MULTIPLE TARGET TRACKING WITH A PIXELIZED SENSOR

Mark R. Morelande, Christopher M. Kreucher*[◇†] *and Keith Kastella*[†]

ABSTRACT

This paper describes a computationally efficient method for tracking multiple moving targets. The method is predicated on estimation of the joint multitarget probability density (JMPD), which is a single probabilistic entity capturing uncertainty in both the number of targets and the states of the individual targets. The non-Gaussian/nonlinear measurement model adopted here does not permit exact computation of the JMPD so some method of approximation is required. A novel particle filtering algorithm for recursive estimation of the JMPD is proposed which provides computational tractability by automatically factoring the high dimensional multitarget state when applicable and using the optimal importance density with Rao-Blackwellisization on each of the factors. The efficiency of the proposed algorithm is shown via simulation results using real, recorded target trajectories.

1. INTRODUCTION

In this paper we investigate the problem of tracking multiple moving targets from a collection of noisy, ambiguous sensor measurements. The approach is based on estimation of the joint multitarget conditional probability density (JMPD) [8]. The JMPD is a single probabilistic entity that simultaneously captures uncertainty in the number of targets and the states of the individual targets although here we restrict our attention to the case where the number of targets is known and fixed. The algorithm developed here forms the basis of an algorithm for the more general case of an unknown and time varying target number [13]. Our work is preceded by many theoretical works including [7][10][16][12] and many implementational approaches including [9][6][14][11].

The measurement model adopted here divides the observation region into cells or pixels with returns in each cell determined probabilistically by the number of targets in the cell and the signal-to-noise ratio. Under such conditions the JMPD cannot be computed exactly so that numerical techniques are required for its representation and propagation through time. The high dimensionality of the JMPD dictates the use of sophisticated numerical procedures. We advocate a particle filtering approach [3], which is a numerical method of solving nonlinear filtering equations. In this approach, the JMPD is represented by a set of samples (particles) and associated weights. Particles are propagated through time via an importance density and particle weights are updated via a weighting equation.

The main contribution of this paper is a multitarget tracking algorithm that recursively estimates the JMPD using a particle filter with a carefully designed importance density. The method takes advantage of the fact that when groups of targets are well separated in sensor space there is no measurement ambiguity and therefore the groups may be treated independently. Additionally, proposals for each target group are done using the exact optimal importance density (OID) with Rao-Blackwellisization of the unobserved state space parameters. These two features combine to result in a computationally efficient multitarget tracking algorithm.

We demonstrate via simulation that this carefully designed importance density allows for reliable tracking of ten targets using sample sizes as small as 50 particles for high SNRs. In the simulations target trajectories are taken from a set of actual targets recorded during an army battle drill.

The paper is organized as follows. The target dynamic and measurement models are given in Section 2. A brief review of the particle filter approximation to the JMPD is given in Section 3. Section 4 contains the proposed algorithm and Section 5 provides a brief performance analysis. Conclusions are given in Section 6.

2. NOTATION AND MODELING

Consider the presence of r targets with the state of the ith target at time k denoted as $\mathbf{x}_i^k \in \mathbb{R}^{n_x}$, $i = 1, \ldots, r$. The multitarget state at time k is defined as the concatenation of the individual target states, $\mathbf{X}^k = (\mathbf{x}_1^{k\prime}, \ldots, \mathbf{x}_r^{k\prime})'$. It is assumed that each target moves independently in a plane with the individual target states composed of position and velocity in each direction. The position elements of the ith target are collected into $\boldsymbol{\rho}_i^k$ and the velocity elements are collected into $\mathbf{v}_i^k$ so that $\mathbf{x}_i^k = (\boldsymbol{\rho}_i^{k\prime}, \mathbf{v}_i^{k\prime})'$. The individual target states evolve according to

$$\mathbf{x}_i^k | \mathbf{x}_i^{k-1} \sim N(\mathbf{F}\mathbf{x}_i^{k-1}, \mathbf{Q}_i^k) \tag{1}$$

where $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, $\mathbf{Q}_i^k$ is the covariance matrix for the process noise of the ith target and

$$\mathbf{F} = \begin{pmatrix} 1 & T \\ 0 & 1 \end{pmatrix} \otimes \mathbf{I}_2$$

with $\mathbf{I}_m$ denoting the $m \times m$ identity matrix, $\otimes$ denoting the Kronecker product and T denoting the sampling period.

