Multitarget Detection and Tracking Using Multisensor Passive Acoustic Data

Chris Kreucher and Ben Shapo

***Abstract*—This paper describes a Bayesian approach to detecting and tracking multiple moving targets using acoustic data from multiple passive arrays. We describe a surveillance application, where a collection of fixed-location passive acoustic arrays is charged with monitoring a predefined spatial region. Our approach combines a unique hybrid discrete-grid/particle approximation to the posterior with a dynamic density factorization. This results in a novel 2-D (X/Y) multisensor multitarget tracker that uses bearing measurements only. The efficacy of the algorithm is illustrated both in simulation and on collected at-sea data.**

***Index Terms*—Fuse-before-track, fusion, nonlinear filtering, passive acoustics, tracking.**

I. Introduction

THIS paper describes a Bayesian approach to detecting and tracking multiple moving targets using acoustic data from multiple passive arrays. We focus on a surveillance application, where a collection of passive acoustic arrays monitors a predefined spatial region to detect and track moving targets in 2-D. This regime presents two main challenges. First, bearing measurements provide incomplete information about target state and couple nonlinearly to the target state. Second, passive acoustic sensors typically receive low signal levels. Therefore, detect-before-track approaches which declare target detections by thresholding received signals may have an unacceptable false alarm/detection tradeoff [1], [2].

Traditional tracking methods [3]–[7] are based on linear (or linearized) Kalman filters and use detections (i.e., threshold exceedances) as input. These suboptimal methods are used for a number of good reasons in different applications. First, in some applications (e.g., RADAR) sensor measurements occur at a very high rate, requiring the tracker to execute very quickly. Second, communication channels between the sensor and the processing unit may be limited (and this is exacerbated by the high-data rate), meaning that only summary information can be sent to the processor and not the raw measurements. For these reasons, Kalman techniques which are implemented by a series of matrix operations are very useful.

In the passive sonar surveillance scenario we focus on here, these constraints are relaxed. First, sensor measurements are aggregated at a low rate, on the order of one Hertz. Second, in our application the arrays are directly cabled to a processing center, allowing raw measurements to be transmitted faithfully. Finally, the processing center can be equipped with high-power computers. For these reasons, techniques which are more computationally burdensome and require higher communication requirements are viable in this application. Therefore, given the nonlinear measurement modality and utility of using non-thresholded data, we advocate a nonlinear track-before-detect approach which fully models the nonlinear measurement to target state coupling and operates with raw measurements rather than just threshold exceedances.

In the multisensor case, conventional methods [8]–[11] develop tracks at each sensor, associate tracks between nodes, and then fuse. In contrast, the present application has the bandwidth to allow data fusion from multiple nodes at the measurement level, and can employ hardware with sufficient computational power to do so. Combined with the challenges of low signal levels and nonlinear measurement to target coupling, this again suggests measurement level nonlinear filtering. A fuse-before-track approach can increase algorithm performance over conventional track-and-then-fuse methods by delaying hard decisions about target existence and state until all of the sensor data has been incorporated.

A number of track-before-detect approaches have been studied in the literature. The "Unified Data Fusion" work of Stone [12], [13], the "JMPD" approach of Kastella [14], and the dim target tracking work of Boers [15] are all examples of Bayesian approaches which do not require thresholding, explicit measurement to track association, linear models, or Gaussian statistics. They also may be used with multiple sensors without using track fusion by incorporating multisensor measurements through the likelihood function. Many others have done important related work [16]–[25]. In particular, Bethel and Shapo [26] apply track-before-detect to single-node passive acoustic tracking.

The method presented in this work is a combination and extension of these and other existing techniques. It gives a novel, bearings-only 2-D (X/Y) multisensor multitarget tracker. In particular, we present the following main contributions. First, we give a unique implementation which approximates the required nonparametric probability density using a fixed grid (discrete) method while the density is diffuse and an adaptive grid (particle) scheme once the density is well localized. Second, we apply a dynamic probability density function factorization method that accommodates closely-spaced targets, modeling coupling in target state uncertainty when necessary, but not wasting computations when such modeling is not required. Finally, we illustrate the efficacy of the algorithm on both simulated and real, collected at-sea data.

Manuscript received March 09, 2010; revised January 10, 2011; accepted February 16, 2011. Date of publication May 12, 2011; date of current version May 27, 2011.

Associate Editor: K. Wage.

The authors are with the Integrity Applications Incorporated, Fusion Group, Ann Arbor, MI 48108 USA (e-mail: ckreuche@umich.edu; bshapo@integrity-apps.com).

Digital Object Identifier 10.1109/JOE.2011.2118630

0364-9059/$26.00 © 2011 IEEE

The paper proceeds as follows. Section II describes the multisensor passive sonar setting and our application. Sections III and IV describe single-target Bayesian detection and tracking and our implementation, which is based on a novel combination of discrete grid detection and particle filter tracking methods. Sections V and VI describe multitarget filtering and our implementation which utilizes a unique adaptive factorization to capture important multitarget target couplings while not requiring full estimation of the high dimensionality joint target density for all targets. Section VII illustrates the value of the approach on simulated and real collected passive acoustic data from a sea test. Section VIII concludes the paper.

II. The Multisensor Passive Sonar Setting

The application we focus on here is surveillance of a large spatial region using multiple passive acoustic arrays. This section describes the nominal sensor layout and develops a statistical signal model for the energy received by the arrays, which is the input to our tracking algorithm.

A. The Surveillance Problem

A region of large spatial extent is to be monitored by a collection of fixed (e.g., bottom mounted) linear passive acoustic arrays. There are $\mathcal{I}$ hydrophone arrays where the position of the center element is denoted $\{\mathbf{x}_{(i)}, \mathbf{y}_{(i)}\}$, for array $i = \{1 \ldots \mathcal{I}\}$. The surveillance region is defined by its extent $\{\mathbf{x}^{min}, \mathbf{x}^{max}, \mathbf{y}^{min}, \mathbf{y}^{max}\}$. Ideally, the arrays are located conveniently with respect to the surveillance region, e.g., along the boundaries and at right angles to each other.

Undersea and surface targets emit acoustic energy which is received at the passive arrays. We are interested in detecting and tracking targets in 2-D from this received energy. For loud targets, conventional signal processing methods use this energy to estimate target bearing relative to each array. With well separated arrays (with respect to target range) triangulation can be used to compute target range and 2-D position. However, our application is to low energy targets, which are not amenable to detect before track (i.e., thresholding approaches) and so more sophisticated methods of estimating target state are required.

B. Signal Modeling

Each of the $\mathcal{I}$ physically separated arrays consists of $\mathcal{M}$ hydrophones. Energy impinges on each element m in array i from acoustic sources, and the time series of received energy is recorded as $s^m_{(i)}(t)$. This data is processed to generate estimates of target energy as a function of bearing from the array.

Processing is typically done according to the standard delay-and-sum beamforming method [27]. The spatial separation of individual elements causes the propagation time from a source to each element to differ. Denote the propagation time delay from a source at bearing θ to hydrophone m by $\Psi^m(\theta)$. For a linear array with elements separated by Δ, $\Psi^m(\theta) = m\Delta\cos(\theta)/c$. For a candidate arrival angle θ, the beamformer implements appropriate delays at each element to force any signals arriving from that direction to add coherently, i.e., the coherent sum at array i for arrival direction θ

$$S_{(i),\theta}(t) = \sum_{m=1}^{M} s^m_{(i)}(t - \Psi^m(\theta)) \tag{1}$$

is the total energy in array i at bearing θ as a function of time.

Computation is typically done in the frequency domain with short (on the order of 1-second) time blocks of data as

$$\tilde{S}_{(i),\theta}(f) = \sum_{m=1}^{M} \tilde{s}^m_{(i)}(f) e^{-j2\pi f \Psi^m(\theta)}. \tag{2}$$

In practice, the beamformer performs this computation for a set of candidate arrival directions (beams) numbered $j = \{1 \ldots \mathcal{J}\}$. It then noncoherently aggregated across frequency for each beam to capture all the energy in that direction, yielding what we refer to as the measurement from array i in beam j, $z_{(i),j}$.

Under certain conditions on element bandwidth and snapshot time, the beamformer values at individual frequencies are independent random Gaussian variables [28]. Therefore, we model the statistics of $z_{(i),j}$ as such. The parameters of the Gaussian random variable depend on whether or not there is a target at direction θ_j.

III. Single-Target Detection and Tracking

This section describes the Bayesian approach to single target detection and tracking using data from multiple passive arrays. For the purposes of this section, we assume there is at most one target present. This assumption is removed in Section V.

Notation

We denote the state of a single target at time k as x^k, which for this work refers to the target 2-D position and velocity, i.e., $x^k = [\mathbf{x} \quad \dot{\mathbf{x}} \quad \mathbf{y} \quad \dot{\mathbf{y}}]$. Additionally, let H^k_0 denote the hypothesis that no target is present at time k, and let H^k_1 denote the hypothesis that a single target is present.

The following notation describes the measurements: $z^k_{(i),j}$ continues to denote the measurement received by array i in bearing beam j at time k; $z^k_{(i)}$ denotes the vector of all measurements received by array i at time k, i.e., $z^k_{(i)} = \left\{z^k_{(i),1} \ldots z^k_{(i),\mathcal{J}}\right\}$; z^k denotes measurements received by all arrays at time k, i.e., $z^k = \left\{z^k_{(1)} \ldots z^k_{(\mathcal{I})}\right\}$; finally, Z^k denotes the collection of all measurements received by all arrays up to and including time k, i.e., $Z^k = \{z^1 \ldots z^k\}$.

The Bayesian method is to estimate the joint probability a target is present (H^k_1 is true) at each state x^k given the measurements. Mathematically, this means we wish to estimate the hybrid continuous-discrete probability density function (pdf)

$$p\left(x^k, H^k_1 | Z^k\right) \tag{3}$$

for all x^k, as well as the discrete probability

$$p\left(H_0^k|Z^k\right) \tag{4}$$

which is simply $1-\int p\left(x^k, H_1^k|Z^k\right) dx^k$.

