

Markov Chain Monte Carlo Data Association for General Multiple Target Tracking Problems

Songhwai Oh, Stuart Russell, Shankar Sastry

Abstract—In this paper, we consider the general multiple target tracking problem in which an unknown number of targets appears and disappears at random times and the goal is to find the tracks of targets from noisy observations. We propose an efficient real-time algorithm that solves the data association problem and is capable of initiating and terminating a varying number of tracks. We take the data-oriented, combinatorial optimization approach to the data association problem but avoid the enumeration of tracks by applying a sampling method called Markov chain Monte Carlo (MCMC). The MCMC data association algorithm can be considered as a deferred logic since its decision about forming a track is based on the current and past observations. But, at the same time, it can be considered as an approximation to the optimal Bayesian filter. The algorithm shows remarkable performance compared to the greedy algorithm and the multiple hypothesis tracker (MHT) under the extreme conditions, such as a large number of targets in a dense environment, low detection probabilities, and a large number of false alarms.

I. INTRODUCTION

The multiple target tracking plays an important role in many areas of engineering such as surveillance, computer vision, and signal processing [1], [4]. Under the most general setup, a varying number of indistinguishable targets is moving around in a region with continuous motions and the positions of moving targets are sampled at random intervals. The measurements about the positions are noisy, with detection probability less than one, and there is a noise background of spurious position reports, i.e., false alarms. Targets arise at random in space and time. Each target persists independently for a random length of time and ceases to exist. A track of a target is defined as a path in space-time traveled by the target. The essence of the multiple target tracking problem is to find tracks from the noisy observations and it requires solutions to both data association and state estimation problems [16].

The data association problem in multiple target tracking is described as a problem of finding a partition of observations such that each element of a partition is a collection of observations generated by a single target or clutter [16]. However, due to the noises in state transitions and observations, we cannot expect to find the exact solution. This data-oriented view of data association has been applied and extended by many authors [9], [17], [15], [8], [5], [14]. The most successful multiple target tracking algorithm

based on this view is the multiple hypothesis tracker (MHT) [15]. In MHT, each hypothesis associates past observations with a target and, as a new set of observations arrives, a new set of hypotheses is formed from the previous hypotheses. Each hypothesis is scored by its posterior and the algorithm returns a hypothesis with the highest score as a solution. MHT is categorized as a deferred logic [14] in which the decision about forming a new track or removing an existing track is delayed until enough observations are collected. Hence, MHT is capable of initiating and terminating a varying number of tracks and suitable for surveillance applications in which an autonomous tracker is required. However, the construction of new hypotheses requires an enumeration of all possibilities and the size of hypotheses grows exponentially. The initial implementation and later extensions proposed several heuristics, such as pruning, gating, clustering and N -scan-back logic, to reduce the complexity of the problem [15], [8]. However, the heuristics are used at the expense of the optimality and the algorithm can still suffer in a dense environment. Furthermore, the running time at each step of the algorithm cannot be bounded easily, making it difficult to be deployed in a real-time surveillance system. As a method of pruning, an efficient method of finding k -best hypotheses based on the algorithm by Murty [10] is developed in [5].

A different approach to the data association problem is the joint probabilistic data association filter (JPDAF) [1]. JPDAF is a suboptimal single-stage approximation to the optimal Bayesian filter. Given a fixed number of targets, JPDAF enumerates all possible associations between the *latest* set of observations and the known tracks and clutter and computes each association weight. For each association, the conditional expectation of the state of a target is estimated by a filtering algorithm. Then, the state of a target is estimated by summing over the conditional expectations weighted by association weights. JPDAF is a sequential tracker in which the associations between the known targets and the latest observations are made sequentially and the associations made in the past are not reversible [14]. Since only the current set of observations is considered, JPDAF cannot initiate or terminate tracks. Also JPDAF assumes a fixed number of targets and requires a good initial state for each target. There are restricted extensions to JPDAF to allow the formation of a new track (see [4] and references within). Other multiple target tracking algorithms include the multisensor multitarget mixture reduction (MTMR) [12] and the probabilistic multi-hypothesis tracker (PMHT) [18]. But they also assume a fixed number of targets and cannot

This work was supported by DARPA F30602-00-2-0538 and ARO MURI DAAD 19-OZ-1-0383.

The authors are with the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, {sho,russell,sastry}@eecs.berkeley.edu.

initiate or terminate tracks. Recently, a Bayesian model-based approach for tracking an unknown number of targets which can initiate and terminate tracks is presented in [11].

The sequential trackers are more efficient than deferred logic trackers such as MHT but they are prone to make erroneous associations [14]. In addition, it is inferred that the exact calculation at each stage is NP-hard [3] since the related problem of finding the permanent of a 0-1 matrix is #P-complete [19]. In [6], a single-stage data association problem is considered and a leave-one-out heuristic is developed to avoid the enumeration of all possible associations. Later, the approach is extended to a multi-stage data association problem using Markov chain Monte Carlo (MCMC) [13].

