Multitarget Detection and Tracking using Multi-Sensor Passive Acoustic Data

Christopher M. Kreucher[1], Benjamin Shapo[1], and Roy Bethel[2]
[1]Integrity Applications Incorporated, 900 Victors Way, Suite 220, Ann Arbor, MI 48108, 734-997-7436, x16, x14
[2]The MITRE Corp, 7515 Colshire Drive, McLean, VA, 22102, 703-983-5455
ckreuche@umich.edu, bshapo@integrity-apps.com, rbethel@mitre.org

Abstract—This paper describes a Bayesian approach to detecting and tracking multiple moving targets using acoustic data from multiple passive arrays. Traditional undersea acoustic systems develop tracks at the single array level, requiring track association between nodes with nonlinear projections from measurement space to target space. In contrast, our nonlinear filtering approach fuses data at the measurement level and operates directly in the target state space. As such, this approach directly addresses both the nonlinear sensor to target state coupling as well as the ambiguities caused by bearings-only nature of the passive regime. In particular, our method better addresses these challenges by combining high-fidelity physics-based sensor statistical modeling, an innovative nonlinear Bayesian filter, and a unique method of handing the computational implementation. [12]

TABLE OF CONTENTS

[1] 978-1-4244-2622-5/09/$25.00 ©2009 IEEE
[2] IEEEAC paper #1636, Version 1, Updated January 5, 2009

1. INTRODUCTION

This paper describes a principled Bayesian approach to detecting and tracking multiple moving targets using acoustic data from multiple passive arrays. Traditional undersea acoustic systems focus on signal processing and track development at the single array level, resulting in reliable but very intensive methods for finding targets that are located in densely populated shipping lanes. This process requires track association between nodes, including nonlinear projections from measurement space to target space. In contrast, our approach fuses data from multiple nodes at the measurement (rather than track) level and operates directly in the target state space (rather than measurement state space). There exist a number of challenges in the problem. First, sensor data couples to the unknown target states in a non-linear fashion. Second, the multi-array data is plagued by ambiguities. There are two main ambiguities present in the problem: a left-right ambiguity at each sensor node inherent to any linear array, and a multi-target ambiguity caused by additional beam intersections due to the multiple targets and multiple sensors. In this paper, we describe a novel approach to combat these challenges by combining high-fidelity physics-based sensor statistical modeling, an innovative nonlinear Bayesian filter, and a unique method of handing the computational implementation.

Nonlinear filtering approaches to multitarget tracking have been studied extensively in the literature. The "PDF Tracker" work by Bethel [2-3] gives a Bayesian nonlinear filtering approach and provides a strong theoretical basis for its viability. The seminal "likelihood ratio tracker" work of Stone [4] also describes a Bayesian multitarget tracking method and is implemented using a discrete grid. The "JMPD" approach of Kastella [5, 6] advocates estimating the joint multitarget state as a single hybrid probability density – i.e., constructing a single probability density over both the number of targets and the states of each. Orton and Fitzgerald [9] and Doucet [10] also study the multitarget problem from the Bayesian perspective. Recently, Streit [8] has described a multitarget filter that utilizes Poisson Point Process parameterizations. The work of Bethel & Shapo [1, 7] shows how these ideas can be adapted to single node passive acoustic tracking using real sensor data.

This work differs from earlier works as it uses the Bayes-optimal nonlinear filtering approach to perform multitarget detection and tracking in a 2D (X/Y) state space ("target

space") using simultaneous measurements from multiple passive arrays.

In this environment, there are several challenges that require the nonlinear filter's generality. First, sensor measurements couple nonlinearly to the desired target state (i.e., bearing measurements provided by passive acoustic sensors are related in a nonlinear manner to target XY position). Second, using multiple bearings only sensors in a multitarget environment leads to ambiguous beam intersections and therefore a multiply peaked conditional density. Additionally, this environment provides several implementational challenges. First, due to the physics, sensor resolution in X/Y is a function of target to sensor range. This makes numerical density estimation difficult. Second, the tactical nature of the problem calls for a large surveillance region requiring accurate density estimation over a large state space.

Our proposed approach has a number of benefits over conventional methods. Most importantly, we do not employ separate filtering at each sensor node. Rather, a single Bayesian multitarget filter operates, combining all of the sensor measurements. This eliminates the need for track fusion, or association between tracks at different nodes. A cost of this approach is the bandwidth required to transfer measurements to a central fusion point. In addition, the approach fuses data at the measurement level, which avoids the SNR losses involved in thresholding operations. These features together allow for improved detection and tracking of low SNR targets.

This approach requires increased computation as compared to conventional linear approaches. In contrast to linear approaches, the nonlinear approach estimates an entire non-parametric probability density (although it is locally truncated) whereas a linear (Kalman) approach instead needs only to estimate the mean and covariance of the density. Both methods scale linearly in target number when the targets are well separated, and superlinearly when targets are close together. Other works [13] study the computational cost for nonlinear filtering in similar settings. In the experiments presented here, we track a few targets with one-second updates in real time using un-optimized MatLab code on a standard Desktop PC.

The rest of this paper proceeds as follows. Section 2 develops a high-fidelity statistical model of sensor output, which is a function of array geometry, location, and spectral characteristics. The model uses element-level physics to generate a description of the final beamformed broadband data that serves as input to the tracker.

Section 3 reviews the Bayes optimal single target detection and tracking paradigm. In this nonlinear approach, we do not assume a parametric form for the probability density, allowing us to directly incorporate sensor bearing measurements through the high fidelity sensor model. In this track-before-detect (actually fuse-before-detect) paradigm, both the probability of target presence and probability of target state conditioned on target presence are estimated from the measurements.

Section 4 describes the numerical approach, which is based on a discrete grid approximation and an efficient alternating direction implicit Fokker-Plank solver. This numerical approach is different than particle filtering [6] which is a stochastic sampling based method of representing the density. The discrete grid approach has a number of benefits in the current setting, most importantly as it gives non-zero probability mass to all locations in the surveillance region. The main deficiency of the fixed grid approach is wasted computations on regions of the state space with very low mass, leading to potentially sharply increased computational demands over the particle filtering approach, particularly in high dimensionality state spaces.

Finally, Sections 5, 6, and 7 show how this approach is extended to the multitarget case. The multitarget filter builds on the Bayes optimal single target detection and tracking approach and employs a unique numerical approach based on multiple spatially overlapping discrete grids and moving tracker grids. Furthermore, we describe one of the interesting features of this domain: ambiguous targets.

2. PASSIVE ACOUSTIC MODELING

To test the multi-node tracker under the most realistic conditions possible, we have generated high-fidelity broadband simulations to model input sensor data. The simulation capability is robust and flexible in terms of input parameters, allowing arbitrary node array geometry (not necessarily linear), non-uniform spectral characteristics (e.g., to model dispersion and/or hydrophone frequency response), and a parameterized range of target kinematics.

In terms of signal physics, the simulator follows the frequency-domain plane-wave approach detailed in [11]. In this model, targets emit plane waves that propagate in the acoustic medium and reach the array sensors. Array geometry induces propagation delays that cause the received plane waves to have different phase characteristics at each hydrophone in the array. Exploiting this phenomenon leads to a beamformer approach to array processing, in which bearing-time displays are typically used to show target motion over time.

Fundamental to the passive sonar paradigm (and the widely used bearing-time display) is the frequency-bearing surface that results from processing an FFT-block of sensor data with a beamformer. Our simulator synthesizes this quantity as a function of target position and array location.

An example of the simulator's processing capabilities begins in Figure 1.

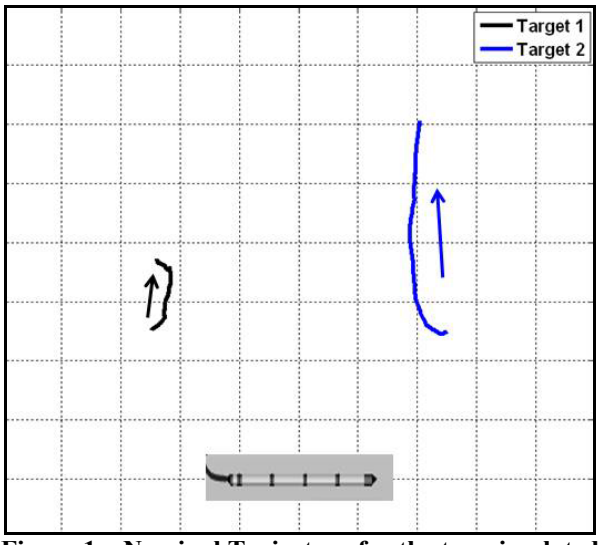

Figure 1 – Nominal Trajectory for the two simulated targets discussed below and a cartoon showing a passive node which receives acoustic energy from each.

A passive node (pictured with an array cartoon) resides along the horizontal axis, centered at (0, 0) and two targets traverse the region. Target trajectories appear as a blue and a black line. From these trajectories, the array location, and sensor geometry we expect the target energy to be spatially separated on a bearing display.

Figure 2 depicts the simulated data at the output of a conventional beamformer ("CBF"). In the frequency-domain CBF used here, phase rotations (equivalent to classical time domain delay-and-sum processing) cause plane wave inputs at each hydrophone to align in a phase-coherent way at the correct arrival angle. The equation that describes the required phase rotations comes from the array manifold vector for the array. In the case of a linear array such as the one simulated here, the manifold vector $\underline{a}$ takes on a simple form. For an array with N sensors, the nth component of the manifold vector is simply

$$a_n = e^{-j\underline{k}\cdot p_n}, \qquad (1)$$

where p_n is the nth sensor location) and the wavenumber $\underline{k} = 2\pi\hat{\underline{\theta}}/\lambda$. In contrast, plane wave signals do not add coherently in directions other than the true arrival angle. This process repeats for every discrete frequency bin in the simulation. In the left hand panel of Figure 2, bearing is on the horizontal axis (in discrete beams) and frequency (in discrete bins) appears on the vertical. This orientation gives rise to the commonly used nomenclature "FRAZ" – short for "FRequency-AZimuth." Bearing extends from zero degrees at the far-left side of the image to 180° (π) at the right-hand end. Targets (physically residing at a single bearing) emit broadband energy across all frequencies, and the results are the vertical funnel-shaped responses indicated for both targets (in blue and red, respectively). The increasing spatial extent of the target energy over bearing as frequency decreases results from the poor spatial discrimination at the large acoustic wavelengths that comprise low-frequencies. In this display, Target 1 is fairly difficult to see because of its low SNR. Note that near the endfire regions (0° and 180°), there is significant distortion.

