SPECSweb Tracking on a Large, Simulated Multistatic Field of the MSTWG

Doug Grimmett
Code 56560
SPAWAR Systems Center Pacific
San Diego, CA, U.S.A.
grimmett@spawar.navy.mil

Abstract - Effective fusion and tracking of multistatic active sonar contacts is challenging, due to high levels of false alarm clutter present on all sonar nodes. Such false alarms often overload the sensor-to-fusion-center communications links and fusion/tracking processes, producing too many false tracks. The Specular-Cued Surveillance Web (SPECSweb) multistatic tracker mitigates these problems by allowing track initiation to occur only when high-strength specular target detections are identified. Using these "specular" cues, and subsequent track state estimates, a selective data retrieval approach is used which significantly reduces the data rate at the input to the fusion/tracking algorithm, and reduces node to fusion-center communication link throughput requirements. This paper provides performance results of this tracking algorithm on a simulated multistatic data set from the Multistatic Tracking Working Group (MSTWG). The data set simulates a large multistatic field, and is characterized by low probability of detection (PD), high false alarm rate (FAR), and high measurement errors. The impacts of these challenging conditions on tracker performance are explained, including the degradation of association gating with large error in bearing measurements.

Keywords: Tracking, data association, Multistatic sonar, Kalman filtering, estimation.

1 Introduction

Distributed multistatic active sonar networks have the potential to increase ASW performance against small, quiet, threat submarines in the harsh, clutter-saturated littoral and deeper ocean environments. This improved performance comes through the expanded geometric diversity of a distributed field of sources and receivers and results in increased probability of detection, area coverage, target tracking, classification, and localization [1].

However, with the increased number of sensors in a multistatic network, come corresponding increases in the data rate, processing, communications requirements and operator loading. Without an effective fusion of the multistatic data, the benefits of such systems will be unrealizable, and thus effective, robust, and automated multi-sensor data fusion and tracking algorithms become an essential part of such systems. Much progress has recently been made in this field [2]; however, overloading due to high false alarm rates is still a major issue. Multistatic fusion algorithms are still challenged to automatically output a sufficiently low false track/alert rate to the operator in these reverberation- and clutter-rich conditions. Communication links may not have the throughput capacity to transfer all of the associated information from the multistatic nodes to a fusion center.

A concept referred to as the "Specular-Cued Surveillance Web (SPECSweb)" is being pursued to address this data rate problem through "specular cueing", directed data retrieval, retrospective tracking, and novel fusion techniques. A "specular" or "glint" detection can occur within a multistatic network when a sonar source pings while located at an angle astern of the target's beam aspect (±90º from the bow), at the same time that a receiver is located with equal angle forward of the target's beam aspect (or vice versa). This occurs when the bistatic aspect angle (bisector of the source and receiver angles) is perpendicular to the target's heading. Such echoes can exhibit an order of magnitude (or more) increase in received level compared to echoes from targets that are not in the specular geometry [3]. When in the specular geometry, there is greatly increased target strength, producing increased echo energy, as indicated by various models and data analyses. SPECSweb exploits specular echoes received within a multistatic network to achieve significant improvements in target detection, classification, and tracking performance.

2 SPECSweb Algorithm Description

Detailed descriptions of the SPECSweb multistatic tracking algorithm and specular cueing approach are found in [4-5]. A summary of the approach is provided here.

Multistatic processing provides the following measurements which relate to target kinematics: bistatic time-of-arrival, bearing, and bistatic range-rate (if Doppler-sensitive waveforms are used). Monostatic sonar measurements may be converted into x-y coordinates and

debiased for use in a Kalman Filter, as described in [6]. A generalization of this approach for bistatic sonar measurements is given in [7] with bias corrections given in [8]. Here, time-of-arrival, bearing, and source/receiver positions are used to calculate converted x-y geographic positions of the detections (and their uncertainties) using the non-linear bistatic mapping/transformation.

Each sonar node self-searches each processed (and locally stored) scan for detection contacts which are likely to correspond to loud specular echoes. This is done either by determining those that exceed a high SNR threshold setting (HTH), or by automatically identifying large positive changes in echo level statistics (via a rise threshold: RTH) from ping-to-ping on each source-receiver node [9]. These are initially sent over the communication link to the multistatic fusion center for potential track initiation. Specular targets also have zero range-rate, where Doppler-sensitive waveforms (eg. pulsed continuous wave, CWs) are overcome by reverberation. Therefore, only frequency modulated (FM) or impulsive echoes are considered as specular cues. Such an approach minimizes the number of spurious tracks initiated by lower-level false alarms, while tracks for specular targets are successfully initiated.