The data association problem of multiple target tracking formulated under the data-oriented view is also known to be NP-hard [14]. Hence, we cannot expect to find an optimal solution in polynomial time unless $P = NP$. An optimization approach to data association has been applied as a 0-1 integer programming problem [9] and as a multidimensional assignment problem [14]. In both cases one needs to find a feasible set of tracks from all possible tracks to prevent the exponential explosion and compute the cost of each feasible track, such as the negative log likelihood. Then the optimization routine finds a subset from the feasible tracks such that the combined costs are minimized while satisfying the constraints, i.e., each track has at most one observation at each time and no two tracks share the same observation. The gating method similar to the ones described in [17], [15] is used to find a feasible set of tracks. However, in a dense environment, the size of the feasible tracks can be very large and the complexity of the optimization routine increases dramatically, since the number of parameters in the optimization routine depends on the number of feasible tracks.

The main contribution of this paper is the development of an efficient real-time algorithm that solves the data association problem and is capable of initiating and terminating a varying number of tracks. We take the data-oriented, combinatorial optimization approach to the data association problem but avoid the enumeration of tracks by applying a sampling method called Markov chain Monte Carlo (MCMC). The immediate benefit of using MCMC is the low memory requirement. The MCMC data association (MCMCDA) algorithm can be considered as a deferred logic since its decision about forming a track is based on the current and past observations. But, at the same time, it can be considered as an approximation to the optimal Bayesian filter if it is used to approximate the association probabilities or expectations such as the average link travel time as done in [13]. So, from the Bayesian point of view, the algorithm can be considered as a generalization of [13] to handle an unknown number of objects, missing observations and false alarms. MCMCDA shows remarkable performance compared to the greedy algorithm and MHT under the extreme conditions such as a large number of

targets in a dense environment, low detection probabilities, and a large number of false alarms.

The remainder of this paper is structured as follows. We formally state the (discrete-time) general multiple target tracking problem in Section II. In Section III, we present a general purpose MCMCDA algorithm for multiple target tracking. The algorithm is applied in simulation to extreme situations and its performance is compared with the greedy algorithm and MHT in Section IV.

II. 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 $[t_i^k, t_f^k] \subset [1, T]$ 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_i^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

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

where $w_t^k \in \mathbb{R}^d$ are white noise processes. The noisy observation about the state of the object is measured with the 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 which includes 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.

Under the data-oriented approach, the multiple target tracking problem is to partition the observations such that

the posterior is maximized, i.e., the maximum a posteriori (MAP) estimate. Under the Bayesian approach, if we are given a function defined on Ω , the collection of all partitions of observations (see below for its definition), we seek the expected value of the function given the observations. The MAP estimate found under the data-oriented approach may not be robust in the Bayesian sense. But it is sometimes more convenient when estimating parameters whose dimension is dependent on the number of tracks, such as the states of targets. Since the size of $x_t = (x_t^1, \dots, x_t^K)^T$ depends on the number of tracks K , the estimation of x_t without fixing the number of tracks is not meaningful. Hence, under the Bayesian approach, if a single set of state estimation is required, we might first estimate the most likely number of targets and then estimate the expected values of states given the estimated number of targets.

B. Probabilistic Model

Let us first specify the dynamic and measurement models. Here we use the usual linear system model but the method can be easily extended to non-linear models coupled with a non-linear regression algorithm. If an object is observed k times at t_1, t_2, \dots, t_k , its dynamic and measurement models can be expressed as:

$$\begin{aligned} x_{t_{i+1}} &= A(t_{i+1} - t_i)x_{t_i} + G(t_{i+1} - t_i)w_{t_i} \\ y_{t_i} &= Cx_{t_i} + v_{t_i} \quad \text{for } i = 1, \dots, k, \end{aligned} \quad (1)$$

where w_{t_i} and v_{t_i} are white Gaussian noises with zero mean and covariance Q and R , respectively. $A(\cdot)$, $G(\cdot)$, and C are matrices with appropriate sizes. The entries of the matrix $A(t_{i+1} - t_i)$ and $G(t_{i+1} - t_i)$ are determined by the sampling interval $t_{i+1} - t_i$ for each i . For clarity, the subsequence notation for the time index is suppressed for now. Let \bar{x}_t be the expected value of x_t given y_1, \dots, y_{t-1} ; \bar{P}_t be the covariance of x_t given y_1, \dots, y_{t-1} ; \hat{x}_t be the expected value of x_t given y_1, \dots, y_t ; and \hat{P}_t be the covariance of x_t given y_1, \dots, y_t .

