

The GM-PHD Filter Multiple Target Tracker

Daniel E. Clark
ECE EPS
Heriot-Watt University
Riccarton, Edinburgh, U.K.
dec1@hw.ac.uk

Kusha Panta
CSSIP
Dept. of Electrical & Electronic Engineering
University of Melbourne
Parkville, VIC 3010, Australia
kusha@ee.unimelb.edu.au

Ba-Ngu Vo
Dept. of Electrical & Electronic Engineering
University of Melbourne
Parkville, VIC 3010, Australia
bv@ee.unimelb.edu.au

Abstract - *The Gaussian Mixture Probability Hypothesis Density Filter (GM-PHD Filter) was proposed recently for jointly estimating the time-varying number of targets and their states from a noisy sequence of sets of measurements which may have missed detections and false alarms. The initial implementation of the GM-PHD filter provided estimates for the set of target states at each point in time but did not ensure continuity of the individual target tracks. It is shown here that the trajectories of the targets can be determined directly from the evolution of the Gaussian mixture and that single Gaussians within this mixture accurately track the correct targets. Furthermore, the technique is demonstrated to be successful in estimating the correct number of targets and their trajectories in high clutter density and shows better performance than the MHT filter.*

Keywords: Tracking, data association, filtering, PHD filter, random sets.

1 Introduction

The Gaussian Mixture Probability Hypothesis Density Filter (GM-PHD Filter) provided a closed form solution to the PHD filter recursion for multiple target tracking [1]. The posterior intensity function is estimated by a sum of weighted Gaussian components whose means, weights and covariances can be propagated analytically in time. In particular, the means and covariances are propagated by the Kalman filter.

The original Gaussian Mixture PHD filter algorithm provided a means of estimating the number of targets and their states at each point in time. The method for determining the targets simply used the weights of the Gaussian components and did not take into account temporal continuity. We show that if a target is not detected at each iteration, the Gaussian components can still track the targets in the presence of some missed detections. Furthermore, the trajectory of the target in the past, before it has been detected,

can also be determined by keeping the trajectories of each of the Gaussian components.

The original formulation of the GM PHD filter allowed targets to be spawned from existing targets. For simplicity, we have removed this functionality, although it is anticipated that the algorithm presented here could be extended to incorporate this scenario.

2 The PHD Filter

In single target tracking problems, the constant gain Kalman filter provides the computationally fastest solution for approximate filtering which propagates the first-order moment of the posterior distribution. The PHD filter was proposed to provide an analogous solution in multiple target tracking problems. The first-order statistical moment of the multiple target posterior distribution, known as the PHD, is propagated instead of the posterior. The integral of the PHD over the state space provides an estimate of the number of targets and the target states can be estimated by determining the peaks of this distribution.

Sequential Monte Carlo implementations of the PHD filter were developed to provide a practical solution to the PHD recursion [2] using a particle filter approach. Practical applications of the filter have been developed and implemented on a range of systems including multiple target tracking in forward-looking sonar [3], tracking vehicles observed by humans in different terrains [4], tracking targets located on an ellipse in passive radar [5] and tracking feature points in image sequences [6].

2.1 PHD Recursion

The multiple target framework adopted here represents the multi-target states and multi-target observations as finite sets X_k and Z_k which contain the individual target positions and measurements respectively. The optimal multi-target Bayes filter propagates the multi-

target posterior density $p_k(\cdot|Z_{1:k})$, which is conditional on the sets of observations up to time k , $Z_{1:k}$, with the following recursion

$$p_{k|k-1}(X_k|Z_{1:k-1}) = \int f_{k|k-1}(X_k|X)p_{k-1}(X|Z_{1:k-1})\mu_s(dX), \quad (1)$$

$$p_k(X_k|Z_{1:k}) = \frac{g_k(Z_k|X_k)p_{k|k-1}(X_k|Z_{1:k-1})}{\int g_k(Z_k|X)p_{k|k-1}(X|Z_{1:k-1})\mu_s(dX)}, \quad (2)$$

where the dynamic model is governed by the transition density $f_{k|k-1}(X_k|X_{k-1})$ and likelihood $g_k(Z_k|X_k)$ which can be derived using finite set statistics [7] and μ_s plays the role of the Lebesgue measure, as described in [2]. Methods for implementing this recursion become intractable due to the combinatorial complexity of the densities. The PHD filter was developed to provide a computationally tractable sub-optimal alternative to this recursion by propagating the first-order statistical moment, or PHD, instead [7].

