

Network-centric Angle Only Tracking

Jason Yosinski^a, Nick Coult^a, and Randy Paffenroth^a

^aNumerica Corporation, 4850 Hahns Peak Drive Suite 200, Loveland, CO 80538, USA

ABSTRACT

The coordinated use of multiple distributed sensors by network communication has the potential to substantially improve track state estimates even in the presence of enemy countermeasures. In the modern electronic warfare environment, a network-centric tracking system must function in a variety of jamming scenarios. In some scenarios hostile electronic countermeasures (ECM) will endeavor to deny range and range rate information, leaving friendly sensors to depend on passive angle information for tracking. In these cases the detrimental effects of ECM can be at least partially ameliorated through the use of multiple networked sensors, due to the inability of the ECM to deny angle measurements and the geometric diversity provided by having sensors in distributed locations. Herein we demonstrate algorithms for initiating and maintaining tracks in such hostile operating environments with a focus on maximum likelihood estimators and provide Cramer-Rao bounds on the performance one can expect to achieve. Cramer-rao bound, maximum likelihood estimation, angle only tracking, passive tracking

1. INTRODUCTION

The coordinated use of multiple distributed sensors by network communication has the potential to substantially improve track state estimates, associated feature and attribute estimates, and combat identification (CID) of targets, even in the presence of enemy countermeasures. The improvement is primarily due to geometric diversity, complementary sensor information, and different coverage areas and these effects are precisely those that are particularly desirable for tracking problems involving passive angle measurements.

In the modern electronic warfare environment a network-centric tracking system will have to function in a variety of jamming scenarios. In some scenarios, hostile electronic countermeasures (ECM) will endeavor to deny range and range rate information leaving friendly sensors to depend on passive angle information for tracking. The *lower dimensionality* of the passive angle information causes difficulty in initiating tracks, inability in correctly assigning measurements to current tracks, and failure in maintaining tracks. On the other hand, the efficacy of ECM can be reduced in multi-sensor network-centric architectures through the inability of the ECM to deny angle measurements and the geometric diversity provided by having sensors in distributed locations.

Throughout this paper we will use the terms *lower dimensional* and *low dimensional* to indicate measurements and track states whose positional dimension is lower than three. Similarly, we will use the term *high dimensional* to refer to measurements and tracks whose positional dimension is three. This distinction is of critical importance to all that follows. A classic example of a high dimensional measurement is a range-bearing-elevation measurement, while a canonical example of a low dimensional measurement is a bearing-elevation measurement, where perhaps the range has been denied by electronic counter measures. Another useful distinction to make in passive angle tracking is between two-dimensional *angle only* tracking and one-dimensional *bearing only* tracking. In two-dimensional angle only tracking, bearing and elevation are available and only range has been denied to the sensor, while in one-dimensional bearing only tracking the elevation of the target is not available. Generally bearing only tracking is significantly more difficult than angle only tracking, but both are important and realistic problems. Finally, one often talks about *composite* tracks, that are comprised of measurements from several sensors, and *local* tracks that are comprised of measurements from only a single sensor. Throughout this text the

Further author information: (Send correspondence to R.P.)

J.Y.: E-mail: jason.yosinski@numerica.us, Telephone: +1 970 461 2000 x 256

N.C.: E-mail: nick.coult@numerica.us, Telephone: +1 970 461 2000 x 252

R.P.: E-mail: randy.paffenroth@numerica.us, Telephone: +1 970 461 2000 x 233

assumption is that composite tracks are always high dimensional, unless explicitly stated otherwise, and local tracks are of whatever dimension is convenient for that particular local sensor.

The determination of higher dimensional states from lower dimensional measurements has a long and storied history going at least as far back as Carl Friedrich Gauss. In particular, Laplace and Gauss both worked on the problem of determining orbits from angle only data in an astrodynamics context.¹ A rich and active extant literature exists for these problems with methods that roughly divide into two categories.² Namely, we have batch parameter estimate methods, such as various flavors of least squares processing,³⁻⁵ and recursive stochastic estimation, such as the Extended Kalman Filter and its like.⁴⁻⁶ One factor that permeates much of the literature is the key idea of *observability* and the Cramér-Rao bound which gives a lower bound on mean squared error among all unbiased estimators.^{5,7,8} The authors would like to call out³ for special attention as that article motivated much of our original work in this area.

With respect to tracking, multiple target tracking methods divide into two broad classes, namely single frame and multiple frame methods. The most successful of the multiple frame methods are Multiple Hypotheses Tracking (MHT)⁴ and Multiple Frame Assignment (MFA).^{9,10} The performance advantage of the multiple frame methods over the single frame methods follows from the ability to hold difficult decisions in abeyance until more information is available or, equivalently, the opportunity to change past decisions to improve current decisions. These methods offer improved track accuracy, attribute/feature association, and covariance consistency, as well as reduced number of track switches, track breaks, and missed targets.

