

A Hierarchical Multiple-Target Tracking Algorithm for Sensor Networks

Songhwai Oh, Luca Schenato and Shankar Sastry

Department of Electrical Engineering and Computer Sciences

University of California, Berkeley, CA 94720, U.S.A.

{sho, lusche, sastry}@eecs.berkeley.edu

Abstract—Multiple-target tracking is a canonical application of sensor networks as it exhibits different aspects of sensor networks such as event detection, sensor information fusion, multi-hop communication, sensor management and decision making. The task of tracking multiple objects in a sensor network is challenging due to constraints on a sensor node such as short communication and sensing ranges, a limited amount of memory and limited computational power. In addition, since a sensor network surveillance system needs to operate autonomously without human operators, it requires an autonomous tracking algorithm which can track an unknown number of targets. In this paper, we develop a scalable hierarchical multiple-target tracking algorithm that is autonomous and robust against transmission failures, communication delays and sensor localization error.

Index Terms—Sensor networks, multiple-target tracking, Markov chain Monte Carlo, data association

I. INTRODUCTION

In wireless ad-hoc sensor networks, many inexpensive and small sensor-rich devices are deployed to monitor and control our environment. It is envisioned that the sensor networks will connect us to the physical world in a pervasive manner [6], [8]. Each device, called a sensor node, is capable of sensing, computation and communication. Sensor nodes form a wireless ad-hoc network for communication. The limited supply of power and other constraints, such as manufacturing costs and limited package sizes, limit the capabilities of each sensor node. For example, a typical sensor node has short communication and sensing ranges, a limited amount of memory and limited computational power. However, the abundant number of spatially spread sensors will enable us to monitor changes in our environment accurately despite of inaccuracy of each sensor node.

Multiple-target tracking is a canonical application of sensor networks as it exhibits different aspects of sensor networks such as event detection, sensor information fusion, communication, sensor management, and decision making. Each sensor node has a limited supply of power and operates in the low SNR regime, leading to low detection probability. If the detection is done by a threshold test, one can increase the detection probability by decreasing the threshold level. However, as we decrease the threshold level, the false alarm rate gets increased. The presence of false alarms and missing observations due to the low detection probability complicate the problems

of track initiation and track termination. These important issues are ignored by many tracking algorithm designed for sensor networks. For example, when the false alarm rate is high, the naive track initiation routine will overflow the network with spurious tracks. Hence, an algorithm for sensor networks must be robust against the low detection probability and high false alarm rate.

In sensor networks, we seek for an autonomous tracking algorithm which does not require a continuous monitoring by a human operator. The localization of sensor nodes in an ad-hoc wireless sensor network, without expensive hardware such as the global positioning system (GPS), is a challenging problem (see [15] and references therein). Since the position of a target is reported with respect to the location of the reporting sensor, the algorithm must be robust against the sensor localization error. We also need to consider the following constraints on sensor networks. Due to the limited supply of power, the multi-hop wireless ad-hoc communication is used in sensor networks. In many cases, the communication bandwidth is low and the communication links are not reliable, causing transmission failures. In addition, due to the low communication bandwidth and a limited amount of memory, communication delays can occur frequently. It is well known that communication is costlier than computation in sensor networks in terms of power usage [7]. Hence, it is essential to fuse local observations before the transmission. Since the data association problem is NP-hard [5], [17], we cannot expect to solve it with only local information. But, at the same time, we cannot afford to have a centralized algorithm since such solution cannot be scalable. In summary, we need a simple and efficient tracking algorithm that is robust against the low detection probability and high false alarm rates; capable of initiating and terminating tracks; uses less memory; combines local information to reduce the communication load; and is scalable. Also it must be robust against transmission failures, communication delays and sensor localization error. But at the same time we want an algorithm that can provide a good solution and improve its solution toward the optimal solution given an enough computation time.

In [16], an efficient real-time algorithm that solves the data association problem and is capable of initiating and terminating a varying number of tracks, Markov chain Monte Carlo data association (MCMCDA), is presented. MCMCDA is an approximation to the optimal Bayesian

filter [16]. It has been shown that MCMCDA is computationally efficient compared to the multiple hypothesis tracker (MHT) [18] and outperforms MHT under extreme conditions, such as a large number of targets in a dense environment, low detection probabilities, and high false alarm rates [16]. MCMCDA is suitable for sensor networks since it can autonomously initiate and terminate tracks. Since transmission failure is another form of a missing observation, MCMCDA is robust against transmission failures. MCMCDA performs data association based both current and past observations, so delayed observations can be easily combined with previously arrived observations to improve the accuracy of estimates. Furthermore, MCMCDA requires less memory as it maintains only the current hypothesis and the hypothesis with the highest posterior. It does not require the enumeration of all or some of hypothesis as in [10], [18]. In this paper, we extend the MCMCDA algorithm to sensor networks in a hierarchical manner so that the algorithm becomes scalable and we show the robustness of the algorithm against transmission failures, communication delays and sensor localization error in simulations. To our knowledge, the algorithm presented in this paper is the first general multiple-target tracking algorithm for sensor networks which can systematically track an unknown number of targets in the presence of false alarms and missing observations and is robust against transmission failures, communication delays and sensor localization error.