In general, it takes longer to detect a specular target than a non-specular target, so this approach is best suited to surveillance applications. Modeling of specular occurrence indicates that though the specular geometry is rare, for large multistatic fields, there will be a sufficient opportunities for it to occur [10] and inititiate tracks. Once a track is initiated, reverse-time tracking is performed to capture any available historical tracking and detection information. This contextual data can be exploited to further inform the operator, even as new data is received and tracking is performed in the forward-time direction.

Cues are mapped to x-y positions in Cartesian coordinates, and these positions with their associated error covariances are sent as snippet requests to other nodes. These nodes calculate the appropriate snippet boundaries in their respective measurement spaces within which data association would be possible, according to a specified gating parameter. Any contacts found within these snippet boundaries, and above a standard low-threshold (LTH), are sent over the communication links to the fusion center for further processing. As track estimates are obtained, they themselves are used to generate snippet requests for selective data retrieval on prior (for reverse-time tracking) and subsequent (for forward-time tracking) scans stored on any of the nodes. A diagram of the concept is shown in Figure 1.

Additional elements of the SPECSweb tracker implementation include the following:

- A logic-based track initiation (M/N) and termination (K) scheme
- Target motion, modeled using a 2-d *nearly constant velocity* motion model [4-5].
- Nearest neighbor data association, with a 2-d or 3-d (if Doppler measurements are available) ellipsoidal association gate [11].
- An Extended Kalman Filter (EKF), using de-biased [8], converted measurements for positional measurements [7,12] and non-linear bistatic range-rate measurements [5,13].

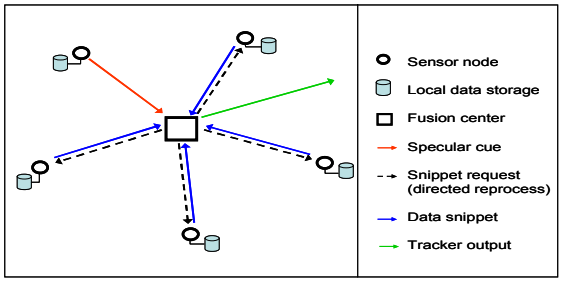

Figure1. Diagram of the SPECSweb cueing concept.

3 Metron Simulated Data Set Description

Metron, Inc. has produced multistatic sonar simulated data sets for use by the members of the Multistatic Tracking Working Group (MSTWG). This working group is organized under the auspices of the International Society of Information Fusion (ISIF), which is addressing the development and testing of multistatic fusion and tracking algorithms [2]. This paper only addresses the first of five simulated data sets which Metron provided. This is the only one for which target-truth information was also provided at the time this work was conducted. This data set simulates a much larger field than has previously been evaluated by the MSTWG.

The simulation scenario is as shown in Figure 2. The field consists of 4 sources (red) and 25 receivers (blue and red) distributed over a square area. Four of the receivers are collocated with the sources (as monostatic nodes). Four targets are present, each executing four cycles of a "box" trajectory over a ten-hour period. Targets 1 (red) and 4 (magenta) proceed in a clockwise direction while targets 2 (green) and 3 (cyan) travel counter-clockwise. Targets 1 and 2 travel with a speed of 6 m/s while targets 3 and 4 travel with a speed of 3 m/s. An aspect-dependent target strength was simulated, with beam aspect echoes slightly (about 12 dB) larger than for non-beam aspects. Source transmissions were cycled amongst the four sources, with a ping repetition interval of three minutes. Each individual source alternates between CW and FM waveforms, creating a cycle time of 24 minutes to return to the same source-waveform pair.

The data set consists of 200 pings. There are 5000 measurement scans, each one representing a source-receiver-waveform triple. For each ping, there are 25 scans received on each of the 25 receivers. Measurement

errors were provided with the data set description, (also validated through analysis) to be Gaussian zero-mean with:

echo timing errors of σ_τ= 0.4 sec for both FM and CW waveforms, bistatic range-rate errors of σ_{RR}= 0.5 m/s, and bearing errors of σ_θ= 8º for all receivers. As will be shown subsequently, the impact of bearing errors of this magnitude degrades tracking performance.

For this data set, identification (truth) tags for contacts originating from targets 2 and 3 were provided, though with the identical labels (not distinguishable from one another). Through separate analysis, we have re-tagged the contacts originating from targets 2 and 3 such that they are distinguishable from one another. An input Receiver Operating Curve (ROC) for target 3 is shown in Figure 3. The data set's probability of detection (PD) is the number of target-originated detections divided the number of detection opportunities. We see that over a range of (SNR) thresholds, the data set's PD (average per scan over the entire scenario) is quite low, always less than 0.13. The false alarm rate (FAR) is shown to be about 35 contacts per scan at the highest level of PD.

