Random Set Tracker Experiment on a Road Constrained Network with Resource Management

James Witkoskie, Walter Kuklinski, Stephen Theophanis, and Michael Otero
MITRE Corporation
M/S E050
202 Burlington Road
Bedford, MA 01730-1420 U.S.A.
jwitkoskie@mitre.org

Abstract – This paper describes the application of finite set statistics (FISST) to real-time multiple target road constrained tracking problems. We studied specific test problems where multiple modality wireless sensor networks monitored road networks of interest. Acoustic and radar detections updated a global density that tracked the number and positions of targets. The global density determines "information states" that form the basis of a closed-loop Markov Decision Process resource management procedure that controls sensor operation.

Keywords: Tracking, data association, estimation, resource allocation, netted-sensors.

1 Introduction

Recent interest in netted-sensors technology is partially motivated by the potential of obtaining information about targets and their environments on length and time scales that are not easily accessed by traditional standoff sensors. Obtaining this information requires efficient methods to fuse measurements from multiple in situ sensors, often operating over a wide range of sensing modalities. Among applications of interest are dense road networks, where target state vectors can change rapidly due to target maneuvers at intersections [1]. Wireless sensor networks with in situ sensors at intersections would be advantageous in these tracking applications.

Traditional Kalman Filter target tracking methods have difficultly tracking targets in dense road networks, where several roads are in a sensor's field of view. Measurements can correspond to several different roads and lead to an association ambiguity between the measurement and the target position [2]. Random set approaches avoid the association ambiguity by statistically weighing all possible hypotheses and associations [3,4].

While a large number of publications describe random set tracker performance in simulation studies, only a limited number of real-time implementations using field data have been reported [5]. In this paper, we use a random set tracker to fuse data from a wireless network of in situ sensors deployed within road networks of interest. An additional consideration is the desire to minimize sensor energy expenditure and correspondingly maximize network operational lifetime [6]. Energy conservation is enhanced through statistically based resource management algorithms that determine an optimal subset of sensor measurements to achieve monitoring objectives by combining the tracker's knowledge about target locations with sensor characteristics, including energy consumption.

This paper is organized as follows: Section 2 is an overview of the applications of the FISST to target tracking in dense road networks. Section 3 discusses the scenarios examined in this paper. This discussion includes sensor modalities and the geometry of the road networks. Section 4 outlines the motion models and update equations for the global density fusion algorithm while Section 5 gives numerical approximations to the algorithms. Section 6 discusses deriving "information states" from the global density, which forms the basis of a closed-loop Markov Decision Process resource management procedure to control sensor operation. Section 7 discusses results from simulations and live scenarios before concluding in Section 8.

2 FISST

Traditional Kalman filters assume that state variables, $\mathbf{x}$, and measurements, $\mathbf{y}$, are fixed length random vectors [3]. For traditional tracking applications, $\mathbf{x}$ represents the targets' geokinetic variables, and $\mathbf{y}$ represents the measurements related to the geokinetic variables. The state vector motion is linear with Gaussian white noise, $\mathbf{x}_{t+1} = \mathbf{\Omega}\mathbf{x}_t + d\boldsymbol{\lambda}_t$. Additionally, the measurements depend linearly on the state vectors with additive Gaussian white noise, $\mathbf{y}_t = \mathbf{W}\mathbf{x}_t + d\boldsymbol{\delta}_t$. The Kalman filter can be extended to non-Gaussian noise, non-linear measurements, and non-linear motion models through the Bayesian filter [3]. The Bayesian filter consists of two steps. 1) Starting with a probability density conditioned on previous measurements, $f(\mathbf{x}_{t-1} \mid \mathbf{y}_{1:t-1})$, a prediction step estimates the probability density of the vector at time step t,

$$f(\mathbf{x}_t \mid \mathbf{y}_{1:t-1}) = \int d\mathbf{x}_{t-1} f(\mathbf{x}_t \mid \mathbf{x}_{t-1}) \times f(\mathbf{x}_{t-1} \mid \mathbf{y}_{1:t-1}). \quad (1)$$

This new probability density is combined with a new set of measurements, $\mathbf{y}_t$, during the update step to determine the new estimate of the probability density,

$$f(\mathbf{x}_t \mid \mathbf{y}_{1:t}) = N^{-1} f(\mathbf{y}_t \mid \mathbf{x}_t) \times f(\mathbf{x}_t \mid \mathbf{y}_{1:t-1}), \quad (2)$$

where N is a normalization factor. The number of state variables and measurements is fixed and the mapping between the state variables and measurements is explicit.

Tracking multiple targets presents difficulties because association of measurements with existing targets is often ambiguous [3, 4]. Even if the state variables are linearly related to measurements with Gaussian errors, the ambiguity in associations produces non-Gaussian effects [3]. Missed detections, false alarms or clutter, and the birth and death of targets complicate the scenario. In addition to their numerical values, the number of targets and measurements are also random variables, and the vectors $\mathbf{x}$ and $\mathbf{y}$ must be replaced with random sets $\{x\}$ and $\{y\}$. The geokinetic state, $\{x\}$, may take on the values $\{x\} = \{\phi\}, \{x^1, v^1\}, \{x^1, v^1, x^2, v^2\}, \ldots$ where x^i and v^i denote the positions and velocities of target i. The set of measurements, $\{y\}$, are estimates of the geokinetic variables recorded from various sensors and include clutter returns and missed detections.