The observation region is divided into C cells with the measurement vector $\mathbf{z}^k = (z_1^k, \ldots, z_C^k)'$ containing the measurements obtained in each cell. The position elements of the targets are collected into $\mathbf{P}^k = (\boldsymbol{\rho}_1^{k\prime}, \ldots, \boldsymbol{\rho}_r^{k\prime})'$. Measurements are made independently in each cell with the distribution of the measurement in the jth cell depending on the number of targets residing in the cell. Let $o_j(\mathbf{P}^k)$ denote the number of targets occupying the jth cell. Then,

$$p(\mathbf{z}^k | \mathbf{P}^k) = \prod_{j=1}^{C} l_{o_j(\mathbf{P}^k)}(z_j^k) \tag{2}$$

In the case of non-thresholded measurements, $z_j^k \in \mathbb{R}$ and $z_j^k | \mathbf{P}^k$ is Rayleigh distributed with parameter $1 + o_j(\mathbf{P}^k)\lambda$ where λ is the

* CRC for Sensor, Signal and Information Processing, The University of Melbourne, Australia, m.morelande@ee.mu.oz.au, ◇ University of Michigan, Ann Arbor, Michigan, USA, † General Dynamics, Ann Arbor, Michigan, USA

signal-to-noise ratio (SNR). For thresholded measurements, $z_j^k \in \{0,1\}$ with $z_j^k = 1$ corresponding to a target detection in the jth cell. The threshold is set so that $\mathsf{P}(z_j^k = 1 | o_j(\mathbf{P}^k) = 0) = P_{FA}$. We then have, for $j = 1, \ldots, C$,

$$l_{o_j(\mathbf{P}^k)}(z_j^k) = \begin{cases} P_{FA}^{1/(1+o_j(\mathbf{P}^k)\lambda)}, & z_j^k = 1, \\ 1 - P_{FA}^{1/(1+o_j(\mathbf{P}^k)\lambda)}, & z_j^k = 0. \end{cases}$$

3. APPROXIMATION OF THE JOINT MULTITARGET PROBABILITY DENSITY FUNCTION

Solution of the multitarget tracking problem involves recursive computation of the JMPD $p(\mathbf{X}^k|\mathbf{Z}^k)$ where $\mathbf{Z}^k = \{\mathbf{z}^1, \ldots, \mathbf{z}^k\}$ is the measurement history [8]. Assuming that $p(\mathbf{X}^{k-1}|\mathbf{Z}^{k-1})$ is available, the JMPD at time k can be computed via the Chapman-Kolmogorov-Bayes recursion:

$$p(\mathbf{X}^k|\mathbf{Z}^k) = \frac{p(\mathbf{z}^k|\mathbf{X}^k)}{p(\mathbf{z}^k|\mathbf{Z}^{k-1})} \int p(\mathbf{X}^k|\mathbf{X}^{k-1}) p(\mathbf{X}^{k-1}|\mathbf{Z}^{k-1})\, d\mathbf{X}^{k-1} \quad (3)$$

Closed-form solution of (3) is generally not possible so it is necessary to use an approximation. Early work in this area used a deterministic grid approximation which is computationally feasible only for simple problems involving a small number of targets moving in one dimension [8]. More recently, particle filters (PFs) have been used to approximate the JMPD in realistic scenarios involving large numbers of targets moving in two-dimensions [9]. A similiar approach will be used in this paper although the technique will be refined to better handle situations in which several targets are in close proximity.

