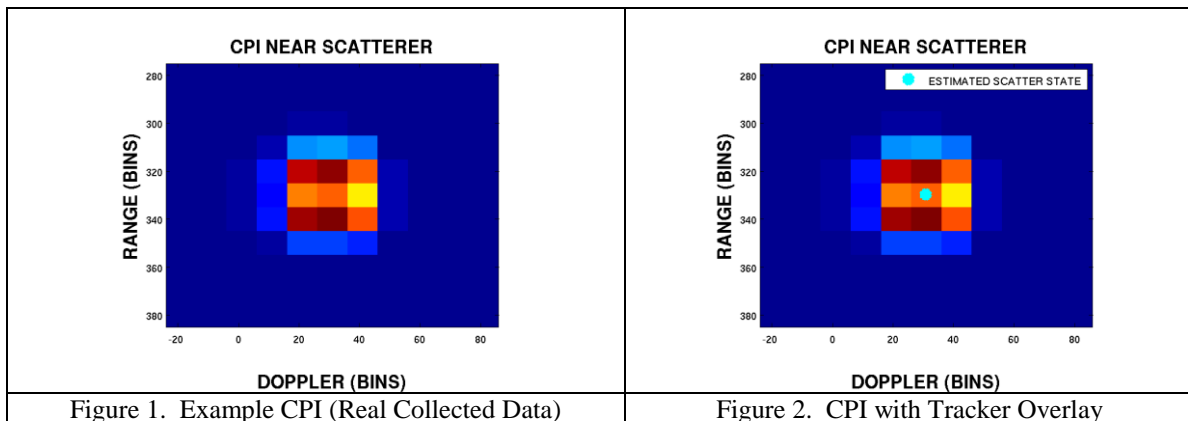
Nonlinear methods are often used in low-SNR regimes, and lead to non-Gaussian probability densities which are often approximated using techniques including grid filters and particle filters. Both of these propagate probability densities that describe target state (typically position and velocity in 2D or 3D) and often do not require explicit target detections or track associations, but rather incorporate all measured data to estimate target state [6].

In contrast, in the current problem, we can select to focus on strong scatterers, meaning the scatterer detection probability is very high. In our setting, we expect dozens to hundreds of scatterers in the scene, but only need range histories from on the order of ten scatterers to effectively estimate the platform trajectory. Therefore, we can elect to use those scatterers that are most easily tracked.

In the work described here, we process collected SAR data by compressing in both range and Doppler, and the input to the tracker described below is a sequence of Coherent Processing Intervals (CPIs). Further details on SAR processing are available in [7]. Figure 1 provides an example of one such surface, with Doppler (in discrete bins, roughly corresponding to target radial velocity with respect to the platform) on the horizontal axis and range (in discrete bins) on the vertical axis. We cropped the figure near a strong scatterer, such as the ones we track in this paper, for illustration.

Since the scatterers we seek to detect and track are very strong, it is appropriate to use a tracker that operates on detections. While we have not yet described the tracking process discussed in this paper, we show an example of the data and the tracker state that we estimate in the sections below. Figure 2 reproduces the CPI from Figure 1, and adds the estimated scatterer state as a cyan symbol overlay. In this example, it is clear that the tracker has estimated the scatterer state (in both range and Doppler) in a reasonable way.

The next section describes our approach to utilizing data surfaces like that in Figure 1 to perform the detection, tracking, and estimation tasks needed to perform motion compensation for SAR systems without GPS for fine platform path estimation and then explains our geometric inversion procedure for estimating platform motion and scatterer locations.



3. TECHNICAL APPROACH

Our technical approach combines an innovative method for high-fidelity estimation of strong scatterer range histories with an advanced geometric inversion. The geometric step uses the range history estimates to construct a three

dimensional object estimate by enforcing rigid scatterer connections. Section 3.1 first describes our SAR data signal processing, scatterer detection process, and two-input Kalman Filter method for strong scatterer range history estimation. Section 3.2 briefly reviews our method of using these range estimates to construct a three-dimensional object estimate.