Notice that we can write

$$p\left(x^k, H_1^k|Z^k\right) = p\left(H_1^k|Z^k\right) p\left(x^k|H_1^k, Z^k\right) \tag{5}$$

i.e., the density is the product of the *target present probability* $p\left(H_1^k|Z^k\right)$ and the *target state probability* $p\left(x^k|H_1^k, Z^k\right)$. Both conceptually and in implementation, we treat the problem as separate (but coupled) tasks of estimating the target present probability and the estimating target state probability.

In the Bayesian approach, we (1) assume an initial or prior estimate of the desired probabilities is present (perhaps completely uninformative), and (2) generate a recursive formula that relates probabilities at one time step with those at the next. This is done in two steps, analogous to the Kalman Filter: the temporal update, which predicts the probability distribution at time k from that at time $k-1$, and the measurement update which corrects the predicted probability distribution at time k given the measurements received at time k.

A. Temporal Update

The first step in recursive Bayesian filtering is to predict the relevant probability distributions forward in time using statistical models on target kinematics. The temporal update of the target present density is

$$p\left(H_1^k|Z^{k-1}\right) = \sum_{i=0}^{1} p\left(H_1^k|Z^{k-1}, H_i^{k-1}\right) p\left(H_i^{k-1}|Z^{k-1}\right) \tag{6}$$

where the quantity $p\left(H_1^k|Z^{k-1}, H_i^{k-1}\right)$ is a statistical model of how targets arrive and exit the surveillance region, to be specified by studying the target arrival properties.

Similarly, the time-predicted target state density is based on a model of how targets move

$$\begin{aligned} p\left(x^k|H_1^k, Z^{k-1}\right) &= \frac{p\left(H_1^{k-1}|Z^{k-1}\right)}{p\left(H_1^k|Z^{k-1}\right)} \\ &\times \int p\left(x^k, H_1^k|x^{k-1}, H_1^{k-1}\right) \\ &\times p\left(x^{k-1}|H_1^{k-1}, Z^{k-1}\right) dx^{k-1} \end{aligned} \tag{7}$$

where the density $p\left(x^k, H_1^k|x^{k-1}, H_1^{k-1}\right)$ is a statistical model on target kinematics specified in the particular implementation. The normalizing term $p\left(H_1^{k-1}|Z^{k-1}\right)/p\left(H_1^k|Z^{k-1}\right)$ does not need to be evaluated, as the density can be forced to integrate to 1. In this work, we assume the nearly constant velocity (NCV) model for the target. Other models, or even multiple models are admissible under the Bayesian framework [29].

B. Measurement Update

The second step in Bayesian filtering is to incorporate the measured data into the probability estimate. The measured data comes into the picture through the likelihood ratio

$$\lambda\left(z^k|x^k, H_1^k\right) \doteq \frac{p\left(z^k|x^k, H_1^k\right)}{p\left(z^k|H_0^k\right)} \tag{8}$$

where the functional form of $\lambda\left(z^k|x^k, H_1^k\right)$ is a model specified by sensor physics. Recall that in this approach, measurements z are not just threshold exceedances, but rather the full measurement set consisting of beamformer outputs from all sensors at all beams.

With this definition, the target present and target absent detection probabilities are measurement-updated using the law of total probability and Bayes' rule, yielding

$$\begin{aligned} p\left(H_1^k|Z^k\right) &= \frac{p\left(H_1^k|Z^{k-1}\right) p\left(z^k|H_1^k, Z^{k-1}\right)}{p\left(z^k|Z^{k-1}\right)} \\ &= p\left(H_1^k|Z^{k-1}\right) \int \lambda\left(z^k|x^k, H_1^k\right) \\ &\quad \times p\left(x^k|H_1^k, Z^{k-1}\right) dx^k \\ &\quad \times \frac{p\left(z^k|H_0^k\right)}{p\left(z^k|Z^{k-1}\right)} \end{aligned} \tag{9}$$

and

$$p\left(H_0^k|Z^k\right) = p\left(H_0^k|Z^{k-1}\right) \times \frac{p\left(z^k|H_0^k\right)}{p\left(z^k|Z^{k-1}\right)}. \tag{10}$$

These equations express the current target present and absent hypothesis probabilities in terms of the target present, target absent, and target state probabilities predicted from the previous time step and the conditional likelihood of the incoming measurements. The normalization constant $p\left(z^k|H_0^k\right)/p\left(z^k|Z^{k-1}\right)$ does not need to be computed since $p\left(H_1^k|Z^k\right) + p\left(H_0^k|Z^k\right) = 1$.

The target state probability is updated in a similar manner

$$\begin{aligned} &p\left(x^k|H_1^k, Z^k\right) \\ &\quad = p\left(x^k|H_1^k, Z^{k-1}\right) \lambda\left(z^k|x^k, H_1^k\right) \frac{p\left(z^k|H_0^k\right)}{p\left(z^k|H_1^k, Z^{k-1}\right)} \end{aligned} \tag{11}$$

again $p\left(z^k|H_0^k\right)/p\left(z^k|H_1^k, Z^{k-1}\right)$ is independent of x^k and does not need to be computed since the probability density integrates to 1.

In our multisensor passive array application, the likelihood ratio $\lambda\left(z^k|x^k, H_1^k\right)$ can be expressed as follows. We assume measurements from different arrays are independent conditioned on the target state, i.e.

$$p\left(z^k|x^k, H_1^k\right) = \prod_{i=1}^{\mathcal{I}} p\left(z_{(i)}^k|x^k, H_1^k\right) \tag{12}$$

and that the measurements a particular array takes in each beam are independent conditioned on the target state, i.e.

$$p\left(z_{(i)}^k|x^k, H_1^k\right) = \prod_{j=1}^{\mathcal{J}} p\left(z_{(i),j}^k|x^k, H_1^k\right). \tag{13}$$

For each array i, the state $x^k = [\,\mathbf{x} \quad \dot{\mathbf{x}} \quad \mathbf{y} \quad \dot{\mathbf{y}}\,]$ corresponds to bearing θ_{x^k}, which maps to a beam we will denote $\hat{j}_{(i)}$. With this notation, we write

$$\begin{aligned} p\left(z_{(i)}^k|x^k, H_1^k\right) &= p_1\left(z_{(i),\hat{j}_{(i)}}^k\right) \prod_{j\neq\hat{j}_{(i)}} p_0\left(z_{(i),j}^k\right) \\ &= \frac{p_1\left(z_{(i),\hat{j}_{(i)}}^k\right)}{p_0\left(z_{(i),\hat{j}_{(i)}}^k\right)} \prod_{j=1}^{\mathcal{J}} p_0\left(z_{(i),j}^k\right) \end{aligned} \tag{14}$$

where $p_1(z)$ is the probability density on the received energy in beams where the target exists (the target present density) and $p_0(z)$ is the probability density on the received energy in beams where no targets exist (the target absent density). Thus, this equation says that the likelihood of a scan of data $z_{(i)}^k$ given target state x^k is proportional to the ratio of the target present probability in the beam x^k maps to divided by the target absent probabilty in that same beam. We have ignored extended target effects such as sidelobes in this simplification. If we were to consider these, the state x^k would map to a set of beams rather than a single beam and (14) would then include a term that depended on all of the cells in the mapping.

Combining (12) and (14), we see that fusing over multiple arrays gives the multisensor likelihood ratio

$$\lambda\left(z^k|x^k, H_1^k\right) = \prod_{i=1}^{\mathcal{I}} \frac{p_1(z_{(i),\hat{j}_{(i)}})}{p_0(z_{(i),\hat{j}_{(i)}})}. \tag{15}$$

IV. Single-Target Implementation

If the probability density of interest $p\left(x^k|H_1^k, Z^k\right)$ is well approximated by a Gaussian or sum-of-Gaussians, techniques such as the extended Kalman filter, unscented Kalman filter, or Gaussian sum filter are preferred. In the multisensor passive acoustic case, however, the density is poorly approximated by such parameterizations. We instead rely on two nonlinear filtering approaches, the discrete grid and particle filter.

A. The Merits of Discrete Grid and Particle Representations

The numerics involved in a discrete grid approximation to a pdf is a well studied area [29]–[31]. In the discrete grid approach, a pdf of interest is typically represented on a fixed, evenly spaced, multidimensional grid. The main benefit of this method is it generates probability estimates over a defined region of state space, only assuming the pdf is zero outside of the grid boundaries. Conversely, the main deficiencies are that the fixed discrete grid approach spends computational effort updating grid cells with near-zero probability and that the number of grid cells grows exponentially with the dimension of the state space.

The particle filter approach, on the other hand, uses an adaptive grid where the tie-points are computed online via importance sampling [32], [33]. For certain classes of problems [34] and a well localized initial pdf, the particle approach provides excellent estimation performance at a fraction of the cost of a discrete grid approach. Thus, its main strengths are that computational effort is only used in areas of high pdf probability, and the grid adaptively changes size. This is important in the present application because the shape of the spatial uncertainty about target location varies dramatically with its position relative to the sensors. The main deficiency is this sparse sampling approach is ill-suited to represent very broad pdfs, such as a pdf which has uniform uncertainty over a large spatial region.

For these reasons, we advocate a hybrid approach which inherits the best features of the two methods. The pdf is first approximated using a discrete grid. At some point, the target present probability exceeds a threshold, indicating a target is present and it has been well localized. At this point, the pdf approximation is transitioned to a particle filter. This allows sustained tracking of this target with a grid that adapts to the pdf's shape and extent.

B. The Discrete Grid Representation

The details of the discrete grid method [29], [30], [35] are briefly reviewed here.

1) Density Representation: The pdf of x is discretized onto a 4D grid (corresponding to the four dimensional state vector x^k) of $N_x \times N_{\dot{x}} \times N_y \times N_{\dot{y}}$ cells. This approximation is appropriate here, given we wish to perform surveillance over a region of fixed spatial extent. The effect of spatial truncation of the probability distribution will be discussed later.

2) Kinematic and Measurement Updates: The NCV model we adopt leads to the Fokker–Plank equation [30], [36]–[38] which says that the rate of change of the probabilty density with respect to time can be expressed in terms of the rates of change with respect to the parameter values x, y, $\dot{x}$, and $\dot{y}$ as

$$\frac{\partial p}{\partial t} = -\dot{x}\frac{\partial p}{\partial x} - \dot{y}\frac{\partial p}{\partial y} + \frac{\sigma_{\dot{x}}^2}{2}\frac{\partial^2 p}{\partial \dot{x}^2} + \frac{\sigma_{\dot{y}}^2}{2}\frac{\partial^2 p}{\partial \dot{y}^2} \tag{16}$$

where $\sigma_{\dot{x}}$ and $\sigma_{\dot{y}}$ are diffusion constants that model the kinematics of the target.