Let $y_t = \{y_t^j : j = 1, \dots, n_t\}$ and $Y = \bigcup_{t \in \{1, \dots, T\}} y_t$. 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 $k = 1, \dots, K$ and $t = 1, \dots, T$; 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. We denote by $\tau_k(t)$ the observation at time t that is assigned to the track τ_k . Notice that $\tau_k(t)$ can be empty if it is a missing observation. The fourth requirement says that a track can have at most one observation at each time, but, in the case of multiple 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. Once a partition $\omega \in \Omega$ is chosen, the tracks $\tau_1, \dots, \tau_K \in \omega$ and a set of false alarms $\tau_0 \in \omega$ are completely determined. Hence, for each track, we can estimate the states of an object independently since each object moves independently from the other objects. For each track $\tau \in \omega$, we apply the Kalman filter to estimate the states $\bar{x}_t(\tau)$ and covariances $B_t(\tau)$, where $B_t(\tau) = C\bar{P}_t(\tau)C^T + R$ is the conditional observation covariance at time t for the track τ .

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:

$$\begin{aligned} P(\omega|Y) &= \frac{1}{Z} \prod_{t=1}^T p_z^{z_t} (1 - p_z)^{c_t} p_d^{d_t} (1 - p_d)^{u_t} \lambda_b^{a_t} \lambda_f^{f_t} \\ &\times \prod_{\tau \in \omega \setminus \{\tau_0\}} \prod_{i=1}^{|\tau|-1} \mathcal{N}(\tau(t_{i+1}) | \bar{x}_{t_{i+1}}(\tau), B_{t_{i+1}}(\tau)), \end{aligned} \quad (2)$$

where Z is a normalizing constant and $\mathcal{N}(\cdot | \mu, \Sigma)$ is the Gaussian density function with mean μ and covariance matrix Σ . Now under the data-oriented, combinatorial optimization approach, our goal is to find a partition of observations such that $P(\omega|Y)$ is maximized.

III. MCMC DATA ASSOCIATION ALGORITHM

In this section, we develop an MCMC sampler to solve the multiple target tracking problem. Solving complex problems by sampling methods such as Markov chain Monte Carlo (MCMC) has become more tractable, due to the increased computational power. 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 which finds a good approximate solution to a complex problem in polynomial time [7]. MCMC techniques have been applied to the complex probability distribution integration problems, counting problems such as #P-complete problems, and combinatorial optimization problems [7], [2]. The MCMC approach applied to the combinatorial optimization problems is generally known as simulated annealing.

The set Ω becomes a state space of the MCMC sampler and we sample from Ω such that its stationary distribution is $P(\omega|Y)$. 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

$$A(\omega, \omega') = \min \left(1, \frac{P(\omega'|Y)q(\omega', \omega)}{P(\omega|Y)q(\omega, \omega')} \right), \quad (3)$$

otherwise the sampler stays at ω , so that the detailed balance is satisfied. If we make sure that the chain is irreducible and aperiodic, then the chain converges to its stationary distribution. The sampler 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.

We index each move by an integer such that $m = 1$ for a birth move, $m = 2$ for a death move and so on. The move m is chosen randomly from the distribution $\xi_K(m)$ where K is the number of tracks of the current partition ω . When there is no track, we can only propose a birth move, so we set $\xi_0(m = 1) = 1$ and 0 for all other moves. When there is only a single target, we cannot propose a merge or track switch move, so $\xi_1(m = 4) = \xi_1(m = 8) = 0$. For other values of K and m , we assume $\xi_K(m) > 0$. The MCMC data association (MCMCDA) algorithm is described in Algorithm 1. The inputs are the set of all observations Y , the number of samples n_{mc} , and the initial state ω_{init} . At each step of the algorithm, ω is the current state of the Markov chain. The acceptance probability $A(\omega, \omega')$ is defined in (3) where the posterior (2) is used.

Algorithm 1 (MCMC Data Association (MCMCDA)):

Input: Y, n_{mc}, ω_{init}
Output: $\hat{\omega} = \arg \max p(\omega(n)|Y)$

```

 $\omega \leftarrow \omega_{init}$ 
for  $n = 1$  to  $n_{mc}$ 
  sample  $m$  from  $\xi_K(\cdot)$ 
  propose  $\omega'$  based on  $m$  and  $\omega$  (described below)
  sample  $U$  from  $\text{Unif}[0, 1]$ 
  if  $U < A(\omega, \omega')$ 
     $\omega \leftarrow \omega'$ 
  end
   $\omega(n) \leftarrow \omega$ 
end

```

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 can simply keep the partition with the maximum posterior. If the algorithm is used to estimate $\mathbb{E}_{P(\omega|Y)} f(\omega)$ for some bounded function f , we will need burn-in samples and need to maintain the sufficient statistics for the desired expectation.

In order to make the algorithm more efficient, we make two additional assumptions: (1) the maximal directional speed of any target in \mathcal{R} is less than \bar{v} ; and (2) the number of consecutive missing observations of any track is less than \bar{d} . The first assumption is reasonable in a surveillance scenario since, in many cases, the maximal speed of a vehicle is generally known based on its type and terrain conditions. The second assumption is a user defined parameter and it can be used as one of the criteria to distinguish an event of a new object's appearance from an event of a continuation of an existing object. We will now assume that these two

new conditions are added to the definition of Ω so each element $\omega \in \Omega$ satisfies these two additional assumptions.