Several methods address variable numbers of targets and detections with ambiguous associations, including joint integrated probabilistic data association (JIPDA) and jump Markov models (JMM) [7, 8]. In this paper, we explore applications of finite set statistics (FISST) to road constrained multiple target tracking. FISST is a generalization of the Bayesian equations, Eq. (1) and (2), to sets. The probability density, called the global density, is defined on the possible number and locations of targets. For road networks, the global density has the form

$$f_g = \begin{cases} f_0 \\ f_r(x,v) \overset{\Delta}{=} f_r(\{x^{(1)}\}) \\ f_{r_1 r_2}(x^1, v^1, x^2, v^2) \overset{\Delta}{=} f_{r_1 r_2}(\{x^{(2)}\}) \\ \vdots \end{cases}, \quad (3)$$

where f_0 is the probability of no targets, $f_{r_1 \ldots r_n}$ is the probability density for n targets on roads $r_1 \ldots r_n$ at $x^1, v^2, \ldots, x^n, v^n = \{x^{(n)}\}$, respectively. The density is normalized with respect to a set integral

$$\sum_{n=0}^{\infty} \frac{1}{n!} \int f_{r_1 \ldots r_n}(x^1, v^1, \ldots, x^n, v^n) dx^1 dv^1 \ldots dx^n dv^n = \sum_{n=0}^{\infty} \frac{1}{n!} \int f_{r_1 \ldots r_n}(\{x^{(n)}\}) d\{x^{(n)}\} = 1, \quad (4)$$

where the summation is over the number and identity of possible roads. Similar to the Bayesian filter one can define conditional probabilities, motion models, and measurement models to develop a set of recursive update equations [3, 4]. The prediction step includes propagation and birth and death processes so that the conditional expectation of the global density on previous measurements has the form

$$f_{r_1 \ldots r_n}(\{x_t^{(n)}\} \mid \{y_{1:t-1}\}) = \sum_{m=0}^{\infty} \frac{1}{m!} \int d\{x_{t-1}\} f_{r_1 \ldots r_n \mid r_1' \ldots r_m'}(\{x_t^{(n)}\} \mid \{x_{t-1}^{(m)}\}) \times f_{r_1' \ldots r_m'}(\{x_{t-1}^{(m)}\} \mid \{y_{1:t-1}\}) \quad (5)$$

where the sum is over all possible roads, $r_1' \ldots r_m'$. The predicted density is then updated with the measurements at time step t,

$$f_{r_1 \ldots r_n}(\{x_t^{(n)}\} \mid \{y_{1:t}\}) = N^{-1} \times f_{r_1 \ldots r_n}(\{y_t^{(n)}\} \mid \{x_t^{(n)}\}) \times f_{r_1 \ldots r_n}(\{x_t^{(n)}\} \mid \{y_{1:t-1}\}). \quad (6)$$

Summation over the associations makes these expressions more complicated than Eq. (1) and (2). For calculations, we assume $f_{r_1 \ldots r_n} = 0$ for $n > N$, which places a cutoff in the possible number of targets. The exact forms of the expressions for the sensors used in this experiment will be outlined in Section 3 after introducing the particular road network scenarios.

3 Road Network Scenarios

The road network geometries depicted in Fig. (1) and (2) were examined using simulation and live experiments. Each road network is decomposed into six road segments. Sensors probing the road networks include monostatic range radars (red circles) and multiple element acoustic arrays (green circles). The performances of these sensors are outlined in Section 3.1 and 3.2. The targets were commercial motor vehicles instrumented with GPS receivers allowing comparison of the tracker's predicted vehicle locations with ground truth.

3.1 Range Radars

The range radars were built by Multispectral Solutions, Inc. (Germantown, MD) and have a detection range of 512 ft (156 m) and a 24 degree beam pattern. Detections are recorded by determining background clutter statistics and detecting statistically significant blocks of range cells that exceed the background clutter. The detection

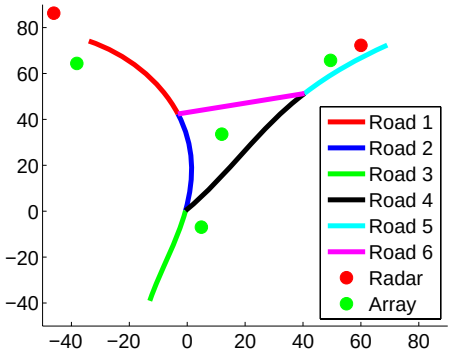

Figure 1: Triangle road network geometry. The axes are topocentric with distances in meters.

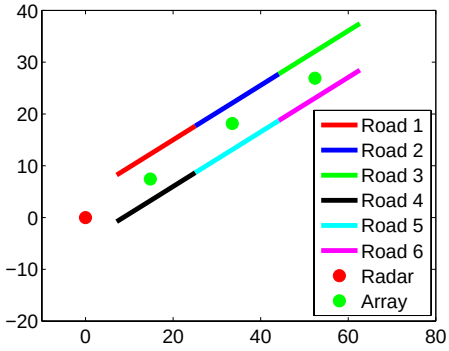

Figure 2: The straight line road network geometry. The axes are topocentric with distances in meters.