Computationally, the state probability is discretized onto the grid and the update is computed from time $k-1$ to k using an implicit Euler method. This approach is always stable and has acceptable accuracy in δt and δx. We use Thomas' algorithm as a fast tridiagonal solver leading to computation linear in the number of grid cells. For more details, see [29], [39].

The temporal evolution of the target present probability $p\left(H_1^k|Z^{k-1}\right)$ assumes constant target arrival/removal, i.e.

$$\begin{aligned} &p\left(H_1^k|Z^{k-1}\right) \\ &\quad = qp\left(H_1^{k-1}|Z^{k-1}\right) + (1-q)p\left(H_0^{k-1}|Z^{k-1}\right) \end{aligned} \tag{17}$$

where q is the arrival rate parameter. As this has been cast as a standard hypothesis test between the two point hypotheses representing target present and target absent, we declare a target present when the probability of target present hypothesis $p\left(H_1^k|Z^k\right)$ exceeds a threshold in accordance with the standard

Neyman-Pearson approach. The Bayes update of the probability density (11) is approximated when using the discrete grid representation by a pointwise multiplication of each cell in the discrete representation by the corresponding data likelihood ratio (15).

3) Effect of the Finite Grid Extent: The problem of interest is surveillance over a fixed spatial region. As such, the spatially limited discrete grid pdf approximation proposed here is a natural approach. However, (nuisance) targets which are located off of the grid do contribute energy and this phenomenon must be accounted for in the formulation.

The likelihood ratio derived above implicitly assumes the probability density on x^k is estimated over an unlimited spatial extent. In practice, in this application we are only interested in estimation over a finite spatial region. Hence the H_1 and H_0 hypotheses correspond to the event that a target is present in the region of interest, and the event that a target is not present in the region of interest, respectively. Therefore, the definition of the likelihood ratio λ must be modified to reflect the finite size region over which we wish to estimate target presence by properly accounting for energy emitted by off-grid targets.

We therefore generalize the likelihood ratio to account for the modified definitions of H_1 and H_0. In particular, the target absent hypothesis H_0 is now the composite hypothesis that either a target is not present, or that a target is present but outside the finite grid. For the two array case and a state x^k that is hypothesized to not contain a target, two things can happen: the target absent statistics are present in both arrays, or there is a real target outside of the detector's spatial region that concidentally puts energy into one of the beams (but not the other). Therefore, the target absent hypothesis is supported when both arrays receive energy from the target absent density, or one array receives energy from the target absent density and one receives energy from the target present density. The target present hypothesis is supported only when both arrays receive energy from the target present density. Concretely, we define λ in the example of two arrays using the generalized likelihood ratio as

$$\lambda\left(z^k|x^k, H_1^k\right) = \frac{p_1(z_{(1),\hat{j}_{(1)}})p_1(z_{(2),\hat{j}_{(2)}})}{\max\begin{pmatrix} p_0(z_{(1),\hat{j}_{(1)}})p_0(z_{(2),\hat{j}_{(2)}}), \\ p_0(z_{(1),\hat{j}_{(1)}})p_1(z_{(2),\hat{j}_{(2)}}), \\ p_1(z_{(1),\hat{j}_{(1)}})p_0(z_{(2),\hat{j}_{(2)}}) \end{pmatrix}} \tag{18}$$

C. The Particle Filter Representation

An alternative method of representing the target state probability $p(x^k|H_1^k, Z^k)$ is via a particle filter. Particle filtering is an adaptive grid method of representing a pdf and numerically updating it temporally and with measurements [32]. The details are briefly reviewed here.

1) Density Representation: In a single target particle filter, the density of interest is approximated by a set of N_{part} weighted samples (particles)

$$p(x|Z) \approx \sum_{p=1}^{N_{\text{part}}} w_p \delta_D(x - x_p) \tag{19}$$

where δ_D represents the usual Dirac delta function.

2) Kinematic and Measurement Updates: The model update and the measurement update are simulated by the following three step recursion. First, the particle locations at time k are generated using the particle locations x_p at time $k-1$ and the current measurements z^k by sampling from an importance density, denoted $q(x^k|x^{k-1}, z^k)$. The design of the importance density is a well studied area [40], as the choice of the importance density can have a dramatic effect of the efficiency of the particle filter algorithm. It is known that the optimal importance density (OID) is given by $p(x^k|x^{k-1}, z^k)$, but this density is typically prohibitively difficult to sample from. In practice, oftentimes the importance density is chosen just to be the kinematic prior $p(x^k|x^{k-1})$, which is what we do here. A more sophisticated choice of importance density may lead to better results for a fixed computational burden, but that is not studied here.

Particle weights are updated according to the weight equation, which involves the likelihood, the kinematic model, and the importance density [32]

$$w_p^k = w_p^{k-1} \frac{p\left(z^k|x_p^k\right) p\left(x_p^k|x_p^{k-1}\right)}{q\left(x_p^k|x_p^{k-1}, z^k\right)}. \tag{20}$$

When using the kinematic prior as the importance density, the weight equation reduces to $w_p^k = w_p^{k-1} * p\left(z^k|x_p^k\right)$. Finally, a resampling step is used to prevent particle degeneracy. We have selected to resample when the number of effective particles [32] falls below a threshold, here chosen as half the number of particles.

D. Combined Approximation

In our approach, we represent the surveillance region via a discrete grid at onset. Once the target present probability exceeds a threshold, indicating the density is well localized, the approximation is transitioned to a particle representation. This hybrid approach allows good performance when the pdf is broad (at onset) as well as good tracking performance once a target is found. This transition is accomplished by sampling N_{part} particles from the $N_x \times N_{\dot{x}} \times N_y \times N_{\dot{y}}$ discrete grid.

V. Multitarget Detection and Tracking

In this section, we show the generalization of the single target tracking algorithm to the general case where there are an unknown and time varying number of targets.

Notation

In multitarget detection and tracking, we wish to estimate the hybrid continuous-discrete density

$$p\left(x_1^k \ldots x_{T^k}^k, T^k|Z^k\right) \tag{21}$$

for all T and $x_1, \ldots, x_T$, where T is the number of targets ($T = 0, 1 \ldots$) and $x_1 \ldots x_T$ are the states of the individual targets.

For notational convenience, we define

$$X = [x_1 \ldots x_T] \tag{22}$$

i.e., X denotes the multitarget state vector, where the cardinality will be clear by context.

As in the single target case (5), the joint multitarget density can be expressed as the product of the target number density and the target state density as

$$p(X^k, T^k|Z^k) = p(T^k|Z^k)p(X^k|T^k, Z^k) \tag{23}$$

A. Temporal Update

The target number temporal update is given by

$$p(T^k|Z^{k-1}) = \sum_{T^{k-1}=0}^{\infty} p(T^k|T^{k-1}, Z^{k-1})p(T^{k-1}|Z^{k-1}). \tag{24}$$

And the target state temporal update can be expressed

$$\begin{aligned} p(X^k|T^k, Z^{k-1}) &= \left(\frac{p(T^{k-1}|Z^{k-1})}{p(T^k|Z^{k-1})}\right) \\ &\times \sum_{T^{k-1}} \int p(X^k, T^k|X^{k-1}, T^{k-1}) \\ &\times p(X^{k-1}|T^{k-1}, Z^{k-1})dX^{k-1} \end{aligned} \tag{25}$$

where the density $p(X^k, T^k|X^{k-1}, T^{k-1})$ is a statistical model on target kinematics and the integral over X^{k-1} is to be interpreted as performing the T required integrations over the domain of x.

B. Measurement Update

In extension of the definition given earlier (8), we define the multitarget likelihood ratio as

$$\lambda(z^k|X^k, T^k) \doteq \frac{p(z^k|X^k, T^k)}{p(z^k|T^k = 0)}. \tag{26}$$

The recursive update of the target number probability is then given in a form analogous to the single target case (9) as

$$\begin{aligned} p(T^k|Z^k) &= \frac{p(T^k|Z^{k-1})p(z^k|T^k, Z^{k-1})}{p(z^k|Z^{k-1})} \\ &= p(T^k|Z^{k-1}) \int \lambda(z^k|X^k, T^k) \\ &\times p(X^k|T^k, Z^{k-1})dX^k \\ &\times \frac{p(z^k|T^k = 0)}{p(z^k|Z^{k-1})} \end{aligned} \tag{27}$$

where the constant term $p(z^k|T^k = 0)/p(z^k|Z^{k-1})$ does not need to be computed since the probability mass function sums to 1.

In the multiple target situation, it is possible that more than one target projects into the same bearing cell. T targets with states $x_1 \ldots x_T$ will project into bearing cells $\hat{j}_{1,(i)} \cdots \hat{j}_{T,(i)}$. Let $S_{(i)}$ denote the set of all beams the hypothesized targets occupy in sensor i, i.e., $S_{(i)} = \hat{j}_{1,(i)} \cup \hat{j}_{2,(i)} \ldots \cup \hat{j}_{T,(i)}$. Let $O(\hat{j})$ denote the occupation number (i.e., the number of targets that are hypothesized to exist in cell $\hat{j}$), and $p_O(j)$ be the probability density on the measurement received in bearing cell j when there are O targets predicted in to be in that cell. Then we can write

$$\begin{aligned} &p\left(z^k_{(i)}|x^k_1 \cdots x^k_T, H^k_{1,1} \cdots H^k_{1,T}\right) \\ &\quad = \prod_{\hat{j} \in S_{(i)}} \frac{p_{O(\hat{j})}\left(z^k_{(i),\hat{j}}\right)}{p_0\left(z^k_{(i),\hat{j}}\right)} \prod_{j=1}^{\mathcal{J}} p_0\left(z^k_{(i),j}\right). \end{aligned} \tag{28}$$

Using reasoning similar to the single target case, we therefore have

$$\lambda(z^k|X^k, T^k) = \prod_{i=1}^{\mathcal{I}} \prod_{\hat{j} \in S_{(i)}} \frac{p_{O(\hat{j})}\left(z^k_{(i),\hat{j}}\right)}{p_0\left(z^k_{(i),\hat{j}}\right)}. \tag{29}$$

The target state update can then be written using Bayes' rule as

$$\begin{aligned} p(X^k|T^k, Z^k) &= p(X^k|T^k, Z^{k-1})\lambda(z^k|X^k, T^k) \\ &\times \frac{p(z^k|T^k = 0)}{p(z^k|T^k, Z^{k-1})} \end{aligned} \tag{30}$$

where the constant term $p(z^k|T^k = 0)/p(z^k|T^k, Z^{k-1})$ does not need to be computed since the pdf integrates to 1.