We consider a simple shortest-path routing scheme on a sensor network. The transmission failures and communication delays of the network are characterized probabilistically. We assume the availability of a small number of special nodes, *supernodes*, that are more capable than regular nodes in terms of computational power and communication range. Each node is assigned to its nearest supernode and nodes are grouped by supernodes. We call the group of sensor nodes formed around a supernode as a “tracking group”. When a node detects a possible target, it communicates with its neighbors and observations from the neighboring sensors are fused and sent to its supernode. Each supernode receives the fused observations from its tracking group and executes the tracking algorithm. Each supernode communicates with neighboring supernodes when a target moves away from its range. Lastly, the tracks estimated by supernodes are combined hierarchically.

The remainder of this paper is structured as follows. In Section III, the multiple-target tracking problem and its probabilistic model are described. The MCMCDA algorithm for multiple-target tracking is presented in Section IV. The sensor network model is described in Section V and the hierarchical MCMCDA method is given in Section VI. The simulation results are shown in Section VII.

II. RELATED WORK

The traditional multiple-target tracking algorithms such as the joint probabilistic data association filter (JPDAF)

[1] and multiple hypothesis tracker (MHT) [18] are robust against the low detection probability and high false alarm rate. But they are not suitable for sensor networks since the track initiation and termination is difficult with JPDAF and both JPDAF and MHT require large memory and computation cycles. Since MHT can initiate and terminate tracks, the tracking task can be easily distributed in a network of sensors. In [3], a distributed tracking algorithm based on MHT is developed for multiple sensors. But the approach is not suitable for sensor networks since it demands large computational power and large amount of memory on each sensor.

In [11], the authors propose to use a classification algorithm to disambiguate closely located targets. But signals received from targets are correlated and we cannot recover the uncorrelated signals in all cases. Since we do not know in advance the number of targets around each sensor, the problem is ill-posed and very challenging even for a high-end computer. In [13], the distributed track initiation and maintenance methods are described. By electing a leader among the sensors by which a target is detected, unnecessary communications are reduced. But considering the complexity of the data association problem, the approach will suffer from incorrect associations when there are many targets crossing or moving close to each other. In addition, when the false alarm rate is high, the proposed approach will overflow the network with spurious tracks and it is unclear how the missing observations are handled.

Many of newly proposed multiple-target tracking algorithms for sensor networks try to solve the identity management problem [12], [19]. They assume the availability of a classification algorithm as in [11] but the disambiguation is delayed until targets are sufficiently separated. As assumed in simulations of [12], when the targets are of different classes, a target can be classified by the signature of its class. But, if all targets are of the same class, a target cannot be easily classified by its signature and, in the absence of reliable classification information, the proposed methods will behave like the naive nearest neighbor tracker. Our algorithm can complement the identity management algorithms when tracking targets with the same class or reliable classification information is not available.

A distributed particle filtering algorithm for sensor networks is presented in [4] and used to track a single maneuvering target. The paper assumes the availability of supernodes and a hierarchical topology similar to ours. The paper also assumes the availability of sensors which can measure an angle and distance to a target. But we are not aware of sensors with such capability available for sensor networks. The most widely used and realistic sensor model is based on the signal strength and this is the model we use in this paper.

III. GENERAL MULTIPLE-TARGET TRACKING

A. Problem

Let $T \in \mathbb{Z}^+$ be the duration of surveillance. Let K be the unknown number of objects moving around the

surveillance region \mathcal{R} for some duration $[t_1^k, t_f^k] \subset [1, T]$ for $k = 1, \dots, K$. Let V be the volume of \mathcal{R} . Each object arises at a random position in \mathcal{R} at t_1^k , moves independently around \mathcal{R} until t_f^k and disappears. At each time, an existing target persists with probability $1 - p_z$ and disappears with probability p_z . The number of objects arising at each time over \mathcal{R} has a Poisson distribution with a parameter $\lambda_b V$ where λ_b is the birth rate of new objects per unit time, per unit volume. The initial position of a new object is uniformly distributed over \mathcal{R} .