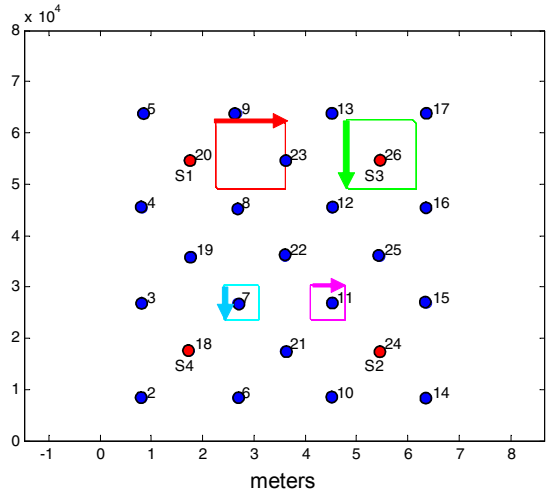

Figure 2. The Metron simulated scenario; 25 receivers (blue and red dots), sources (red dots), and four targets executing box geometries (red, green, cyan, and magenta).

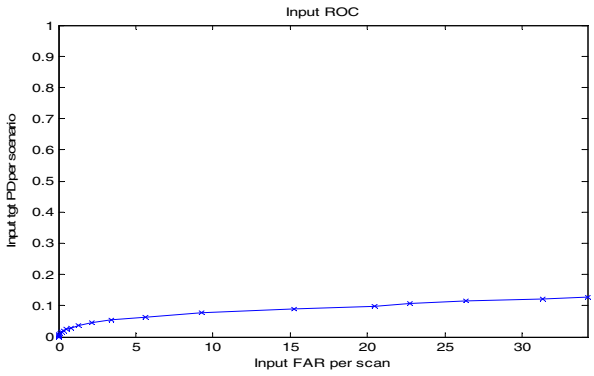

Figure 3. Probability of detection vs. FAR, per scan, over the entire simulated scenario (target #3).

4 Tracking Results on Metron data set

First, the tagged contacts originating from target 3 were extracted and input into the SPECSweb tracker. In this configuration, no non-target false alarms were input to the tracker. The results are shown in Figure 4. We see that the target is successfully tracked. However, the output track localization is somewhat poor, and following the target through turns is particularly challenging. This is due to the large measurement errors, which not only produce poor localization estimates, but also contribute to "missed" detections. Some target detections are not used in the tracker update because they don't fall within the association gate. These "missed" target detections are ones with large bearing errors, and particular source-receiver geometries, where the mapping linearization is invalid. For the Metron data set, target #3 provided 641 contacts out of 5000 detection opportunities, resulting in an input PD of 0.13. In the target-only tracking, 571 contacts originating from target #3 were successfully associated to the track. 70 contacts (11% of the total contacts originating from target #3) were not associated, or "missed", due to large bearing errors. However, taking these losses into account over the entire scenario, the input PD only drops from 0.13 to 0.11, which is not a significant impact. It is possible for "missed" target detections to initiate additional tracks, however. In addition, we expect further degradation in tracking when the non-target contacts are also processed.

Figure 5 shows the cumulative output of the SPECSweb tracker after the entire scenario has been processed, with input parameters as listed in Table 1. A total of twenty-eight tracks were confirmed and output, with each shown in a different color (target trajectories are shown in yellow). Nine of these tracks correspond to target trajectories, at least over a portion of their reported duration. Selected single output tracks for each target are plotted individually and can be seen in a less-cluttered zoom detail, in Figures 6-9. It can be seen that the tracker has difficulty following the target through the turns, and in several cases, the tracks wander away from truth by finding and associating false contacts which feed the continuation of the track. This is most evident for the tracks corresponding to the faster (6 m/s) targets (Figures 6 and 7). The false alarm density is significant enough that false alarms overwhelm and confuse the tracker because few target-originated detections are available. Better performance is seen for the targets 3 and 4. Sixteen of the output tracks are determined to be completely false; six of these are very short-lived tracks and the others initiate on strong false alarm contacts and continue tracking as they are fed by false alarms. This corresponds to a false track rate of 1.6 per hour.

Track fragmentation appears also to be an issue, as the tracker loses the target and dies, or continues as fed by false contacts. Nevertheless, tracks are successfully

initiated three times for each of the four targets, and the targets are held over one or more legs of the box geometries. The tracker performance and target holding (output track PD) for the combination of all track fragments is shown in Table 2. Good tracking performance was achieved on the slower (3 m/s) targets with track holding better than 85%. Performance for the fast targets (6 m/s) was poorer, around 20%. Better performance than this was not achieved, despite trying a range of input parameters, due to the low input PD, high FAR, and large bearing errors. We now consider some of the issues that challenge tracker performance for this data set.