VI. Multitarget Implementation

The dimension of the state space required to directly estimate the joint multitarget probability grows exponentially with the number of targets. This "curse of dimensionality" makes it impractical to directly estimate the full multitarget density when there are more than 1 or 2 targets in the surveillance region of interest [34], [41], [42].

This reality is addressed in the literature in two ways. In conventional Kalman multitarget methods [3], [4], the approach is to run a bank of single target trackers (one for each target) and use data association to determine which measurements to give to each tracker. This implicitly assumes the input data is thresholded and can be treated by linear or linearized methods. Other approaches [41]–[46] fully model the joint multitarget density, but use procedures which amount to factoring the joint density into a product of smaller dimensionality densities. Sophisticated approaches do this adaptively at each time step, selecting which targets are "close" and performing more intense processing to account for the coupling of measurements on these targets.

As discussed below, we treat the problem as separate but coupled detection and tracking stages. In the detection stage, we employ a collection of single target fuse-before-track and track-before-detect filters. Once a detector has accumulated enough evidence to initiate a tracker, we use an adaptive factorization method which determines which targets are close in sensor space and treats these pairs (or triplets, etc.) jointly as necessary. In this case, the individual particle filter approximation to the multiple targets is combined into a single multiple target density and this density is time and measurement updated jointly. The main effect is that the sensor modeling of (28) now becomes relevant. In particular, the measurement likelihood for a pair of targets where both targets project into the same beam takes a different functional form from when they project into

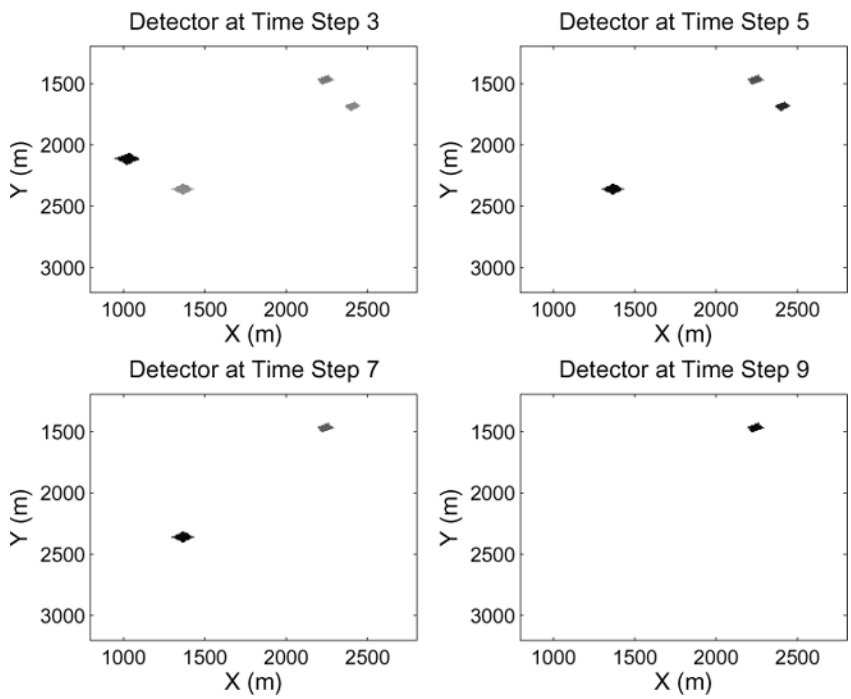

Fig. 1. Detector with four targets. At time Step 3, the detector has four peaks corresponding to the XY location of the four targets. A tracker is initiated about the highest peak, leaving three targets in the detector at time Step 5. This continues until all targets have been initiated by time Step 10.

different beams. In practice, this high fidelity modeling prevents track coalescence onto the stronger track and/or dropping of the weaker track.

A. Detection Stage

The surveillance region is divided up into $\mathcal{D}_r \times \mathcal{D}_c$ static spatial subregions, each of which is called a "detector." Each detector $d = 1 \ldots \mathcal{D}_r \mathcal{D}_c$ is defined by its extent $\left\{\mathbf{x}_d^{min}, \mathbf{x}_d^{max}, \mathbf{y}_d^{min}, \mathbf{y}_d^{max}\right\}$ and is a represented by a single target discrete grid. Detectors thus consist of a probability density for each state x that maps to the detector and a target present probability $p(H_{1,d})$ corresponding to the detector. Each state x in the overall surveillance region maps to a unique detector denoted $\mathfrak{D}_x$.

Each detector is updated temporally and with measurements according to the single target method described earlier. There are competing desires for detector size. On the one hand, each detector should be small to contain either 0 or 1 targets, so the problem exactly breaks down into a collection of single target problems. Conversely, it is also important for each region to be large enough to allow sufficient track-before-detect updates on the target to integrate SNR over multiple time steps before the target moves into a different subregion.

With well separated targets, the factorization

$$p(x_1, x_2 \ldots x_T, T|Z) = \prod_{t=1}^{T} p(x_t, H_{1,\mathfrak{D}_{x_t}}|Z) \tag{31}$$

holds. More critically, the multitarget likelihood also factors in this situation. By assumption, measurements in different beams are independent conditioned on the (multitarget) state, i.e.

$$p(z_{(i)}|x_1 \ldots x_t \ldots x_T) \propto \prod_{t=1}^{T} p(z_{(i),\hat{j}_{t,(i)}}) \tag{32}$$

where $\hat{j}_{t,(i)}$ is again the beam target t maps to in array i. This is valid when the $\hat{j}_{t,(i)}$ are distinct (i.e., the targets are well separated in measurement space and map to different beams). Therefore, with well separated targets

$$p(z_{(i)}|x_i \ldots x_T) \propto \prod_{t=1}^{T} p(z_{(i)}|x_t). \tag{33}$$

B. Tracking Stage

When the detection statistic for detector d, $P(H_{1,d})$ exceeds a threshold a tentative target is declared and a "tracker" is initiated around the target. A tracker is implemented by a particle filter. The tracker follows the tentative target and continues to update the estimate of target state and target present probability. We continue to allow for the possibility new targets arrive in the original detector subregion.

C. Closely Spaced Targets

Although targets are typically well separated, there is the possibility that multiple targets are present in a single detector, or there is spatial overlap between a detector and a tracker, or between trackers. The consequences of this mismatch between assumptions and reality must be explored.

First, each detector assumes there is either 0 or 1 target in its spatial extent. It is possible that more than 1 target is actually present on the detector. In this case, the likelihood ratio (8) will become peaked at multiple points, rather than at a single point. The detector pdf will correspondingly peak at multiple points. This leads to the target present hypothesis computed by (9) to grow faster than if a single target was present in the detector. In this situation, the detector's target present hypothesis crosses a threshold and a tracker is first initiated around the largest peak by sampling an area spatially close to that peak. The detector's

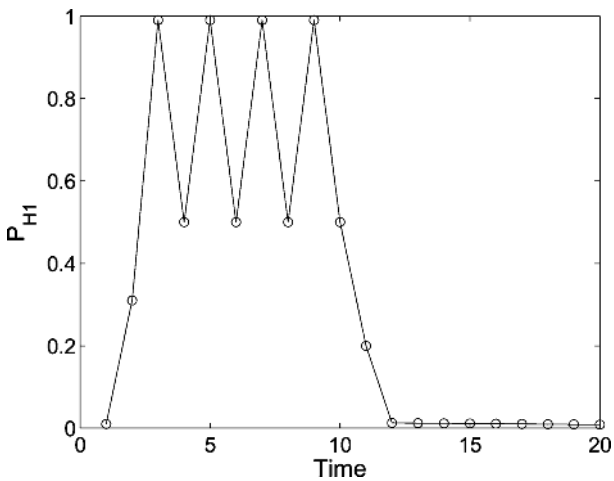

Fig. 2. The detection statistic verbs time for the example of Fig. 1. Targets are initiated at time step 3, 5, 7, and 9.

target present hypothesis again exceeds the threshold at a subsequent time step and a second target is initiated around the second peak, and so on. Figs. 1 and 2 illustrate what happens when there are four targets in a single detector using synthetic data. First, the detector pdf develops four peaks. In our implementation, a tracker is launched first around the strongest peak. The target existence probability remains high even after this target is initiated since there are multiple targets in that detector. Therefore, shortly later a second target is initiated, and so on. As explained below, if targets are very close together, there may be some delay in initiating the second target since we prevent measurement sharing between trackers and detectors. However, the approach typically performs well in practice.

Second, since a mobile tracker moves its position in x and y to follow a target, there may be spatial overlap between the fixed detectors and the mobile trackers. In this case, our approach is to prevent both a detector and a colocated tracker from simultaneously updating their single target pdf with the same measurements. This is done by determining when entities overlap, and modifying the likelihood ratio computation for the detector to omit the measurements made in the overlapping region. A target state x^k is said to be included in a tracker if the convex hull defined by the tracker's particles includes x^k. With this, we then define the indicator function

$$I_r(x^k) = \begin{cases} 1, & x^k \text{ maps to tracker } r \\ 0, & \text{otherwise} \end{cases} \tag{34}$$

and construct the logical disjunction over the R trackers as

$$I(x^k) = \bigvee_{r=1}^{R} I_r(x^k) \tag{35}$$

which is 1 if x^k maps to any tracker. Then the modified detector likelihood, accounting for closely spaced targets, is

$$\lambda'\left(z^k|x^k, H_1^k\right) = (1 - I(x^k))\lambda\left(z^k|x^k, H_1^k\right) + I(x^k). \tag{36}$$

Finally, two trackers may be close together in measurement space. In this case, the pdfs must be updated jointly. As discussed earlier, we compute the separation between trackers and cluster them into groups to determine which trackers must be updated jointly and which may be updated independently. The clustering is performed as follows. Each track is characterized by its state estimate, projected into each sensor's bearing beam. Those tracks that are close in beamspace are clustered together and treated jointly. In the ideal situation where all targets are well separated, each cluster contains one target. In the case where multiple targets are close, a cluster may contain 2, 3, or more targets. In this instance, we form the joint multitarget density on the cluster (in 8 or 12 dimensions as appropriate) and treat the multiple targets as a single unit for temporal and measurement update. For computational purposes, we use the coupled partition method of [42] to propose and weight new joint particles. If tracks are incorrectly clustered together when they, in fact, correspond to well separated targets the algorithm runs slower than it otherwise would. However, the density estimation is still done correctly. Conversely, if tracks are incorrectly deemed separate when they are in fact close, the coupling is not correctly preserved. Therefore, we choose this threshold conservatively to err on the side of over clustering.