Let $F^k : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be the discrete-time dynamics of the object k , where d is the dimension of the state variable, and let $x_t^k \in \mathbb{R}^d$ be the state of the object k at time t for $k = 1, 2, \dots, K$. The object k moves according to

$$x_{t+1}^k = F^k(x_t^k) + w_t^k, \quad \text{for } t = t_1^k, \dots, t_f^k - 1,$$

where $w_t^k \in \mathbb{R}^d$ are white noise processes. The noisy observation of the state of the object is measured with a detection probability p_d which is less than unity. There are also false alarms and the number of false alarms has a Poisson distribution with a parameter $\lambda_f V$ where λ_f is the false alarm rate per unit time, per unit volume. Let n_t be the number of observations at time t , including both noisy observations and false alarms. Let $y_t^j \in \mathbb{R}^m$ be the j -th observation at time t for $j = 1, \dots, n_t$, where m is the dimensionality of each observation vector. Each object generates a unique observation at each sampling time if it is detected. Let $H^j : \mathbb{R}^d \rightarrow \mathbb{R}^m$ be the observation model. Then the observations are generated as follows:

$$y_t^j = \begin{cases} H^j(x_t^k) + v_t^j & \text{if } j\text{-th observation is from } x_t^k \\ u_t & \text{otherwise,} \end{cases}$$

where $v_t^j \in \mathbb{R}^m$ are white noise processes and $u_t \sim \text{Unif}(\mathcal{R})$ is a random process for false alarms. Notice that, with probability $1 - p_d$, the object is not detected and we call this a missing observation. We assume that targets are indistinguishable in this paper. But, if observations include target type or attribute information, the state variable can be extended to include target type information.

B. Probabilistic Model

Let $y_t = \{y_t^j : j = 1, \dots, n_t\}$ and $Y = \{y_t\}_1^T$. Let Ω be a collection of partitions of Y such that, for $\omega \in \Omega$,

- 1) $\omega = \{\tau_0, \tau_1, \dots, \tau_K\}$;
- 2) $\bigcup_{k=0}^K \tau_k = Y$ and $\tau_i \cap \tau_j = \emptyset$ for $i \neq j$;
- 3) τ_0 is a set of false alarms;
- 4) $|\tau_k \cap y_t| \leq 1$ for $k = 1, \dots, K$ & $t = 1, \dots, T$; and
- 5) $|\tau_k| > 1$ for $k = 1, \dots, K$.

Here, K is the number of tracks for the given partition $\omega \in \Omega$. We call τ_k a track when there is no confusion although the actual track is the set of estimated states from the observations τ_k . However, we assume there is a deterministic function that returns a set of estimated states given a set of observations, so no distinction is required. The fourth requirement says that a track can have at most one observation at each time, but, in the case of multiple




Fig. 1. (a) An example of observations Y (each circle represents an observation and numbers represent observation times); (b) an example of a partition ω of Y

sensors, we can easily relax this requirement to allow multiple observations per track. A track is assumed to contain at least two observations since we cannot distinguish a track with a single observation from a false alarm. An example of a partition is shown in Fig. 1.

Let e_t be the number of targets from time $t-1$ and a_t be the number of new targets at time t . Let z_t be the number of targets terminated at time t and $c_t = e_t - z_t$. Let d_t be the number of detections at time t and $u_t = e_t - z_t + a_t - d_t$ be the number of undetected targets. Finally, let $f_t = n_t - d_t$ be the number of false alarms. It can be shown that the posterior of ω is:

$$P(\omega|Y) \propto \prod_{t=1}^T p_z^{z_t} (1 - p_z)^{c_t} p_d^{d_t} (1 - p_d)^{u_t} \lambda_b^{a_t} \lambda_f^{f_t} P(Y|\omega) \quad (1)$$

where $P(Y|\omega)$ is the likelihood of observations Y given ω , which can be computed based on the chosen dynamic and measurement models. Our goal is to find a partition of observations such that $P(\omega|Y)$ is maximized.

IV. MCMC DATA ASSOCIATION ALGORITHM

In this section, we develop an MCMC sampler to solve the multiple-target tracking problem. MCMC-based algorithms play a significant role in many fields such as physics, statistics, economics, and engineering [2]. In some cases, MCMC is the only known general algorithm that finds a good approximate solution to a complex problem in polynomial time [9]. MCMC techniques have been applied to complex probability distribution integration problems, counting problems such as #P-complete problems, and combinatorial optimization problems [2], [9]. The MCMC approach applied to combinatorial optimization problems is generally known as simulated annealing.