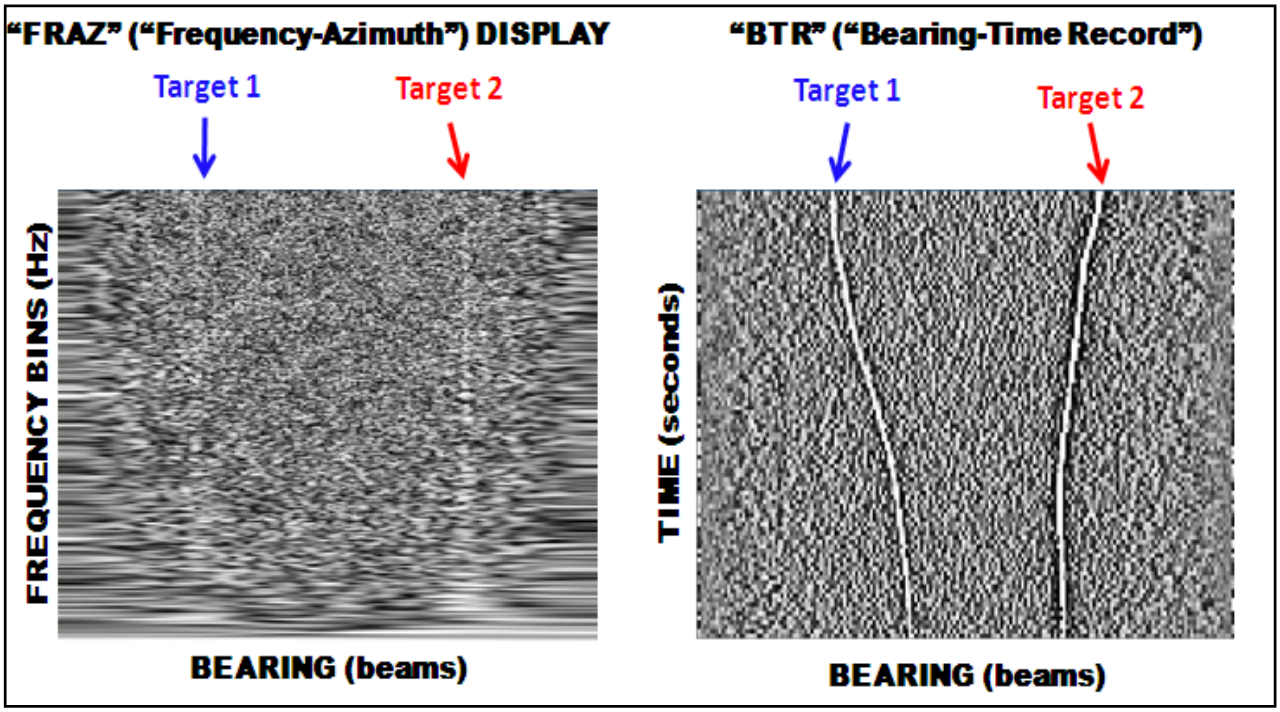

Figure 2 – Left: A single time snapshot of the Frequency-Azimuth (FRAZ) display for two targets; Right: The associated Bearing-Time Record (BTR).

This phenomenon results from the intrinsically poor spatial resolution of linear arrays near endfire (due to lack of physical aperture in these directions). The effect would be less pronounced on a cosine(bearing) display, but here we keep the horizontal axis in true bearing space for more flexibility with nonlinear array geometries and the 2D multiple array scenarios that are the focus of this work.

Omnidirectional hydrophones actually provide full 360° (2π) coverage of the area surrounding the node. However, the 0°-180° region and the 180°-360° are indistinguishable from the perspective of linear sensors like the one modeled here. Therefore, we have eliminated that portion of the display since it adds no new information. This phenomenon is different from the one discussed in Section 6, where ambiguities arise because true passive target bearings from multiple targets intersect in multiple places.

Because FRAZ images do not provide integration gain over frequency or an easy way for operators to follow contacts over time on a 2D display, bearing-time record ("BTR") displays like the one in the right hand panel of Figure 2 enjoy widespread use. A BTR is formed by integrating over all frequency bins. This results in easier target detection, and the reduced dimensionality of the process enables a display of target bearing over time. In contrast to the situation in the left panel of Figure 2, both targets are easily and clearly visible. Furthermore, this display provides time coverage absent in the FRAZ display.

In the following sections, we assume multiple passive arrays, each of which records acoustic energy. The energy is processed into BTR surfaces like this one at each array and these are used as input for our 2D detection and tracking algorithms. Therefore, our algorithm requires only the BTR data (and not the data from the individual array elements). The BTR can be characterized statistically, in terms of expected signal given $0, 1, \text{or } T$ targets by modeling the energy in the FRAZ and appealing to the statics that appear when summing this energy.

3. SINGLE TARGET DETECTION AND TRACKING THEORY

This section explains the mathematics of the Bayesian approach to detecting and tracking a single moving target. Section 4 describes the numerical implementation. This methodology is expanded to the multiple target situation in Section 5.

Notation

The notation that will be used throughout this paper is as follows. The state of a single target at time k will be denoted x^k. In this paper, our state model refers to the targets' 2D position and velocity as $x^k = [x\ \dot{x}\ y\ \dot{y}]$, although the approach generalizes to higher order state models. Furthermore, let H_o^k denote the hypothesis that no target is present at time k, and let H_1^k denote the hypothesis that a target is present. Finally, we use the following notation to describe the measurements: z_i^k denotes the measurements taken by sensor i at time k (e.g., a collection of energy arriving at different bearings taken from the BTR); z^k denotes the measurements taken by all sensors at time k (i.e., $z^k = \{z_1^k, \dots, z_S^k\}$); and Z^k denotes the collection of all measurements taken by all sensors up to and including time k (i.e., $Z^k = \{z^0, z^1, \cdots, z^k\}$).

Fundamentally, we wish to estimate the joint probability that a target is present (i.e., H_1^k is true) and its location is x^k given the measurements. Mathematically, the means we wish to estimate the hybrid continuous-discrete density

$$p(x^k, H_1^k | Z^k) \tag{2}$$

for all x^k. Notice that this quantity can be expressed as

$$p(x^k, H_1^k | Z^k) = p(H_1^k | Z^k) p(x^k | H_1^k, Z^k), \tag{3}$$

i.e., as the product of the *target present probability* $p(H_1^k | Z^k)$ and the *target state probability* $p(x^k | H_1^k, Z^k)$. Both conceptually and implementationally, we treat the problem as separate (but coupled) tasks of estimating the target present probability ("detection") and the estimating target state probability ("tracking"). The two portions will be covered in the following subsections.

Bayes Optimal Detection

In the Bayesian approach, we (i) assume that an initial or prior estimate of the desired probabilities is present (perhaps completely uninformative), and (ii) generate a recursive formula that relates the probabilities at one time step with those at the next. This is then used at each time step to update the desired probability.

The target present $p(H_1^k | Z^k)$ and target absent $p(H_0^k | Z^k)$ detection probabilities are computed recursively using the law of total probability and Bayes' rule, yielding:

$$\begin{aligned} p(H_1^k | Z^k) &= \int p(x^k, H_1^k | Z^k) dx^k \\ &= p(H_1^k | Z^{k-1}) \times \left(\frac{p(z^k | H_0^k)}{p(z^k | Z^{k-1})} \right) \times \\ &\int \left(\frac{p(z^k | H_1^k, x^k)}{p(z^k | H_0^k)} \right) p(x^k | H_1^k, Z^{k-1}) dx^k, \end{aligned} \tag{4}$$

and:

$$p(H_0^k | Z^k) = p(H_0^k | Z^{k-1}) \times \left(\frac{p(z^k | H_0^k)}{p(z^k | Z^{k-1})} \right). \tag{5}$$

These equations express the current target present and

absent hypothesis probabilities in terms of the target present and target state probabilities predicted from the previous time step (e.g., $p(H_1^k|Z^{k-1})$ and $p(x^k|H_1^k, Z^{k-1})$) and the conditional likelihood of the incoming measurements (i.e., $p(z^k|H_1^k, x^k)/p(z^k|H_0^k)$).

The predicted densities are formed in the standard Bayesian manner: a model on target appearance and disappearance generates predicted target present and absent probabilities $p(H_1^k|Z^{k-1})$ and $p(H_0^k|Z^{k-1})$, and a target kinematic model predicts the target state density $p(x^k|H_1^k, Z^{k-1})$. We discuss these processes in more detail below.

The fundamental data dependent quantity needed to construct the detection update is the likelihood ratio $p(z^k|x^k, H_1^k)/p(z^k|H_0^k)$, which is a function of the sensor measurement statistics and the measured data. Notice that the normalizing constant $p(z^k|H_0^k)/p(z^k|Z^{k-1})$ present in both updates need not actually be computed directly since we can determine its value using the relation

$$p(H_1^k|Z^k) + p(H_0^k|Z^k) = 1. \tag{6}$$

Bayes Optimal Tracking

The target state probability is computed recursively in a manner similar to the above:

$$p(x^k|H_1^k, Z^k) = p(x^k|H_1^k, Z^{k-1}) \left(\frac{p(z^k|x^k, H_1^k)}{p(z^k|H_0^k)} \right) \times \left(\frac{p(z^k|H_0^k)}{p(z^k|H_1^k, Z^{k-1})} \right), \tag{7}$$

where the constant term $p(z^k|H_0^k)/p(z^k|H_1^k, Z^{k-1})$ does not need to be computed since the probability density integrates to 1. Again, the fundamental data dependent quantity needed to construct this update is the likelihood ratio. Also required again is the prediction density $p(x^k|H_1^k, Z^{k-1})$.

Therefore, in principle the update for each possible x^k proceeds by predicting its probability forward in time (according to a temporal target kinematic model) and then updating using the likelihood ratio $p(z^k|x^k, H_1^k)/p(z^k|H_0^k)$. Obviously, since x^k is drawn from the continuum, some discrete representation or parameterization must be employed for computation. We defer this discussion to Section 4. The following subsections describe the temporal prediction and measurement update steps in detail.