In the absence of system latency (e.g., communications and processing) a centralized architecture is provably optimal in that it is capable of producing the best track quality (e.g., purity and accuracy) and a consistent air picture. However, the centralized tracker is unacceptable for several reasons, notably communication overloads, system latency, and single-point-failure. Thus, one must turn to a distributed architecture for both estimation/fusion¹¹ and data association. Network-centric MFA architectures have been developed¹²⁻¹⁴ that achieve tracking performance comparable to that of a centralized tracker, while maintaining SIAP. The achievement of SIAP is based on a Cooperative Engagement Capability (CEC)-like rule that requires that “each platform be in charge of associating its own measurements to the network tracks.”

Throughout this paper we have endeavored to analyze important component parts of a distributed tracking system in the context of passive angle tracking and the relationship between low dimensional local tracks and high dimensional composite tracks. Our focus has been on advanced methods for properly treating the inherent non-linearity between the various spaces of interest (e.g. Cartesian tracking spaces vs. cylindrical measurement spaces) within both track initiation and extension. Also, the network infrastructure in which these algorithms function is of critical importance for realistic systems, and considerations of such a distributed system permeate all of our work.

This paper is organized as follows. In Section 2 we address of the challenge of track initiation with low-dimensional measurements by deriving initiation schemes using maximum likelihood estimation. The core idea here is to carefully analyze the covariances generated by the initiation schemes for consistency. In Section 3, we use these error estimates as a means to examine the efficiency of different sensor and protected asset configurations. These ideas allow one to compute optimal sensor placements based upon given threat scenarios. In Section 4, we explain our current network architecture and describe the additional architectural components necessary to include passive angle local tracks with the higher dimensional composite tracks.

2. INITIATION OF HIGH-DIMENSIONAL TRACKS FROM LOW DIMENSIONAL MEASUREMENTS

In this section, we consider the problem of estimating the initial high-dimensional state of a single target using a sequence of low-dimensional measurements. In particular, we study the mathematical aspects of initiating a two-dimensional track using bearing-only measurements from two or more sensors.

Given a sequence of noisy bearing-only measurements, an *estimator* will produce an estimate of the target state as well as an estimate of the covariance. Our goal is not only to find an algorithm for such an estimator, but also to evaluate the covariance of the estimator relative to the smallest possible covariance that could be achieved given the noise in the measurements. We can accomplish this by utilizing the Cramér-Rao lower bound (CRLB),

which provides a lower bound on the covariance of all possible estimators. Any estimator which achieves this covariance is said to be *efficient*.

2.1 Maximum Likelihood Estimator and the Cramér-Rao lower bound

Let $Z = \{z_i\}_{i=1}^N$ be a set of measurements at times t_i , and $f(t_i, \theta)$ be the function that maps from the high-dimensional state θ into the measurement space. If $p(Z|\theta)$ is the conditional probability of Z given the state θ , a *maximum likelihood estimator* is an estimator which seeks to maximize $p(Z|\theta)$ as a function of θ . Furthermore, if the individual measurements z_i have independent Gaussian noise with equal variances, then the MLE equivalently minimizes the non-linear function

$$S(\theta) = \sum_{i=1}^N (z_i - f(t_i, \theta))^2. \quad (1)$$

This means that for such noise models, the MLE is equivalent to a non-linear least-squares problem (NLLS). A variety of methods including Levenberg-Marquardt, Gauss-Newton, etc., can be used for solving NLLS problems; for our purposes it suffices that there exist well-known techniques for solving the problem.

Given a maximum-likelihood estimator $\hat{\theta}$, we can compute its covariance $Var(\hat{\theta})$. This gives a measure of the uncertainty of the estimator. Since the measurements contain noise, one cannot expect the covariance of any estimator to be zero. Nevertheless, we would like it to be as small as possible. The CRLB¹⁵ provides a lower bound on the covariance; under the assumption that the individual measurements z_i have independent Gaussian noise with equal variances σ^2 , the CRLB reduces to

$$Var(\hat{\theta}) \geq ((Df)^T \sigma^{-2} Df)^{-1} \quad (2)$$

where Df is the Jacobian of f . If the inequality in Equation 2 is replaced by strict equality for a particular estimator, then we call that estimator *efficient*. Note in particular that if f is linear and full rank, then one can easily show that the estimator is efficient.

2.2 Choice of motion model and state space

The function f in Equation 1 provides a mapping from state space to measurement space. Both the assumed motion of the target and the parameterization of the state space are represented by this function. For the sake of simplicity, we assume that during initiation the target has constant velocity.

We consider the problem of estimating this state based on bearing-only measurements at multiple times from sensors at two distinct locations. Since there are four degrees of freedom for 2D constant velocity motion, simple dimension-counting shows that bearing measurements from at least two different times are required in order to have enough degrees of freedom in the measurements to fully recover the state.

Likewise, any parameterization of this state space must have four degrees of freedom. Different choices will result in different versions of f in Equation 1, but *all* such choices (under reasonable assumptions) will result in equivalent solutions of the NLLS problem, in that a solution using one parameterization can be mapped to a solution using another parameterization. The only difference between parameterizations is the lower bound provided by the CRLB. In some spaces, the CRLB may be sharp, while in others, it might not be sharp. If the CRLB is sharp in one particular space and the MLE achieves the CRLB in that space, then we know that in fact the performance of the MLE for *all* parameterizations is as good as it can be.