MCMC is a general method to generate samples from a distribution π by constructing a Markov chain \mathcal{M} whose states are ω and whose stationary distribution is $\pi(\omega)$. If we are at state $\omega \in \Omega$, we propose $\omega' \in \Omega$ following the proposal distribution $q(\omega, \omega')$. The move is accepted with an acceptance probability $A(\omega, \omega')$ where

$$A(\omega, \omega') = \min \left(1, \frac{\pi(\omega')q(\omega', \omega)}{\pi(\omega)q(\omega, \omega')} \right), \quad (2)$$

otherwise the sampler stays at ω , so that the detailed balance is satisfied. If we make sure that \mathcal{M} is irreducible and aperiodic, then \mathcal{M} converges to its stationary distribution by the ergodic theorem.

The MCMC data association (MCMCDA) algorithm is described in Algorithm 1. MCMCDA is an MCMC

Algorithm 1 (MCMC Data Association):

Input: Y, n_{mc}, ω_{init}
Output: $\hat{\omega}$

```

 $\omega \leftarrow \omega_{init}; \hat{\omega} \leftarrow \omega_{init}$ 
for  $n = 1$  to  $n_{mc}$ 
  propose  $\omega'$  based on  $\omega$  (see [16])
  sample  $U$  from  $\text{Unif}[0, 1]$ 
   $\omega \leftarrow \omega'$  if  $U < A(\omega, \omega')$ 
   $\hat{\omega} \leftarrow \omega$  if  $p(\omega|Y)/p(\hat{\omega}|Y) > 1$ 
end

```

algorithm whose state space is Ω described in Section III-B and whose stationary distribution is the posterior (1). The proposal distribution for MCMCDA consists of five types of moves. They are (1) birth/death move pair; (2) split/merge move pair; (3) extension/reduction move pair; (4) track update move; and (5) track switch move. The MCMCDA moves are graphically illustrated in Fig. 2. For detail description of each move, see [16]. The inputs for MCMCDA are the set of all observations Y , the number of samples n_{mc} , and the initial state ω_{init} . The acceptance probability $A(\omega, \omega')$ is defined in (2) where $\pi(\omega) = P(\omega|Y)$ from (1). In Algorithm 1, we use MCMC to find a solution to a combinatorial optimization problem. So it can be considered as simulated annealing at a constant temperature. No burn-in samples are used since we are simply looking for a partition which maximizes the posterior. In addition, the memory requirement of the algorithm is at its bare minimum. Instead of keeping all $\{\omega(n)\}_{n=1}^{n_{mc}}$, we simply keep the partition with the maximum posterior, $\hat{\omega}$. Notice that, in MCMC, the construction of ω' is done on fly according to the proposal distribution $q(\omega, \omega')$ and there is no need to store previously visited states.

The Markov chain designed by Algorithm 1 is irreducible (Theorem 1 in [16]) and aperiodic [16]. In addition, the transitions described in Algorithm 1 satisfy the detailed balance condition since it uses the Metropolis-Hastings kernel (2). Hence, by the ergodic theorem, the chain converges to its stationary distribution.

V. SENSOR NETWORK MODEL

In this section, we describe the sensor network and sensor model used for simulations in Section VII. Let N_s be the number of sensor nodes, including both supernodes and regular nodes, deployed over the surveillance region $\mathcal{R} \subset \mathbb{R}^2$. We assume that each supernode can communicate with its neighboring supernodes. Let $s_i \in \mathcal{R}$ be the location of the i -th sensor node and let $S = \{s_i : 1 \leq i \leq N_s\}$. Let $R_t \in \mathbb{R}$ be the transmission range of a regular sensor node. A pair of sensor nodes i and j can communicate to each other if the Euclidean distance $\|s_i - s_j\| \leq R_t$. Let $G = (S, E)$ be a communication graph such that $(s_i, s_j) \in E$ if and only if $\|s_i - s_j\| \leq R_t$. Let $N_{ss} \ll N_s$ be the number of supernodes and let $s_j^s \in S$ be the position of the j -th supernode, for $j = 1, \dots, N_{ss}$. Let $g : \{1, \dots, N_s\} \rightarrow \{1, \dots, N_{ss}\}$ be the assignment of each sensor to its nearest supernode such that $g(i) = j$




Fig. 2. Graphical illustration of MCMCDA moves (associations are indicated by dotted lines and hollow circles are false alarms)

if $\|s_i - s_j^s\| = \min_{k=1, \dots, N_{ss}} \|s_i - s_k^s\|$. For a node i , if $g(i) = j$, then the shortest path from s_i to s_j^s in G is denoted by $sp(i)$.