Temporal update

Both the detection and tracking update stages perform temporal evolution on the relevant probability distributions. For the target present and absent probabilities $p(H_1^k|Z^{k-1})$ and $p(H_0^k|Z^{k-1})$, we use a simple mixing matrix approach, where the probabilities at the previous time are used to predict the current probabilities, i.e.,

$$\begin{bmatrix} p(H_0^k|Z^{k-1}) \\ p(H_1^k|Z^{k-1}) \end{bmatrix} = \begin{bmatrix} 1-q & q \\ q & 1-q \end{bmatrix} \begin{bmatrix} p(H_0^{k-1}|Z^{k-1}) \\ p(H_1^{k-1}|Z^{k-1}) \end{bmatrix}. \tag{8}$$

This corresponds to a fixed target arrival/removal probability. Our approach allows more complicated models with little impact on computational requirements.

Similarly, a model on target kinematics $p(x^k|H_1^k, x^{k-1})$ performs the temporal update of the target state probability. This relationship is expressed in discrete time as:

$$p(x^k|H_1^k, Z^{k-1}) = \int p(x^{k-1}|H_1^k, Z^{k-1}) p(x^k|H_1^k, x^{k-1}) dx^{k-1}. \tag{9}$$

This step is the most computationally challenging of the algorithm and requires a numerical approach able to efficiently compute the required probabilities. The approach we've selected is discussed in more detail in the implementation section below.

The Multi-Sensor Likelihood ratio

The fundamental data dependent quantity needed to update all densities is the likelihood ratio. In this section, we describe the computation in our model problem.

For convenience, we specialize here to two bearings-only sensors interrogating a region. The method extends to more sensors in a similar manner. We further assume each sensor makes measurements in bearing "cells" (referred to hereafter as beams) with some known distribution when the target is present $p_1(z)$ and some other distribution when the target is absent $p_0(z)$. For notational brevity, we assume $p_1(z)$ and $p_0(z)$ are sensor independent, although this is also not required.

The measurements z^k are energy values corresponding to discrete beams θ. Assuming independence among the beams, the likelihood ratio can be rewritten as

$$\begin{aligned} \frac{p(z^k|x^k, H_1^k)}{p(z^k|H_0^k)} &= \frac{\prod_{\theta=1}^{N_\theta} p(z^{\theta,k}|x^k, H_1^k)}{\prod_{\theta=1}^{N_\theta} p(z^{\theta,k}|H_0^k)} \\ &= \frac{p(z^{\theta_x,k}|x^k, H_1^k) \prod_{\substack{\theta=1 \\ \theta \neq \theta_x}}^{N_\theta} p(z^{\theta,k}|H_0^k)}{p(z^{\theta_x,k}|H_0^k) \prod_{\substack{\theta=1 \\ \theta \neq \theta_x}}^{N_\theta} p(z^{\theta,k}|H_0^k)} \\ &= \frac{p(z^{\theta_x,k}|x^k, H_1^k)}{p(z^{\theta_x,k}|H_0^k)}, \end{aligned} \tag{10}$$

where the notation $z^{\theta_x,k}$ is used to indicate the measurement in the bearing cell θ corresponding to x^k. The likelihood of the two-node measurement set $z^{\theta_x,k} = \left\{z_1^{\theta_{x1},k}, z_2^{\theta_{x2},k}\right\}$ conditioned on the target present hypothesis (the numerator of the likelihood ratio) is simply the product of the two target present probabilities since target presence is more likely when the data follows the target present distribution in both sensors. In contrast, the H_0 (target absent) hypothesis is the composite hypothesis that either (a) both measurements are drawn from the target absent density $p_0(z)$, or (b) one measurement is drawn from the target present hypothesis and one from the target absent hypothesis. Therefore, in following with the GLRT, the denominator of the likelihood ratio is computed as the maximum over the three atomic hypotheses: p_0 from both sensors, p_0 from sensor 1 and p_1 from sensor 2, and p_1 from sensor 1 and p_0 from sensor 2.

4. Single Target Detection and Tracking Implementation

This section describes how the mathematics of Section 3 is implemented. The approach we take here utilizes a discrete grid representation of the posterior. In contrast, other work [5] has employed a particle approximation to the relevant posteriors. As discussed earlier, there are tradeoffs involved in the two approaches.

Representation of target state probability

We represent the target state probability $p(x^k|H_1^k, Z^k)$ on a 4D discrete grid (corresponding to the four dimensional state vector x^k) of $N_x \, x \, N_{\dot{x}} \, x \, N_y \, x \, N_{\dot{y}}$ cells. The spatial extent of this grid dictates the overall region where targets may be detected. The cell resolution must be carefully chosen based on the sensors and geometry to allow sufficient accuracy. Figure 3 gives an example of marginal PDFs as represented on a discrete grid.

The tradeoffs involved in grid resolution include the following: (i) Algorithm speed is roughly linear in the number of grid cells. The number of the grid cells is the product of cells in each of the spatial and each of the velocity dimensions. This fact rewards minimizing of the number of cells either by limiting the overall region upon which the PDF is approximated or making the discretization coarse. (ii) Single-target tracking performance degrades as the grid cells become coarser. In particular, since the sensor to target state mapping is nonlinear, performance breaks down substantially above certain cell resolution sizes (depending on the sensor geometry). Both of these phenomena will be illustrated later.

Temporal update of the tracking probability

The temporal evolution of the probability density on x^k can be expressed in continuous time using a partial differential equation. We wish to compute $p(x, \dot{x}, y, \dot{y}, t + \Delta t)$ from $p(x, \dot{x}, y, \dot{y}, t)$. The relation between these two densities can be expressed as

$$\begin{aligned} &p(x, \dot{x}, y, \dot{y}, t + \Delta t) \\ &= \int p(x - \Delta x, \dot{x} - \Delta \dot{x}, y - \Delta y, \dot{y} - \Delta \dot{y}, t) \\ &\times p(\Delta x, \Delta \dot{x}, \Delta y, \Delta \dot{y}, \Delta t) \, d\Delta x \, d\Delta \dot{x} \, d\Delta y \, d\Delta \dot{y}. \end{aligned} \tag{11}$$

Using (i) a second order Taylor series approximation to $p(x - \Delta x, \dot{x} - \Delta \dot{x}, y - \Delta y, \dot{y} - \Delta \dot{y}, t)$, (ii) assuming the "nearly constant velocity" model for the transition density $p(\Delta x, \Delta \dot{x}, \Delta y, \Delta \dot{y}, \Delta t)$, and (iii) assuming small Δt, we find the Fokker-Planck equation

$$\frac{\partial p}{\partial t} = -\dot{x}\frac{\partial p}{\partial x} - \dot{y}\frac{\partial p}{\partial y} + \frac{\sigma_{\dot{x}}^2}{2}\frac{\partial^2 p}{\partial \dot{x}^2} + \frac{\sigma_{\dot{y}}^2}{2}\frac{\partial^2 p}{\partial \dot{y}^2}. \tag{12}$$

This derivation is given in more detail in the Appendix.

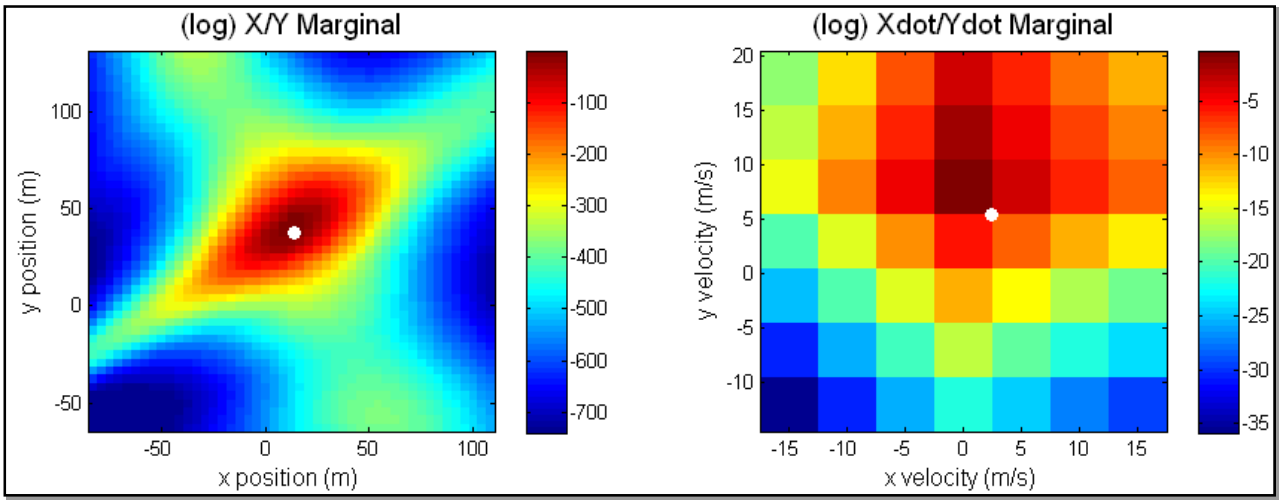

Figure 3– Grid-based representation of the target state density. Left: The 4D posterior marginalized to give the x/y probability density. Right: The $\dot{x}/\dot{y}$ marginal.

In the following discussion we use this transition model. Other models can be incorporated similarly, including models which use roadway constraints and higher order motion terms.

In the experiments we discuss later, this model is valid (i.e., the diffusivity of the target is captured statistically by the diffusivity coefficient in the model). In cases where robust modeling of the target dynamics is not possible, one typically appeals to a multiple model approach [13] which adds robustness to the filtering by allowing it to select from a number of candidate motion models.

The state probability density is represented on a discrete grid and so this differential equation must be used to update that representation. This discrete update is computed from time k to time $k+1$ using a backward Euler method as

$$\frac{p_{i,j,l,m}^{k+1} - p_{i,j,l,m}^{k}}{\Delta t} = \frac{\sigma_x^2}{2\Delta x^2}\left[p_{i,j+1,l,m}^{k+1} - 2p_{i,j,l,m}^{k+1} + p_{i,j-1,l,m}^{k+1}\right] + \cdots \qquad (13)$$

This approach has nice stability properties in both Δt and Δx. For more details, see [12].