3.1 Scatterer Detection and Tracking

We employ a two-stage approach to scatterer range history estimation. The first step in our tracking algorithm is target detection. We use an automated detection scheme based on the local signal-to-clutter ratio (SCR) and a cell-averaging CFAR criterion. Figure 3 gives a pictorial illustration in one dimension. The detector, which in practice operates in 2D, partitions a group of cells into a “cell under test” (CUT – blue in the diagram), buffer cells (yellow), and clutter (red) cells. The buffer cells prevent signal energy from the CUT from leaking into the clutter cells and falsely lowering SCR. Extension to two dimensional images is straightforward, with buffer cells and clutter cells comprising rectangular cell regions instead of 1D cell sequences.


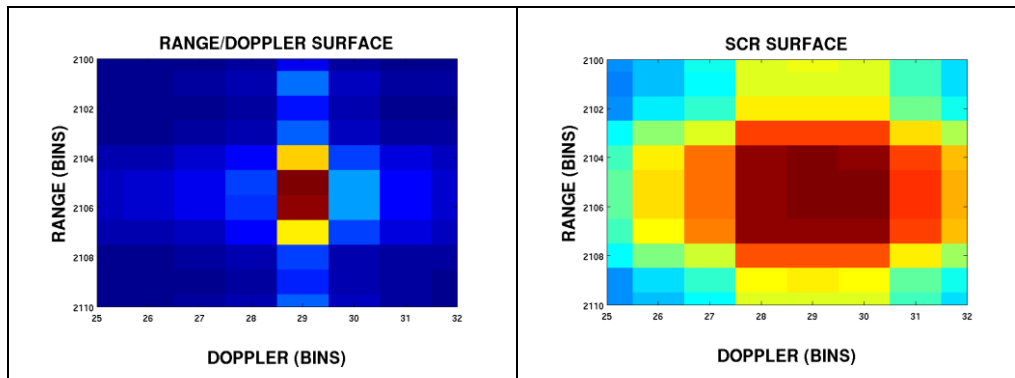


Figure 3. Cell-Averaging CFAR Detector

An example with simulated strong scatterers demonstrates our scatterer detection algorithm. Figure 4 shows the process for a single scatterer. All four panels of the image show the same set of pixels from a range/Doppler surface (Doppler on the horizontal axis and range on the vertical). The top-left panel is a cropped amplitude map of the 2D compressed range/Doppler image. We see a single, strong scatterer. The top-right panel shows the SCR of each pixel from the original range/Doppler image, estimated using the CFAR detector illustrated in Figure 3. The bottom-left panel shows the results of applying a detection threshold to the SCR surface. Here, only pixels with SCR above a threshold survive; all others are set to zero. In practice, we set the threshold to accept only the number of scatterers we wish to track for the geometric inversion process. This process results in an SCR surface with all sub-threshold pixels set to zero. Finally the bottom-right panel shows that we group all contiguous super-threshold pixels into a single “target” via a recursive algorithm. In this case, all green pixels represent a single target.



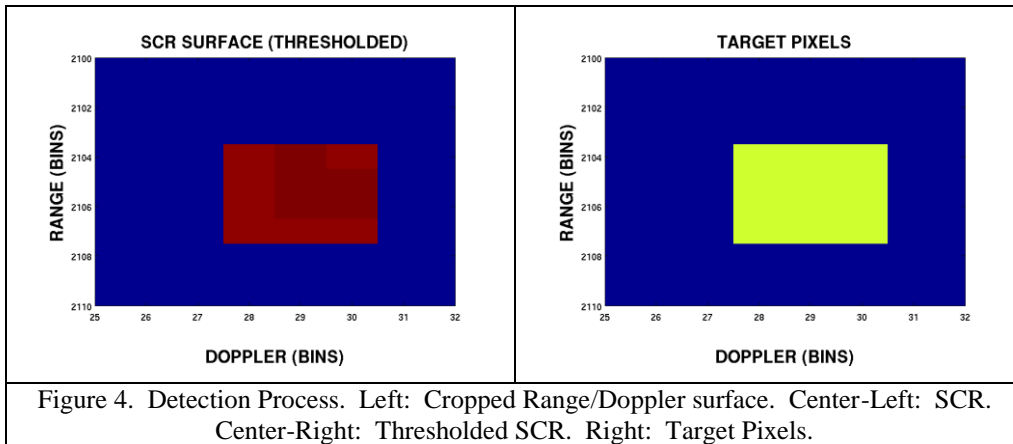


Figure 4. Detection Process. Left: Cropped Range/Doppler surface. Center-Left: SCR. Center-Right: Thresholded SCR. Right: Target Pixels.