Let $R_s \in \mathbb{R}$ be the sensing range. If there is an object at $x \in \mathcal{R}$, a sensor can detect the presence of the object. Each sensor records the sensor's signal strength,

$$z_i = \begin{cases} \frac{\beta}{1 + \gamma \|s_i - x\|^\alpha} + w_i, & \text{if an object is present at } x \\ w_i, & \text{if no object is present,} \end{cases} \quad (3)$$

where α , β and γ are constants specific to the sensor type and they are normalized such that w_i has the standard Gaussian distribution. This signal-strength based sensor model (3) is general for sensors available in sensor networks, such as acoustic and magnetic sensors, and has been used frequently [12]–[15]. For each i , if $z_i \geq \eta$, where η is a threshold set for appropriate values of detection and false-positive probabilities, the node transmits z_i to its neighboring nodes, which are at most $2R_s$ away from s_i , and listens to incoming messages from its $2R_s$ neighborhood. Note that this approach is similar to the leader election scheme in [13] and we assume that $R_t \geq 2R_s$. However, this approach may cause some missing observations if there is more than one object in this disk of radius $2R_s$. A better approach to fuse local data is required and we will address this issue in our future work. For the node i , if z_i is the larger than all incoming messages, $z_{i_1}, \dots, z_{i_{k-1}}$, and $z_{i_k} = z_i$, then the position of an object is estimated as

$$\hat{z}_i = \frac{\sum_{j=1}^k z_{i_j} s_{i_j}}{\sum_{j=1}^k z_{i_j}}. \quad (4)$$

Then \hat{z}_i is transmitted to the supernode $g(i)$ via the shortest path $sp(i)$. If z_i is not the largest compared to the incoming

messages, the node i does nothing and goes back to the sensing mode. Although each sensor cannot give an accurate estimate of object's position, as more sensors collaborate, the accuracy of estimates improves as shown in Fig. 3 (left). The collaboration of sensors makes the system more robust against node failures and we can increase the detection probability and decrease the false alarm rate by collaboration.

A transmission along the edge (s_i, s_j) fails independently with probability p_{te} and the message never reaches a supernode. So we can consider transmission failure as another form of a missing observation. If k is the number of hops required to relay data from a sensor node to its supernode, the probability of successful transmission decays exponentially as k increases. To overcome this problem, we use k independent paths to relay data if the reporting sensor node is k hops away from its supernode. The probability of successful communication from the reporting node i to its supernode $g(i)$ can be computed as $1 - (1 - (1 - p_{te})^k)^k$, where $k = |sp(i)|$.

The (additional) communication delay is modeled by the negative binomial distribution. We assume each node has the same probability p_{de} of delaying a message. If d_i is the number of delays occurred on the message originating from the sensor i , d_i is distributed as

$$p(d_i = d) = \binom{|sp(i)| + d - 1}{d} (1 - p_{de})^{|sp(i)|} (p_{de})^d. \quad (5)$$

If the network is heavily loaded, the independence assumptions on transmission failure and communication delay may not hold. However, the model is realistic under the moderate conditions and we have chosen it for its simplicity.

VI. HIERARCHICAL MCMCDA

We use the online MCMCDA algorithm with a sliding window of size w_s [16]. The supernodes maintains a set of observations $Y = \{y_t^j : t_{curr} - w_s < t \leq t_{curr}, 1 \leq j \leq n_t\}$, where t_{curr} is the current time. Each y_t^j is a fused observation \hat{z}_i from some sensor i . At time $t_{curr} + 1$, the observations at time $t_{curr} - w_s$ are removed from Y and a new set of observations is appended to Y . Any delayed observations are appended to appropriate slots. Then each supernode initializes the Markov chain with the previously estimated tracks and executes Algorithm 1 on Y . Once tracks are found, the next state of each track is predicted. If the predicted next state belongs to the surveillance area of another supernode, track information is passed to the corresponding supernode. The newly received M_t tracks are incorporated into the initial state of the Markov chain for the next time step.

Since each supernode maintains its own set of tracks, there can be multiple tracks from a single object maintained by different supernodes. To make the algorithm fully hierarchical, we do track-level data association to combine tracks from different supernodes. Let ω_j be the set of tracks maintained by supernode $j \in \{1, \dots, N_{ss}\}$. Let $Y = \{\tau_i(t) \in \omega_j : 1 \leq t \leq T, 1 \leq i \leq |\omega_j|, 1 \leq j \leq N_{ss}\}$ be the combined observations only from the

established tracks. We form a new set of tracks ω_{init} from $\{\tau_i \in \omega_j : 1 \leq i \leq |\omega_j|, 1 \leq j \leq N_{ss}\}$ while making sure that constraints defined in Section III-B are satisfied. Then we run Algorithm 1 on this combined observation set Y with the initial state ω_{init} .