This is illustrated graphically in Figure 4, where the left hand side represents the posterior from the k^{th} time step and the right hand side represents the prediction for the $k+1^{th}$ time step. The white dot represents the true target position. In this example, the target is moving in the positive x direction, causing the probability mass to evolve preferentially along the positive x direction.

This computation method is linear in the number of grid cells used in the discretization. Linear dependence on grid cell number is achieved using the backward Euler method above, yielding a tridiagonal system of equations which is efficiently solved using Thomas' algorithm [12]. Figure 5 illustrates this dependency from empirical tests.

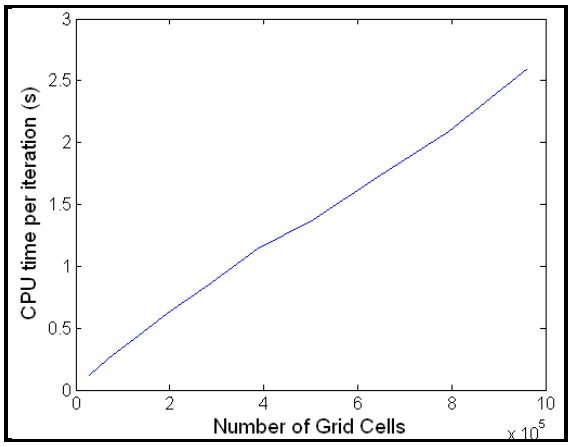

Figure 5– Computation is linear in number of grid cells.

Temporal update of the target present probability

The target present and absent probabilities are updated according to the simple mixing model described in eq. (8) corresponding physically to targets arriving and leaving at a constant rate. Since this is a discrete PMF with only two possible events (target present and target absent) it is implemented by simply storing a single floating point number corresponding to the target present probability.

Measurement update of the tracking probability

As discussed above, evaluating the measurement likelihood for each cell x^k performs the target state probability update:

$$p(x^k|H_1^k, Z^k) \propto p(x^k|H_1^k, Z^{k-1})\left(\frac{p(z^k|x^k, H_1^k)}{p(z^k|H_0^k)}\right) \qquad (14)$$

This simply requires computation of the likelihood ratio $p(z^k|x^k, H_1^k)/p(z^k|H_0^k)$ at each cell x^k. Specializing again to the two-node, bearings only situation, at each time step, each sensor node makes measurements at all (discretized) bearings (beams θ). The sensor model describes statistically the likelihood of each measurement conditioned on target

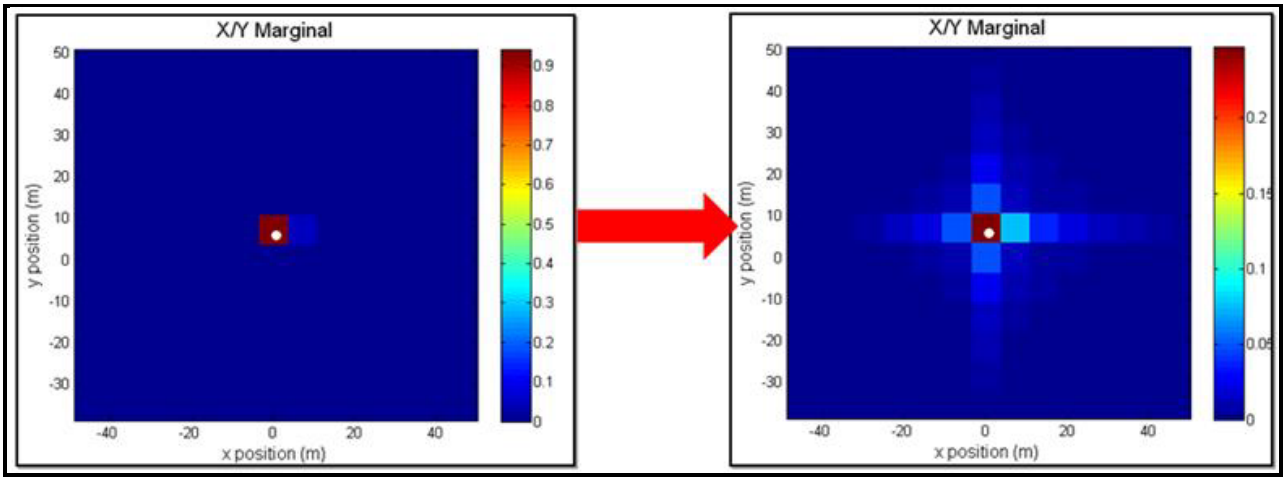

Figure 4– The temporal update of the target state probability density.

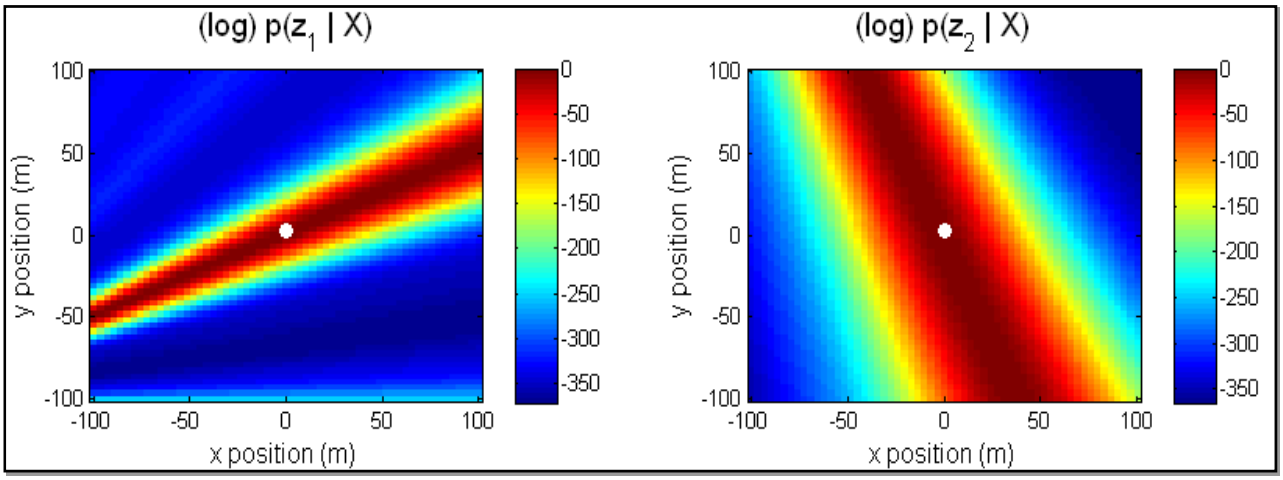

Figure 6– Single Sensor conditional likelihoods.

present, as illustrated in Figure 6. Note that different node locations yield different beam widths and spatial information.

By assumption, the joint likelihood of the measurements conditioned on target present (the numerator) is the product of the two individual likelihoods (i.e., the noise is independent from sensor to sensor):

$$p(z^k|x^k, H_1^k) = p(z_1^k|x^k, H_1^k)p(z_2^k|x^k, H_1^k), \tag{15}$$

where from the end of Section 3 we know

$$p(z_1^k|x^k, H_1^k) \propto p\left(z_1^{\theta_{x1},k}|x^k, H_1^k\right). \tag{16}$$

Here, θ_{x1} is the bearing cell (beam) into which x projects at sensor node 1. Therefore, for each grid cell in the discrete representation of x^k, we compute the H_1^k hypothesis probability as proportional to:

$$\begin{aligned} &p(z^k|x^k, H_1^k) \\ &\propto p\left(z_1^{\theta_{x1},k}|x^k, H_1^k\right)p\left(z_2^{\theta_{x2},k}|x^k, H_1^k\right). \end{aligned} \tag{17}$$

As mentioned above, the H_0^k hypothesis (the denominator of the likelihood ratio) corresponds to the composite hypothesis that either (i) neither sensor has target energy in the target bearing, or (ii) one of the two sensors (but not both) has energy in the target bearing. Therefore, the probability of the H_0^k hypothesis is evaluated by appealing to the GLRT as

$$\begin{aligned} &p(z^k|x^k, H_1^k) \\ &\propto max\begin{Bmatrix} p\left(z_1^{\theta_{x1},k}|H_0^k\right)p\left(z_2^{\theta_{x2},k}|H_0^k\right), \\ p\left(z_1^{\theta_{x1},k}|x^k, H_1^k\right)p\left(z_2^{\theta_{x2},k}|H_0^k\right), \\ p\left(z_1^{\theta_{x1},k}|H_0^k\right)p\left(z_2^{\theta_{x2},k}|x^k, H_1^k\right) \end{Bmatrix}. \end{aligned} \tag{18}$$

Finally, the measurement update for cell x^k is computed as:

$$\begin{aligned} &\frac{p(z^k|x^k, H_1^k)}{p(z^k|H_0^k)} \\ &= \frac{p\left(z_1^{\theta_{x1},k}|x^k, H_1^k\right)p\left(z_2^{\theta_{x2},k}|x^k, H_1^k\right)}{\max\begin{Bmatrix} p\left(z_1^{\theta_{x1},k}|H_0^k\right)p\left(z_2^{\theta_{x2},k}|H_0^k\right), \\ p\left(z_1^{\theta_{x1},k}|x^k, H_1^k\right)p\left(z_2^{\theta_{x2},k}|H_0^k\right), \\ p\left(z_1^{\theta_{x1},k}|H_0^k\right)p\left(z_2^{\theta_{x2},k}|x^k, H_1^k\right) \end{Bmatrix}}. \end{aligned} \tag{19}$$

Measurement update of the target present probability

As discussed in Section 3, the measurement update of the target present probability is performed as

$$\begin{aligned} &p(H_1^k|Z^k) \\ &\propto p(H_1^k|Z^{k-1}) \int \left(\frac{p(z^k|H_1^k, x^k)}{p(z^k|H_0^k)}\right) p(x^k, H_1^k|Z^{k-1})dx^k. \end{aligned} \tag{20}$$

By comparison, it is seen that the inner term in this integral is the measurement update done on the target state probability density (eq. (14)). Therefore, the update of the target present and absent probabilities simply requires summing the (non-normalized) target state probability after the measurement update is performed, followed by a normalization step which forces $p(H_1^k|Z^k) + p(H_0^k|Z^k) = 1$.