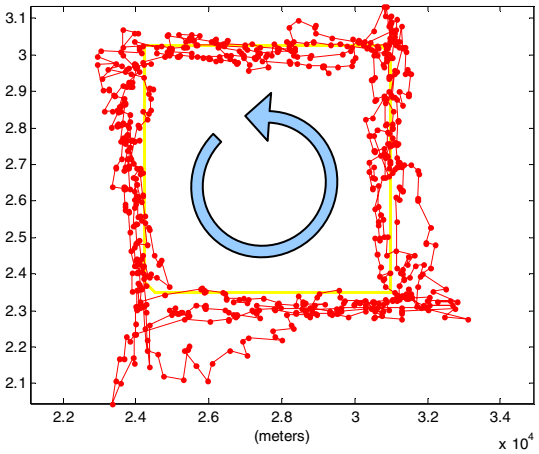

Figure 4. Tracking result for the simulated data, using only the tagged contacts originating from target #3; target's true trajectory (yellow); track estimates (red).

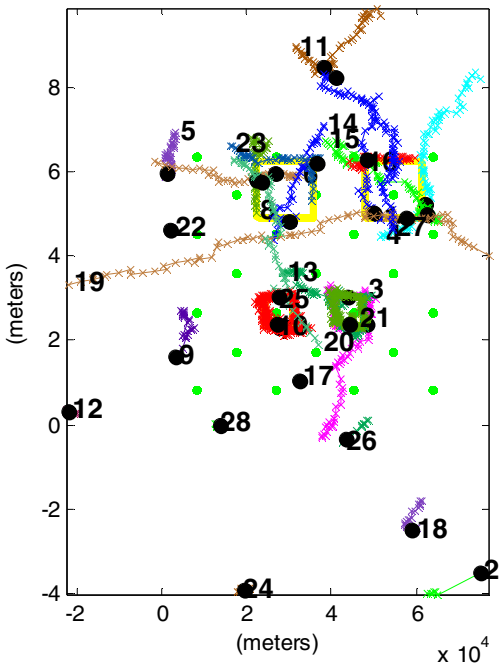

Figure 5. Cumulative SPECSweb tracking result of the simulated data with the input of all contacts (true trajectories shown in yellow).

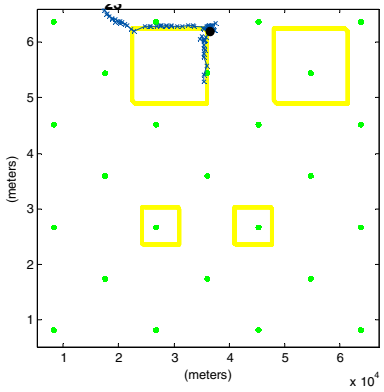

Figure 6. An output track corresponding to target #1.

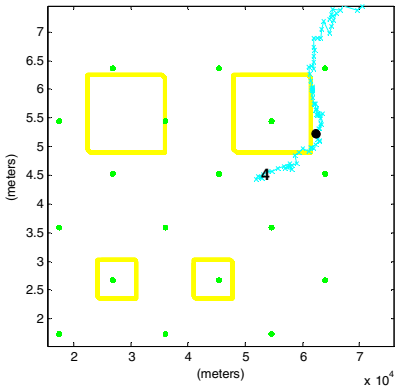

Figure 7. An output track corresponding to target #2.

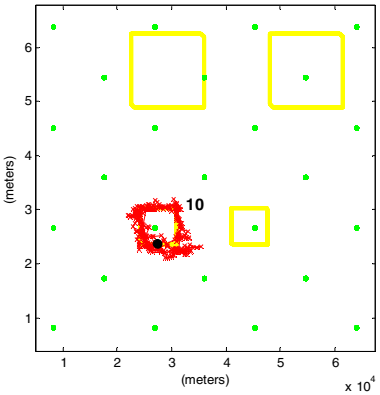

Figure 8. An output track corresponding to target #3.

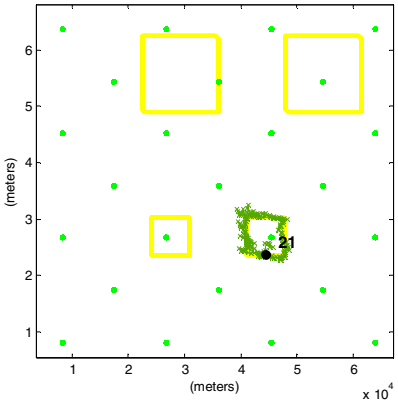

Figure 9. An output track corresponding to target #4.

Table 1. Tracker Input Parameters.