Probabilities are calculated from the signal, S, using the Swerling I model,

$$P(S>T)=\exp(-T/(1+\sigma_0^2(\theta)r_0^4/r^4\,e^{-r/r'})),$$

where T is a threshold, r is the range, r_0 is a reference range, and $\sigma_0^2(\theta)$ is the mean signal returned at the reference range and angle θ with respect to the aim of the radar beam. The $e^{-r/r'}$ term is a correction for deficiencies in the detector. In the center of the beam, the detection probability is around 0.75 at 150 ft (45 m) . The errors in the reported ranges are approximately Normally distributed with a standard deviation of around 2m for simple targets.

3.2 Acoustic Arrays

The acoustic arrays consist of four microphones spaced 8 inches apart in a square planar configuration. Angles are reported when the sound level exceeds a threshold, T defined relative to the background. The time delays between the different microphones determine the incident angle sines and cosines of the sound source,

$$\begin{bmatrix}\cos(\theta)\\ \sin(\theta)\end{bmatrix}=-\frac{1}{\Delta x/v_s}\left[\lambda+A^TA\right]^{-1}A^T\Delta\tau, \qquad (7)$$

where $\Delta\tau$ are the time delays calculated by maximizing the cross-correlation between the signals from the microphones,

$\Delta\tau_{i+4(j-1)}=\arg\max_{\Delta\tau}\int dtS_i(t)S_j(t+\Delta\tau)$, Δx is the spacing, v_s is the velocity of sound, A is a matrix of the displacements, and λ is a Lagrange multiplier that ensures $\cos^2(\theta)+\sin^2(\theta)=1$. Experiments show a Normally distributed error with 5^o accuracy and a detection range of 10-15 meters. The estimated bearing angle generally points to the loudest (and often closest) vehicle, although more complicated behaviors occasionally occur. The detection probability model is a passive version of the Swerling I model,

$$P(S>T)=\exp(-T/(1+\sigma_0^2(\theta)r_0^2/r^2\,e^{-r/r'})),$$

where σ_0 is velocity dependent, but only weakly vehicle dependent for non-diesel passenger motor vehicles. However, only the loudest object is reported so the probability of the angle from source i being returned is

$$P(s_i>T,s_{j\neq i})=\int_T^\infty dt\frac{1}{(1+\sigma_{i0}^2(\theta_i)r_{i0}^2/r_i^2)}\times e^{-t/(1+\sigma_{i0}^2(\theta_i)r_{i0}^2/r_i^2e^{-r/r'_{i0}})}\prod_{j\neq i}(1-e^{-t/(1+\sigma_{j0}^2(\theta_i)r_{j0}^2/r_j^2e^{-r/r'_{i0}})}) \qquad (8)$$

where s stands for the source, which may be clutter, and $\sigma_{i0}^2(\theta)$, r'_{i0}, and r_{i0} are source dependent. The probability of no return is $P(T>s_i)=1-\sum_{s_i}P(s_i>T,s_{j\neq i})$. Because the acoustic measurements require multi-target information, the probability hypothesis density PHD approximation to the FISST cannot be implemented on this system, and a full random set implementation is necessary [9].

4 Global Density Calculation

The prediction and update equations for the road constrained network are complicated by the association ambiguity of the sensors. In this section, we outline the expressions for these equations.

4.1 Prediction

The prediction step includes propagation along a road, road switching, and birth and death processes. The centerline of each road, i, is defined by a parametric curve, $x_i(s)$ and $y_i(s)$, with $\sqrt{(dx_i/ds)^2+(dy_i/ds)^2}=1$. The assumed independent motion of targets along the arc length of a road is linear with Gaussian white noise acceleration ,

$$\begin{aligned}\dot{s}^i&=v^i\\ \dot{v}^i&=-\gamma_i v^i+d\Lambda_i\end{aligned}. \qquad (9)$$

Over a small time interval, the expected variance in the integral of the acceleration is $\left\langle\left(\int_0^{\Delta t}d\Lambda_i\right)^2\right\rangle=\sigma_\Lambda^2\Delta t$.

The model's steady state root mean square velocity is $\sigma_{v_{ss}}=\sqrt{\sigma_\Lambda^2/2\gamma_i}$. Since the roads have a finite length, the propagated density will extend beyond the end of the road. The overhanging density is truncated from the original road and equipartitioned among the connected roads (See Fig. (3)). The velocity profile of the equipartitioned distribution is the same as the original density, with velocities pointing towards the intersection on the original road pointing away from the intersection on the new road.