For computational purposes, it is desirable that the target state grid contain the fewest number of cells that allow robust estimation. Cells which contain zero probability mass are useless computation and should be avoided. However, since we are tracking moving targets, the locations of the grid cells needed to estimate the target state density change over time. Therefore, we use a moving grid, which constantly re-centers around the estimated state of the target every update. In practice, we only allow the grid to translate a small number of cells at each time step, which minimizes the chance that a small number of bad measurements can shift the target off the grid.

On Grid Resolution

As mentioned to earlier, the resolution (spacing) of the target detection and tracking spatial grid cells is of critical importance. In our method, we choose to allow each cell center to represent the entire cell. An alternative approach would treat each cell as containing many sub-cells and requiring a weighted update. This has a similar computational burden as making the grid cell finer, and provides a benefit in the single target tracking case. However, in the multitarget tracking case discussed later it introduces other problems. A discussion of these problems is deferred until later.

In the single target tracking case where cell centers represent the cell, the main concern is to ensure the cell spacing is fine enough. This avoids cases where the conditional likelihood has energy in large parts of a cell but does not overlap with the cell center. Figure 7 illustrates this. In this figure, cell boundaries are indicated by white lines, and cell centers by black dots. The true target location is given by the green dot. The conditional probability density is indicated by the color scale showing the intersection of the two sensor node beams. Because the grid resolution is insufficient, the cell center in the cell the target actually occupies does not correspond to the peak of the conditional as it should. This mismatch leads to the target containing cell incorrectly receiving low likelihood in the measurement update.

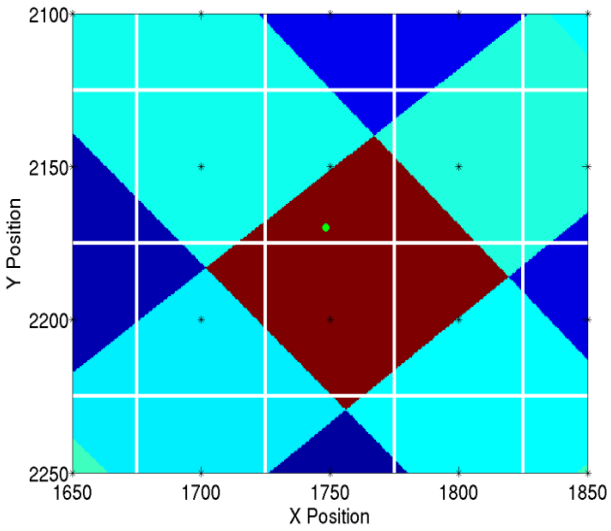

Figure 7 - Grid cells must be spaced finely enough to avoid degenerate cases where no cell center corresponds to the peak of the likelihood function.

The remedy for this issue is to make the cells more finely spaced. Also note that unless the grid cells are spaced grossly inadequately (as they are in the example given Figure 11 later), this problem does typically does not persist from time step to time step as the target is moving. The most catastrophic consequence of poor grid cell resolution is track fragmentation. In the case of overly coarse cell discretization, the track existence probability will be artificially driven lower resulting in (incorrectly) removing the track. A new track would then be instituted very shortly.

On Computational Requirements

The dominant properties that effect computations are (i) the number of grid cells and (ii) the number of targets. Section 5 discusses algorithm scaling with the number of targets. Here we focus on the single target case and note that the algorithm scales linearly with number of grid cells (which is the product $N_x \; x \; N_{\dot{x}} \; x \; N_y \; x \; N_{\dot{y}}$).

Figure 8 illustrates empirically the tradeoff between cell resolution, tracking, and algorithm run time in the single target detection and tracking case. As the grid cell resolution decreases (i.e., the number of cells used to represent the probability density increase), the tracking error decreases. It reaches an asymptote which is dictated by the sensor resolution. Furthermore, as the grid cell resolution decreases (i.e., the number of grid cells increases), the computation time increases.

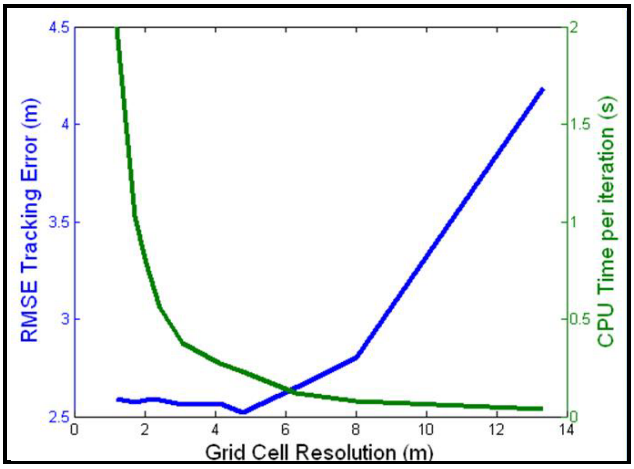

Figure 8 – For single target tracking, performance improves as grid cells become more finely spaced. This is at the cost of increased computation time.

A second way to deal with the degeneracy caused by large grid cells is to compute the cell likelihood function by weighting the likelihood of the beams into which the cell projects. In other words, instead of using the cell center to represent the entire cell, treat the cell as the continuum of points it represents. Doing this exactly is computationally prohibitive, but approximate methods such as averaging discrete points in a cell are feasible. This is a net computation savings over simply making the grid cells smaller because it avoids increasing the number of cells that require temporal update.

In the multitarget case, the problem is more complicated. Since many cells will correspond to the peak of the conditional likelihood, blindly shrinking cell size or interpolating the conditional will lead to false (double) initializations of targets. We will discuss this further in Sections 6 and 7.

5. MULTITARGET DETECTION AND TRACKING THEORY

In principle, multiple target detection and tracking requires estimation of the joint multitarget probability density

$$p(x_1^k, x_2^k, \dots, x_T^k, T^k | Z^k) = p(x_1^k, x_2^k, \dots, x_T^k | T^k, Z^k) p(T^k | Z^k), \quad (21)$$

where T^k denotes the number of targets present at time k $(0, 1, \dots)$ and $x_1^k, x_2^k, \dots, x_T^k$ are the individual states of those targets.

The state space of this joint multitarget probability density grows exponentially with the number of targets and hence precise computation grows exponentially as well. Brute-force multitarget discrete grid representations of this high dimensional posterior become intractable with more than two or three targets in a four dimensional state space.

Fortunately, high fidelity modeling of the joint coupling is only necessary when targets are close together, i.e., widely spaced targets can be treated nearly optimally by solving multiple single target detection and tracking problems. In the (unrealistic) limiting case, where all targets are well separated in measurement space, numerical estimation of the joint density grows linearly with the number of targets rather than exponentially.

A more sophisticated approach to this problem is to automatically factorize the joint density into small groups of targets which must be treated jointly and develop a computational solution which is a compromise between the exponential growth of the joint computation and the linear growth of the fully factored computation. A detailed discussion of a joint density adaptive factorization approach and the precise algorithmic details in a related environment are discussed at length in [5] and [6].

In our present implementation, we have chosen to simply treat the multiple target situation as a collection of single target detection and tracking problems. To account for the sub-optimality of this approach when targets are nearby in measurement space, we employ a data-censoring algorithm which operates when targets are close - i.e., we make the approximation

$$p(x_1^k, x_2^k, \dots, x_T^k, T^k | Z^k) \approx \prod_i p(x_i^k, H_{1(i)}^k | Z^k), \quad (22)$$

and censor some of the data from nearby trackers to prevent improper evaluation of the conditional likelihood.

Future work will extend this to high fidelity sensor modeling and joint density calculation for closely spaced targets when necessary. A sketch of this extension, analogous to the approach in [5], is as follows. First, the target state and relevant uncertainty will be estimated for each target. Then those targets that are close together will be treated in clusters. The (now joint) probability density for the cluster of targets will be temporally and measurement updated as a group. This computation is superlinear in the number of targets, but will only operate on targets in that cluster, rather than the entire target set. This method allows for careful physics-based modeling of the sensor returns when targets are close together (e.g., including the expected coherent sums of energy from nearby targets and accurately modeling sidelobe interference) and will allow for more effective handling of crossing targets and convoy movements.

6. AMBIGUOUS TARGETS

In addition to left-right ambiguity arising from linear arrays, in the present setting there are intersection ambiguities. In the multisensor, multitarget, bearings-only environment we study here there are ambiguities arising from the (persistent) intersection of bearing measurements across sensors from e.g., (sensor 1, target 1) and (sensor 2, target 2).

Figure 9 gives an example of this phenomenon, by showing the multisensory conditional likelihood surface in a high-SNR and low-SNR case, respectively. In this example, there are two passive sensors, one located at the northeast and one at the southeast of the surveillance region. There are two real targets indicated by white circles. Each passive sensor receives high energy at the bearings corresponding to the true targets as expected. This results in (correct) intersections at the true locations of the targets. However, there are also false intersections, which correspond to the mismatched beams (i.e., a beam from sensor 1, target 1 intersection with a beam from sensor 2, target 2).

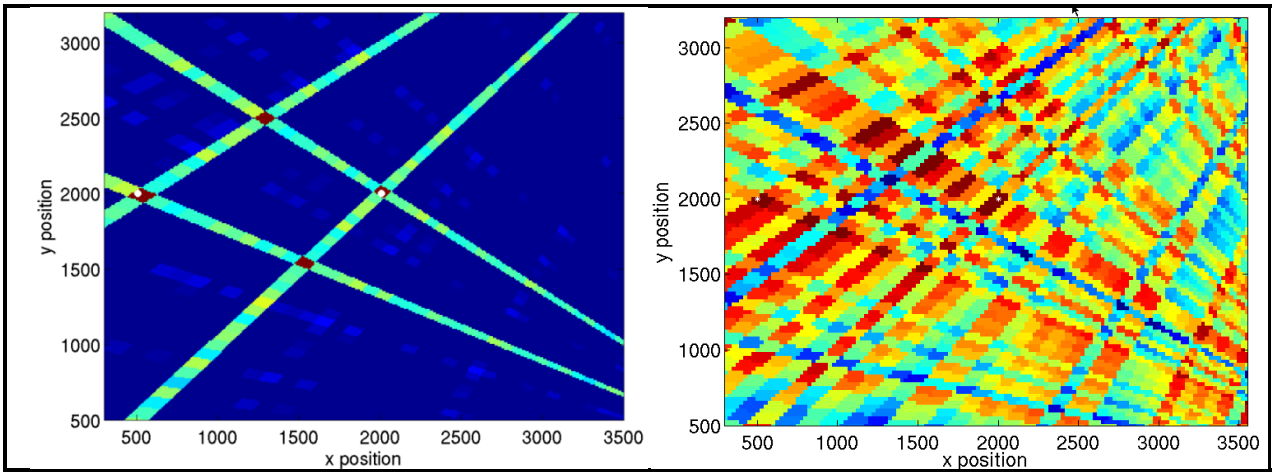

Figure 9 – Bearings intersections corresponding to true targets (white circles) and ambiguous intersections. There are two sensors, located NE and SE of the image, respectively. Each sensor has high energy returns at the bearings corresponding to the true targets, yielding intersections at the true target locations. However, there are also false intersections between, for example, (sensor 1, target 1) and (sensor 2, target 2).