PFs provide a recursive stochastic grid approximation to the exact solution of Bayesian state estimation problems. The idea is that, by simulating the assumed dynamic and measurement models many times, a set of particles will be obtained which move of their own volition to the desired parts of the state-space. PFs are often implemented using the technique of sequential importance sampling. Under this scheme, given a set of particles $\mathbf{X}_1^{k-1}, \ldots, \mathbf{X}_n^{k-1}$ with weights $w_1^{k-1}, \ldots, w_n^{k-1}$ which approximates the JMPD at time $k-1$, the PF approximation to the JMPD at time k is found by performing the following steps for $t = 1, \ldots, n$

$$\mathbf{X}_t^k \sim q(\cdot|\mathbf{X}_t^0, \ldots, \mathbf{X}_t^{k-1}, \mathbf{Z}^k) \quad (4)$$

$$w_t^k \propto w_t^{k-1} \frac{p(\mathbf{z}^k|\mathbf{X}_t^k, \mathbf{Z}^{k-1}) p(\mathbf{X}_t^k|\mathbf{X}_t^{k-1})}{q(\mathbf{X}_t^k|\mathbf{X}_t^0, \ldots, \mathbf{X}_t^{k-1}, \mathbf{Z}^k)} \quad (5)$$

where q is the importance density. In order to ensure that particles remain approximately evenly weighted, a necessity for accurate approximation, the particle set should be resampled at regular intervals [5]. This material is covered in greater depth in [1].

4. AN EFFICIENT PARTICLE FILTERING ALGORITHM

It is relatively straightforward to develop a particle filtering scheme which, given a sufficient sample size, will provide an accurate approximation of the JMPD. However a solution of minimal computational expense requires careful design of the particle filter by using the inherent structure of the tracking model to reduce the amount of numerical simulation the particle filter is required to perform. In the multitarget tracking problem this is achieved through exploiting the approximate marginalization of the measurement likelihood for well-separated target clusters, the use of joint measurement-directed proposals for each target cluster and Rao-Blackwellization. The dynamic equation (1) for the ith target can then be written as

$$\boldsymbol{\rho}_i^k = \boldsymbol{\rho}_i^{k-1} + T\mathbf{v}_i^{k-1} + \boldsymbol{\epsilon}_i^k \quad (6)$$

$$\mathbf{v}_i^k = \mathbf{v}_i^{k-1} + \boldsymbol{\eta}_i^k \quad (7)$$

where

$$\begin{pmatrix} \boldsymbol{\epsilon}_i^k \\ \boldsymbol{\eta}_i^k \end{pmatrix} \sim N\left(\mathbf{0}, \begin{pmatrix} \mathbf{Q}_{i,\rho}^k & \boldsymbol{\Lambda}_i^{k\prime} \\ \boldsymbol{\Lambda}_i^k & \mathbf{Q}_{i,v}^k \end{pmatrix}\right) = N(\mathbf{0}, \mathbf{Q}_i^k) \quad (8)$$

Since (6) and (7) form a linear/Gaussian system of equations with (7) the "process" equation and (6) the "measurement" equation, the distribution of $\mathbf{v}_i^k$ conditional upon the position trajectory of the ith target is Gaussian and can be found exactly using the Kalman filter. This suggests the following decomposition of the JMPD of $\mathbf{X}^0, \ldots, \mathbf{X}^k$:

$$p(\mathbf{X}^0, \ldots, \mathbf{X}^k|\mathbf{Z}^k) = p(\mathbf{P}^0, \ldots, \mathbf{P}^k|\mathbf{Z}^k) \times \prod_{i=1}^r p(\mathbf{v}_i^0, \ldots, \mathbf{v}_i^k|\boldsymbol{\rho}_i^0, \ldots, \boldsymbol{\rho}_i^k, \mathbf{Z}^k)$$

The densities $p(\mathbf{v}_i^0, \ldots, \mathbf{v}_i^k|\boldsymbol{\rho}_i^0, \ldots, \boldsymbol{\rho}_i^k, \mathbf{Z}^k)$, $i = 1, \ldots, r$ can be computed using the Kalman filter (KF) and $p(\mathbf{P}^0, \ldots, \mathbf{P}^k|\mathbf{Z}^k)$ can be approximated using a PF.