Death processes correspond to density that overhangs at dead ends in the road network such as roads $1,3$, and 5 in Fig. (1). Death processes result in marginalizing over the overhanging profile and adding the remaining density profile to the corresponding hypothesis. For example, if Fig. (3) (b) corresponds to the projection of $f_{r_1r_2}$ onto the r_1 road segment, the overhang would be

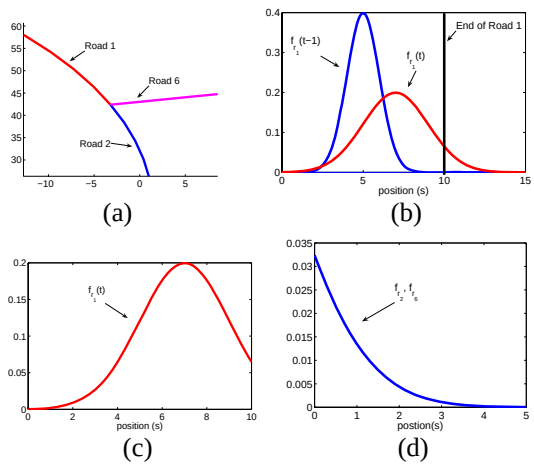

Figure 3: Road switching in the network. (a) The intersection of r_1, r_2 and r_6 in the upper left intersection in Fig. (1). (b) Starting from a density on r_1, the propagation step results in the density overhanging r_1. (c) The new profile for r_1 is created by truncating the overhanging portion of the density. (d) The overhanging density is equipartitioned between r_2 and r_6.

truncated and $\int_{x^1>10} dx^1 dv^1 f_{r_1 r_2}(x^1, v^1, x^2, v^2)$ is added to f_{r_2}.

Several birth processes are possible, such as coupling birth processes to measurements not associated with any previously detected targets [4]. Resource management of sensors requires knowledge of the rate of decay of information about the global density, which is greatly affected by birth processes.

To ensure that the birth processes reflect the decay in information, the birth processes correspond to adding targets to a road segment by a Poisson process with parameter, $\lambda_{r_i} << 1$. The added target has a Gaussian position profile with a mean located near the end of the road and a road dependent standard deviation. The velocity profile is a truncation of the steady state velocity profile that ensures the velocity points onto the road so that targets do not immediately depart upon entry.

4.2 Update

As mentioned above, summation over the associations complicates updating the probability density. For example, if a radar sensor reports the range measurement, $\{\rho_1 \ldots \rho_m\}$, the likelihood function becomes

$$f_{r_1 \ldots r_n}(\{\rho_1 \ldots \rho_m\} \mid \{x^{(n)}\}) = \sum_{\substack{\rho_1 = (clutter, 1 \ldots n) \\ \vdots \\ \rho_m = (clutter, 1 \ldots n)}} \left[\prod_{\substack{b \neq \rho_a \\ \rho_a = clutter}} (1 - P_c(b)) \prod_{\rho_a = clutter} P_c(\rho_a) \prod_{t_k \neq \rho_a} (1 - P_t(t_k)) \right. \times \left. \prod_{t_k = \rho_a} P_t(t_k) g_r(\rho_a - h_r(x^k(s^k), y^k(s^k))) \right]. \quad (10)$$

In this expression, ρ_a is the measurement (in the product and summation) and the association (in the indices). The probability of a clutter return from bin b, $P_c(b)$, or from the target, P_t, is spatially dependent. The measurement model, g_r, is a Normally distributed error around the point in the measurement space, $h_r(x(s), y(s)) = \sqrt{x(s)^2 + y(s)^2}$, which is a non-linear map between the state space and measurement space.

Association ambiguities lead to non-Gaussian effects and make the Kalman filtering equations invalid even if h_r was linear, g_r is Normally distributed, and the detection probabilities are position independent constants. The non-linear nature of the measurements and detection probabilities increases these complications. Various linear approximations, pruning, and measurement weighting strategies attempt to avoid this difficulty, but they may lead to poor tracking performance [3, 7].

Since the acoustic array only reports a single measurement, there are fewer associations to sum over, but the detection probabilities depends on all coordinates Eq. (8), and the likelihood of detecting an angle α is

$$\begin{aligned} f_{r_1 \ldots r_n}(\{\phi\} \mid \{x^{(n)}\}) &= P(T > clutter, s_j) \\ f_{r_1 \ldots r_n}(\{\alpha\} \mid \{x^{(n)}\}) &= P(clutter > T, s_j) + \\ &\sum_{s_i = t_1 \ldots t_n} P(s_i > clutter, T, s_{j \neq i}) g_\alpha(\alpha - h_\alpha(x^i)) \end{aligned} \quad (11)$$

Similar to the radar model, g_α is Normally distributed and the relation between roads positions and angle, $h_\alpha(x(s), y(s)) = \tan^{-1}(y(s)/x(s))$, is highly non-linear.