We now introduce a data structure which is used to propose a new partition ω' in Algorithm 1. We define a neighborhood tree of observations as

$$L_d(y_t^j) = \{y_{t+d}^k \in Y_{t+d} : \|y_t^j - y_{t+d}^k\| \leq d \cdot \bar{v}\}$$

for $d = 1, \dots, \bar{d}$, $j = 1, \dots, n_t$ and $t = 1, \dots, T - 1$. Here $\|\cdot\|$ is the usual Euclidean distance. This neighborhood tree groups temporally separated observations based on their distances. The parameter d allows missing observations. The use of this neighborhood tree makes the algorithm more scalable since distant observations will be considered separately and makes the computations of the proposal distribution easier. It is similar to the clustering technique used in MHT but $L_d(\cdot)$ is fixed for a given set of observations Y .

We now describe each move of the sampler in detail. First, let $\zeta(d)$ be a distribution of a random variable d taking values from $\{1, 2, \dots, \bar{d}\}$. We assume the current state of the chain is $\omega = \omega^0 \cup \omega^1 \in \Omega$, where $\omega^0 = \{\tau_0\}$ and $\omega^1 = \{\tau_1, \dots, \tau_K\}$. The proposed partition is denoted by $\omega' = \omega'^0 \cup \omega'^1 \in \Omega$. Note the abuse of notation below with indexing of time, i.e., when we say $\tau(t_i)$, t_i means the time at which a target corresponding to the track τ is observed i times.

A. Birth and Death Moves

For a birth move, we increase the number of tracks from K to $K' = K + 1$ and select t_1 uniformly at random (u.a.r.) from $\{1, \dots, T - 1\}$ as an appearance time of a new track. Let $\tau_{K'}$ be the track of this new object. Then we choose d_1 from the distribution ζ . Let $L_{d_1}^1 = \{y_{t_1}^j : L_{d_1}(y_{t_1}^j) \neq \emptyset, y_{t_1}^j \notin \tau_k(t_1), j = 1, \dots, n_{t_1}, k = 1, \dots, K\}$. $L_{d_1}^1$ is a set of observations at t_1 such that, for any $y \in L_{d_1}^1$, y does not belong to other tracks and y has at least one descendant in $L_{d_1}(y)$. We choose $\tau_{K'}(t_1)$ u.a.r. from $L_{d_1}^1$. If $L_{d_1}^1$ is empty, the move is rejected since the move is not reversible. Once the initial observation is chosen, we then choose the subsequent observations for the track $\tau_{K'}$. For $i = 2, 3, \dots$, we choose d_i from ζ and choose $\tau_{K'}(t_i)$ u.a.r. from $L_{d_i}(\tau_{K'}(t_{i-1})) \setminus \{\tau_k(t_{i-1} + d_i) : k = 1, \dots, K\}$ unless this set is empty. But, for $i = 3, 4, \dots$, the process of adding observations to $\tau_{K'}$ terminates with probability γ , where $0 < \gamma < 1$. If $|\tau_{K'}| \leq 1$, the move is rejected. We then propose this modified partition where $\omega'^1 = \omega^1 \cup \{\tau_{K'}\}$ and $\omega'^0 = \{\tau_0 \setminus \tau_{K'}\}$. For a death move, we simply choose k u.a.r. from $\{1, \dots, K\}$ and delete the k -th track and propose a new partition where $\omega'^1 = \omega^1 \setminus \{\tau_k\}$ and $\omega'^0 = \{\tau_0 \cup \tau_k\}$.

B. Split and Merge Moves

For a split move, we select $\tau_s(t_r)$ u.a.r. from $\{\tau_k(t_i) : |\tau_k| \geq 4, i = 2, \dots, |\tau_k| - 2, k = 1, \dots, K\}$. Then we split the track τ_s into τ_{s_1} and τ_{s_2} such that $\tau_{s_1} = \{\tau_s(t_i) : i = 1, \dots, r\}$ and $\tau_{s_2} = \{\tau_s(t_i) : i = r + 1, \dots, |\tau_s|\}$. The modified track partition becomes $\omega'^1 = (\omega^1 \setminus \{\tau_s\}) \cup$

$\{\tau_{s_1}\} \cup \{\tau_{s_2}\}$ and the false alarm partition ω^0 is updated accordingly. For a merge move, we consider the set

$$M = \{(\tau_{k_1}(t_f), \tau_{k_2}(t_1)) : \tau_{k_2}(t_1) \in L_{t_1-t_f}(\tau_{k_1}(t_f)), \\ f = |\tau_{k_1}| \text{ for } k_1 \neq k_2, 1 \leq k_1, k_2 \leq K\}.$$

We select a pair $(\tau_{s_1}(t_f), \tau_{s_2}(t_1))$ u.a.r. from M . The tracks are combined into a single track $\tau_s = \tau_{s_1} \cup \tau_{s_2}$. Then we propose a new partition where $\omega^{t_1} = (\omega^1 \setminus (\{\tau_{s_1}\} \cup \{\tau_{s_2}\})) \cup \{\tau_s\}$ and ω^0 with appropriate rearrangements.