VII. Results

This section illustrates the proposed technique on simulated and real collected passive acoustic data. Section VII-A starts with a synthetic experiment which illustrates the performance of the algorithm by directly comparing it to (simulated) truth. Section VII-B shows results of the tracking approach when applied to a set of real collected data.

A. Results on Synthetic Data

The following simulation illustrates the utility of the approach. Three targets were simulated to move according to the NCV model of Section IV. The surveillance region (shown in Fig. 3) covered $x = -1200 \cdots 1200$ and $y = 500 \cdots 2900$. The initial location of the three targets is (900,1000), (1200,500), and $(-1000,1700)$. The sensor arrays are located at (0,0) and (700,0). All units are nominally in meters.

The surveillance region was implemented with a 4×4 set of overlapping detector grids. Overlapping grids are used to allow quicker detection of targets that start close to the edge of one detector and then move onto another.

Synthetic input data was created as follows. Simulated bearing time records (BTRs), which are profiles of the quantized target bearing versus time, were created according to the model of Section II-B. The BTR surfaces, shown in Fig. 4, are the input to the tracker. The simulation used 1 degree bearing beams. The sensor model we employ to describe how the measured data couples to the target state under estimation is the nonlinear pixelated model described in (14).

The method described in this paper was evaluated using this input data. Each detector grid was chosen to have $26 \times 21 \times 26 \times 21$ cells with spatial and velocity resolutions of 25 m and 1 m/s, respectively. Each single target detector was measurement and time updated according to the method of Section IV. When a detectors' estimate of the target present probability $p(H_1|Z)$ exceeds the threshold (here set at $\tau = 0.8$), the target is transitioned to a particle filter representation according to the description in Section IV-C.

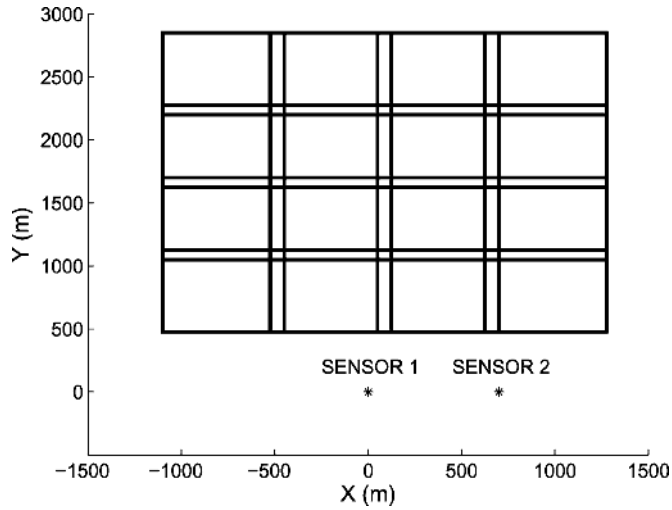

Fig. 3. Surveillance region for the synthetic experiment.

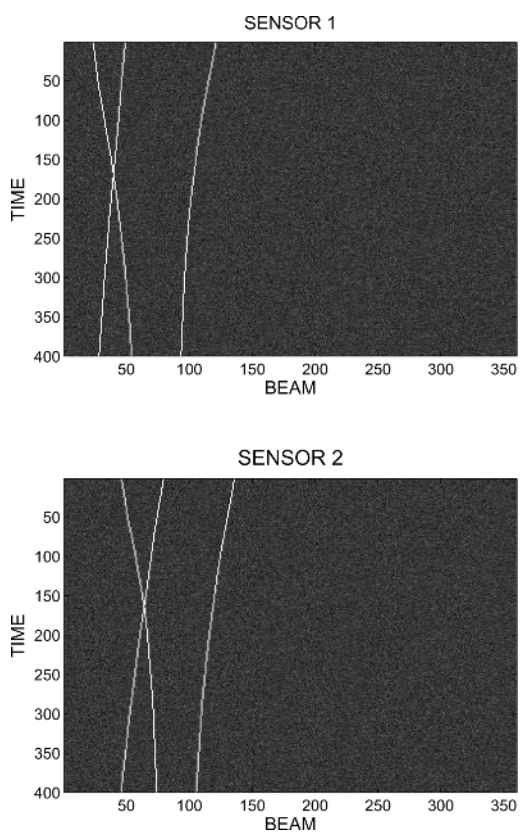

Fig. 4. Synthetic BTRS from two sensors with three targets.

Detector thresholds were crossed (resulting in a target transition from fixed grid to particle filter) at time steps 4, 6, and 9. Target initializations at time step 4 and 9 came from the same detector. Two targets were close together for a portion of the tracking and were clustered and treated jointly by the algorithm between time steps 146 and 188.

At each time step, this particle filter representation of the probability density was used to compute the minimum mean square error estimate of target state. Fig. 5 shows the position component of this state estimate over time in comparison to the truth.

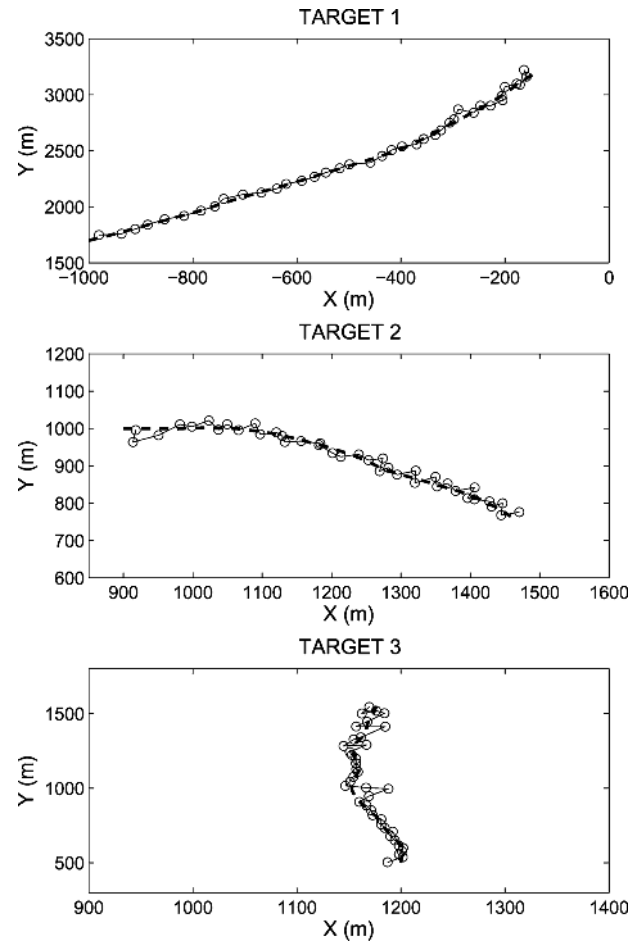

Fig. 5. XY point estimates (circles) produced by the method, as compared to truth (dashed line) for the three targets in this simulated scenario. Tracker point estimates are only shown every 10 s for clarity.

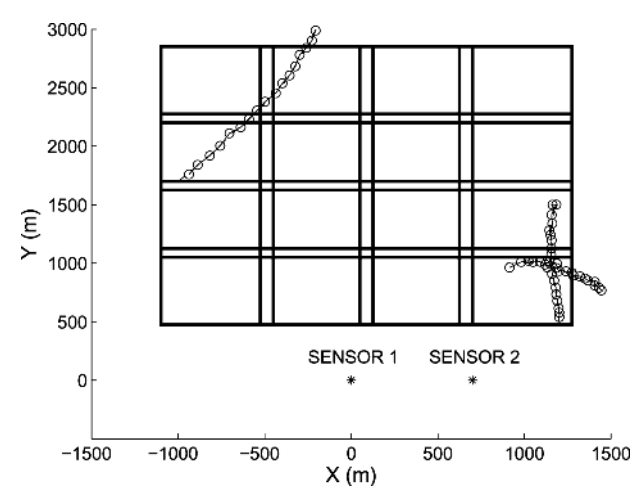

Fig. 6. Surveillance region for the synthetic experiment, with true target tracks (dashed lines) and estimated target tracks (circles) drawn in.

The tracker estimates and target truth are shown with respect to each other, the surveillance region, and the sensors in Fig. 6.

Fig. 7 shows the estimated tracks projected back on to the BTRs. Note the tracker operates in XY but these XY tracks can be reprojected onto the input surfaces to show algorithm efficacy, as is done here.