5 Numerical Implementation

We used a Gaussian Mixture approximation to implement a computationally tractable random set fusion algorithm. The Gaussian mixture model for this road network scenario also allows a scalable Gaussian sum particle filter (GSPF) representation [10]. The GSPF is similar to particle filter sampling methods, but the delta function kernel associated with each particle is replaced by a variable dimensional Gaussian. The covariance and mean of each particle is propagated instead of simply the position of each particle. Each term in the global density is represented by a finite number of Gaussians,

$$f_{r_1 \ldots r_n}(\{x^{(n)}\}) = \sum_i a_i G(x^{(n)} - \mu_i^{(n)}, C_i^{(n)}) \quad (12)$$

where G is a multidimensional Gaussian distribution, $G(x, C) = \sqrt{Det(2\pi C)}^{-1} \exp(x^T C^{-1} x)$. The Gaussian mixture representation requires several approximations: 1) After combining detections with the Gaussian, the probability of detect is based on the mean value of each Gaussian component. 2) The non-linear maps between the coordinates and the measurements are Taylor expanded around the mean values, $h \approx h(x(s_0)) + (\partial h(x(s_0))/\partial s)(s - s_0)$. Gaussian

displacements transverse to the road centerline are also considered. 3) The motion model is still linear. 4) Instead of truncating overhanging density, road switching and death processes correspond to the mean of the mixture component overhanging the end of the road. The entire Gaussian component either switches roads or is marginalized. 5) Birth processes correspond to adding mixture components to the ends of roads with a mean velocity pointing onto the roads and a fixed initial width.

All of these approximations maintain the Gaussian mixture representation of the global density. The validity of these approximations depends on the Gaussian mixture components' variances being much smaller than variations in the detection probabilities, the variations in h, and the distance from the ends of the roads. We ensure the validity of these approximations by replacing components with large variances with several Gaussian components of smaller variances through a Kullback-Leibler measure.

Although these approximations ensure that the global density maintains a Gaussian mixture functional form, the associations of the measurements result in a geometric explosion in the number of mixture components. To avoid the geometric growth, the mixture components are recombined based on the Kullback-Leibler metric. The KL metric compares $f_{r_1 \ldots r_n}$ with a new Gaussian mixture, $\tilde{f}_{r_1 \ldots r_n}$, that contains two components with the same mean and variance [11],

$$\tilde{f}_{r_1 \ldots r_n}(\{x^{(n)}\}) = \sum_k b_k G(x^{(n)} - m_k^{(n)}, R_k^{(n)}) = \sum_{k=i,j} a_k G(x^{(n)} - \tilde{\mu}^{(n)}, \tilde{C}^{(n)}) + \sum_{k \neq i,j} a_k G(x^{(n)} - \mu_k^{(n)}, C_k^{(n)}). \quad (13)$$

The KL metric gives

$$D_{f|\tilde{f}} = \sum_k \int d\{x^{(n)}\} a_k G(x^{(n)} - \mu_k^{(n)}, C_k^{(n)}) \times \ln\left(\frac{a_k G(x^{(n)} - \mu_k^{(n)}, C_k^{(n)})}{b_k G(x^{(n)} - m_k^{(n)}, R_k^{(n)})}\right) = = \sum_{k=i,j} \int d\{x^{(n)}\} a_k G(x^{(n)} - \mu_k^{(n)}, C_k^{(n)}) \times \ln\left(\frac{G(x^{(n)} - \mu_k^{(n)}, C_k^{(n)})}{G(x^{(n)} - \tilde{\mu}^{(n)}, \tilde{C}^{(n)})}\right). \quad (14)$$

where $\tilde{\mu}^{(n)}$ and $\tilde{C}^{(n)}$ are the optimal mean and covariance for combining the two mixtures. The combined mixture components that result in the smallest KL metric are combined until the number of mixtures is smaller than a set maximum number and the KL metric for combining two mixtures is greater than a tolerance [11].

If the final distribution was reduced to a single Gaussian at each time step, the resulting filter would be similar to the JIPDA filter, which sums over associations to determine the best fitting Gaussian distribution [7]. Maintaining a finite number of mixture components gives greater flexibility and hopefully improves performance. Each mixture component is similar to a single hypothesis in a multiple hypothesis tracker (MHT) [12]. Unlike the MHT that performs branching and pruning procedures, the FISST tracker combines branches of the track estimates [11]. This procedure reduces the amount of information lost at each update step since a hypothesis with low probability may contain a significant amount of information, but makes the definition of a track and track lifetime ambiguous since tracks with different lifetimes may be combined.

6 Resource Management

Sensor network efficiency is optimized with a Markov decision process resource management algorithm. The objective of the sensor network is to maintain knowledge of the surveillance area while maximizing the sensor's battery lifetime. Excessive use can also compromise the sensor network by increasing the probability of hostile entities detecting and disabling the sensor network.

The resource management is based on considering each of the road segments as tasks. The global density determines the "information state" of each task, which reflect our confidence in a set of mutually exclusive and exhaustive hypotheses. For a maximum of two targets per road segment the hypotheses are listed in Table (1).

Table 1: The states calculated from the global density and used to determine resource management.

S_{-1}	Not sure of any other hypothesis
S_0	>95% positive no targets on road segment
S_1	>95% positive 1 target on road segment, but $\sigma > \sigma_0$
S_2	>95% positive 1 target on road segment and $\sigma < \sigma_0$
S_3	>95% positive 2 targets on road segment, but $\sigma > \sigma_0$ for at least one target
S_4	>95% positive 2 targets on road segment, and $\sigma < \sigma_0$ for all targets

Defining these states allows us to cast the resource management algorithms into the restless bandit formulation [13]. Each task area can be actively probed by a set of sensors. A probed task is active ,"a", and an inactive task is passive, "p". An active or passive task i that is in state, S_{ij}, results in a reward $R_{ij}^{(a)}$ or $R_{ij}^{(p)}$ and a transition between states denoted by Markov transition matrix $P_i^{(a)}$ or $P_i^{(p)}$. Actively working on tasks results in transitions from states of relative ignorance, S_{-1}, S_1, and S_3, to states of relative knowledge S_0, S_2, and S_4. Tasks that remain passive produce transitions that degrade knowledge. The Markovian transition probabilities are not exact, since the real transitions depend on f_g, which has many more degrees of

freedom, but these transitions can be approximated from average behaviors of the global density.