C. Extension and Reduction Moves

In a track extension move, we select a track τ u.a.r. from K available tracks in ω . We reassign observations for τ after the disappearance time $t_{|\tau|}$ as done in the track birth move. For a track reduction move, we select a track τ u.a.r. from K available tracks in ω and r u.a.r. from $\{2, \dots, |\tau| - 1\}$. We shorten the track τ to $\{\tau(t_1), \dots, \tau(t_r)\}$ by removing the observations assigned to τ after the time t_{r+1} .

D. Track Update Move

In a track update move, we select a track τ u.a.r. from K available tracks in ω . Then we pick r u.a.r. from $\{1, 2, \dots, |\tau|\}$ and reassign observations for τ after the time t_r as done in the track birth move.

E. Track Switch Move

For a track switch move, we select a pair of observations $(\tau_{k_1}(t_p), \tau_{k_2}(t_q))$ from two different tracks such that, $\tau_{k_1}(t_{p+1}) \in L_d(\tau_{k_2}(t_q))$ and $\tau_{k_2}(t_{q+1}) \in L_{d'}(\tau_{k_1}(t_p))$, where $d = t_{p+1} - t_q$, $d' = t_{q+1} - t_p$ and $0 < d, d' \leq \bar{d}$. Then we let

$$\tau_{k_1} = \{\tau_{k_1}(t_1), \dots, \tau_{k_1}(t_p), \tau_{k_2}(t_{q+1}), \dots, \tau_{k_2}(t_{|\tau_{k_2}|})\} \\ \tau_{k_2} = \{\tau_{k_2}(t_1), \dots, \tau_{k_2}(t_q), \tau_{k_1}(t_{p+1}), \dots, \tau_{k_1}(t_{|\tau_{k_1}|})\}.$$

Theorem 1: Suppose that $0 < p_z, p_d < 1$ and $\lambda_b, \lambda_f > 0$. If $\zeta(d) > 0$, for all $d \in \{1, \dots, \bar{d}\}$, then the Markov chain designed by Algorithm 1 is irreducible.

Proof: The birth and death moves are sufficient to illustrate the irreducibility of the chain. Since $0 < p_z, p_d < 1$ and $\lambda_b, \lambda_f > 0$, $P(\omega|Y) > 0$ for all $\omega \in \Omega$. Take an arbitrary partition $\omega \in \Omega$, say $\omega = \{\tau_0, \tau_1, \dots, \tau_K\}$. Now consider the partition $\omega' \in \Omega$, such that $\omega' = \{\tau'_0\}$, i.e., ω' assigns all observations as false alarms. Since ω is arbitrary, the chain is irreducible if the chain can move from ω' to ω and from ω to ω' . For the move from ω' to ω , consider K consecutive birth moves: $\omega_0 = \omega', \omega_1 = \{\{\tau'_0 \setminus \tau_1\}, \tau_1, \dots, \omega_K = \{\{\tau'_0 \setminus \{\cup_{i=1}^K \tau_i\}\}, \tau_1, \dots, \tau_K\} = \omega$. Since $\omega \in \Omega$, all tracks τ_k are legal, i.e., τ_k satisfies the constraints described in Section II-B and, for $i = 1, \dots, |\tau_k| - 1$, $\tau_k(t_{i+1}) \in L_d(\tau_k(t_i))$ for $1 \leq d = t_{i+1} - t_i \leq \bar{d}$. Thus, $\omega_k \in \Omega$ for all k . Because $\zeta(d) > 0$ and all tracks τ_k are legal, the probability of proposing τ_k at ω_{k-1} by the birth move is positive and $q(\omega_k, \omega_{k+1}) > 0$. For the move from ω to ω' , consider K consecutive death moves: $\omega_K = \omega, \omega_{K-1}, \dots, \omega_0 = \omega'$. The probability of removing the track τ_k at ω_k by the death move is positive and $q(\omega_{k+1}, \omega_k) > 0$. Since $P(\omega_k|Y) > 0$

for all k , the chain can move from ω' to ω and from ω to ω' . Hence, the chain is irreducible. ■

The Markov chain designed by Algorithm 1 is irreducible (Theorem 1) and aperiodic since there is always a positive probability of staying at the current state in the track update move. In addition, the transitions described in Algorithm 1 satisfy the detailed balance. Hence, by the ergodic theorem, the chain converges to its stationary distribution. Notice that the other moves are designed to improve the performance of the algorithm.