Computation of the posterior density of the velocity elements can be performed using well-known recursions after allowing for the dependence between $\boldsymbol{\epsilon}_i^k$ and $\boldsymbol{\eta}_i^k$ in (8) [2]. The details are omitted for the sake of brevity. The posterior mean and covariance matrix for the ith target at time k are denoted as $\mathbf{v}_i^{k|k}$ and $\boldsymbol{\Sigma}_i^{k|k}$, respectively.

The JMPD of the target positions at time $k-1$ is represented by the particle set $\{\mathbf{P}_t^{k-1}, w_t^{k-1}\}_{t=1}^n$. Note that we must have $w_t^{k-1} = 1/n$ for reasons which will be explained below. This particle set can be considered to approximate the JMPD as

$$\hat{p}(\mathbf{P}^{k-1}|\mathbf{Z}^{k-1}) = 1/n \sum_{t=1}^n \delta(\mathbf{P}^{k-1} - \mathbf{P}_t^{k-1})$$

where δ is Dirac's delta function. The particle filter approximation to the JMPD at time k can then be found as

$$\hat{p}(\mathbf{P}^k|\mathbf{Z}^k) \propto p(\mathbf{z}^k|\mathbf{P}^k)/n \sum_{t=1}^n p(\mathbf{P}^k|\mathbf{P}_t^{k-1}) \quad (9)$$

The PF seeks a set of samples from (9). In the case of multitarget tracking it is desired to increase the efficiency of the sampling process by taking advantage of the approximate marginalization of the JMPD for well-separated targets. The targets are separated into $s \le r$ clusters $C_1, \ldots, C_s$ such that $\bigcup_{l=1}^s C_l = \{1, \ldots, r\}$ and $\forall l \in \{1, \ldots, s\}, \forall i \in C_l$,

$$|\hat{\boldsymbol{\rho}}_i^k - \hat{\boldsymbol{\rho}}_j^k| \le \Gamma \Rightarrow j \in C_l$$

where $\hat{\boldsymbol{\rho}}_i^k$ is the predicted position of the ith target and Γ is a threshold. The positions of the targets in the lth cluster are collected into $\mathbf{c}_l^k$. For a sufficiently large value of Γ the likelihood can be written as

$$p(\mathbf{z}^k|\mathbf{X}^k) \approx \prod_{l=1}^s \pi(\mathbf{z}^k|\mathbf{c}_l^k) \quad (10)$$

where $\pi(\mathbf{z}|\mathbf{c})$ is the density of the measurement vector $\mathbf{z}$ conditional on the target cluster $\mathbf{c}$ under the assumption that only this target cluster exists. The π notation will be used to indicate densities evaluated in this manner. The approximate marginalisation of the likelihood allows particles at time k to be constructed from target clusters which belonged to different particles at time $k-1$. When using a measurement-directed proposal density for each cluster, this greatly reduces computational expense, since the expense of jointly drawing samples for a group of targets increases exponentially with the number of targets, and improves performance, since it is easier to separately select several good clusters than it is to jointly select several good clusters.

The process of constructing particles by collecting clusters from several particles is equivalent to sampling from the distribution

$$p(\mathbf{c}_1^k, d_1, \ldots, \mathbf{c}_s^k, d_s|\mathbf{Z}^k) \propto p(\mathbf{z}^k|\mathbf{P}^k) \prod_{l=1}^{s} p(\mathbf{c}_l^k|\mathbf{c}_{l,d_l}^k) \tag{11}$$

where $d_l \in \{1, \ldots, n\}$, $l = 1, \ldots, s$ is the index of the particle from which the lth cluster will be selected. This is similiar in spirit to the formulation of the auxiliary particle filter [15].