We compare two different parameterizations of the state space.

- A straightforward choice for parametrizing the state space of all such targets in 2D is

$$\theta_C = [x_0, y_0, \dot{x}, \dot{y}] \quad (3)$$

We call this parametrization the *Cartesian parametrization*; it has an associated mapping f_C which we use in Equation 1.

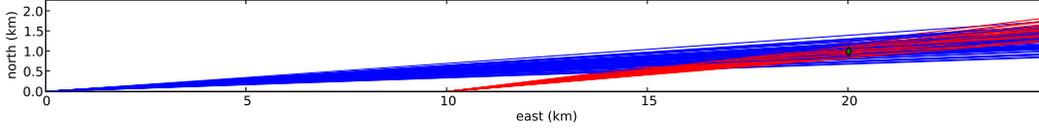


Figure 1. Sample of bearing-only measurements for two sensors and one target. The diamond is the true target position.

- Another choice of parametrization is

$$\theta_A = [\alpha_1, \beta_1, \alpha_N, \beta_N] \quad (4)$$

where α_i is the bearing from sensor 1 to the target at time t_i , and β_i is the bearing from sensor 2 to the target at time t_i . The state is therefore parameterized by first and last bearings to the target. This parametrization also has a mapping function f_A which is used in Equation 1. We call this the *angle-based* parametrization.

To illustrate the differences between these parametrizations, consider two bearing-only sensors with an 8 second update rate. Sensor 1 located at $(0, 0)$ and sensor 2 is located at $(10, 0)$ (10 km east of Sensor 1). Both sensors have a measurement error with 0.5 degree standard deviation. A target starts at $(20, 1)$ and is traveling at 200 kph west and 200 kph north. Figure 1 shows the geometry of the sensors, the target, and a sample of measurements at $t = 0$. As can be seen in the figure, the geometry is fairly extreme; even though enough data exists in principle to recover the target state, the observability is poor.

Figure 2 shows the distribution produced by the MLE over 50,000 Monte Carlo runs. Since these results on the surface appear to be of fairly low quality, we might ask *is this the best we can do?* In fact, we find that the covariance of the MLE in the Monte Carlo simulation is about 30% higher than the covariance calculated by the CRLB. This gap could indicate that an estimator exists with a smaller covariance than the MLE, or it could be that the CRLB is simply not sharp in this particular case. Using only the data provided so far, we have no way to tell whether the former or the latter is true.

Again, we propose that it is not the estimator but rather the *state space parameterization* that is limiting the sharpness of the CRLB. The angle parametrization, while producing the same set of solutions to the NLLS problem, will generally have a different CRLB.

In the case where $N = 2$, the mapping f_A from state space to measurement space is actually just the identity, and the CRLB is easily shown to be sharp. Figure 3 shows the same Monte Carlo results as Figure 2, but mapped into the angle-based parametrization. Since the CRLB is sharp for the angle-based parametrization, and the MLE for the Cartesian parametrization is just a remapping of the MLE for the angle-based parametrization, we can see that the MLE results shown in Figure 2 are in fact the best we can do (in the sense that the estimator has the smallest possible covariance). Finally, Figure 4 shows the comparison of the covariances of the Cartesian and angle-based parametrizations with an ideal χ^2 distribution. As expected, the angle-based parametrization is a closer match.

For the same sensor and target configuration, we can also consider increasing the number of measurements. If $N > 2$, it is no longer the case that $f = I$, and we cannot easily show that the CRLB is sharp for the