Algorithm 2 (Greedy Multiple Target Tracking):

Input: Y, σ (threshold function)

Output: $\omega = \omega^0 \cup \omega^1$

```

 $\omega^1 \leftarrow \emptyset$ 
for  $t = 1$  to  $T - 1$ 
  repeat
     $\omega_t \leftarrow \emptyset$ 
     $G \leftarrow \{(y_t^i, y_t^j) : 1 \leq s - t \leq \bar{d} \text{ and } y_t^i, y_t^j \notin \tau \in \omega^1\}$ 
    foreach  $(y_t^i, y_t^j)$  in  $G$ 
       $\tau_t \leftarrow \emptyset$ 
      estimate an initial state from  $(y_t^i, y_t^j)$ 
       $\tau_t(t_1) \leftarrow y_t^i$  and  $\tau_t(t_2) \leftarrow y_t^j$ 
       $r \leftarrow s$ 
      while  $r < T$ 
        for  $d = 1$  to  $\bar{d}$ 
           $B \leftarrow \{y \in L_d(\tau_t(r)) : y \notin \tau \in \omega^1\}$ 
          if  $B \neq \emptyset$ 
             $\tau_t(r+d) \leftarrow \arg \min_{y \in B} \|y - \bar{x}_{r+d}(\tau_t)\|$ 
            break
          end
        end
         $r \leftarrow r + d$ 
      end
      if  $p(\tau_t|Y) \geq \sigma(|\tau_t|)$ 
         $\omega_t \leftarrow \omega_t \cup \{\tau_t\}$ 
      end
    end
     $\omega^1 \leftarrow \omega^1 \cup \{\arg \max_{\tau \in \omega_t} p(\tau|Y)\}$ 
  until  $\omega_t \neq \emptyset$ 
end
 $\omega^0 \leftarrow \{Y \setminus (\cup_{\tau \in \omega^1} \tau)\}$ 

```

IV. SIMULATION RESULTS

For the simulations we consider the surveillance over a rectangular region on a plane, $\mathcal{R} = [0, L] \times [0, L] \subset \mathbb{R}^2$. The state vector is $x = [x, y, \dot{x}, \dot{y}]^T$ where (x, y) is a position on \mathcal{R} along the usual x and y axes and (\dot{x}, \dot{y}) is a velocity vector. The linear system model (1) is used where δ is an interval between observations and

$$A(\delta) = \begin{bmatrix} 1 & 0 & \delta & 0 \\ 0 & 1 & 0 & \delta \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} G(\delta) = \begin{bmatrix} \frac{\delta^2}{2} & 0 \\ 0 & \frac{\delta^2}{2} \\ \delta & 0 \\ 0 & \delta \end{bmatrix} C = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}^T$$

The covariance matrices are $Q = \text{diag}(100, 100)$ and $R = \text{diag}(25, 25)$.

The complexity of the multiple target tracking problems can be measured by several metrics: (1) the intensity of the false alarm rate λ_f ; (2) the detection probability p_d ; (3) the number of tracks K ; and (4) the density of tracks.




Fig. 1. NCA (left), ICAR (middle), and the estimated number of tracks (right) as functions of a number of tracks

The problem gets more challenging with increasing λ_f , decreasing p_d , increasing K , and increasing density of tracks. The number of tracks itself may not make the problem more difficult if they are scattered apart. But the difficulty arises when there are many tracks that are moving closely and crossing each other. This is when the ambiguity of data association is greater. Hence, we only consider situations in which tracks move very closely so we can control the density of tracks by the number of tracks. We study the performance of the MCMCDA algorithm against the greedy tracker (see Algorithm 2) and MHT by varying the parameters listed above.

Based on our model described above, we have generated different scenarios. In particular, in all cases, except for the online tracking, a half of new objects appears from the left bottom quadrant of \mathcal{R} and the other half of new objects appears from the right bottom quadrant. The actual initial positions are chosen randomly from each quadrant. They all move diagonally so each group of tracks crosses the other group in the middle of \mathcal{R} . Also targets move very close to each other and there are also crossovers within each group. The situations we have used for simulations below include very extreme cases and, in our opinion, such complex situations have not appeared in the multiple target tracking literatures.

Since the number of targets is not fixed, it is difficult to compare algorithms using a standard criterion such as the residual mean square error. Hence, we introduce two new metrics to measure the effectiveness of each data association algorithm. Let ω^* be the true partition with which the test case was generated. For $\omega \in \Omega$, we represent the set of all associations in ω as $\text{SA}(\omega) = \{(\tau, t_i^T, t_{i+1}^T) : i = 1, \dots, |\tau| - 1, \tau \in \omega\}$, where t_i^T is the time at which the track τ is observed i times. Let $\text{CA}(\omega) = \{(\tau, t, s) \in \text{SA}(\omega) : \tau(t) = \tau^*(t), \tau(s) = \tau^*(s), \tau^* \in \omega^*\}$ be the set of correct associations in ω relative to ω^* . The two new metrics we will be using are the normalized correct associations (NCA) and incorrect-to-correct association ratio (ICAR):

$$\text{NCA}(\omega) = \frac{|\text{CA}(\omega)|}{|\text{SA}(\omega^*)|} \quad (4)$$

$$\text{ICAR}(\omega) = \frac{|\text{SA}(\omega)| - |\text{CA}(\omega)|}{|\text{CA}(\omega)|}. \quad (5)$$




Fig. 2. Average running time vs. number of tracks

NCA measures the ratio between the number of correct associations and the number of associations in the true partition while ICAR measures the number of incorrect associations per correct association. We measure the performance of each algorithm by NCA, ICAR, the estimated number of tracks and the running time of the algorithm.