Ambiguous targets often move physically for some time and for that time are indistinguishable from real targets. This short term phenomenon is not a problem with the tracker but is a fundamental issue of physics. However, over time, the ambiguous targets can be distinguished from real targets in a tracking environment as they will move in a non-physical manner. Typically this non-physical motion takes the form of a jump movement as the ambiguities approach a line of symmetry defined by the nodes.

Figure 10 provides an illustration of this behavior. For the first 500 time steps (shown at left), both the true targets (green) and the ambiguous targets (red) move in a manner that is plausible physically. However, as shown at right, after time step 500 the ambiguous targets move in a dramatically non-physical manner. This behavior is common and in this case is correlated with the ambiguity position moving through a sensor line of symmetry.

7. MULTITARGET DETECTION & TRACKING IMPLEMENTATION

As discussed earlier, we have chosen to factorize the joint density into a collection of single target densities. This approach is optimal when targets are well separated but is inappropriate as targets become close in sensor space. Future work includes an adaptive approach which will appropriately treat these closely spaced targets jointly. In the present work however, we employ an engineering approach to cope with this sub optimality which prevents measurement sharing among closely spaced trackers. These algorithmic modifications mean that the final product is not simply multiple single target detection and tracking algorithms running in parallel.

First, we partition the surveillance region into multiple overlapping static detection grids. Each detection grid is a single target detector as described above for detecting

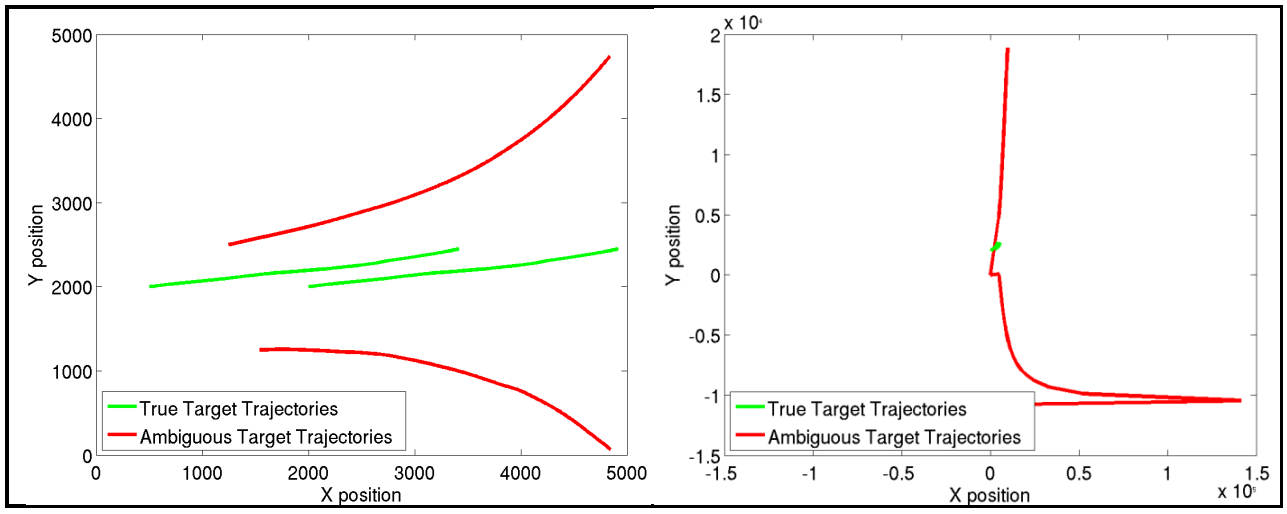

Figure 10– Left: true and ambiguous trajectories for the first 500 time steps of the simulation. Right: After time step 500, the ambiguous targets make large non-physical jumps.

targets in its sub-region. We compute the target present hypothesis independently for each detector grid. It is desirable that each detector grid is small in total spatial extent to allow detection of closely spaced targets since each individual detection grid is only capable of accurately modeling the target absent/present hypothesis when there is either a single target present or no targets present. In contrast, it is desirable that each grid is large enough (in total spatial extent) to capture multiple time steps of measurements on a moving target to allow accurate computation of the target present hypothesis.

Second, the single target detector/tracker equations (discussed above) update each detector grid. The temporal and measurement updates proceed as if there were a single target present on the grid, updating the target present hypothesis with the new data. In this manner, each detector grid performs the Bayes-optimal single target detection and tracking algorithm for its spatial region.

If the target present hypothesis associated with a detector exceeds a threshold, the algorithm declares a new target and initializes a tracking grid to follow the target. This target grid is mobile and continually re-centered on the predicted target location. It is desirable that this tracking grid is as small as feasible for both computational reasons and also to allow multiple closely spaced targets to be tracked on their own grids. However, the grid must be large enough in spatial extent to account for temporal uncertainties in target motion and measurement error.

The individual target temporal updates proceed exactly as in the single target case. The measurement updates of the detection and tracking grids use a measurement-censoring step not present in the single target tracker. This measurement censoring step is executed in lieu of fully estimating the joint multitarget density, and should be looked upon as an engineering method for dealing with closely spaced targets that is less costly than fully estimating the joint density. In experiments with real data it has been found that this method often provides sufficient accuracy to perform adequate tracking. However, it is our plan to look at joint density estimation as in [5] in future work.

To elaborate, measurements that fall into the spatial extent of any tracker are censored from the detectors. Second, trackers compete for measurements based on their prior probabilities. These steps prevent multiple targets from being incorrectly detected at the same location, and also prevent multiple nearby trackers from simply following the strongest target.

On Grid Resolution

Like the case of single target tracking, it is important that grid cell spatial resolution be chosen judiciously. If grid cells are too coarse, it is possible no cell centers will project into the maximum of the conditional likelihood (refer back to Figure 7). However, even if the grid size too large, the behavior of the tracker may not be catastrophic. Typical behavior is that a detector initiates a track; the tracker follows the target for some period of time and terminates the track; then the detector reinitiates a tracker on the same target. Figure 11 illustrates the effect of grid cell resolution on tracking performance. The left panel shows performance when the grid is too coarsely spaced. The red dots show track termination points, and illustrate that a single target track is routinely broken and restarted when the grid spacing is too coarse. The right panel shows performance when grid spacing is appropriately selected. All tracks are followed throughout the entire vignette with no track fragmentation.

Unlike the case of single target tracking, the multitarget case exhibits problems when the grid resolution is too small. In particular, since the multitarget detectors constantly seek new targets, small grid cell resolution can have the unintended consequence of allowing energy from the conditional update to bleed onto both a tracker (correctly) and the underlying detector (incorrectly) and thus generate

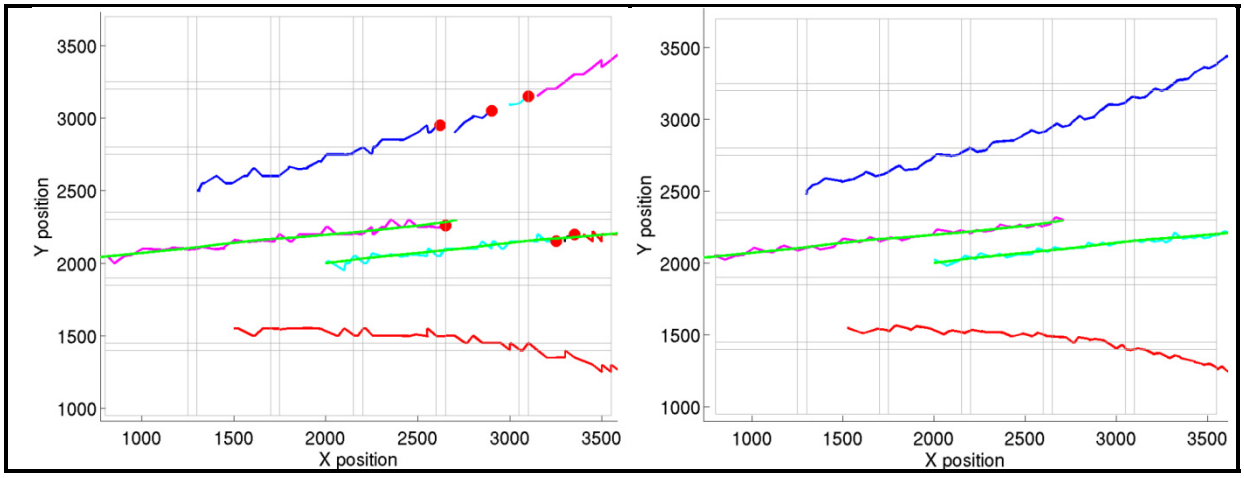

Figure 11 – Left: Multitarget tracking with (too) coarse grid resolution. Right: Multitarget tracking with appropriately selected resolution. In both panels, red dots show track termination points, green lines show true target trajectories, and the other colored lines show the track estimates.

false (double) targets.

Figure 12 illustrates this situation. The conditional update correctly updates the tracker which is tasked with following the target. However, the detector which is partially spatially coincident with the tracker also receives energy from the conditional update. This can lead the detector to initiate (falsely) a second target nearby the first target.