At this point in our processing, we have identified this scatterer (and similarly we have identified all high-SCR scatterers). We wish to obtain an initial estimate of scatterer state in range and Doppler, and we accomplish this task by taking the 2D centroid over all “target-present” pixels (e.g., the green pixels in Figure 4’s right-hand panel). Figure 5 demonstrates. The left-hand panel reproduces the range/Doppler surface from the left-hand panel in Figure 4, and the right hand panel shows the 2D centroid as a green overlay. Visual inspection confirms that this initial estimate of target location (in range/Doppler) has small error compared to the system’s point spread function.

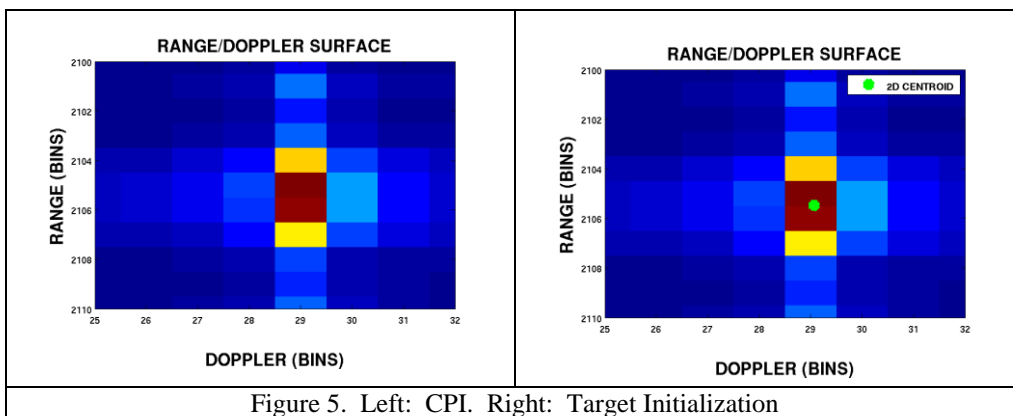


Figure 5. Left: CPI. Right: Target Initialization

In the second stage of our range history estimation, we use new measurements to update our estimate of target range and Doppler. As indicated above, the high-SCR nature of the measurements at each CPI justifies an approach that associates detections with existing targets. Many traditional trackers use each detection’s range estimate and feed these to the tracker as data. By considering data that spans a long time period, the tracker estimates scatterer velocity as well as position.

However, Doppler information is available to the system. We exploit this information to improve the estimation. We thus use a Kalman Filter that uses both range and Doppler as inputs to estimate scatterer state. This approach has significant advantages over traditional methods that utilize only range measurements because it exploits the available data more effectively. Using Doppler is particularly effective because it integrates phase over coherent processing intervals (CPIs) and is less susceptible to amplitude variations.

The following table summarizes the model equations:

$x = \begin{bmatrix} r \\ \dot{r} \end{bmatrix}$	Target state
$x_{k+1} = Fx_k + w_k$	Target kinematic model. Typically, $F = \begin{bmatrix} 1 & PRI \\ 0 & 1 \end{bmatrix}$
$z_k = Hx_k + v_k$	Measurement model

In the table, k is the time index, F is the system dynamics matrix, w is process noise, z is the measurement, H is the measurement matrix (that relates data to target state), and v is measurement noise. As is typically done, we assume the noise is Gaussian and independent.