VII. SIMULATION RESULTS

For simulations below, we consider the surveillance over a rectangular region on a plane, $\mathcal{R} = [0, 100]^2$. The state vector is $x = [x, y, \dot{x}, \dot{y}]^T$ where (x, y) is a position in \mathcal{R} along the usual x and y axes and (\dot{x}, \dot{y}) is a velocity vector. The following linear dynamic and measurement models are used

$$\begin{aligned} x_{t+\delta} &= A_\delta x_t + G_\delta w_t \\ y_t &= C x_t + v_t, \end{aligned} \quad (6)$$

where δ is a sampling interval, w_t and v_t are white Gaussian noises with zero mean and covariance $Q = \text{diag}(.15^2, .15^2)$ and R (set according to Fig. 3 (left)), respectively, and

$$A_\delta = \begin{bmatrix} 1 & 0 & \delta & 0 \\ 0 & 1 & 0 & \delta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} G_\delta = \begin{bmatrix} \frac{\delta^2}{2} & 0 \\ 0 & \frac{\delta^2}{2} \\ \delta & 0 \\ 0 & \delta \end{bmatrix} C = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}^T.$$

We assume a 100×100 sensor grid, in which the separation between sensors is normalized to 1. So the unit length in simulation is the length of the sensor separation. In all simulations, $R_t = 10$, $n_{mc} = 1000$, $w_s = 10$ and the window is forward by a single step. For the sensor model, we use $\alpha = 2$, $\gamma = 1$, $\eta = 2$, and $\beta = 3(1 + \gamma R_s^\alpha)$.

Since the number of targets is not fixed, it is difficult to measure the performance of an algorithm using a standard criterion such as the mean square error. Hence, we use two separate metrics to measure performance: the estimation error in the number of targets ϵ_K and the estimation error in position ϵ_X . Let K_t^* be the number of targets at time t and K_t be the estimated number of targets at time t . We define

$$\epsilon_K = \frac{1}{T} \sum_{t=1}^T |K_t - K_t^*|. \quad (7)$$

The computation of ϵ_X is done when it makes sense. At any t , there can be at most $M_t = \min(K_t, K_t^*)$ common tracks. We find M_t matches between true tracks and estimated tracks based on positions at $t-1, t, t+1$. For each match i , let $x_t^*(i)$ and $x_t(i)$ be the position of the true track and the estimated track at t , respectively. We define

$$\epsilon_X^2 = \frac{1}{\sum M_t} \sum_{t=1}^T \sum_{i=1}^{M_t} \|x_t(i) - x_t^*(i)\|^2. \quad (8)$$

We first evaluate the effect of the sensing range and empirically find that there is an optimal value at which the estimation error is minimized. Then we illustrate the robustness of our algorithm against sensor localization error, transmission failures and communication delays. We then give an example of surveillance with sensor networks and demonstrate how the hierarchical MCMCDA algorithm works.




Fig. 3. (left) estimation error ϵ_X (single target, Monte Carlo simulation of 1000 samples); (middle) estimation error ϵ_K (10 targets); (right) estimation error ϵ_M (10 targets) - as functions of sensing range R_s

A. Sensing Range

When localizing a single target, we can minimize the localization error by allowing more sensors to collaborate, which is equivalent to increasing R_s as shown in Fig 3 (left). But when there is more than one target, this is no longer true, since observations from different targets can collide, giving missing observations and observations away from target positions. Fig. 3 (middle) and (right) show the estimation errors ϵ_K and ϵ_X when 10 targets appear and disappear at random times and $T = 50$. The speeds \dot{x} and \dot{y} of slow-speed vehicles are between 0 and 1 unit length per sampling period while $\dot{x}, \dot{y} \in [1, 2]$ for medium-speed vehicles and $\dot{x}, \dot{y} \in [2, 5]$ for high-speed vehicles. For each vehicle type, we used five different scenarios and an example is shown in Fig. 4 (left). When $R_s = .5$, the sensors do not cover the surveillance region \mathcal{R} and do not detect a target at all times, hence, the estimation error is higher. As we increase R_s , estimation errors increase, since there are more collisions among observations of different targets. The estimation errors are low for high-speed vehicles since it is easier to disambiguate crossing targets. We find that $R_s = 1.5$ is a good range for all types of vehicles and it is used in simulations below. We can also interpret this result in terms of sensor density for a fixed value of R_s . Hence, once the surveillance region is fully covered by sensors, a further increase in density does not improve the estimation error.