This effect can be countered a number of ways. First, we can adjust the speed at which the tracker re-centers itself. The double initialization phenomenon occurs when the PDF peaks near the edge of the tracker grid. However, this method has the side effect of potentially allowing probability to fall off of the grid in low SNR environments, causing track loss. Of course, if the SNR is low enough or measurement outages occur tracks will be dropped. Second, a guardband around the tracker that does not allow any detector sufficiently near the tracker to receive reinforcement via the conditional density can mitigate the double target problem. However, this has the side effect of preventing detection of closely spaced targets. Third, increasing the spatial extent of the tracker has a similar effect as the using a guardband. It does require increased computation, but generates a better representation of the posterior.

There are several engineering tradeoffs. The first is that large tracker grids (or large guard bands) prevent falsely detecting new targets because of conditional probability spill over. However, if applied too aggressively, this will prevent correctly detecting closely spaced targets. Second, quick tracker grid translation correctly centers the target mass, again preventing spillover into nearby detectors. However, overly liberal tracker repositioning may in fact move trackers to spurious energy locations and drop true targets off of the finite grid.

On Ambiguous Targets

As discussed earlier, ambiguous targets will eventually move non-physically and this will cause the tracker to remove them via its natural prediction and update process.

Figure 13 illustrates this phenomenon. There are two real targets that create two persistent ambiguities. All four are detected and tracked automatically. The ambiguous targets, however, eventually move non-physically due to their reliance on the node bearing angles. The tracker automatically penalizes the non-physical motion and the targets' present hypothesis decrease quickly over time.

Ambiguous target removal is done automatically in the Bayesian framework as follows. The PDF on target state is predicted forward in time according to the kinematic model. True targets will have behavior consistent with the kinematic model (note the kinematic model is a statistical model so it is predicting a range of possibilities for the future target state). Ambiguous targets may behave consistently with this model for a period of time, but eventually they will appear to perform a non-physical maneuver (these epochs typically come when the ambiguous target crosses a line of symmetry in the sensor). At this point, the predicted target position will be in strong disagreement with the incoming measurements on that target. This mismatch in predicted target position and measurements leads to a decrease in the target present hypothesis as calculated in eq. (4). Before long, only true targets remain.

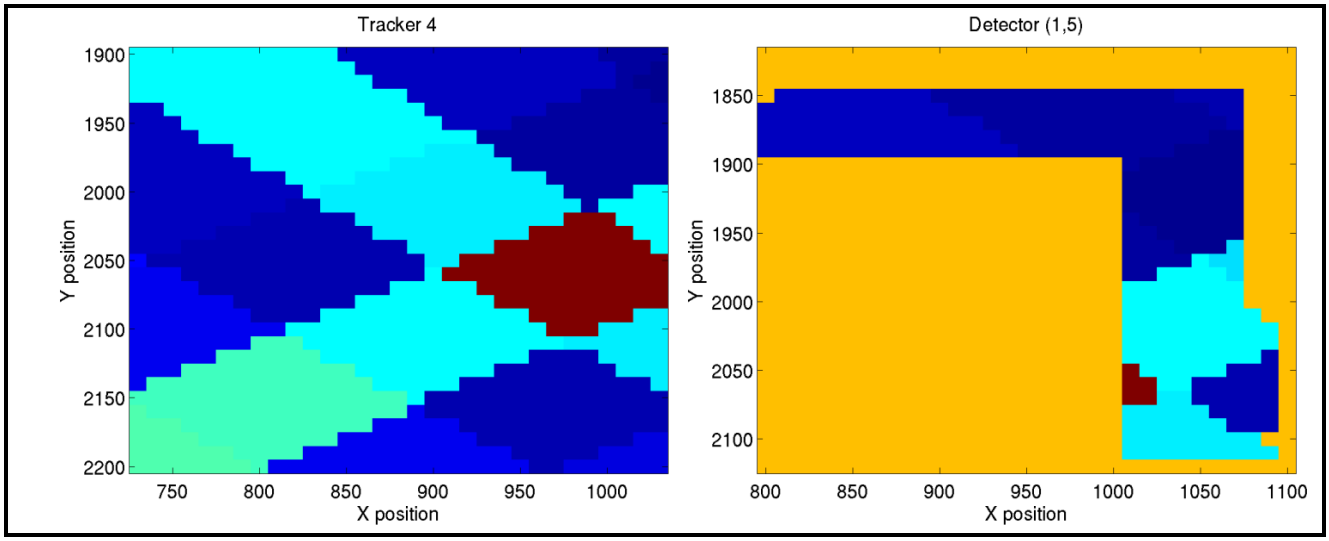

Figure 12 – Improper selection of grid resolution leads to multiple initializations on the same target. Left: Measurement update of a Tracker (red=highest likelihood, blue=lowest). Right: Measurement update of a detector which lies near the Tracker. Since the tracker size has been improperly chosen, some energy from the measurements of a single target leaks on to the detector. This can lead to false double-initializations.

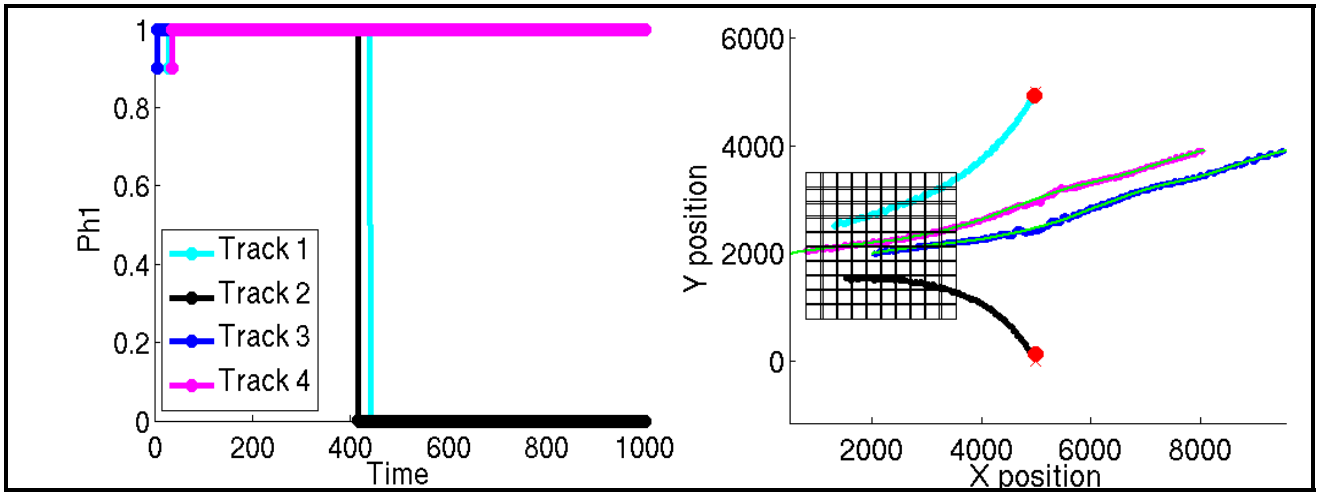

Figure 13– Left: P_{h1} over time for four targets, two of which are real and two of which are ambiguous. Although the ambiguous intersections are persistent, eventually the false targets have non-physical motion. The target present hypothesis quickly goes to zero for these targets and they are eliminated. Right: the tracker estimate of target position and red circles indicating the removal point for the false targets.

8. Conclusion

This paper has described a Bayesian approach to detecting and tracking multiple moving targets using acoustic data from multiple passive arrays. In contrast to traditional undersea acoustic systems, which develop tracks at the single array level and require track association, our approach fuses data at the measurement level and operates directly in the target state space.

We have detailed a well known nonlinear filtering approach to single target detection and tracking [1, 4] and described our computationally efficient finite-grid approach to the required density estimation. We have furthermore extended this to the multiple target case by employing a bank of single target detector trackers and approximation methods that adjust for closely spaced targets. This approximate approach avoids fully treating the computationally complex joint multitarget problem.

Future work includes modified approaches to posterior estimation including dynamic grid extent, dynamic grid resolution, and particle filtering. It is anticipated that adaptive sampling of the posterior will lead to computational savings. Furthermore, future work includes more detailed modeling and estimation of closely spaced targets allowing a more accurate representation of the joint target density. Naively implemented, this implies exponential growth (in the number of targets) for the probability state space being estimated. However, recent work in a related tracking domain on adaptive density factorization [5] and stochastic sampling (particle filtering) [6] provide methods that mitigate this computation growth when the full joint density is treated.

Appendix

This section discusses the details of how the single target probability density is time evolved on a discrete grid. This discussion is similar to that found elsewhere [15, 14, 13].

We wish to compute the single target probability density at time $t + \Delta t$, $p(x, \dot{x}, y, \dot{y}, t + \Delta t)$, from the density at time t, $p(x, \dot{x}, y, \dot{y}, t)$. The relation between these two densities can be expressed using the law of total probability as

$$\begin{aligned} &p(x, \dot{x}, y, \dot{y}, t + \Delta t) \\ &= \int p(x - \Delta x, \dot{x} - \Delta \dot{x}, y - \Delta y, \dot{y} - \Delta \dot{y}, t) \\ &\times p(\Delta x, \Delta \dot{x}, \Delta y, \Delta \dot{y}, \Delta t)\, d\Delta x\, d\Delta \dot{x}\, d\Delta y\, d\Delta \dot{y}. \end{aligned} \tag{23}$$

We expand $p(x - \Delta x, \dot{x} - \Delta \dot{x}, y - \Delta y, \dot{y} - \Delta \dot{y}, t)$ using a second order Taylor series as

$$\begin{aligned} &p(x - \Delta x, \dot{x} - \Delta \dot{x}, y - \Delta y, \dot{y} - \Delta \dot{y}, t) \\ &\approx p(x, \dot{x}, y, \dot{y}, t) - (Dp)^T [\Delta x\ \ \Delta \dot{x}\ \ \Delta y\ \ \Delta \dot{y}] \\ &+ \frac{1}{2} [\Delta x\ \ \Delta \dot{x}\ \ \Delta y\ \ \Delta \dot{y}]^T (D^2 p)^T [\Delta x\ \ \Delta \dot{x}\ \ \Delta y\ \ \Delta \dot{y}], \end{aligned} \tag{24}$$

where (Dp) is the vector of partial derivatives, i.e., $(Dp) = \left[\frac{\delta p}{\delta x}\ \ \frac{\delta p}{\delta \dot{x}}\ \ \frac{\delta p}{\delta y}\ \ \frac{\delta p}{\delta \dot{y}}\right]$, and $(D^2 p)$ is the matrix of second order partial derivatives.