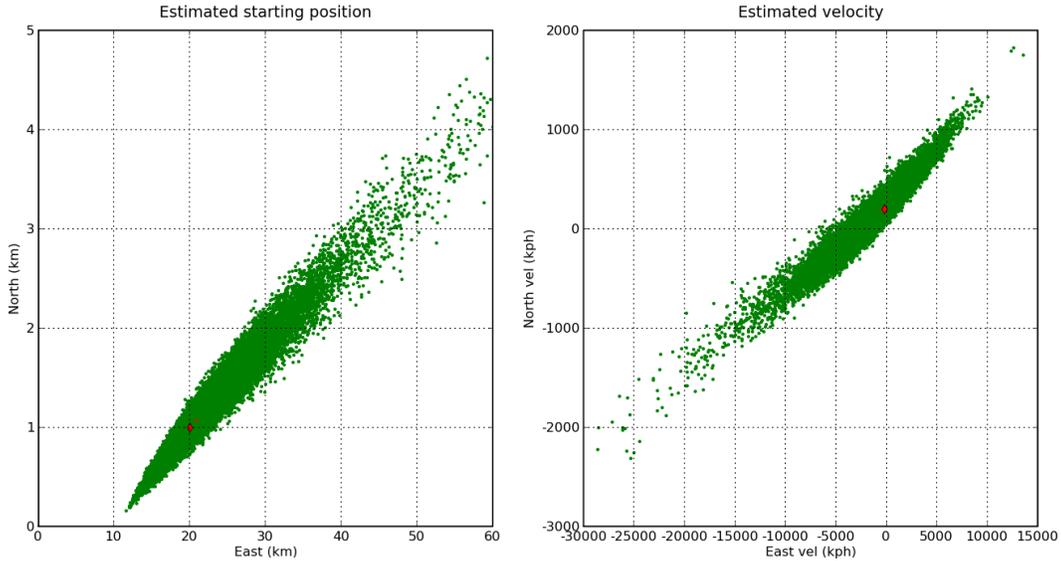


Figure 2. MLE solution for Monte Carlo simulation with $n = 50,000$ for the Cartesian parametrization. Truth is indicated by the diamond, and the mean of the 50,000 runs is indicated by the '+' symbol. Not only is the distribution non-symmetric, the MLE appears to be biased as well.

angle-based parametrization. However, Monte-Carlo simulation shows that even if $N = 5$, the Cartesian-based parametrization has a covariance that is about 10% larger than the CRLB covariance, whereas the angle-based parametrization has a covariance that exceeds the CRLB covariance by only about 2%. In other words, for $N = 5$ the MLE using the Cartesian-based parametrization is still has close to the smallest possible covariance even though the CRLB might suggest otherwise.

(We also note that as $N \rightarrow \infty$, it has been shown previously¹⁵ that the covariance of the MLE approaches the CRLB for *any* parametrization. This result is not particularly useful for the small values of N that typically are used in realistic scenarios).

2.3 Accuracy maps

In Section 2.2 we studied one particular example of sensor/target configuration. Given that we know the MLE is nearly efficient, we now wish to better understand what effect the target and sensor geometry has on the accuracy that can be achieved by the MLE.

In order to produce meaningful results in a graphical form, we consider our measure of accuracy to be the trace of the covariance matrix returned by the MLE. Since the trace of a matrix is the sum of its eigenvalues, this boils the “size” of the covariance matrix down to one number. Larger traces mean that the eigenvalues are larger which means a larger uncertainty in the state estimate.

For a given sensor configuration, we consider a set of targets starting at points (x_0, y_0) which lie on a regular grid and traveling with constant velocity motion towards the origin.

For each target, we then construct a set of bearing-only measurements on that target, and use the MLE to estimate the state and covariance of the state estimate. We can then reduce this to a single number (the trace of the covariance) for each target starting location in the regular grid.

Figure 5 shows the target starting locations and velocity vectors for a three-sensor configuration, and also the trace of the covariance for each target. The accuracy falls off primarily as a function of distance, and does not appear to have a strong directional dependence.

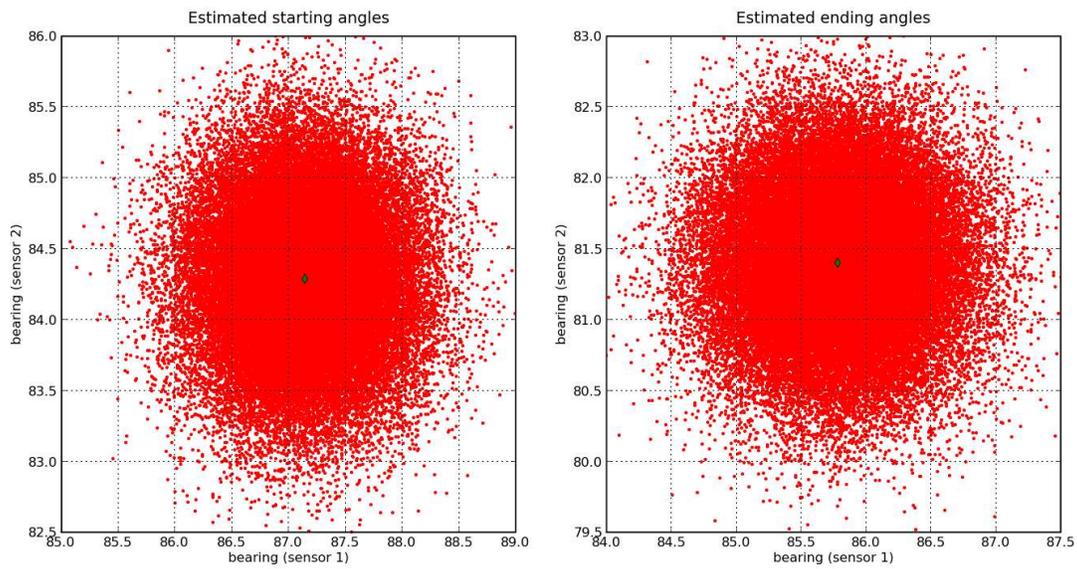


Figure 3. MLE solution for Monte Carlo simulation with $n = 50,000$ for the angle-based parametrization. Truth is indicated by the diamond, and the mean of the 50,000 runs is indicated by the '+' symbol.