B. Sensor Localization Error

The localization of sensor nodes in an ad-hoc wireless sensor network, without expensive hardware such as the global positioning system (GPS), is a challenging problem [15]. Hence, an algorithm which utilizes sensor positions needs to be robust against the sensor localization error. Suppose that the true position of sensor node i is $s_i^* = s_i + w_i$, where w_i are Gaussian noises with zero mean and covariance $\Sigma = \text{diag}(\sigma^2, \sigma^2)$. Fig 4 (middle) and (right) show the estimation errors from tracking 10 targets as functions of the sensor localization error σ . It shows that the algorithm is robust against the sensor localization error and, for $\sigma \leq .5$, the algorithm performs as if there is no sensor localization error. Notice that ϵ_K is always under 1.8, so the algorithm finds most tracks for all σ . But ϵ_X gets larger at high σ , since the target position estimation was

based on incorrect node positions. However, considering the fact that ϵ_X is computed from the norm of a vector in \mathbb{R}^2 , ϵ_X is mostly due to the sensor localization error.

C. Transmission Failures

To assess the effects of transmission failures alone, we assume that there are no delayed observations, no false alarms, and no missing detections. A single supernode is placed at the center. As mentioned earlier, transmission failures are missing observations and Fig. 5 (left) shows the ratio between the number of lost packets and the number of total packets as a function of the transmission failure rate p_{te} . As p_{te} increases, we lose more packets and, at $p_{te} \approx .9$, we lose all packets. Fig. 5 (middle) and (right) show the estimation errors and the algorithm performs well for $p_{te} \leq .4$. Notice that when $p_{te} = .4$ more than 50% of packets are lost. It shows that our algorithm is very robust against transmission failures. The estimation error ϵ_X is low at high p_{te} since the algorithm loses the most of tracks at high p_{te} and ϵ_X is computed with a small number of samples.

D. Communication Delays

As in the previous section, we assume that there are no transmission failures, no false alarms, and no missing observations. Fig. 6 (left) shows the ratio between the number of delayed packets and the number of packets as a function of the communication delay rate p_{de} . As p_{de} increases to 1, all packets are delayed. Since $w_s = 10$, we do not receive all the delayed packets and the ratio between the number of delayed packets that are eventually received and the number of packets is shown as a dotted line in Fig. 6 (left). The estimations errors are shown in Fig. 6 (middle) and (right). It shows a good performance for $p_{de} \leq .6$; this is when the most of delayed packets are received. Clearly, the performance can be improved if we increase w_s .

E. An Example of Surveillance with Sensor Networks

In this section, we give an example of surveillance with sensor networks. The surveillance region \mathcal{R} is divided into four quadrants and sensors in each quadrant form a tracking group, where a supernode is placed at the center of each quadrant. We used $p_{te} = .3$, $p_{de} = .3$, and $\eta = 2$.




Fig. 4. (left) a scenario used in simulation (numbers are target appearance and disappearance times, initial positions are marked by circles); (middle) estimation error ϵ_K as a function of σ ; (right) ϵ_X as a function of σ




Fig. 5. (left) ratio between no. of lost packets and no. of total packets; (middle) estimation error ϵ_K ; (right) ϵ_X - as functions of p_{te}




Fig. 6. (left) ratio between no. of delayed packets and no. of total packets; (middle) estimation error ϵ_K ; (right) ϵ_X - as functions of p_{de}

The surveillance duration is increased to $T = 100$. The scenario is shown in Fig. 8 (left). Fig. 7 (left) shows the accumulated fused observations at the sensor level. There were a total of 1174 observations and 603 observations were false alarms. Fig. 7 (middle) shows the observations received by supernodes and the delayed observations are circled in Fig. 7 (right). Notice that we solve the multiple-target tracking problem with observations shown in Fig. 7 (middle), not those in Fig. 7 (left). A total of 319 packets out of 1174 packets were lost due to transmission failures and 449 packets out of 855 received packets were delayed. The tracks estimated by the algorithm are shown in Fig. 8 (middle) and Fig. 8 (right). Fig. 8 (middle) shows the tracks estimated by supernodes while Fig. 8 (right) shows the tracks estimated by the track-level data association step. Fig. 8 (right) shows that the track-level data association step corrects mistakes made by supernodes due to missing observations. The ability to correct mistakes made by a lower-level agent is another strength of our algorithm. The

algorithm is written in C++ and MATLAB and run on PC with a 2.6-GHz Intel Pentium 4 processor. It takes less than 0.06 seconds per supernode, per simulation time step.

VIII. CONCLUSIONS

In this paper, a scalable hierarchical multiple-target tracking algorithm for sensor networks is presented. The algorithm is based on the efficient MCMC data association algorithm and it is suitable for autonomous surveillance in sensor networks. This new multiple-target tracking algorithm can initiate and terminate tracks and requires a small amount of memory. The algorithm is also robust against transmission failures, communication delays and sensor localization error. In order to reduce the communication overhead, observations are first locally fused and then transmitted to its supernode. The task of tracking is done hierarchically by forming a tracking group around a supernode and later combining tracks from different supernodes. The algorithm also features an ability to correct mistakes made by a lower-level agent. The simulation




Fig. 7. (left) accumulated observations at the sensor level from $t = 1$ to $t = T$; (middle) accumulated observations received by supernodes; (right) accumulated observations received by supernodes with delayed observations circled




Fig. 8. (left) a scenario used in Section VII-E (numbers are target appearance and disappearance times, initial positions are marked by circles); (middle) tracks estimated by supernodes superimposed; (right) tracks estimated by the track-level data association step of hierarchical MCMCDA

results show that the algorithm is well suited for sensor networks where transmission failures and communication delays are frequent.