Sampling from (11) is performed through an importance density which factorises as

$$q(\mathbf{c}_1^k, d_1, \ldots, \mathbf{c}_s^k, d_s|\mathbf{Z}^k) = \prod_{l=1}^{s} q(\mathbf{c}_l^k, d_l|\mathbf{Z}^k) \tag{12}$$

The importance density for the lth cluster can be written as

$$q(\mathbf{c}_l^k, d_l|\mathbf{Z}^k) = \psi_{l,d_l} \pi(\mathbf{c}_l^k|\mathbf{c}_{l,d_l}^0, \ldots, \mathbf{c}_{l,d_l}^{k-1}, \mathbf{Z}^k) \tag{13}$$

where $\psi_{l,t}$ is the probability of selecting the lth cluster from the tth particle. The trajectory $\mathbf{c}_l^0, \ldots, \mathbf{c}_l^{k-1}$ should be interpreted to mean the collection of position trajectories of targets in the cluster C_l at time k. Eq. (13) therefore involves a slight abuse of notation, made to reduce notational complexity, since the targets in the lth cluster will almost certainly not be the same at all time steps. The weights $\psi_{l,1}, \ldots, \psi_{l,n}$ can be chosen arbitrarily although, as discussed above, proper setting of these weights is necessary to take full advantage of the approximate marginalization. The most sensible choice seems to be the usual OID weight update [4]

$$\psi_{l,t} \propto \pi(\mathbf{z}^k|\mathbf{c}_{l,t}^{k-1}), \qquad t = 1, \ldots, n. \tag{14}$$

The suitability of this choice of weighting will become apparent when the weighting update for the reconstructed particles is derived at the end of the section.

Expressions for the JOID and the weights (14) will now be given for the case of threshold measurements. Similar expressions can be obtained for the non-thresholded case. The working has been omitted for the sake of brevity. Assume a cluster of q targets and let $\mathbf{c}^k = (\boldsymbol{\rho}_1^{k\prime}, \ldots, \boldsymbol{\rho}_q^{k\prime})'$ denote the collection of target positions and $\mathbf{c}^0, \ldots, \mathbf{c}^{k-1}$ denote a trajectory of positions for these targets from time 0 to time $k-1$. Let M_i, $i = 1, \ldots, q$ denote measurement cells in the neighbourhood of the ith target and V_j, $j = 1, \ldots, C$ denote the region of measurement space occupied by the jth measurement cell. According to this notation the JOID can be written as

$$\pi(\mathbf{c}_l^k|\mathbf{c}_l^0, \ldots, \mathbf{c}_l^{k-1}, \mathbf{Z}^k) = \sum_{j_1 \in M_1} \cdots \sum_{j_q \in M_q} \beta_{j_1,\ldots,j_q} \phi_{j_1,\ldots,j_q}(\mathbf{c}_l^k) \tag{15}$$

where

$$\beta_{j_1,\ldots,j_q} = \alpha_{j_1,\ldots,j_q} \Bigg/ \sum_{e_1 \in M_1} \cdots \sum_{e_q \in M_q} \alpha_{e_1,\ldots,e_q}$$

$$\alpha_{j_1,\ldots,j_q} = \prod_{i=1}^{q} \gamma_{i,j_i} \prod_{u=1}^{\bar{q}} \left(P_{FA}^{-m_u/(1+m_u\lambda)}\right)^{z_{\bar{j}_u}^k} \times \left(\frac{1 - P_{FA}^{1/(1+m_u\lambda)}}{1 - P_{FA}}\right)^{1 - z_{\bar{j}_u}^k}$$

$$\phi_{j_1,\ldots,j_q}(\mathbf{c}_l^k) = \prod_{i=1}^{q} \left\{ \chi_{V_{j_i}}(\boldsymbol{\rho}_i^k) N(\boldsymbol{\rho}_i^k; \hat{\boldsymbol{\rho}}_i^k, \boldsymbol{\Psi}_i^k)/\gamma_{i,j_i} \right\}$$

with $\bar{q}$ the number of distinct cell indices in $\{j_1, \ldots, j_q\}$, $\bar{j}_1, \ldots, \bar{j}_{\bar{q}}$ the distinct cell indices and $m_1, \ldots, m_{\bar{q}}$ the multiplicities of the distinct cells, $\chi_A(z) = 1$ if $z \in A$ and zero otherwise and