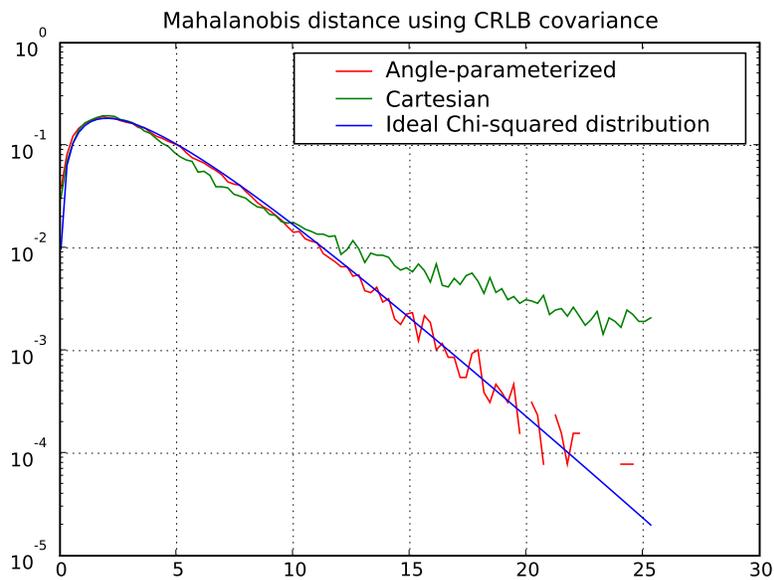


Figure 4. Estimated PDF of Mahalanobis distances for MLE solution of 50,000 Monte Carlo simulations. The angle-based parametrization matches the ideal χ^2 distribution fairly well, while the Cartesian parametrization shows significant deviation from it.

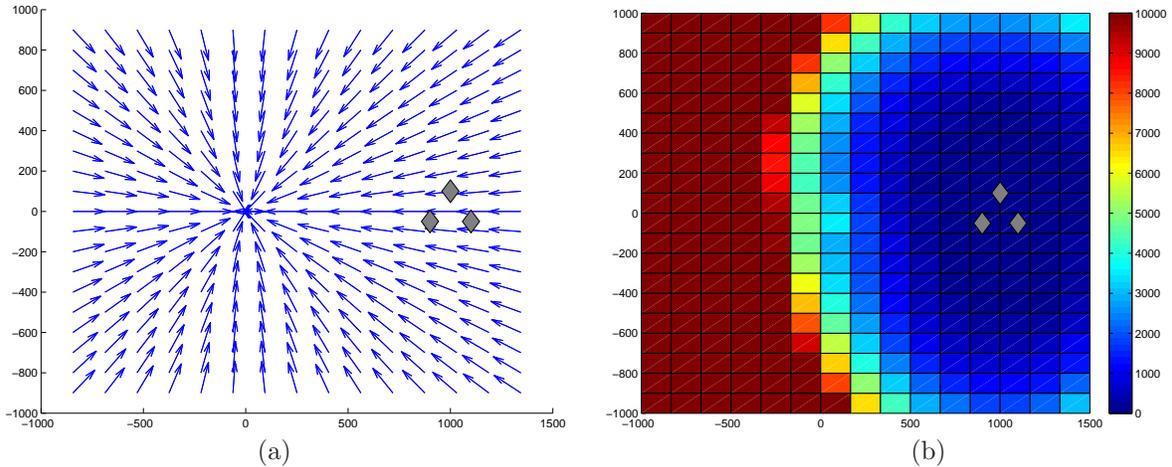


Figure 5. Trace covariance map for target direction to the left 1000 from the sensors. (a) shows target direction. (b) shows the trace covariance map. The accuracy falls off as a function of distance from the sensors.

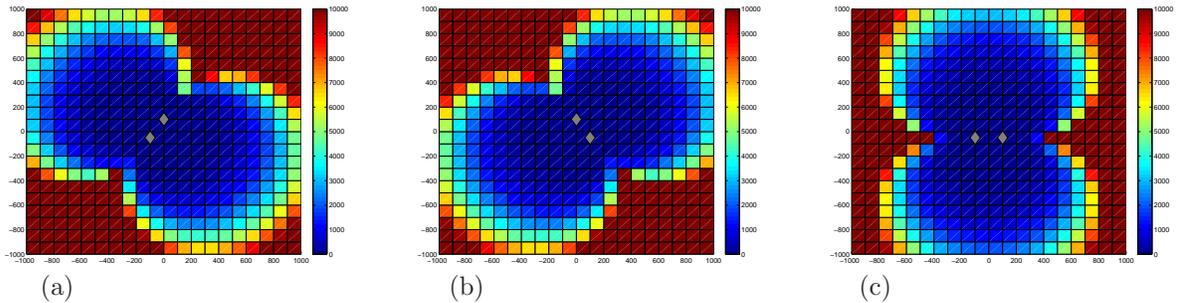


Figure 6. Trace covariance maps for three combinations for two of the three sensors. The target can be seen clearly in certain directions from the sensors, but it cannot be seen as clearly when it forms a line with the sensors.