In our approach, we use Doppler measurements as estimates of range rate. Under this measurement model, our

measurement matrix becomes $H = \begin{bmatrix} \frac{1}{r_{bin}} & 0 \\ 0 & \frac{-4\pi(PRI)}{\lambda_0} \end{bmatrix}$, where r_{bin} is the range bin resolution. The (1,1) entry

relates the range to the range measurement because range is in units of meters and the measurement is in units of range bins. Thus, the range-cell size (in meters) is the constant relating scatterer state to measurement. The (2,2) entry comes from the interpretation of phase changes as a measure of radial velocity. Under this interpretation, scatterer phase

change for a displacement of Δr is $-2\frac{2\pi}{\lambda_0}(\Delta r)$, which yields the expression

$$\Delta\phi = -2\frac{2\pi(PRI)}{\lambda_0}\left(\frac{\Delta r}{PRI}\right) = \frac{-4\pi(PRI)}{\lambda_0}\dot{r}.$$

Thus, the (2, 2) entry of the H matrix relates the phase change at the scatterer to the radial velocity of the scatterer.

Under these conditions, $z_k = \begin{bmatrix} \frac{1}{r_{bin}} & 0 \\ 0 & \frac{-4\pi(PRI)}{\lambda_0} \end{bmatrix} \begin{bmatrix} r \\ \dot{r} \end{bmatrix} + \begin{bmatrix} v_A \\ v_\phi \end{bmatrix}$, where v_A and v_ϕ are process noise terms for the

amplitude and phase measurements, respectively. Here, the data consist of a measurement based on the amplitude's peak location (z_A), and a range-rate measurement based on its phase (z_ϕ). Noise corrupts each of these two inputs.

To verify that the approach is effective, we have run a set of simulations using a real airborne platform motion history with simulated scatterers on the ground. Figure 6 shows the test scenario, plotting the XY component of the airplane's path (black line) and five scatterers (blue asterisks). The platform trajectory is derived from an actual circular SAR airborne collection. The circular collection is a good choice because it provides differing radial motion across the scatterers.

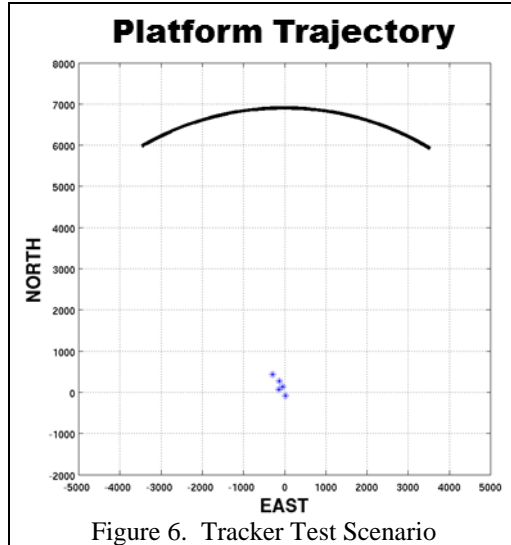


Figure 6. Tracker Test Scenario

Figure 7 illustrates tracking results. The x-axis represents the pulse index, spanning 500,000 pulses (about four minutes). The y-axis denotes range, in meters, from the platform. The blue lines show true target ranges, and the red overlays (virtually indistinguishable from the blue truth curves) show the tracker outputs for all five scatterers. The range profiles of each scatterer change in different ways over time. It is important to note that the tracker can handle each scatterer's range history with minimal error.

Because the error is difficult to visualize in Figure 7, we provide Figure 8, which shows tracker **error** for each scatterer in a different color. Once startup transients have diminished, the figure reveals clearly that tracker error is less than 1mm for all scatterers at each pulse.

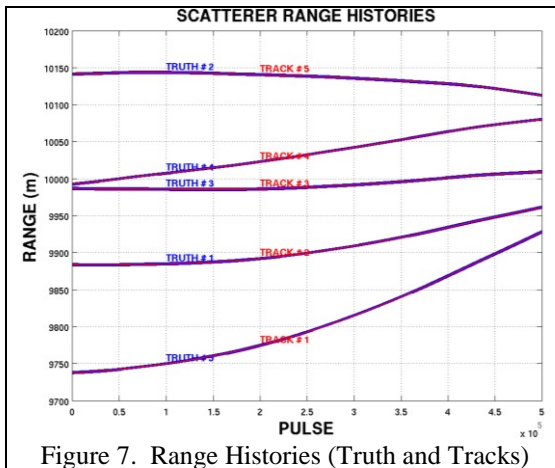


Figure 7. Range Histories (Truth and Tracks)