The major barrier to the restless bandit formulation is the determination of the appropriate rewards structure. The rewards are chosen to obtain desired behavior from the system, including discovery of new targets and tracking known targets. The rewards are obtained from the guidelines listed below. 1) States of high knowledge, S_0, S_2, and S_4, do not require much additional information so the active rewards are low for these states and the passive rewards are high. 2) States of high ignorance, S_{-1}, require a large amount of additional information, so the active rewards are high while the passive rewards are low.

A policy that adaptively controls the sensors consists of a precomputed set of performance measures and a heuristic rule that uses the performance measures to determine the optimal set of actions given the current state of the system. The performance measures are determined by solving for the maximum expected infinite time horizon discounted reward through a linear programming solution [13]. The original restless bandit heuristic described in [14] selects the best m tasks to be active at each discrete time step since the number of active tasks is fixed in that application.

While this heuristic is applicable to a limited resource with the capability to activate only m tasks at a time, the sensor network application has a flexible number of active tasks at each time step. The netted sensors scenario better corresponds to the linear programming solution for the performance measures, which assumes a fixed discounted time averaged number of active tasks that prevents excessive use of any single set of sensors. In the netted sensor heuristic, the active tasks are selected by comparing the performance measure with a threshold at each time step.

7 Result

To analyze the tracking performance, we instrumented a vehicle with a GPS. This vehicle traversed the linear road network in Fig. (2) while targets of opportunity enter the road network and act as decoys. We examined the ability of the fusion algorithm to track the true target while missed detections and false alarms from the clutter and the decoys create association ambiguities. The FISST tracker and resource manager were run at a 3 Hz rate. During each iteration, the global density was calculated, which subsequently determined sensor alocation through the resource manager. To aid initial development on the straight road network and reduce computational time, a diffusion approximation was applied to the motion model to eliminate the velocity.

The FISST global density reported tracks determined by: 1) determining the most likely hypothesis through marginalizing over the geokinetic variables,

$$P_{r_1 \ldots r_n} = \frac{1}{n!} \sum \int d\{x^{(n)}\} f_{r_1 \ldots r_n}(\{x^{(n)}\}),$$ where the summation is performed over the permutations of $r_1 \ldots r_n$. Selecting hypotheses by marginalizing over the geokinetic variables avoids incommensurable units. 2) The most likely positions of the most likely hypothesis (the MAP estimate) are reported,

$$\arg\max_{\{x^{(n)}\}} f_{r_1 \ldots r_n}(\{x^{(n)}\}).$$

Fig. (4) compares a typical track on the linear road network for the fusion algorithm with all sensors collecting data at every time step in (a) against a typical track for the fusion algorithm with the resource manager. The measured distance corresponds to the distance from the radar to the target or track. Continuously employing all sensors results in longer track lifetimes at the expense of energy usage and a reduced sensor lifetime. Track lifetime is difficult to define for FISST since the information associated with the true track may still be present in global density even if it may not be the most likely hypothesis. Similar behavior is found in MHT, where the true hypothesis may not be the most likely hypothesis, but it is still being tracked [12].

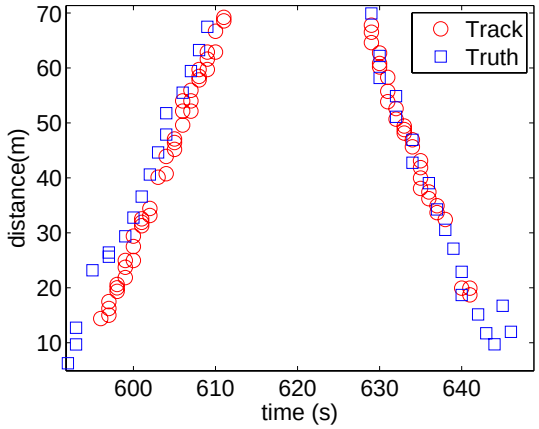

(a)

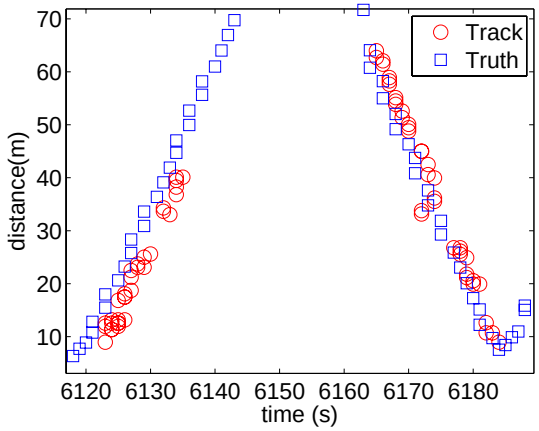

(b)

Figure 4: Tracking on the linear road network in Fig. (2). (a) Tracking with all sensors running at 3 Hertz. (b) Tracking with the resource manager. Distances are measured in meters from the radar (see Fig. (2)). Errors in tracks vary with time and the measurements, but are around 4m in (a) and 6-8m in (b). Truth error estimates are around 3m, and may be systematic, due to GPS.