Figure 6 shows the trace of the covariance for three different orientations of a two-sensor configuration. In this case, there is a strong directional dependence, because targets which are in line with the two sensors are not as observable as targets which are out of line with the two sensors. (Note that Figure 2 falls into the “in line” category.)

3. DISAMBIGUATION OF GHOST HYPOTHESES

One issue that often arises in multi-sensor multi-target bearing only tracking is commonly known as *ghosting*. Figure 7 shows an example of such problem in a two-sensor two-target example. In effect, ghosting arises from multiple target configurations giving rise to the same set of measurement reports. It does not only arise from a particular configuration of targets and sensors, but is instead a general phenomena. Ghosting gives rise to multiple possible assignment solutions for the given data and leads to confusion in the air picture.

Since ghost targets will in general have different kinematics than the true targets, one can use kinematic estimates of the target states to disambiguate true from ghost targets. One way to keep track of this will be to maintain the tracks that contain the ghosts as a separate hypothesis. Each hypothesis contains a set of measurement to target assignments. For two sensors and n targets, there exists a total of n^2 possible target positions. There are then $n!$ hypotheses, or possible groups of targets with assigned sensor measurements. For m sensors and n targets, there exist $(n!)^m$ hypotheses.

We can use the methods from multiple hypotheses tracking. Each hypothesis will have a score consisting of the sum of the likelihood scores for each target. As time goes on, the hypotheses that contain the ghosts

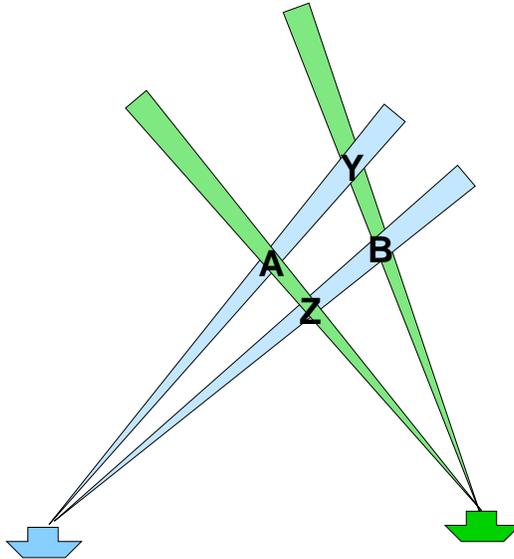


Figure 7. An example of ghosting. We have two configurations of targets that give rise to the *same* set of angle only measurements. In particular, targets at positions A and B will give rise to the same set of measurements as targets at positions Y and Z. Even worse, this is not an issue specific to this configuration. Any smoothly varying set of measurements can be generated by varying either A/B or X/Y, but note the kinematics of A/B and X/Y can be substantially different.

will have worse likelihood scores than those with the true targets, because the ghosts do not follow the expected target motion model. At that point, those hypotheses with ghosts will be discarded, and we will have maintained tracks on the true targets. The existing body of research on multiple hypothesis tracking provides algorithms for maintaining good hypotheses and removing unlikely ones.

The hypotheses in the MHT/MFA solver can be compared to discern whether the most probable hypothesis has an appreciable difference in probability over the other possible hypotheses. In effect, a “winning” hypothesis which is close in score to the other “losing” hypotheses is ambiguous since, under slightly different conditions, one of the other hypotheses might have been the “winner”. Such ambiguity is easy to detect when one uses a MHT/MFA type assignment solver since a number of hypotheses are retained, rather than being pruned away as in an SHT.

For these calculations we now use *two* targets of interest. As before, the targets approach the asset at the origin from many directions. We assume the two targets are flying in parallel, with various distances from each other*. In this section, we particularly focus on searching for the ambiguity of ghost hypotheses as compared to the truth hypothesis.

Each of the following figures represents the ambiguity between the two parallel targets at each point of the map. The ambiguity is determined by how well the true hypothesis compares with the next best hypothesis. Each hypothesis is the set of measurement associations that generate two targets.

The figures show the relative difference in the two best hypotheses scores. With three sensors and two targets, there should be eight possible hypotheses, though in actuality, there are only four because of symmetry. Of the four unique hypotheses, we choose the true hypothesis and the hypothesis with the next lowest score, though in certain cases due to random perturbations, the hypothesis with the lowest score is not necessarily the true hypothesis. The ratio tells us how much information the score will tell us about which hypothesis is true and which is the ghost. We take as our metric the ratio between the true hypothesis, and the hypothesis with the next best score. We define this ratio as “ambiguity”.

*There are many other target formations one might consider. We choose parallel flight since we believe that is one of the more difficult configurations.