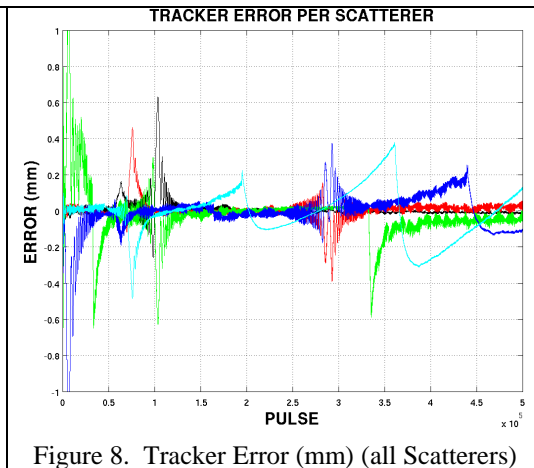


Figure 8. Tracker Error (mm) (all Scatterers)

3.2 Geometric Inversion Estimator

For a radar scene that includes a finite set of reflectors, the reflectors' range histories contain nearly all the information that is present in the radar signals concerning the size and shape of the reflector configuration. They also contain essentially all the information (present in the original radar signals) concerning the relative motions of the sensor and the reflector configuration. So, in the absence of other information, the range histories of the reflected signals provide sufficient statistics for the estimation of the shape, size, and relative motion.

The size and shape information is encoded in the range histories via a set of motion-invariant constraints that the range histories must satisfy [2]. Consequently, at each time sample, the vector of ranges to the reflectors satisfies an invariant equation. When the motion is sufficiently complex, the resulting system of invariant equations uniquely determines the size and shape of the three-dimensional reflector configuration.

The system of invariant equations fails to uniquely determine the size, shape, and relative motions when the direction vectors from the centroid of the reflector configuration to the sensor are all contained in (or near) the surface of some fixed elliptic cone, which consists of a fixed pair of intersecting planes, or a single plane. The algebraic description and derivation of the illumination conditions are given, for the far field case, in [8]. Our research has determined that, in the near field case, the same issues arise.

In practice, such restricted flight trajectories never exactly occur due to minor perturbations from intended flight trajectories. But the most common intended flight trajectories for SAR systems are straight lines or circles. These intended paths are special cases of trajectories that keep the illumination directions on the surface of an elliptical cone. The result is a set of actual trajectories that are close to the ambiguous cases for the proposed algorithm, and thereby make the estimation of the size, shape, and motion sensitive to small errors in the range measurements. This practical situation has motivated our efforts to minimize range error estimates since an important issue with the proposed algorithm is range accuracy. To reliably get accurate estimates of the shape and motion, the random deviations in the range estimates need to be small. The deviations must be smaller than the true flight trajectory departures from the nearest trajectory having illumination directions that generate an elliptical cone.

4. RESULTS

We provide an example result using 10,000 pulses from a real airborne platform. The platform in this case was the General Dynamics Data Collection System (DCS) flying a straight line near a field of corner reflectors near Willow Run Airport in Ypsilanti, Michigan.

Figure 9 shows the range history estimates for 13 of the reflectors. These estimates are outputs of our tracker, similar to the simulated results in Figure 7.