$$\gamma_{i,j} = \int_{V_j} N(\boldsymbol{\rho}_i^k; \hat{\boldsymbol{\rho}}_i^k, \boldsymbol{\Psi}_i^k)\, d\boldsymbol{\rho}_i^k$$

where $\hat{\boldsymbol{\rho}}_i^k = \boldsymbol{\rho}_i^{k-1} + T\hat{\mathbf{v}}_i^{k-1|k-1}$ and $\boldsymbol{\Psi}_i^k = T^2\boldsymbol{\Sigma}_i^{k-1|k-1} + \mathbf{Q}_{i,\rho}^k$. It can be seen from (15) that the JOID is a mixture of truncated Gaussian distributions. Each component can be interpreted as a hypothesis on the cell locations of the q targets in the cluster. A draw from (15) can be made by selecting a mixture component using the probabilities $\beta_{j_1,\ldots,j_q}$ and then drawing each target position from the appropriate truncated Gaussian distribution.

The collection of cells forming the neighbourhood of the ith target can be defined as

$$M_i = \{j \in \{1, \ldots, C\} : \gamma_{i,j} > \Upsilon\}$$

where Υ is a small, pre-defined lower bound.

The weights (14) are given by the normalization factor for the JOID which can be found as

$$\pi(\mathbf{z}|\mathbf{c}_l^{k-1}) = \sum_{j_1 \in M_1} \cdots \sum_{j_q \in M_q} \alpha_{j_1,\ldots,j_q}$$

It remains to compute the weight update for the reconstructed particles. Since reconstructed particles contain target clusters originating from different particles the weight assigned to a particular particle at time $k-1$ will have no connection to the reconstructed particle at time k. It must therefore be assumed that the reconstructed particles are resampled at each time step so that $w_t^k = 1/n$ for $t = 1, \ldots, n$, $k = 0, 1, \ldots$. Let $\mathbf{c}_{l,1}^k, \ldots, \mathbf{c}_{l,n}^k$ denote the collection of particles drawn for the lth target cluster. The weight of the tth particle can be found, by substituting (12), (13), (14) and (11) into (5), as

$$\tilde{w}_t^k \propto p(\mathbf{z}^k|\mathbf{c}_{1,t}^k, \ldots, \mathbf{c}_{s,t}^k) \Bigg/ \prod_{l=1}^{s} \pi(\mathbf{z}^k|\mathbf{c}_{l,t}^k) \tag{16}$$

Resampling based on $\tilde{w}_1^k, \ldots, \tilde{w}_n^k$ is performed in order to obtain an evenly weighted particle set, i.e., $w_t^k = 1/n$ for $t = 1, \ldots, n$. If the threshold Γ used in the clustering is sufficiently large the resampling step will select a large number of distinct particles so that particle duplication is minimised. In fact, in most cases the clustering threshold can be selected so that exactly uniform weights are obtained and resampling of the reconstructed particles is unnecessary. Increasing the clustering threshold above the level at which uniform weights are obtained will decrease performance while further increasing computational expense.

Eq. (16) also motivates the choice of weighting given in (14) since any other choice will not have the property that uniform weights are obtained for a sufficiently large clustering threshold. This implies that other weighting choices do not select the best clusters for reconstruction.

5. PERFORMANCE ANALYSIS

The scenario used here involves ten targets moving in a 5500m x 5500m observation region for 1000 time steps of 1s each. The target trajectories belong to real targets and were obtained from an exercise at the US Army's National Training Centre. For much of the observation period groups of as many as four targets move in close proximity. The observation region is divided into 100m x 100m cells with measurements in each cell generated from the target trajectories according to the thresholded measurement model described in Section 2. Monte Carlo realisations are obtained by generating independent measurement sequences from the same set of target trajectories.