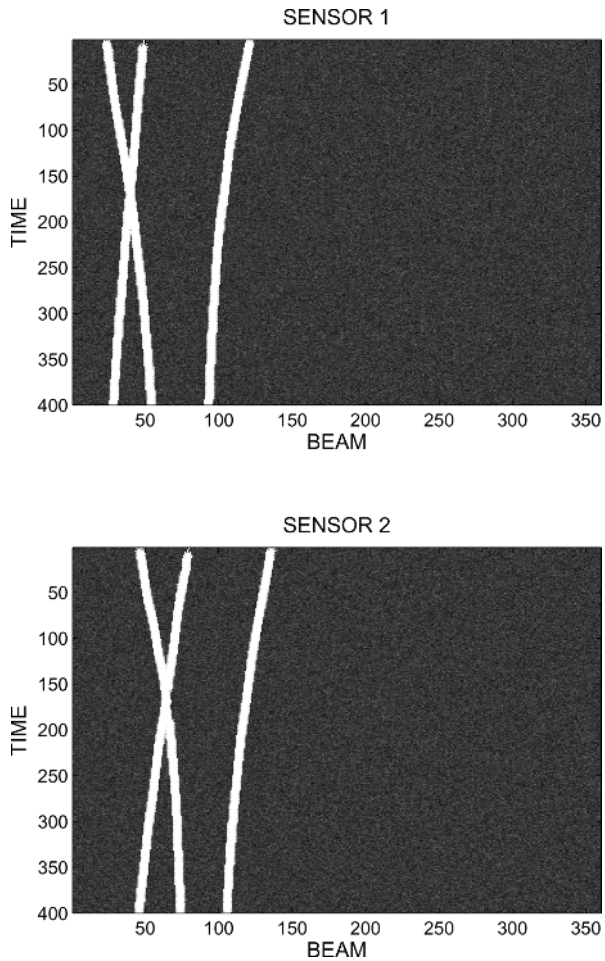

Fig. 7. XY point estimates from the tracker projected back on to the BTRs.

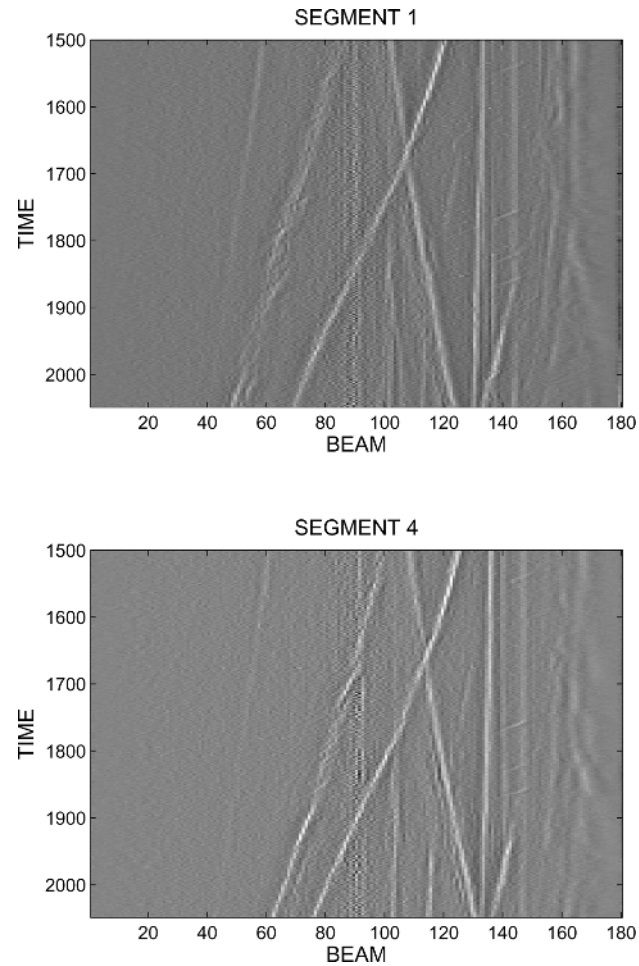

Fig. 8. Input data surfaces (BTRs) that were used to evaluate the method. Top: Segment 1. Bottom: Segment 4.

B. Results on the Shallow Water Array Performance Array

This subsection illustrates the proposed technique on a set of real collected passive acoustic data. The shallow water array performance (SWAP) array is located off the eastern cost of Florida near Ft. Lauderdale. The SWAP array has four linear segments, each of which are approximately 200 m long and contain 125 hydrophones. The hydrophone locations are known with high accuracy. A small number of the hydrophones did not operate during the collection, but these elements are known have been excluded from the beamforming process. The 4 segments are labeled 1, 2, 3, and 4 with segment 1 closest to shore (westernmost) and segment 4 farthest from shore (easternmost). The elements in each segment run approximately west-to-east. Segment 1 is oriented approximately 0.6 ° (relative to east), and Segment 4 is oriented approximately $-0.4°$. All hydrophones are approximately 265 m deep.

The environment has heavy commercial and recreational traffic. The experiments shown here use data collected by segments 1 and 4 on August 9, 2007 starting at 1115 local time. The data was recorded as part of a four-day sea test. We selected a time segment with two targets of opportunity. This data includes a challenging situation where the two targets cross in sensor space, therefore the adaptive algorithm must temporarily treat the joint target state to prevent target coalescence or removal.

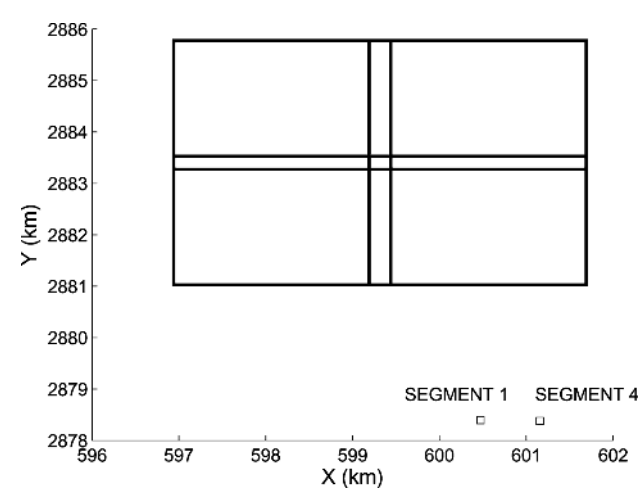

Fig. 9. Surveillance region, and the 2×2 grid of single-target detectors used to tile the region. The four detectors are each $51 \times 13 \times 51 \times 13$ grids updated using the methods of Section IV.

The raw sensor data was prepared according to the process described in Section II. This is briefly summarized as follows. Each hydrophone sampled and recorded raw acoustic data. This raw data was decoded and interpolated to obtain 1000 samples per second at equal intervals, time synchronized at all hy-

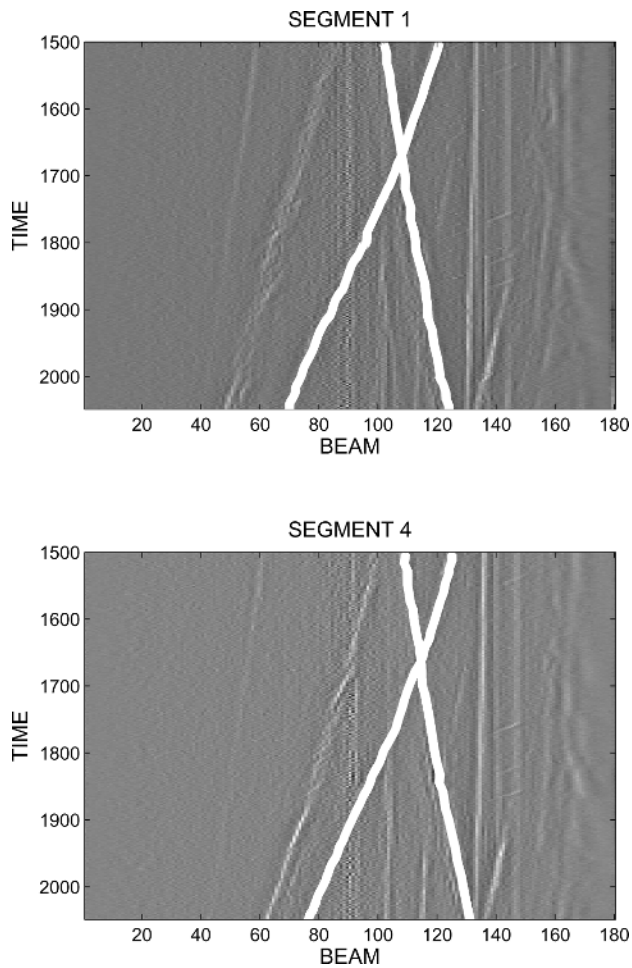

Fig. 10. Input data surfaces, with hand-truthed traces in each segment.

drophones. A conventional beamformer, followed by integration over frequency, was used to produce BTRs at 1 Hz. The sensor model we employ to describe how the measured data couples to the target state under estimation is again the nonlinear pixelated model described in (14). The target present and absent distributions p_1 and p_0 are assumed Gaussian by appealing to the analysis of [28]. The target present probability threshold used to elevate a target to detected status was fixed at $\tau = 0.8$. Fig. 8 shows the input data surfaces.

The surveillance region is defined as shown in Fig. 9. It was implemented using a 2×2 set of overlapping detector grids. Notice that the finite spatial extent of surveillance region means that only that portion of the BTR data corresponding to the surveillance region will be used. For example, measurements in bearing cells directly to the left of the segments do not map to the surveillance region and are thus not used in updating the system at onset.

Each detector grid has $51 \times 13 \times 51 \times 13$ cells with spatial and velocity resolution of 50 m and 2 m/s, respectively. Each single target detector is measurement and time updated according to the method of Section IV. When a detectors' estimate of $p(H_1|Z)$ exceeds the threshold, the target is transitioned to a particle filter representation according to the description in Section IV-C.

Here we illustrate the benefit offered by our target space tracking method (XY) over traditional measurement domain methods, which track on the individual data surfaces and then fuse the results (the track-fusion approach). The track-fusion approach proceeds as follows. First, a measurement domain tracker is used on each BTR to generate bearing estimates. Kinematics are enforced on the data surface, rather than the XY space in which the target actually operates. Next, the tracks are associated between sensors and the targets' XY position are computed using trigonometry. We have simulated that method here by hand-truthing the traces in the BTR as shown in Fig. 10. Note that in the hand-truthing, the beams of the individual traces are known exactly and correctly.

Whether done by hand or with an automated tracker that operates independently at each node, this approach does not take into account kinematic models operating in XY space, but instead only enforces kinematics on the bearings. As a consequence, the XY tracks associated with this hand truthing may be nonphysical in the XY domain. Fig. 11 shows the noisy tracks that result from tracking in bearing and then associating the tracks, characterized by nonphysical jumps in X and Y. Note that the figure also indicates the locations of the two sensor arrays near the bottom center, with labels "SEGMENT 1" and "SEGMENT 4," respectively. The error seen in the target XY estimate is due entirely to the fact that the estimate of the target bearing is the center of the (correct) beam, but it does not have subbeam accuracy. If, for example the target is actually in the center of both beams, the tracker estimate of target XY is perfect. If, on the other hand, it is at the edge of the beam, the error is larger. The wild oscillations are a manifestation of the target moving through a beam and having its bearing poorly estimated, then estimated well, and then poorly estimated again.