In the absence of targets, the resource manager runs sensors at around ½ Hz, which is six times lower than the continuous running rate. When targets are present, the resource manager approximately increases the rate to 1½ to 2 Hz, which is still less than the maximum rate of 3 Hz. The rate changes depending on the quality of the previous

measurements by allowing good tracks to coast for a time step, which conserves the sensor's battery. Using the resource management, a typical track error is around 6-8m compared to 4m for the continuously run sensors.

The resource manager also displays a sentinel behavior. The rate of decay of the information in the global density affects "information states", which alters the resource allocation. The global density information degrades the most rapidly at intersections (due to road switching events), and dead ends (due to birth processes). As a result, the acquisition rates are the largest for sensors near dead end road segments or intersections approached by targets. These behaviors were not explicitly programmed into the resource manager, but are reasonable consequences of the restless bandit rewards structure discussed in Section 6.

To demonstrate that the FISST tracker can be applied to more complicated road networks, we examined simulations on the road network in Fig.(1) (see Fig. (5)). The results from our initial linear road network experiment were used to increase the fidelity of sensor models by adjusting parameters such as false alarm rates, spatial detection probabilities, and measurement accuracy. With the improved sensor models, we were able to address the more challenging road network depicted in Fig. (1). A typical simulation run is depicted in Fig. (5).

The tracker can readily handle the more complicated road network. The resource manager greatly reduced sensor operation while slightly degrading tracking performance. The MAP output temporarily dropped more tracks (~10% versus ~5%), and the MAP output may initially follow the wrong road segments at intersections. However, the FISST tracker does correct itself and resume the correct MAP estimates, while reducing energy usage. Most tracks are only dropped for one time step. The correction is due to the nature of the FISST tracker. The track output is a MAP estimate and the true target positions might not always be the most likely positions, but they are almost always very probable positions. The likelihood of the true position may be the best measure of FISST tracker performance. The error variances are similar to those reported for the straight line network with variable errors around 4m for the continuously run sensors and 6m for the resource managed network.

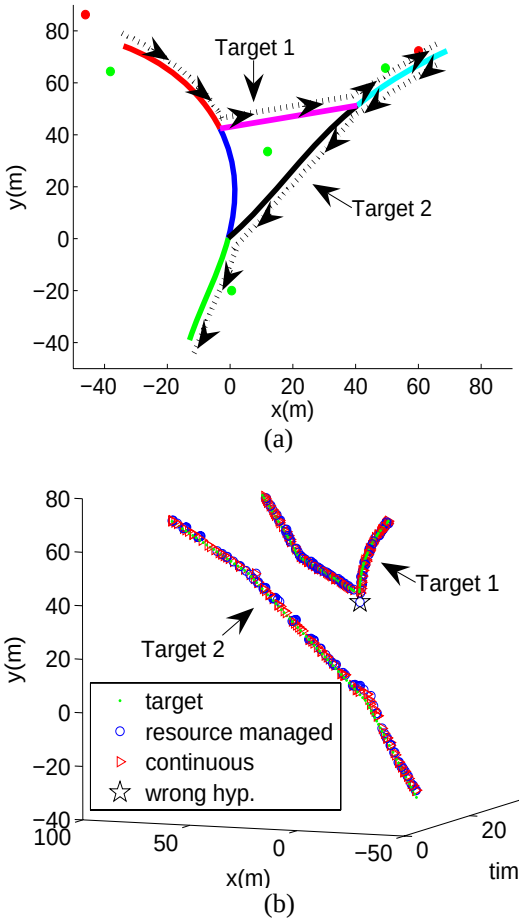

Figure 5: Simulation of the network in Fig. (1). (a) The paths of two targets entering the networks. (b) The true target paths are compared with estimates from continuously running all sensors and running the sensors with a resource manager. The resource managed simulation shows an initial adoption of a wrong hypothesis at a road intersection before correcting itself. Similar to Figure 4, errors vary with time, but are around 4 meters for the continuously run sensors and 6m for the resource managed simulation. The target positions are exact.

8 Conclusions

FISST trackers are ideal for dense road networks. Every intersection creates a bifurcation in the target trajectory and makes simple Kalman filtering impossible. A MHT can track branching processes if intelligent pruning is possible [12], but the FISST tracker examined in this paper introduces hypothesis branching reduction through merging similar hypotheses, which avoids information loss from pruning processes. The multi-body nature of the acoustic sensors prevents use of the PHD tracker, a simplified FISST tracker that neglects multi-body information [9].

The FISST tracker performed well in live experiments using a simple scenario and in our simulation experiments using a more complex scenario. Initial applications to the triangle network and other road networks show successful tracking capabilities, which will be quantified in future publications. This success is encouraging since the tracker is not specifically optimized to a single geometry and the other examined intersections are near major roads and in uncontrolled environments. The application to several experimental scenarios demonstrates that the algorithm does not depend on the specifics of the road network and is extensible.