Track initiation (M of N scans)	4/50
Track termination (K scans)	110
Association Gate Probability	99%
Cue Threshold (HTH)	14 dB
Low Threshold (LTH)	3 dB
Manoeuvrability index (m^2/s^3)	0.001
Initial guess target speed	4 kts
Error of initial target speed	2 kts
Error of source/receiver positions	0.1 m
Error of receiver bearing	8°
Error of receiver timing	0.4 s
Error of speed of sound	1 m/s
Error of specular heading	10°

Table 2. SPECSweb Tracker Performance

Target #	Number Track Segments (Fragments)	Track Hold (Output PD)	Number of False Tracks
1	3	25%	16
2	3	19%	
3	3	85%	
4	3	88%	

5 Evaluation of the degradation of data association due to large bearing errors

The presence of large bearing errors for the measurements in the Metron data set challenges data association in Kalman-based tracking. This is due to the linearization approximation of sonar geometries which are significantly non-linear. The SPECSweb tracker currently uses converted bistatic measurements for positions with a Kalman filter update [4]. The measurement error covariance is obtained using analytical expressions derived using small error assumptions [7]. This section evaluates the degradation in measurement-to-track association as a function of bearing errors for various bistatic configurations. The missed target detections in the filtering of target 3, as seen in section 4, are attributable to this effect.

Figure 10a shows an example from the data analysis where the linearization approximation is valid, and which results in an effective data association. The predicted measurement (and 3-sigma error ellipse) is shown in green, near to the true target position (black circle). The target measurement is shown (in blue) along the bistatic ellipse (magenta), very near to the correct bearing (cyan) from the receiver. This measurement passes the gating probability criteria of 99%, and is determined to be the nearest neighbor contact. It is successfully associated with the track to yield the new estimate (and its 3-sigma error ellipse), shown in red.

Figure 10b shows another case with increased range from the receiver and a more eccentric bistatic equi-time ellipse. Here, there is a significant bearing error (on the order of about 21°, or $2.6\sigma_\theta$) in the measurement, relative to truth (cyan). We see that the measurement error covariance (blue) is not able to capture the curvature effect of the bistatic ellipse (magenta), leaving the predicted target state error too far away and disjoint from the measurement error. Here, the measurement does not pass the gating criteria, and it is not associated to the existing track.

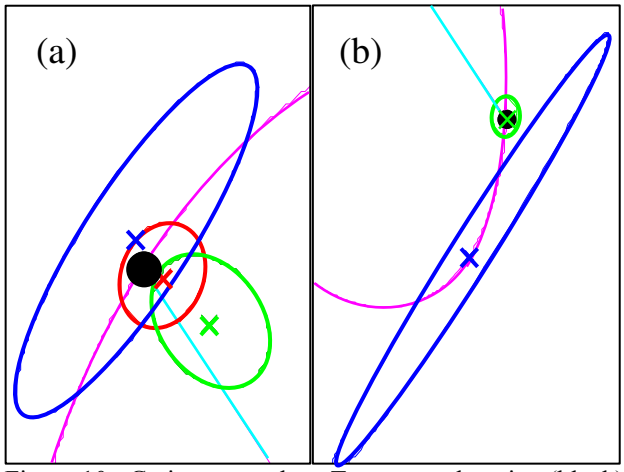

Figure 10. Gating examples. True target location (black); true ranging ellipse (magenta); true bearing (cyan); measurement (blue), predicted target estimate (green); updated estimate (red); (a) Successful gating example, with small bearing error (b) Unsuccessful gating example, with large bearing error.

In order to understand more precisely the point at which bearing errors will degrade the association step within a Kalman filter-based tracker, we consider a set of simulations. We evaluate the gating degradation for a number of monostatic and bistatic source-receiver geometries. For simplicity, we choose to ignore the bistatic range-rate in this analysis and focus on the positional errors which are the primary cause of gating degradation. The gating equation is given as [11]

$$(Z-X(k+1|k))^T(CP(k+1|k)C^T+R)^{-1}(Z-X(k+1|k))<G \quad (1)$$

where $Z=\left[x^m\ y^m\right]^T$ is the x-y converted measurement as obtained from the actual measurements of arrival time and bearing. R is the measurement error covariance, as obtained by analytical, linearization expressions derived assuming small errors [7]. X(k+1|k)= $\left[\bar{x}\quad \bar{y}\right]^T$ is the predicted target state positional estimate, and P is the predicted target state error covariance. C is the measurement matrix, which, for converted measurements,

is just the identity matrix, I_2, and G is the gating value, which can be related to a gating probability. The error analysis is made by simulating a large number (100,000) of measurements, Z, each one derived from bearings drawn randomly from the assumed (Gaussian) distribution of the bearing errors (σ_θ). The measurement covariance (σ_x, σ_y, σ_{xy}) depends on geometry and is calculated using the linearization expressions of [7]. The predicted target position is centered on the point of interest within the geometry, and its positional error covariance is assumed constant with values of $\sigma_x = \sigma_y = 500$ meters, and $\sigma_{xy} = 0$. The gating values for the set of simulated measurements are computed and the percentage of them that are within (smaller than) a gate of 99% is calculated. The results are then compiled as a function of geometry and assumed bearing error.