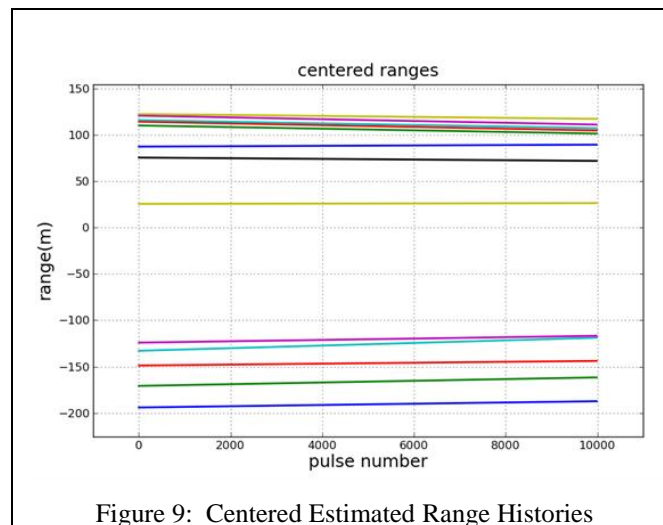
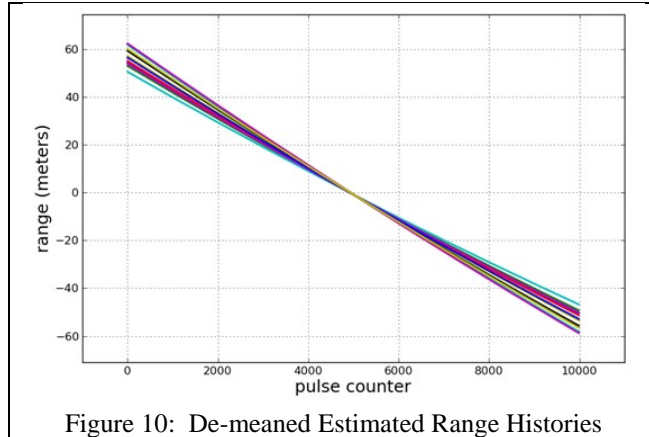


Figure 9: Centered Estimated Range Histories

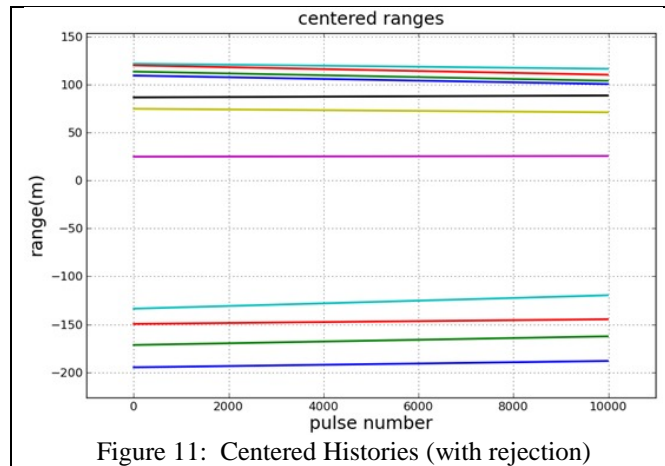
In this display, the range histories appear as nearly straight lines. This is to be expected, because of the limited resolution and extent of the display, in range and time. The display compresses over 250 meters of slant range extent into the figure, so fine details of the range estimates are not apparent.

Since the range histories differ primarily by large constants (range offsets), an improved way to visualize the other characteristics of these estimates involves subtracting each scatterer's average range (over time or pulse number) from its range history. We plot these deviations from the mean values in Figure 10.



The centered range estimates in the figure appear reasonable. As previously noted, the reflectors contribute range history estimates that have various slopes, which are due primarily to their unequal cross-range locations. The tracks appear completely smooth at the resolution of the plot, so there is no immediate indication that the geometric analysis will fail. We performed the scatterer geometric analysis using the approach of Section 3.2. The range history estimates were centered at each pulse time, and projected onto the principal three-dimensional subspace.

Two of the 13 scatterers were outliers due to poor tracking performance. We removed the two bad scatterers and continued processing with the remaining 11. The remaining range history estimates are in Figure 11.



For these edited range estimates, the least-squares solution of the invariant equations yields geometric invariants that correspond to a physically plausible geometric configuration. These coordinates are displayed, from the principal points of view in Figure 12.




Figure 12: Three views of the 3D scatterer Estimates (with outlier rejection)

Relative to the estimated coordinates, the geometric analysis estimates the radar's illumination directions. This is related to the platform's moving position over time. These are illustrated in Figure 13.