Then the relation of (23) is then approximated as

$$
\begin{aligned}
&p(x,\dot{x},y,\dot{y},t+\Delta t)\\
&\quad\approx p(x,\dot{x},y,\dot{y},t)-\frac{\delta p}{\delta x}\mathbb{E}[\Delta x]\\
&\quad-\frac{\delta p}{\delta \dot{x}}\mathbb{E}[\Delta \dot{x}\,]+\frac{1}{2}\frac{\delta^2 p}{\delta x^2}\mathbb{E}[(\Delta x)^2]\\
&\quad+\frac{1}{2}\frac{\delta^2 p}{\delta \dot{x}^2}\mathbb{E}[(\Delta \dot{x})^2]\\
&\quad+\frac{\delta^2 p}{\delta \dot{x}^2}\mathbb{E}[\Delta x\Delta \dot{x}]+\cdots
\end{aligned}
\tag{25}
$$

Where $\mathbb{E}$ denotes the expectation with respect to the transition distribution $p(\Delta x,\Delta\dot{x},\Delta y,\Delta\dot{y},\Delta t)$, and the omitted terms involve similar terms involving Δy and $\Delta\dot{y}$ and cross terms between the x and y coordinates.

We use the nearly constant velocity (NCV) model to specify the transition distribution $p(\Delta x,\Delta\dot{x},\Delta y,\Delta\dot{y},\Delta t)$. This assumption corresponds to one where the target moves at constant velocity except for random jump changes (i.e., nearly constant velocity). This is a plausible model when Δt is small as it is here.

Specifically, the NCV model assumes step changes in target velocity defined by the Ito Equations

$$
\begin{aligned}
d\dot{x}&\sim N(0,dt\;\sigma_{\dot{x}}^2)\\
d\dot{y}&\sim N\big(0,dt\;\sigma_{\dot{y}}^2\big).
\end{aligned}
\tag{26}
$$

This model implies $dx=\dot{x}dt$ (and likewise for dy). It is furthermore assumed that the noise processes in each coordinate are independent.

Under this model, we can evaluate the required terms from (25) as follows:

$$
\begin{aligned}
\mathbb{E}[\Delta x]&=\dot{x}\Delta t\\
\mathbb{E}[\Delta\dot{x}]&=0\\
\mathbb{E}[(\Delta x)^2]&=\dot{x}^2\Delta t^2\\
\mathbb{E}[(\Delta\dot{x})^2]&=\sigma_{\dot{x}}^2\Delta t\\
\mathbb{E}[\Delta x\Delta\dot{x}]&=0.
\end{aligned}
\tag{27}
$$

And likewise for terms involving Δy and $\Delta\dot{y}$. Notice that all cross terms (e.g., $\mathbb{E}[\Delta x\Delta\dot{y}]$) have expectation 0 due to the assumption that the noise process is independent in the two coordinates.

This model simplifies (25) to

$$
\begin{aligned}
&p(x,\dot{x},y,\dot{y},t+\Delta t)\\
&\quad\approx p(x,\dot{x},y,\dot{y},t)-\frac{\delta p}{\delta x}\dot{x}\Delta t\\
&\quad+\frac{1}{2}\frac{\delta^2 p}{\delta x^2}\dot{x}^2\Delta t^2+\frac{1}{2}\frac{\delta^2 p}{\delta \dot{x}^2}\sigma_{\dot{x}}^2\Delta t\\
&\quad+\cdots
\end{aligned}
\tag{28}
$$

where the terms omitted are replicas involving the y coordinate.

Under the assumption that Δt is small, this can be rewritten as

$$
\frac{\delta p}{\delta t}=-\frac{\delta p}{\delta x}\dot{x}+\frac{\sigma_{\dot{x}}^2}{2}\frac{\delta^2 p}{\delta \dot{x}^2}-\frac{\delta p}{\delta y}\dot{y}+\frac{\sigma_{\dot{y}}^2}{2}\frac{\delta^2 p}{\delta \dot{y}^2}.
\tag{29}
$$

For implementation, this is approximated using an implicit Euler scheme [12] where

$$
\frac{\delta p}{\delta t}\approx\frac{p_{i,j,l,m}^{k+1}-p_{i,j,l,m}^{k}}{\Delta t}.
\tag{30}
$$

Where the indices i,j,l,m represent the discrete $x,\dot{x},y,\dot{y}$ locations where the probability mass is captured. Likewise, using forward differencing

$$
\frac{\delta p}{\delta x}\approx\frac{p_{i,j,l,m}^{k+1}-p_{i-1,j,l,m}^{k+1}}{\Delta x}
\tag{31}
$$

and

$$
\frac{\delta^2 p}{\delta \dot{x}^2}\approx\frac{p_{i,j+1,l,m}^{k+1}-2p_{i,j,l,m}^{k+1}+p_{i,j-1,l,m}^{k+1}}{(\Delta\dot{x})^2}
\tag{32}
$$

and similarly for the y coordinate system.

When substituted into (28), this leads to a series of equations of the form

$$
\begin{aligned}
p_{i,j,l,m}^{k+1}&\left[1+\frac{\dot{x}\Delta t}{\Delta x}+\frac{\sigma_{\dot{x}}^2\Delta t}{(\Delta x)^2}+\frac{\dot{y}\Delta t}{\Delta y}+\frac{\sigma_{\dot{y}}^2\Delta t}{(\Delta y)^2}\right]\\
&+p_{i-1,j,l,m}^{k+1}\left[-\frac{\dot{x}\Delta t}{\Delta x}\right]\\
&+p_{i,j+1,l,m}^{k+1}\left[-\frac{\sigma_{\dot{x}}^2\Delta t}{2(\Delta x)^2}\right]\\
&+p_{i,j-1,l,m}^{k+1}\left[-\frac{\sigma_{\dot{x}}^2\Delta t}{2(\Delta x)^2}\right]+\cdots=p_{i,j,l,m}^{k}
\end{aligned}
$$

This series of equations define the probability at each point at time $k+1$. It can be efficiently solved via Thomas' algorithm (rather than simply inverted) as the matrix is tridiagonal.

REFERENCES

[1] Roy E. Bethel, Benjamin Shapo, Christopher M. Kreucher, "PDF Detection and Tracking", under review *IEEE Transactions on Aerospace and Electronic Systems*.

[2] Roy. E. Bethel and G. J. Paras, "A PDF Multitarget Tracker", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 30, no. 2, pp. 386-403, April 1994.

[3] Roy E. Bethel and G. J. Paras, "A PDF Multisensor Multitarget Tracker", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, no. 1, pp. 153-168, January 1998.

[4] L. D. Stone, C. A. Barlow, and T. L Corwin, "Bayesian Multiple Target Tracking". Boston: Artech House, 1999.

[5] C. Kreucher, K. Kastella, and A. Hero, "Multitarget Tracking using the Joint Multitarget Probability Density", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1396-1414, October 2005.

[6] M. Morelande, C. Kreucher, K. Kastella, "A Bayesian Approach to Multiple Target Detection and Tracking", *IEEE Transactions on Signal Processing*, vol. 55, no. 5, pp. 1589-1604, May 2007.

[7] B. Shapo, and R. E. Bethel, "An Overview of the Probability Density Function (PDF) Tracker" *Oceans 2006*, Boston, Sept. 2006.

[8] Roy L. Streit, "Multisensor Multitarget Intensity Filter", *International Conference on Information Fusion*, Cologne, Germany July 2008.

[9] M. Orton and W. Fitzgerald, "A Bayesian approach to tracking multiple targets using sensor arrays and particle filters" IEEE Transactions on Signal Processing, vol. 50, no. 2, pages 216-223, Feb 2002

[10] A. Doucet, B. Vo, C. Andrieu, and M. Davy, "Particle filtering for multi-target tracking and sensor management", IEEE International Conference on Information Fusion, 2002

[11] H. Van Trees, "Detection, Estimation, and Modulation Theory IV: Optimum Array Processing"

[12] J. C. Strikwerda, Finite Difference Schemes and Partial Differential Equations, Chapman & Hall, New York, 1989

[13] K. Kastella and C. Kreucher, "Multiple Model Nonlinear Filtering for Low Signal Ground Target Applications", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 2, April 2005, pp. 549-564.

[14] Z. Tang and Ü. Özgüner, "Sensor Fusion for Target Track Maintenance with Multiple UAVs based on Bayesian Filtering Method and Hospitability Map", *Proceedings of the 42nd IEEE Conference on Decision and Control,* pages 19-24, December 2003.

[15] K. Kastella, "Finite difference methods for nonlinear filtering and automatic target recognition", *Multitarget-Multisensor Tracking: Applications and Advances,* vol. III, pages 233-258, Artech House, 2000.

BIOGRAPHY

Chris Kreucher *received his Ph.D. in Electrical Engineering from the University of Michigan in 2005. He is currently a Senior Systems Engineer at Integrity Applications Incorporated in Ann Arbor, Michigan. From 1998 to 2007, he was a Staff Scientist at General Dynamics Advanced Information Systems' Michigan Research & Development Facility (formerly ERIM). His current research interests include nonlinear filtering (specifically particle filtering), Bayesian methods of multitarget tracking, self localization, information theoretic sensor management, and distributed swarm management.*

Ben Shapo *earned his Ph.D. in Electrical Engineering in 1996 from the University of Michigan. He is currently a Senior Systems Engineer at Integrity Applications Incorporated in Ann Arbor, Michigan. From 2003 to 2008, he was a Lead Engineer at General Dynamics, where he contributed to a number of RF and acoustics signal processing and tracking efforts. Dr. Shapo has 12 years experience in the DoD research community in the areas of detection, tracking, and data fusion, with emphasis on high-fidelity simulations and applying new methods to real data.*

Dr. Roy Bethel *is currently employed at The MITRE Corporation in McLean, VA. He has been actively involved in development, testing, and evaluation of signal processing and detection and tracking systems. In particular, he has developed many systems that have been implemented on United States Navy airborne, surface, and submerged platforms. He is currently engaged in research and development of innovative approaches to multitarget detection and tracking.*

ACKNOWLEDGEMENTS

This work was partially funded by the Office of Naval Research contract N00014-08-C-0275. The authors would like to thank Dr. John Tague for his support, and Mr. Scott Spencer and Dr. Charles Choi for their assistance.