Figure 11 shows a bistatic source-receiver pair, with five possible target positions located around the equi-time ellipse. Figure 12 shows the results of the gating analysis for these five positions, as a function of assumed bearing error. The specified 99% gating probability is drawn as a black line, and therefore, when the probability of gating drops below this line, there will be a corresponding loss in association opportunity, due to the increasing impact of non-linearities. It is seen that for this geometry, that gating and association is still robust up to about 5 degrees of bearing error. Those target locations further away from the receiver (eg. position #1) and with more elliptical curvature exhibit a greater sensitivity to bearing error on gating probability.

Figure 13 shows the gating analysis for the same receiver used in Figure 11, and varies the position of the source, from monostatic (red) to more extremely bistatic (cyan). Position "S3" is the same position as was used previously. A single target position is indicated, which is the same point as case #2 in the previous analysis. Figure 14 shows the results of the simulation. We see that the more eccentric (non-linear) the bistatic ellipse, the worse the gating degradation.

Figure 15 shows the gating analysis for monostatic geometries, where the source and receiver are collocated, and varies the position of the target along a constant line of bearing at different ranges. Figure 16 shows the results of the simulation. We see that for these positions, the gating degrades with increasing range from the receiver.

We see that bearing errors of less that 5º are less sensitive to non-linearities. Bearing errors of 8º may degrade gating (association) performance, depending on the bistatic geometry. In summary, gating degradation due to extreme bearing errors increases with ellipse eccentricity, range from receiver, and for bearings (from the receiver) in the direction of the source.

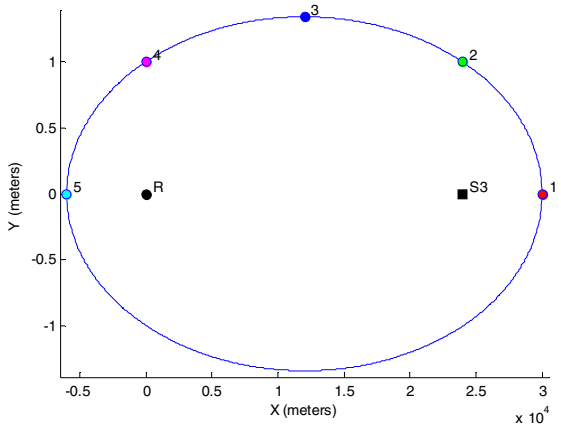

Figure 11. Five positions with equal time delay for a bistatic source and receiver configuration; geometries used for gating results of figure 12.

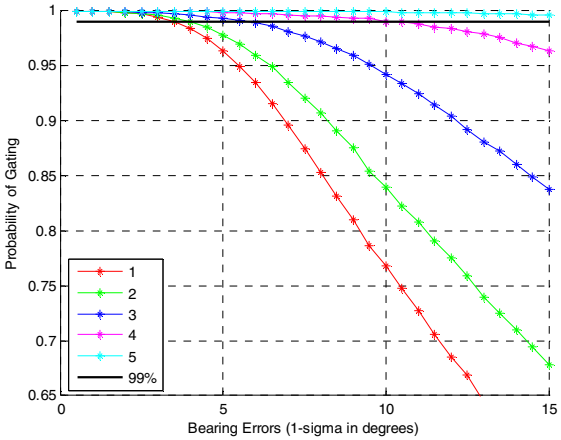

Figure 12. Probability of gating as a function of bearing error for the positions seen in figure 11.

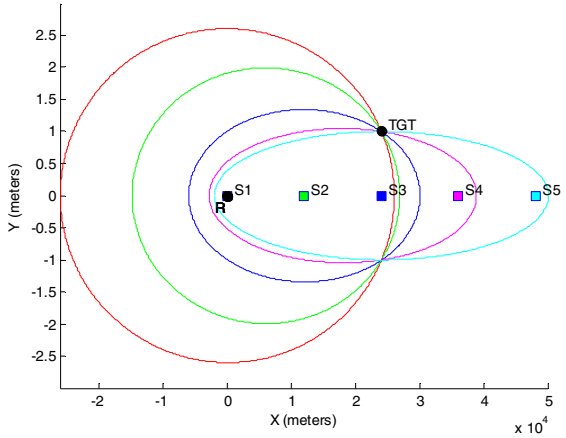

Figure 13. Ellipses with different bistatic eccentricity; geometries used for gating results of figure 14.

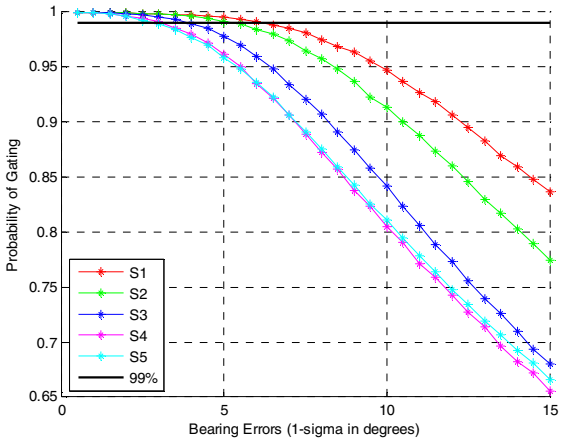

Figure 14. Probability of gating as a function of bistatic ellipse eccentricity and bearing error for the ellipses shown in figure 13.