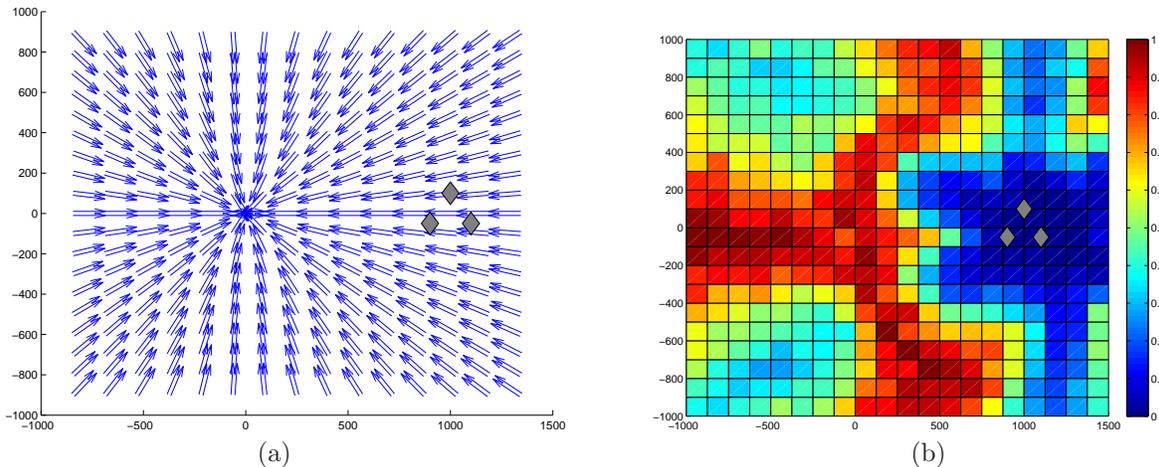


Figure 8. Ambiguity map for the target direction to the left 1000 from the sensors. (a) shows target direction. (b) shows the ambiguity map. Counter-intuitively, there is a region in which the ambiguity is substantially worse closer to the sensors than it is further away. The reason may be that the parallel motion of the targets orthogonal to the sensor direction allows the ghosts to appear to travel in constant velocity similar to the true targets.

For each square, the complete initiation method of Section 2 is computed for each ghost and each true target. The targets are initialized with measurements generated by the two true targets and the sensors. The three sensors are arranged in an equilateral triangle. They use bearing only measurements with measurement noise 0.01 rad. The targets travel with constant velocity of 50/s (dimensionless space units) towards the point of interest. Five measurements are used, each 1s apart.

Figure 8 shows a particular target and sensor configurations, and the ambiguity map computed from this configuration. The figure highlights one of the many non-intuitive properties of angle-only tracking; in this case, some target locations which are *farther away* from the sensors have lower ambiguity than targets which are closer to the sensors.

4. NETWORK ARCHITECTURE

Coordination of multiple distributed sensors through the use of a communications network can substantially improve track state, feature, and attribute estimates. However, in order to fully utilize this rich collection of data, we require advanced algorithms for data fusion. We have developed an architectural concept in which high-dimensional measurement-to-track and low dimensional measurement-to-track tracking can simultaneously coexist and interact. Figure 9 illustrates the general framework, which we also describe below.

Three-dimensional sensors and their associated composite trackers (CT) communicate with one another over the HDT (“High Dimensional Tracking”) network. Track initiation is coordinated among the trackers over this logical network, and associations between measurements and existing tracks are also communicated over this network. Such association messages are called Associated Measurement Reports or AMR’s.

At the same time, low-dimensional sensors send all of their measurements to a corresponding low-dimensional tracker (LDT) which maintains 1D or 2D tracks in e.g. bearing or bearing-elevation space. Associations between low-dimensional measurements and low-dimensional tracks are called LDAMR’s, and these are reported by the LDT over the LDT logical network.

The Low-Dimensional Initiation and Correlation (LDIC) component is a gatekeeper between the LDT and HDT networks. It maintains the same high-dimensional track database that all of the CT’s maintain, but it also attempts to initiate high-dimensional tracks from multiple low-dimensional measurements, and to correlate low-dimensional tracks to high-dimensional tracks. It therefore consumes LDAMR’s on the LDT logical network, but also produces and consumes track initiation reports and AMR’s on the HDT network.

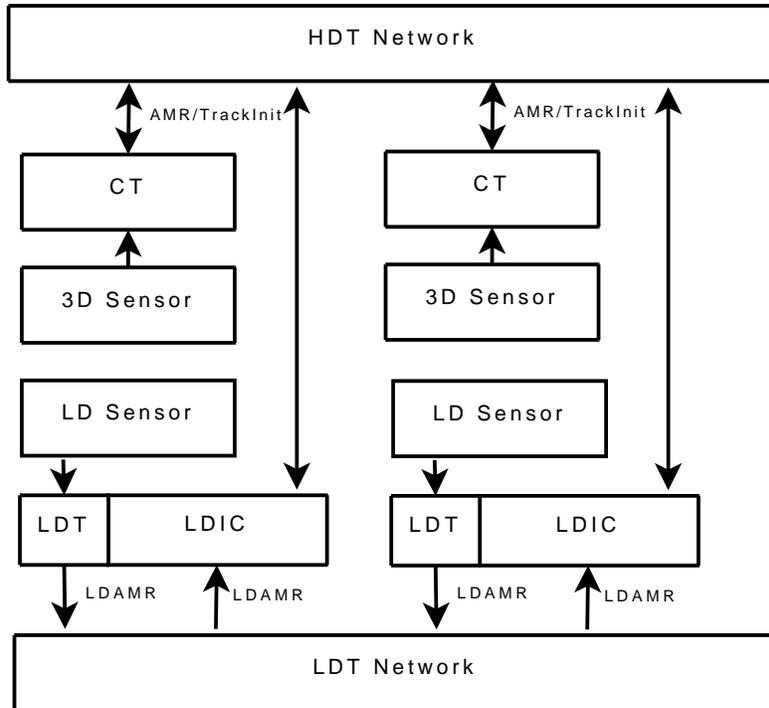


Figure 9. Network hierarchy for mixed high-dimensional and low-dimensional tracking networks.