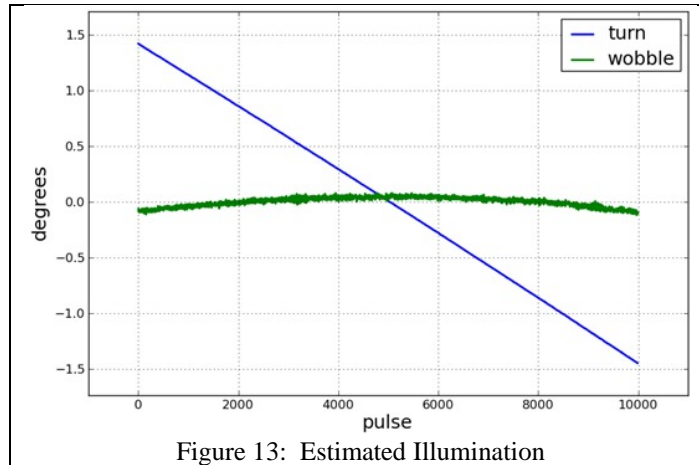


Figure 13: Estimated Illumination

The blue turn curve is approximately linear, indicating that the sensor appears to be progressing about the array at an approximately constant angular rate of about 0.000287° per inter-pulse period. The green curve, which estimates the wobble, shows an approximately parabolic arc, which has about 0.15° deviation from a constant curve. (A constant curve is the wobble-free case, which is desired for standard SAR imaging).

As in the previous subsection, we also calculated the residuals, to examine the apparent rigidity of the estimated range histories. The residual deviations are displayed in Figure 14.

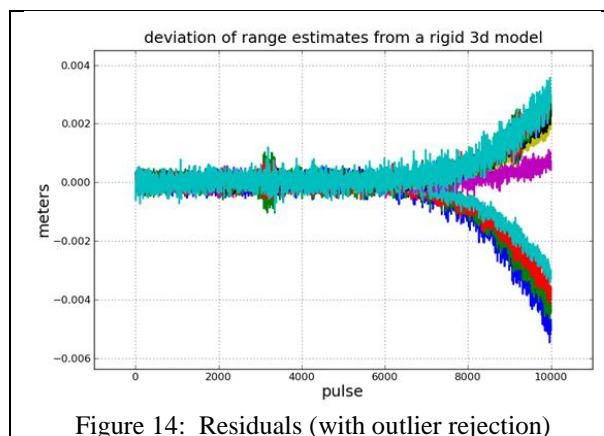


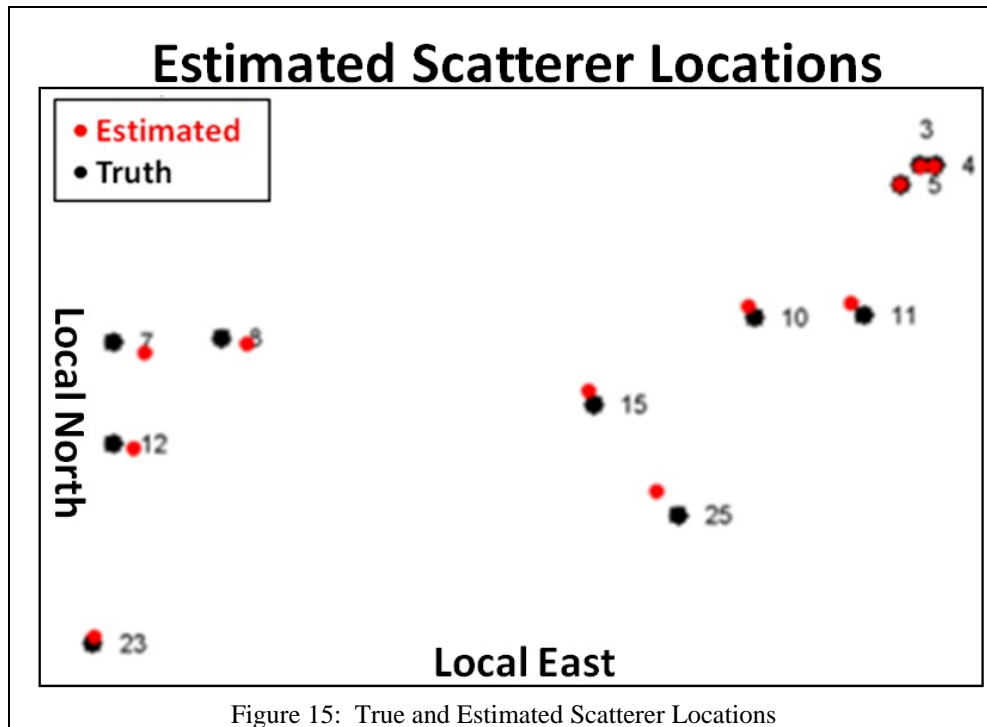
Figure 14: Residuals (with outlier rejection)