In contrast, the tracker proposed here constructs a pdf on the region by combining XY kinematic models with the measurements. Although the input data has the same resolution as the track fusion approach, the performance is much better because the tracker enforces plausible kinematics, and this provides additional information which ensures that the tracks are smooth in XY.

The complete nonparametric target pdf it estimates can be used to produce both point estimates and covariance ellipses, although the internal representation is much richer than this. Fig. 12 shows the tracker point estimates of target XY in this scenario. Note that these targets were automatically detected and initiated according to the methodology of Section III. The trackers are deemed closely spaced from time steps 1639 to 1693 and thus treated jointly by the method during that time period.

Since this was a controlled experiment, we also have latitude and longitude truth sources from the automatic identification system (AIS) for some of the contacts in the collection. One of the targets in our surveillance region during the period of interest was an AIS equipped vessel, and its truth track (dashed gray line) is shown in the plot for comparison to the tracker output. The system runs at about 1.5 s per update on a standard off the shelf 2.83 GHz Linux machine running MatLab™.

As can be seen by comparing the track fusion results (see Fig. 11) with the output of the tracker proposed here (see

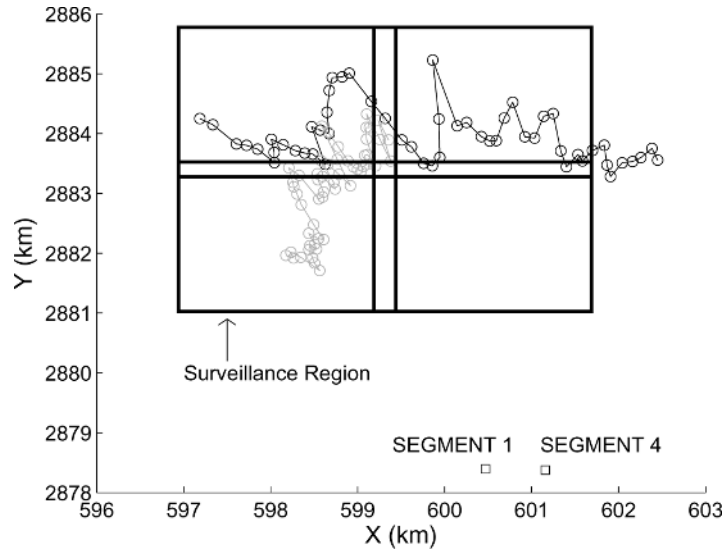

Fig. 11. Nonphysical XY tracks associated with the hand-truthed bearing. Because tracking is done on the data surfaces, the XY tracks are not guaranteed to be physically plausible.

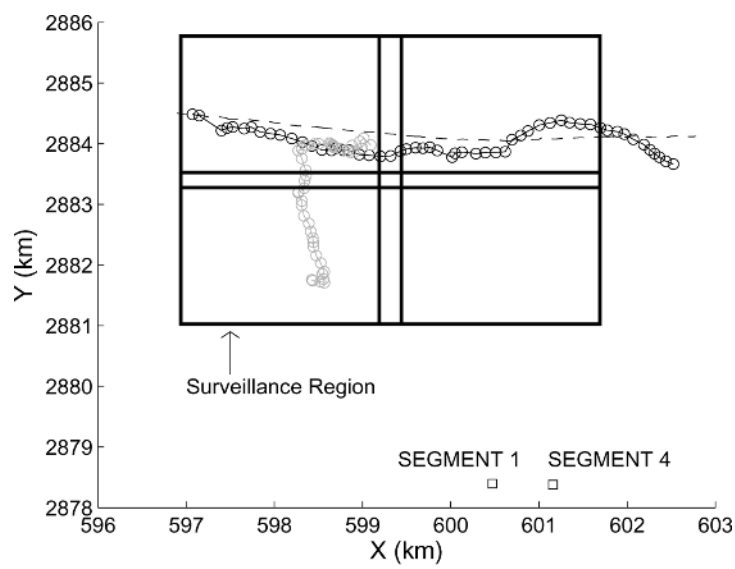

Fig. 12. Tracker target state point estimates. Two targets were automatically detected and tracked. A known truthed target was in the region (dashed). Notice that these track estimates are much more physically plausible then the track-fusion tracks.

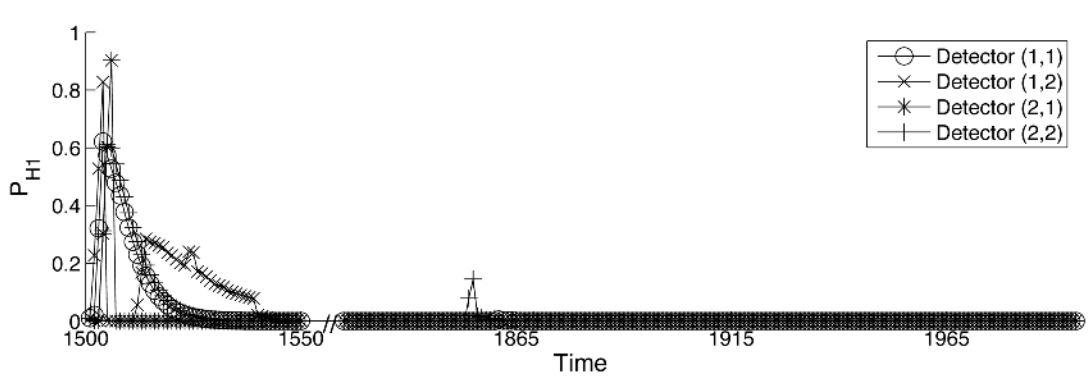

Fig. 13. Detection statistic calculated by each of the four detectors. Two targets are initiated near the beginning of the trial. The detection statistic never rises above the threshold again.

Fig. 12), utilizing physical kinematic models in the target's natural coordinate system provides significant value.

Further detail on the tracker performance is as follows. Two tracks are initiated automatically by the 2×2 detector grid. Fig. 13 shows the detection statistic $p\left(H_1^k|Z^k\right)$ (9) for each of the four detectors over time. Notice that at the beginning of the test, two detectors have detection statistic that grows past the threshold. Trackers are initiated from these detectors. For the rest of the experiment, there are no detectors that have detection statistic which grows past the threshold. There is a brief period

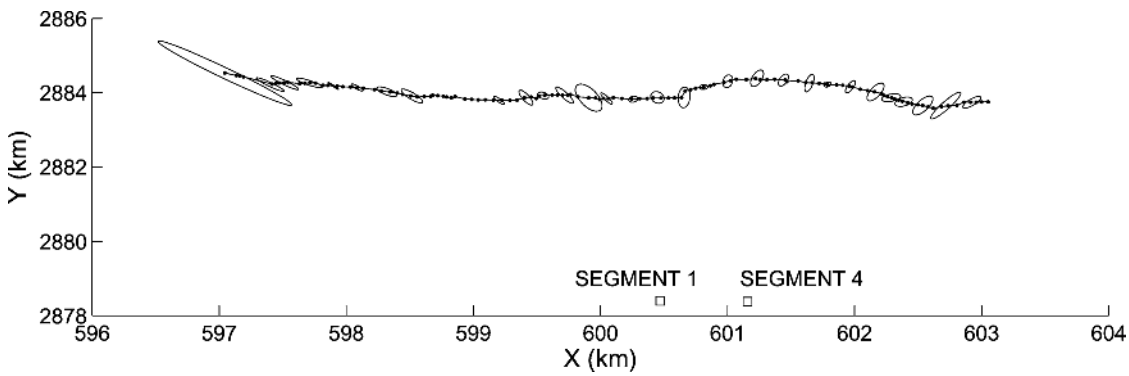

Fig. 14. One track, with tracker estimated covariance ellipses plotted every 16 s.

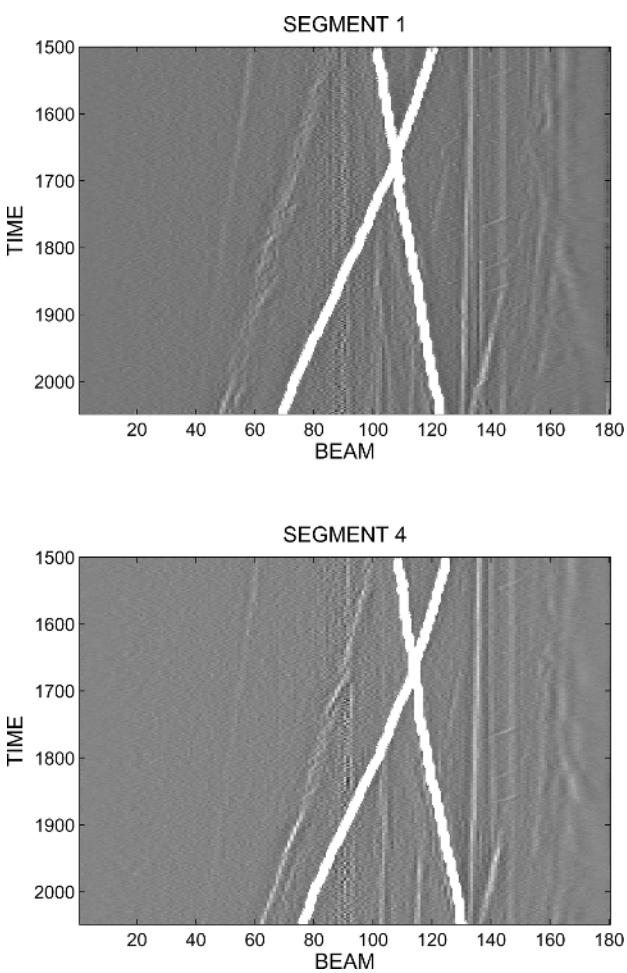

Fig. 15. Input data surfaces, with tracker estimates projected into measurement space.

where one detector begins to accumulate evidence of a target but it never reaches the decision threshold.