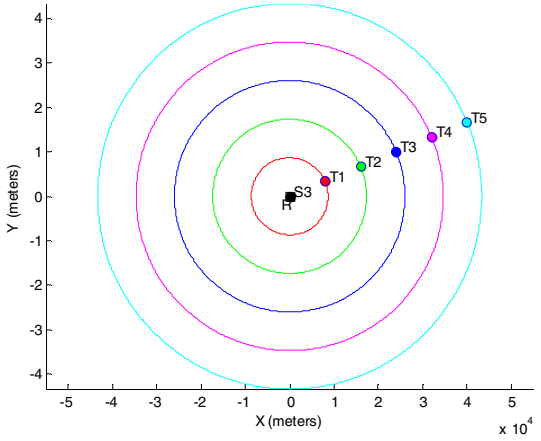

Figure 15. Five positions with increasing range from a monostatic sonar, constant bearing; geometries used for gating results of figure 16.

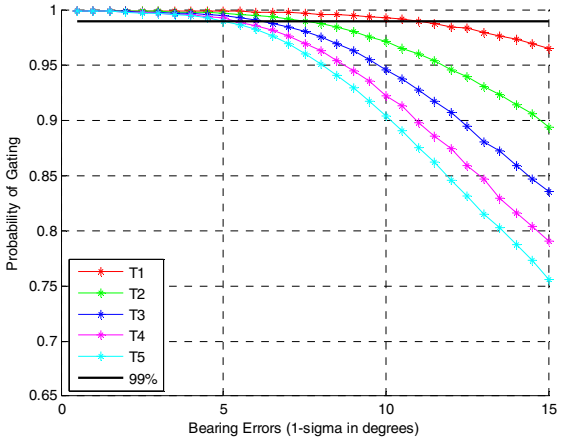

Figure 16. Probability of gating as a function of target range and bearing error, for the target positions seen in figure 15.

6 Ping-fusion preprocessing

An even more significant impact to tracker performance with the Metron data set is the very challenging input PD and FAR statistics. The SPECSweb tracker processes data on a per-scan basis. This means that one update of the Kalman filter is executed for every source-receiver-waveform combination. Since on a per-scan basis the PD is very low (and some of these are lost due to the gating degradation issue), there is a large probability that numerous false contacts will gate and erroneously update the track before a target-originated contact will. This likely explains the loss of target tracking at the turns. While waiting for true contacts to occur, false contacts feed the continuation of the track in an erroneous direction.

Figure 17 shows a ROC (PD vs. FAR) curve for target #3 of the Metron data set on a per-ping basis, where all of the 25 receiver scans for the ping are combined. The PD is calculated by testing if "one or more" target-originated detections exceed the threshold out of all 25 receiver scans. The FAR rate simply combines the number of false contacts for the 25 receiver scans. We see that although the FAR dramatically increases with decreasing threshold, so does the input PD. The PD at the worst FAR has gone from 0.13 to 0.89. This suggests that there may be value in attempting receiver fusion on each ping, prior to performing fusion over time.

We are investigating whether the introduction of a "J out of K receivers" (J/K) fusion criteria, together with Probabilistic Data Association (PDA) would be able to provide increased tracking performance. Assuming an input PD (based on the Metron data) of 0.13, we calculate the output probabilities of detection achievable for such a ping fusion approach. The binomial distribution provides the probabilities shown in Figure 18a, for the case of J detections of 25 receivers. Figure 18b shows the probability for J or more detections out of 25 receivers. We see that by picking an appropriate J (in the range of 1-4) we increase the PD, and at the same time we may be able to reduce the number of track updates produced by randomly occurring false alarms.

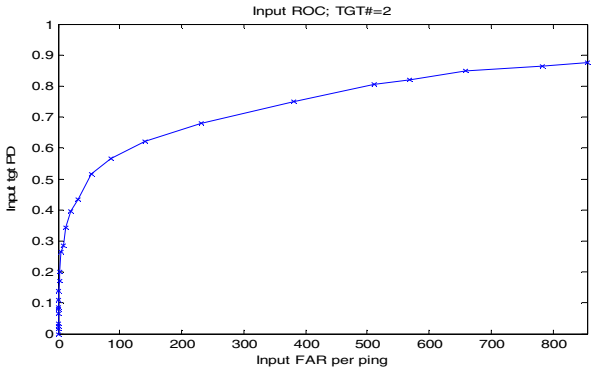

Figure 17. Measured input PD and FAR for target #3 as a function of threshold for fused receivers for target #3: 1 or more target detections out of 25 receiver scans.