The network architecture is therefore a hierarchical composition of two networks, with a initiation and correlation component which acts as a gatekeeper between the two logical networks.

5. CONCLUSIONS

In this paper we have demonstrated that the maximum likelihood estimator is nearly efficient for a model bearing-only track initiation problem, even though a naive analysis using the CRLB would suggest otherwise. We also showed the geometric relationship between the accuracy that can be achieved by this estimator and the configuration of the sensors.

Furthermore, we showed that in multi-target scenarios, ghosting can occur, and the ambiguity in the ghosting problem can be resolved using kinematic considerations. Similarly to the accuracy, the ambiguity depends on the geometric arrangement of targets and sensors; some combinations produce more ambiguity than others in non-intuitive ways.

Finally, we demonstrated a prototype network architecture for integrating low-dimensional and high-dimensional sensors onto a hierarchy of networks.

ACKNOWLEDGMENTS

The authors would like to acknowledge the SIAP JPO SBIR contracts W9113M-06-C-0196, W9113M-08-C-0080, and W9113M-09-C-0172 for their past and continuing support of this work. We also want thank Aubrey Poore, Sean Roberts, Ben Slocumb, Nate Knight, and James Thornbure for their guidance and support of this work.

REFERENCES

- [1] Vallado, D., [*Fundamentals of Astrodynamics and Applications*], McGraw-Hill (1997).

- [2] Blair, W., “Target tracking concepts,” Course material, Georgia Tech Research Institute, Georgia Institute of Technology (May 2007).
- [3] Lin, F., Kirubarajan, T., and Bar-Shalom, Y., “3-d track initiation in clutter using 2-d radar measurements,” *IEEE Transactions on Aerospace and Electronic Systems* **38**, 1434–1441 (October 2002).
- [4] Blackman, S. and Popoli, R., [*Design and Analysis of Modern Tracking Systems*], Artech House, Norwood, MA (1999).
- [5] Bar-Shalom, Y., Li, X.-R., and Kirubarajan, T., [*Estimation with Applications to Tracking and Navigation: Theory, Algorithms, and Software*], J. Wiley (2001).
- [6] Spingarn, K., “Passive position location estimation using the extended kalman filter,” *IEEE Transactions on Aerospace and Electronic Systems* **AES-23**, 558–567 (July 1987).
- [7] Brehard, T. and Cadre, J. L., “Closed-form posteriori cramer-rao bounds for bearing-only tracking,” *IEEE Transactions on Aerospace and Electronic Systems* **42**, 1198–1233 (October 2006).
- [8] Hammel, S. and Aidala, V., “Observability requirements for three-dimensional tracking via angle measurements,” *IEEE Transactions on Aerospace and Electronic Systems* **AES-21**, 200–207 (March 1985).
- [9] Poore, Jr., A. B., “Method and system for tracking multiple regional objects by multi-dimensional relaxation.” US Patent Number 5,537,119 (issued 16 July 1996). (Assignee: Colorado State University Research Foundation, Fort Collins, CO).
- [10] Poore, A. B. and Robertson III, A. J., “A new class of Lagrangian relaxation based algorithms for a class of multidimensional assignment problems,” *Computational Optimization and Applications* **8**(2), 129–150 (1997).
- [11] Chong, C.-Y., Mori, S., and Chang, K.-C., “Distributed multitarget multisensor tracking,” in [*Multitarget-Multisensor Tracking: Advanced Applications*], Bar-Shalom, Y., ed., Artech House, Norwood, MA (1990).
- [12] Poore, A. B., Klusman, M. E., Paffenroth, R., Roberts, S., Novoselov, R., Danford, S., and Chan, S., “Network-centric multiple frame association for distributed multiple target tracking,” Air Force SBIR Phase II Final Report Contract F33615-03-C-1434, Numerica Corporation (2006).
- [13] Lu, S., Poore, A. B., and Suchomel, B., “Network MFA tracking architectures,” in [*Proceedings of the SPIE, Signal and Data Processing of Small Targets*], **4473**, 447–457 (2001).
- [14] Novoselov, R., Gadaleta, S., and Poore, A., “Network-centric MFA tracking architecture based on soft-level data association,” in [*Proceedings of the SPIE, Signal and Data Processing of Small Targets*], **5913** (2005).
- [15] DeGroot, M. and Schervish, M., [*Probability and Statistics*], Addison Wesley, third ed. (2002).