REFERENCES

- [1] Y. Bar-Shalom and T.E. Fortmann. *Tracking and Data Association*. Mathematics in Science and Engineering Series 179 Academic Press, San Diego, CA, 1988.
- [2] I. Beichl and F. Sullivan. The metropolis algorithm. In *Computing in Science and Engineering*, volume 2(1), pages 65–69, 2000.
- [3] C.Y. Chong, S. Mori, and K.C. Chang. Multitarget multisensor tracking. In Y. Bar-Shalom, editor, *Multitarget-Multisensor Tracking: Advanced Applications*, pages 247–295. Artech House: Norwood, MA, 1990.
- [4] Mark Coates. Distributed particle filters for sensor networks. In *Proc. of 3rd workshop on Information Processing in Sensor Networks (IPSN)*, April 2004.
- [5] J.B. Collins and J.K. Uhlmann. Efficient gating in data association with multivariate distributed states. In *IEEE Trans. Aerospace and Electronic Systems*, volume 28(3), 1992.
- [6] David Culler, Deborah Estrin, and Mani Srivastava. Overview of sensor networks. In *IEEE Computer, Special Issue in Sensor Networks*, Aug. 2004.
- [7] L. Doherty, B. A. Warneke, B. Boser, and K. S. J. Pister. Energy and performance considerations for smart dust. In *International Journal of Parallel and Distributed Sensor Networks*, Dec 2001.
- [8] Deborah Estrin, David Culler, Kris Pister, and Gaurav Sukhatme. Connecting the physical world with pervasive networks. In *IEEE Pervasive Computing*, volume 1(1), pages 59–69, 2002.
- [9] M. Jerrum and A. Sinclair. The markov chain monte carlo method: An approach to approximate counting and integration. In Dorit Hochbaum, editor, *Approximations for NP-hard Problems*. PWS Publishing, Boston, MA, 1996.
- [10] Thomas Kurien. Issues in the design of practical multitarget tracking algorithms. In Y. Bar-Shalom, editor, *Multitarget-Multisensor Tracking: Advanced Applications*. Artech House: Norwood, MA, 1990.
- [11] D. Li, K. Wong, Yu Hen Hu, and A. Sayeed. Detection, classification and tracking of targets. In *IEEE Signal Processing Magazine*, volume 17-29, March 2002.
- [12] J.J. Liu, J. Liu, M. Chu, J.E. Reich, and F. Zhao. Distributed state representation for tracking problems in sensor networks. In *Proc. of 3rd workshop on Information Processing in Sensor Networks (IPSN)*, April 2004.
- [13] J.J. Liu, J. Liu, J. Reich, P. Cheung, and F. Zhao. Distributed group management for track initiation and maintenance in target localization applications. In *Proc. of 2nd workshop on Information Processing in Sensor Networks (IPSN)*, April 2003.
- [14] Seapahn Meguerdichian, Farinaz Koushanfar, Gang Qu, and Miodrag Potkonjak. Exposure in wireless ad hoc sensor networks. In *Procs. of 7th Annual International Conference on Mobile Computing and Networking*, pages 139–150, July 2001.
- [15] XuanLong Nguyen, Michael I. Jordan, and Bruno Sinopoli. A kernel-based learning approach to ad hoc sensor network localization. In *AAAI-2004 Workshop on Sensor Networks*, July 2004.
- [16] Songhwai Oh, Stuart Russell, and Shankar Sastry. Markov chain monte carlo data association for general multiple-target tracking problems. In *43rd IEEE Conference on Decision and Control (to appear)*, Paradise Island, Bahamas, Dec. 2004.
- [17] A.B. Poore. Multidimensional assignment and multitarget tracking. In *Partitioning Data Sets. DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, volume 19, pages 169–196, 1995.
- [18] D.B. Reid. An algorithm for tracking multiple targets. In *IEEE Transaction on Automatic Control*, volume 24(6), pages 843–854, December 1979.
- [19] J. Shin, L. Guibas, and F. Zhao. A distributed algorithm for managing multi-target identities in wireless ad-hoc sensor networks. In *Proc. of 2nd workshop on Information Processing in Sensor Networks (IPSN)*, April 2003.