Both MCMCDA and greedy algorithms are written in C++ with Matlab interfaces. We have used the C++ implementation of MHT from [5]¹, which implements pruning, gating, clustering, N -scan-back logic and k -best hypotheses. The parameters for MHT are fine-tuned so that it gives similar performance when there are 10 targets: the maximum number of hypotheses in a group is 1000, the maximum track tree depth is 5, and the maximum Mahalanobis distance is 5.9. All simulations are run on a PC with a 2.6-GHz Intel processor.

A. Experiment 1 (Number of Tracks)

In this experiment, we vary K from 5 to 100 (the actual values of K are 5, 10, 20, 30, 40, 50, 75 and 100). The other parameters are held fixed: $\mathcal{R} = [0, 1000] \times [0, 1000]$, $T = 10$, $\lambda_f V = 1$, $\bar{v} = 130$ unit lengths per unit time. The main focus of this experiment is to test the performance of data association of MCMCDA against other algorithms so the tracks are detected at all times, however, we have

¹<http://www.ee.ucl.ac.uk/~icox/>

set $p_d = .9$ for the prior calculation. We have also set $\bar{d} = 1$. Since all tracks are observed, the number of observations increases as the number of tracks increases. For each value of K , we randomly generated five tests. The results for MCMCDA are the average values over 10 repeated runs and the initial state is initialized with the greedy algorithm and 10,000 samples are used. The average NCAs, ICARs and the estimated numbers of tracks for three different algorithms are shown in Fig. 1. The running times of three algorithms are shown in Fig. 2 (the running time of MCMCDA includes the initialization step). Although the maximum number of hypotheses of 1000 per group is a large number, with increasing number of tracks, the performance of MHT deteriorates as the optimality is compromised by pruning. But both greedy and MCMCDA keep good performance throughout. In addition, the running times of both greedy and MCMCDA are significantly less than that of MHT.

B. Experiment II (False Alarms)

Now the settings are the same as Experiment I but we vary the false alarm rates while the number of tracks is fixed at $K = 10$. The test cases for this experiment are prepared as follows. We first generated five different random scenarios each with 10 tracks. Then, we applied different false alarm rates to generate test cases. The false alarm rates are varied from $\lambda_f V = 1$ to $\lambda_f V = 100$ with an increment of 10. 10,000 samples are used for MCMCDA and the results for MCMCDA are the average values over 10 repeated runs. The average NCAs, ICARs and the estimated numbers of tracks for three different algorithms at different false alarm rates are shown in Fig. 3. It shows remarkable performance of MCMCDA at very high false alarm rates while the other two algorithms perform poorly. The greedy algorithm scores higher in NCA than MCMCDA but poorly in ICAR. In addition, it reports spurious tracks at high false alarm rates. Notice that MHT does not make any correct associations at high false alarm rates, $\lambda_f V \geq 80$, so ICARs for MHT at $\lambda_f V \geq 80$ are not reported.

C. Experiment III (Detection Probability)

In this experiment we vary the detection probability p_d from .3 to .9 with an increment of .1 while keeping the other parameters as the previous experiments except $K = 10$, $\lambda_f V = 1$, $T = 15$ and $\bar{d} = 5$. Now the tracks are not observed all the time. For each value of p_d , five test cases are randomly generated and the average NCAs, ICARs and the estimated numbers of tracks are shown in Fig. 4. For MCMCDA, we present two cases: MCMC(15K) with 15,000 samples and MCMC(150K) with 150,000 samples. It shows that MCMCDA outperforms at low detection probabilities. At high detection probability region, MHT scores higher than MCMCDA but it reports a higher number of tracks, meaning that it fragments tracks.

Although, in theory, MHT gives an optimal solution in the sense of MAP, it performs poorly when the detection

TABLE I
PERFORMANCE OF ONLINE MCMCDA TRACKER
RUNNING TIMES (RT) IN SECONDS

K	Number of samples					
	1,000			5,000		
	NCA	ICAR	RT	NCA	ICAR	RT
100	.95	.19	.06	.98	.13	.28
200	.94	.06	.09	.97	.05	.41
300	.92	.07	.11	.97	.05	.55

probability is low or the false alarm rate is high due to the heuristics such as pruning and N -scan-back techniques used to reduce the complexity. The heuristics are required parts of MHT in practice. Without the pruning and N -scan-back logic, the problem complexity grows exponentially fast even for a small problem. In practice, MHT with heuristics works well when there is a few number of hypotheses which carry the most of weights. When the detection probability is low or the false alarm rate is high, there are many hypotheses with low weights and there is no small set of dominating hypotheses, so MHT cannot perform well. In addition, when the detection probability is high, MHT again suffers from a large number of observations. Another noticeable benefit of the MCMCDA algorithm is that its running time can be regulated by the number of samples and the number of observations but the running time of MHT depends on the complexity of the problem instance and is not predictable in advance.