Fig. 14 shows one of the tracks and the associated covariance ellipses at equally spaced time steps. As can be seen from the figure, the main axis of the uncertainty ellipse is in the range direction with respect to the (bearings-only) sensors. For example, when the target is northwest of the sensors, uncertainty is predominately in the north-west direction. Likewise, when the target is northeast of the sensors, uncertainty is oriented in the north-east direction.

Fig. 15 shows the tracker estimates projected back onto the two original input surfaces. Some contacts visible in the BTRs do not have associated tracks. As discussed earlier, this is because the XY locations corresponding to the bearing traces are outside the specified surveillance region and thus do not map to any detector and are therefore not initiated. The surveillance region size is a choice made by the operator.

VIII. Conclusion

This paper has described a Bayesian approach to multiple target detection and tracking using multiple passive acoustic arrays. It included the development of a novel method of approximating the required multitarget probability density which combines a fixed grid (discrete) method for target detection and an adaptive grid (particle) scheme for tracking and a dynamic factorization method to deal with closely spaced targets. The method was illustrated on real collected multinode passive acoustic data. The method was shown to provide a benefit over track fusion, in terms of the smoothness of the track estimate.

References

[1] M. Morelande and B. Ristic, "Signal-to-noise ratio threshold effect in track before detect," *IET Radar, Sonar, Navig.*, no. 6, pp. 601–608, 2009.

[2] Y. Boers, F. Ehlers, W. Koch, L. D. S. T. Luginbuhl, and R. L. Streit, "Special issue on track before detect algorithms," *EURASIP J. Adv. Signal Process.*, 2008.

[3] S. Blackman, *Multiple-Target Tracking With Radar Applications*. Norwood, MA: Artech House, 1986.

[4] Y. Bar-Shalom, *Multitarget Multisensor Tracking: Advanced Applications*. : Artech House, 1990.

[5] D. Musicki, "Multi-target tracking using multiple passive bearings-only asynchronous sensors," *IEEE Trans. Aerosp. Electron. Syst,*, vol. 44, pp. 1151–1160, 2008.

[6] T. Fortman and Y. Bar-Shalom, "Sonar tracking of multiple targets using JPDA," *IEEE J. Ocean. Eng.*, vol. OE-8, no. 3, pp. 173–184, Jul. 1983.

[7] F. El-Hawary, F. Aminzadeh, and G. Mbamalu, "The generalized kalman filter approach to adaptive underwater target tracking," *IEEE J. Ocean. Eng.*, vol. OE-17, no. 1, pp. 129–137, Jan. 1992.

[8] K. Chang, T. Zhi, and R. Saha, "Performance evaluation of track fusion with information matrix filter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 38, no. 2, pp. 455–466, Apr. 2002.

[9] Y. Bar-Shalom, "On the track-to-track correlation problem," *IEEE Trans. Autom. Control*, vol. 26, no. 2, pp. 571–572, Apr. 1981.

[10] Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association*. New York: Academic , 1988.

[11] S. Coraluppi and C. Carthel, "Recursive track fusion for multisensor surveillance," *Inform. Fus.*, vol. 5, no. 1, pp. 23–33, Mar. 2004.

[12] C. A. Barlow, L. D. Stone, and M. V. Finn, "Unified data fusion," in *Proc. 9th Nat. Symp. Sensor Fus.*, Monterey, CA, Mar. 1996.

[13] L. Stone, C. Barlow, and T. Corwin, *Bayesian Multiple Target Tracking*. Boston, MA: Artech House, 1999.

[14] K. Kastella, "Joint multitarget probabilities for detection and tracking," in *Proc. SPIE Acquisition, Track. Point. XI*, Orlando, FL, 1997, pp. 122–128.

[15] Y. Boers and H. Driessen, "Multitarget particle filter track before detect application," *IEE Proc. Radar, Sonar, Navig.*, vol. 151, pp. 351–357, Dec. 2004.

[16] R. Bethel and G. Paras, "A PDF multitarget tracker," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 30, no. 2, pp. 386–403, Apr. 1994.

[17] R. Bethel and G. Paras, "A PDF multisensor multitarget tracker," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 34, pp. 153–168, Jan. 1998.

[18] A. Srivastava, M. Miller, and U. Grenander, "Jump-diffusion processes for object tracking and direction finding," in *Proc. 29th Allerton Conf. Commun., Contr., Comput.*, Urbana, IL, 1991, pp. 563–570.

[19] E. Kamen, "Multiple target tracking based on symmetric measurement functions," *IEEE Trans. Autom. Control*, vol. 37, no. 3, pp. 371–374, Mar. 1992.

[20] R. Mahler, "A unified foundation for data fusion," in *Proc. 7th Joint Serv. Data Fus. Symp.*, Baltimore, MD, 1994, pp. 154–174.

[21] K. Kastella, "Event-averaged maximum likelihood estimation and mean-field theory in multitarget tracking," *IEEE Trans. Autom. Control*, vol. AC-40, no. 6, pp. 1070–1074, Jun. 1995.

[22] J. Vermaak, S. Maskell, and M. Briers, "A unifying framework for multi-target tracking and existence," in *Proc. 2005 Int. Conf. Inform. Fus.*, Philadelphia, PA, Jul. 2005.

[23] M. Orton and W. Fitzgerald, "A Bayesian approach to tracking multiple targets using sensor arrays and particle filters," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 216–223, Feb. 2002.

[24] A. Doucet, B. Vo, C. Andrieu, and M. Davy, "Particle filtering for multi-target tracking and sensor management," in *Proc. IEEE Int. Conf. Inform. Fus.*, Annapolis, MD, 2002, pp. 474–481.

[25] F. Maurelli, S. Krupinski, Y. Petillot, and J. Salvi, "A particle filter approach for auv localization," in *Proc. OCEANS 2008*, QC, Canada, Sep. 2008, pp. 1–7.

[26] B. Shapo and R. Bethel, "An overview of the probability density function (PDF) tracker," in *Proc. Oceans 2006*, Boston, MA, Sep. 2006.

[27] D. Johnson and D. Dudgeon, *Array Signal Processing*. : Simon and Schuster, 1992.

[28] H. Van Trees, *Detection, Estimation, and Modulation Theory IV: Optimum Array Processing*. New York: Wiley, 2002.

[29] K. Kastella and C. Kreucher, "Multiple model nonlinear filtering for low signal ground target applications," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 2, pp. 549–564, Apr. 2005.

[30] J. C. Strikwerda, *Finite Difference Schemes and Partial Differential Equations*. New York: Chapman and Hall, 1989.

[31] H. J. Kushner, *Probability Methods for Approximations in Stochastic Control and for Elliptic Equations*. New York: Academic, 1977.

[32] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, Feb. 2002.

[33] M. Orton and W. Fitzgerald, "A Bayesian approach to tracking multiple targets using sensor arrays and particle filters," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 216–223, Apr. 2002.

[34] F. Daum and J. Huang, "The mysterious computational complexity of particle filters," in *Proc. SPIE Conf. Signal Data Process. Small Targets*, Orlando, FL, 2002, pp. 418–426.

[35] H. Kushner and P. Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*. Berlin, Germany: Springer-Verlag, 1991.

[36] Y. Bar-Shalom, X. Li, and T. Kirubarajan, *Estimation With Applications to Tracking and Navigation*. New York: Wiley, 2001.

[37] Z. Tang and U. Ozguner, "Sensor fusion for target track maintenance with multiple uavs based on Bayesian filtering method and hospitability map," in *Proc. 42nd IEEE Conf. Decision Contr.*, Maui, HI, Dec. 2003, pp. 19–24.

[38] M. McDonald and B. Balaji, "Track-before-detect using swerling 0, 1, and 3 target models for small manoeuvring maritime targets," *EURASIP J. Adv. Signal Proc.*, vol. 2008, 2008.

[39] C. Kreucher, B. Shapo, and R. Bethel, "Multitarget detection and tracking using multi-sensor passive acoustic data," in *Proc. 13th Annu. IEEE Aerosp. Conf.*, Big Sky, MT, Mar. 2009, pp. 1–16.

[40] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Berlin, Germany: Springer-Verlag, 2001.

[41] J. Vermaak, S. Godsill, and P. Perez, "Monte carlo filtering for multitarget tracking and data association," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 1, pp. 309–332, Jan. 2005.

[42] C. Kreucher, K. Kastella, and A. Hero, "Multitarget tracking using the joint multitarget probability density," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 10, pp. 1396–1414, Oct. 2005.

[43] T. Vercauteren, D. Guo, and X. Wang, "Joint multiple target tracking and classification in collaborative sensor networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 4, pp. 714–723, Apr. 2005.

[44] I. Kyriakides, D. Morrell, and A. Papandreou-Suppappola, "Sequential monte carlo methods for tracking multiple targets with deterministic and stochastic constraints," *IEEE Trans. Signal Process.*, vol. 56, no. 3, pp. 937–948, Mar. 2008.

[45] V. Cevher, F. Shah, R. Velmurugan, and J. McClellan, "A multi target bearing tracking system using random sampling consensus," in *Proc. 2007 Aerosp. Conf.*, Big Sky, MT, Mar. 2007, pp. 1–15.

[46] M. Morelande, C. Kreucher, and K. Kastella, "A Bayesian approach to multiple target detection and tracking," *IEEE Trans. Signal Process.*, vol. 5, no. 5, pp. 1589–1604, May 2007.

Chris Kreucher received the B.Sc., M.Sc., and Ph.D. degrees in electrical engineering from the University of Michigan, Ann Arbor, in 1997, 1998, and 2005, respectively.

He is currently a Principal System Engineer at Integrity Applications Incorporated, Ann Arbor, MI. From 1998 to 2008, he was a Staff Scientist at General Dynamics Advanced Information Systems at Michigan Research and Development Facility (formerly ERIM). His current research interests include nonlinear filtering, Bayesian multitarget tracking, and information theoretic methods for optimal sensor utilization.

Ben Shapo received the B.S.E. degree from Princeton University, Princeton, NJ, in 1986, the M.Sc. degree from the University of Illinois, Urbana, in 1989, and the Ph.D. degree from the University of Michigan, Ann Arbor, in 1996.

He has been a Principal System Engineer at Integrity Applications Incorporated, Ann Arbor, MI, since 2008. His research interests include Bayesian detection and tracking, and data fusion.