Although we only reported MAP estimates, complete target information is retained and tracks can be re-established or corrections can be made for initially taking the wrong branch at an intersection. Quantifying the ability to correct hypotheses will be explored in the future.

The resource manager only slightly degrades tracker performance by reducing the amount of data collected, but substantially reduces sensor operation rates hence greatly increasing network lifetime. The resource manager creates a feedback mechanism, where inadequate information in the global density (areas of high entropy) determine the course of action.

The approximate Gaussian mixture fusion algorithm also lends itself to a Gaussian Sum Particle Filtering implementation (GSPF). GSPF will greatly increase the number of tracked targets[10]. The particle filter can be computationally distributed, since different processors can perform the calculation on different Gaussian particles [15]. If the sensors are independent, the FISST tracker also shows the ability to distribute computation through a hierarchy since the log-likelihood will be a sum of independent terms that can be added through the hierarchy.

Another interesting observation is the similar performances of the full motion model and the diffusion approximation. Since velocity is not directly measured in our system, the velocity simply carries information about past positions. The performance of the diffusion model implies that the sensor measurements frequency and quality makes velocity information unnecessary. This result suggests that velocity information may allow further reduction in sensor utilization at the expense of the higher computational cost associated with solving a more complex model. The trade-off between these two aspects of tracking should be explored.

The live data experiment examined in this paper is an important first step in developing a FISST tracker for netted-sensor monitoring of dense road networks. Our experiments demonstrated the capabilities of the FISST tracker. The fusion algorithm can be combined with resource management to optimally control sensor networks, and a GSPF will distribute computing and insure scalability.

References

[1] T. Kirubarajan, Y. Bar-Shalom, K.R. Pattipati, I. Kardar, B. Abrams, and E. Eadan. "Tracking Ground Targets with Road Constraints Using an IMM Estimator." *IEEE Aerospace Conference Proceedings.* 5:5--12, Mar. 1998.

[2] C.A. Scott and C.R. Drane. "An Optimal Map-Aided Position Estimator For Tracking Motor Vehicles." *IEEE Proceedings of the 6th International Conference on Vehicle Navigation and Information Systems.* 360--367, Jul. 1995.

[3] I.R. Goodman, R.P.S. Mahler, and H.T. Nguyen. *Mathematics of Data Fusion.* Kluwer Academic Publishers. Dordrecht, Netherlands. 1997.

[4] M.R Morelande and S. Challa. "A Multitarget Tracking Algorithms Based on Random Sets." IEEE Proceedings of the 6th International Conference of Information Fusion. 2:807—814, 2003.

[5] B.N. Vo and S. Singh, and W.K. Ma. "Tracking Multiple Speakers Using Random Sets." *IEEE Conference Proceedings on Acoustics, Speech, and Signal Processing*. 2: ii357-ii360, 2004.

[6] R. Evans, V. Krishnamurthy, G. Nair, and L. Sciacca. "Networked Sensor Management and Data Rate Control for Tracking Maneuvering Targets." *IEEE Transactions on Signal Processing.* 53(6):1979-1991, 2005.

[7] D. Musicki and R. Evans. "Joint Integrated Probabilistic Data Association." *IEEE transactions on Aerospace and Electronic Systems*. 40(3): 1093—1099, 2004.

[8] A. Doucet, B.N. Vo, C. Andrieu and M. Davy. "Particle Filtering for Multi-Target Tracking and Sensor Management." *Proceedings of the Fifth International Conference on Information Fusion*. 1:474—481, Jul. 2002.

[9] R.P.S. Mahler. "Multitarget Bayes Filtering via First-Order Multitarget Moments." *IEEE Transactions on Aerospace and Electronic Systems*. 39(4):1152—1178, Oct. 2003.

[10] J.H. Kotecha and P.M Djuric. "Gaussian Sum Particle Filtering." IEEE Transactions on Signal Processing. 51(10):2602—2612, Oct. 2003.

[11] W. Adam, R. Fruhwirth, A. Strandlie, and T. Todorov. "Reconstruction of elections with the Gaussian-sum filter in CMS tracker at LHC." *Journal of Physics G-Nuclear and Particle Physics*. 31:N9-N20, Mar. 2005.

[12] S.S. Blackman. "Multiple Hypothesis Tracking For Multiple Target Tracking." IEEE Aerospace and Electronic Systems Magazine. 19:5—18, 2004.

[13] P. Whittle. "Restless bandits: Activity allocation in a changing world." In *A Celebration of Applied Probability*. J. Gani (Ed.). *Journal of Applied Probability*. 25a: 287—298, 1988.

[14] D. Bertsimas and J. Nino-Mora. "Restless bandits, linear programming relaxations and a primal-dual index heuristic." *Operations Research*. 48(1):80—90, Jan. 2000.

[15] M. Coats. "Distributed Particle Filters for Sensor Networks." In *Proceedings of the 3rd International Symposium on Information Processing in Sensor Networks*. ACM Press. NY, NY. 99-107, Apr. 2004.