D. Online MCMCDA Multiple Target Tracker

The extension of MCMCDA to an online, real-time tracking is a trivial task. As we have seen from the previous experiments, MCMCDA works well even when T is small. Hence, we implement a sliding window of size w_s using Algorithm 1. At each time, we use the previous estimate to initialize MCMCDA and run MCMCDA on the observations belonging to the current window. A total of three test cases are generated: (case 1) 100 tracks, (case 2) 200 tracks and (case 3) 300 tracks. The surveillance duration is increased to $T = 1000$ and the surveillance region is now $\mathcal{R} = [0, 10000] \times [0, 10000]$. The other parameters are: $\lambda_f V = 10$, $p_d = .9$, $\bar{d} = 3$, $\bar{v} = 230$ and $w_s = 10$. The objects appear and disappear at random in time and space so the number of tracks changes in time. These test cases represent instances of the general (discrete-time) multiple target tracking problem. The average NCAs and ICARs over the sliding window and the average execution time per simulation time are shown in Table I. Notice that MCMCDA achieves excellent performance in all cases with less than one second of execution time.

V. CONCLUSIONS

The general (discrete-time) multiple target tracking problem is described and an MCMCDA algorithm is proposed. Our MCMCDA tracker, a data association algorithm capable of initiating and terminating a varying number of tracks, is flexible and easy to incorporate any domain




Fig. 3. NCA (left), ICAR (middle), and the estimated number of tracks (right) as functions of false alarm rate




Fig. 4. NCA (left), ICAR (middle), and the estimated number of tracks (right) as functions of detection probability

specific knowledge to make it more efficient. Instead of searching over the whole solution space, the MCMC algorithm randomly searches over the space where the posterior is concentrated. Our simulation results show remarkable performance of the MCMCDA algorithm under the extreme conditions such as a large number of targets in a dense environment, low detection probabilities, and a large number of false alarms. We have shown that the algorithm can be extended as an online, real-time algorithm with excellent performance.

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] 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.
- [4] I.J. Cox. A review of statistical data association techniques for motion correspondence. In *International Journal of Computer Vision*, volume 10(1), pages 53–66, 1993.
- [5] I.J. Cox and S.L. Hingorani. An efficient implementation of reid’s multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking. In *International Conf. on Pattern Recognition*, pages 437–443, 1994.
- [6] Timothy Huang and Stuart J. Russell. Object identification in a bayesian context. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pages 1276–1283, 1997.
- [7] 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.
- [8] 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.
- [9] C. L. Morefield. Application of 0-1 integer programming to multitarget tracking problems. In *IEEE Trans. on Automatic Control*, volume 22, pages 3:302–312, June 1971.
- [10] K.G. Murty. An algorithm for ranking all the assignments in order of increasing cost. In *Operations Research*, volume 16, pages 682–687, 1968.
- [11] Songhwai Oh, Jin Kim, and Shankar Sastry. A sampling-based approach to nonparametric dynamic system identification and estimation. In *American Control Conference*, Boston, MA, June 2004.
- [12] L.Y. Pao. Multisensor multitarget mixture reduction algorithms for tracking. In *Proc. AIAA Guidance, Navigation, and Control Conf.*, pages 28–37, Monterey, CA, Aug. 1993.
- [13] Hanna Pasula, Stuart J. Russell, Michael Ostland, and Yaacov Ritov. Tracking many objects with many sensors. In *IJCAI-99*, Stockholm, 1999.
- [14] 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.
- [15] D.B. Reid. An algorithm for tracking multiple targets. In *IEEE Transaction on Automatic Control*, volume 24(6), pages 843–854, December 1979.
- [16] R.W. Sittler. An optimal data association problem on surveillance theory. In *IEEE Trans. on Military Electronics*, volume MIL-8 (Apr), pages 125–139, April 1964.
- [17] J.J. Stein and S.S. Blackman. Generalized correlation of multi-target track data. In *IEEE Trans. Aerosp. Electron. Syst.*, volume AES-11(6), pages 1207–1217, November 1975.
- [18] R. Streit and T. Luginbuhl. Maximum likelihood method for probabilistic multi-hypothesis tracking. In *Proc. SPIE*, volume 2235, pages 394–405, April 1994.
- [19] L.G. Valiant. The complexity of computing the permanent. In *Theoretical Computer Science*, volume 8, pages 189–201, 1979.