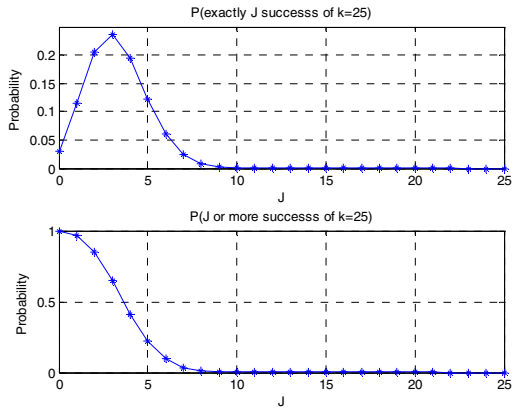

Figure 18. Fused receiver probabilities, based on an input PD of 0.13 per source-receiver scan; (a) Probability of J detections out of K=25 receiver scans for one ping; (b) Probability of J or more detections out of K=25 receivers scans on one ping.

7 Conclusions

The MSTWG Metron data set has been very useful in providing a challenging data set with a large number of multistatic nodes. This has allowed us to understand and identify SPECSweb tracker limitations and deficiencies, and provides insights into possible future enhancements.

The Metron data set is characterized by low PD, high FAR, and large bearing errors. The tracker successfully formed some tracks on the targets though continuous holding was not achieved. Tracker PD ranged from 0.13-0.50. The target tracks were often unable to follow the target through turns. There were more false tracks than desired.

A degradation in the gating of measurements with large bearing errors was observed. This is caused by the linearization of non-linear measurements. The dependence of this effect on various bistatic geometries has been shown, and indicates that geometries with eccentric ellipses (when the target is on the side of the source), will produce serious gating difficulties. However, the frequency of occurrence of these geometries is probably not extraordinary. Bearing errors of 5º or less are less problematic. Bearing errors of 8º (as provided in the Metron data set) are the cause of lost target-to-track associations, depending on geometry. About 10% of the target #3 contacts were lost, reducing the PD for this target by only 2%.

It is hypothesized that a ping fusion process (same ping all receivers), may be advantageous and improve the detections and false alarm statistics enough to improve tracking performance. This is considered a topic for further, future research.

References

[1] D. Grimmett and S. Coraluppi, Multistatic Active Sonar System Interoperability, Data Fusion, and Measures of Performance, NURC Technical Report NURC-FR-2006-004.

[2] S. Coraluppi, D. Grimmett, and P. de Theije, Benchmark Evaluation of Multistatic Trackers, in Proceedings of the 9th International Conference on Information Fusion, July 2006, Florence, Italy.

[3] R. F. Gragg, The BASIS-3D Acoustic Target Strength Model, NRL/FR/7140--07-101052, Washington D.C., U.S.A., Naval Research Laboratory, 2002.

[4] D. Grimmett, Multistatic Target Tracking using Specular Cue Initiation and Directed Data Retrieval, in Proceedings of the 11th International Conference on Information Fusion, July 2008, Cologne, Germany.

[5] D. Grimmett, Specular-Cued Multistatic Sonar Tracking on the SEABAR'07 Dataset, in Proceedings of the 12th International Conference on Information Fusion, July 2009, Seattle, Washington, U.S.A.

[6] D. Lerro and Y. Bar-Shalom, Tracking With Debiased Consistent Converted Measurements Versus EKF, IEEE Transactions on Aerospace and Electronic Systems, Vol. 29, No. 3, July 1993.

[7] S. Coraluppi, Multistatic Sonar Localization, IEEE Journal of Oceanic Engineering, Vol. 31, No. 4, October 2006.

[8] S. Coraluppi and C. Carthel, Progress in Multistatic Sonar Localization and Tracking, SACLANTCEN Report SR-384, December 2003.

[9] D. J. Grimmett, Automatic Identification of Specular Detections in Multistatic Sonar Systems, Proceedings of the MTS/IEEE Oceans'09 Conference, October 2009, Biloxi, Mississippi, U.S.A.

[10] D. Grimmett, S. Sullivan, Sr., and J. Alsup, Modeling Specular Occurrence in Distributed Multistatic Fields, in Proceedings of the IEEE OCEANS'08 Conference, July 2008, Kobe, Japan.

[11] S. Blackman and R. Popoli, Design and Analysis of Modern Tracking Systems, Norwood, MA: Artech House, 1999, ch. 6.3.

[12] S. Coraluppi and D. Grimmett, Multistatic Sonar Tracking, in Proceedings of the SPIE Conference on Signal Processing, Sensor Fusion, and Target Recognition XII, April 2003, Orlando, FL, USA.

[13] S. Coraluppi, C. Carthel, and D. Hughes, Doppler-aided multistatic sonar tracking, NURC Technical Report NURC-FR-2007-024, November 2007.