These residuals illustrate the estimated range history perturbations that would be required to make them consistent with the rigid scattering center configuration. The errors present are small, suggesting a high probability that the scatterer location estimates will be close to the true positions.

We now compare the scatterer location estimates from Figure 12 with Willow Run truth data. Because the shape estimation algorithm is unaware of absolute directions (like local East/North/Up), it provides results in an arbitrary coordinate system. Thus, we performed an alignment optimization step to visualize the accuracy of the results.

Figure 15 shows a result. We display the results in the local East/North plane. In the figure, the black symbols show scatterer truth positions, and they are also labeled with their respective indices from Willow Run truth documentation. The red symbols show the results of a least-squares optimization, which minimizes the squared error between the estimated scatterer positions and the true positions (in 3D). The errors in scatterer location estimates have a maximum value of about 10 meters.

While this accuracy is not sufficient for phase-coherent motion compensation, we have made substantial progress towards a usable product with the work presented here.



5. CONCLUSION

We have described an algorithm for performing SAR motion compensation in the absence of the traditional motion compensation sensors. Our approach is scene-driven and requires the presence of some prominent scatterers in the scene. The method first uses a fine-resolution tracker to estimate scatterer ranges to the platform and then uses a geometric inversion to estimate the platform's position over time. In this paper, we introduced a high-fidelity tracking algorithm to perform this first step. Results with both simulated and real data show tracking with very low errors. The method's second step uses the range estimates and a geometric inversion approach to estimate scatterer locations and platform motion. Results with real data are very promising, geo-locating scatterers in an approximately 1km scene to within about 10m.

The next steps for this research are to improve tracking and geometric inversion fidelity to the point where the estimated platform trajectory can serve as an input to a motion-compensation algorithm.

ACKNOWLEDGMENT

This work was supported by the US Air Force Research Laboratory under contract FA8650-10-C-1704. Approved for public release by Wright Patterson AFB, document number 88 ABW-12-1828

REFERENCES

- [1] C.A. Wiley, "Synthetic Aperture Radars", IEEE Transactions on Aerospace and Electronic Systems, v. AES-21, no. 3, pp 440-443, May 1985
- [2] Stuff, Mark A., Sanchez, Pedro, Biancalana, Martin "Extraction of Three-Dimensional Motion and Geometric Invariants from Range Dependent Signals", Multidimensional Systems and Signal Processing, 14, 161-181, 2003, Kluwer Academic Publishers
- [3] W. Carrara, R. Goodman, and R. Majewski, Spotlight Synthetic Aperture Radar Signal Processing Algorithms. Artech House: Boston, MA, 1995.
- [4] S.S. Blackman and R. Popoli, Design and Analysis of Modern Tracking Systems. Artech House: Boston, MA, 1999.
- [5] L.D. Stone, T.L. Corwin, and C.A. Barlow, Bayesian Multiple Target Tracking. Artech House: Boston, MA, 1999.
- [6] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for on-line non-linear/non-Gaussian Bayesian tracking," IEEE Transactions on Signal Processing, vol. 50, pp. 174–188, Feb. 2002.
- [7] L.A. Gorham and L.J. Moore, "SAR Image Formation Toolbox for MATLAB," Proceedings of the SPIE, Volume 7699, pp. 769906-769906-13 (2010).
- [8] M. A. Stuff, Derivation and Estimation of Euclidean Invariants of Far Field Range Data, PhD Dissertation, University of Michigan, 2002