The performance metric will be the expected number of targets in track at the end of the observation interval. This is computed from 100 Monte Carlo realisations for two values of the SNR. For each SNR the false alarm probability P_{FA} is such that the probability of registering a return in a cell occupied by one target is 0.5, i.e., $P_{FA} = e^{\ln(0.5)(1+\lambda)}$ for an SNR λ. The results given in Table 1 clearly show the efficiency of the proposed algorithm. The larger SNR considered here, $\lambda = 10$, results in a false alarm probability of $P_{FA} = 4.9 \times 10^{-4}$. Under these conditions almost perfect tracking is obtained with only 25 particles. The situation becomes considerably more challenging at an SNR of 5dB for which $P_{FA} = 0.056$. Performance deteriorates considerably for the smaller sample sizes although a sample size of 200 particles is still sufficient to provide excellent performance.

Table 1. Mean number of targets in track for the proposed algorithm across 100 realisations of 1000 time steps each

	Number of particles				
SNR (dB)	25	50	100	200	500
5	7.45	8.84	9.49	9.74	9.85
10	9.93	9.99	9.98	9.98	9.97

6. CONCLUSIONS

An efficient particle filtering algorithm for tracking a known number of moving targets was developed. The efficiency of the algorithm derives from the use of a carefully designed importance density which exploits the structure of the multitarget tracking problem. Successful tracking of ten targets across a range of signal-to-noise ratios was demonstrated using as few as 200 particles.

An important part of the general multitarget tracking problem is the need to allow for uncertainity in the number of targets present. This requires procedures for track initiation and deletion in addition to track maintenance. A forthcoming paper shows how procedures for track initiation and deletion can be incorporated into the proposed algorithm to provide an efficient solution to the general multitarget tracking problem.

7. REFERENCES

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

[2] Y. Bar-Shalom and X.-R. Li, *Estimation and Tracking: Principles, Techniques and Software*. Artech House, 1993.

[3] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Springer Publishing, 2001.

[4] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, pp. 197–208, 2000.

[5] N. Gordon, D. Salmond, and A. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," *IEE Proceedings-F*, vol. 140, no. 2, pp. 107–113, 1993.

[6] M. Isard and J. MacCormick, "BraMBLe: A Bayesian multiple-blob tracker," *Proceedings of the 8th International Conference on Computer Vision*, 2001.

[7] K. Kastella, "Event-averaged maximum liklihood estimation and mean-field theory in multitarget tracking," *IEEE Transactions on Automatic Control*, vol. 40, pp. 1070–1074, 1995.

[8] ——, "Discrimination gain for sensor management in multitarget detection and tracking," in *Proceedings of the IMACS Conference on Computational Engineering in Systems Applications*, Lille, France, 1996, pp. 167–172.

[9] C. M. Kreucher, K. Kastella, and A. O. H. III, "Tracking multiple targets using a particle filter representation of the joint mulitarget probability density," *Proceedings of SPIE Conference on Signal and Data Processing of Small Targets*, 2003.

[10] R. Mahler, "A unified foundation for data fusion," *Proceedings of SPIE*, vol. 124, pp. 325–345, 1996.

[11] S. Maskell, M. Rollason, N. Gordon, and D. Salmond, "Efficient particle filtering for multiple target tracking with application to tracking in structured images," *Procedings of SPIE Conference on Signal and Data Processing of Small Targets*, 2002.

[12] M. I. Miller, A. Srivastava, and U. Grenander, "Conditional mean estimation via jump-diffusion processes in multiple target tracking/recognition," *IEEE Transactions on Signal Processing*, vol. 43, no. 11, pp. 2678–2690, 1995.

[13] M. R. Morelande, C. M. Kreucher, and K. Kastella, "A unified Bayesian approach to multiple target detection and tracking," *IEEE Transactions on Signal Processing*, vol. submitted, 2004.

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

[15] M. Pitt and N. Shephard, "Filtering via simulation: auxiliary particle filters," *Journal of the American Statistical Association*, vol. 94, pp. 590–599, 1999.

[16] L. D. Stone, T. L. Corwin, and C. A. Barlow, *Bayesian Multiple Target Tracking*. Artech House, 1999.