The posterior intensity function v_k is propagated by the following prediction and update recursion,

$$v_{k|k-1}(x) = \int \phi_{k|k-1}(x, \zeta)v_{k-1}(\zeta)d\zeta + \gamma_k(x), \quad (3)$$

$$v_k(x) = [1 - p_{D,k}(x)]v_{k|k-1}(x) + \sum_{z \in Z_k} \frac{\psi_{k,z}(x)v_{k|k-1}(x)}{\kappa_k(z) + \int \psi_{k,z}(\xi)v_{k|k-1}(\xi)d\xi}. \quad (4)$$

In the prediction equation (3), the transition density

$$\phi_{k|k-1}(x, \xi) = p_{S,k}(\xi)f_{k|k-1}(x|\xi) + \beta_{k|k-1}(x|\xi), \quad (5)$$

is determined from $f_{k|k-1}(x_k|x_{k-1})$, the single target transition density, $p_{S,k}$, the probability of target survival, and $\beta_{k|k-1}$, the PHD for spawned target birth from targets at time $k-1$. The intensity γ_k is the PHD for spontaneous birth of new targets at time k .

In the data update equation (4),

$$\psi_{k,z} = p_{D,k}(x)g(z|x), \quad (6)$$

where g is the single target likelihood function and $p_{D,k}$ is the probability of detection, and the intensity of clutter points $\kappa_k(z)$ is given by

$$\kappa_k(z) = \lambda_k c_k(z) \quad (7)$$

where λ_k is the Poisson parameter specifying the expected number of false alarms and c_k is the probability distribution over the measurement space.

It is assumed that each target evolves and generates observations independently of one another, clutter is independent of target-originated measurements, and that the predicted multi-target Random Finite Set (RFS) governed by the multi-target prior density $p_{k|k-1}$ is Poisson. Note that the dependence on the measurements $Z_{1:k}$ is omitted here for simplicity. In our multiple target tracking algorithm, the spawned target PHD, $\beta_{k|k-1}$ has been omitted.

2.2 Target Estimation and Track Continuity

In the Sequential Monte Carlo version of the PHD filter [2], the target estimates needed to be determined from the particle distribution by using clustering techniques such as the EM algorithm and k-means. When the estimated number of targets is not the same as the actual number of particle clusters, this can lead to inaccurate target estimation. The convergence of the empirical distribution to the true density [2] [8] [9] is not affected but the target states cannot always be determined accurately in high levels of clutter.

In addition to estimating the number of targets and their states at each point in time, it is also important in tracking scenarios to know the trajectories of the targets and to be able to distinguish between different targets. Methods for associating the targets between frames have been reported in the literature. The first of these [10] used the PHD filter for pre-filtering the data input to a Multiple Hypothesis Tracker. The second technique [11] represents the PHD in a resolution cell to differentiate the peaks of the PHD posterior, and validation gating was used to determine the weights of the particles. The PHD filter estimated the number and locations of the targets. The results of data association determined the peaks of the PHD.

More recently, two methods were presented independently in [12] and [13]. The first of these considered associating target estimates between iterations, also known as estimate-to-track association. The second method used the partitioning of the particle data to assign labels to the particles within the same cluster and associate the clusters between time frames if there is a large intersection of particles with the same label from the previous time step. This technique directly uses the empirical PHD distribution and is similar to the technique which we use in this paper.

3 The Gaussian Mixture PHD Filter

The linear Gaussian model of the targets is described here before presenting the Gaussian Mixture PHD Multiple Target Tracker.

3.1 The Linear Gaussian Model

Each target follows a linear Gaussian dynamical model,

$$f_{k|k-1}(x|\zeta) = \mathcal{N}(x; F_{k-1}\zeta, Q_{k-1}), \quad (8)$$

$$g_k(z|x) = \mathcal{N}(z; H_k x, R_k), \quad (9)$$

where $\mathcal{N}(\cdot; m, P)$ denotes a Gaussian density with mean m and covariance P , F_{k-1} is the state transition matrix, Q_{k-1} is the process noise covariance, H_k is the observation matrix, and R_k is the observation noise covariance.

The survival and detection probabilities are state independent, $p_{S,k}(x) = p_{S,k}$, and $p_{D,k}(x) = p_{D,k}$. The

intensity of the spontaneous birth is a Gaussian mixture,

$$\gamma_k(x) = \sum_{i=1}^{J_{\gamma,k}} w_{\gamma,k}^{(i)} \mathcal{N}(x; m_{\gamma,k}^{(i)}, P_{\gamma,k}^{(i)}), \quad (10)$$

where $J_{\gamma,k}$, $w_{\gamma,k}^{(i)}$, $m_{\gamma,k}^{(i)}$, $P_{\gamma,k}^{(i)}$, $i = 1, \dots, J_{\gamma,k}$, are given model parameters that determine the shape of the birth intensity.

3.2 The Gaussian Mixture PHD Multiple Target Tracker

The algorithm presented here is initialised in **Step 0** and then iterates through **Steps 1 to 5**:

Step 0: Initialisation

At time $k = 0$, the initial intensity, v_0 , is the sum of J_0 Gaussians,

$$v_0(x) = \sum_{i=1}^{J_0} w_0^{(i)} \mathcal{N}(x; m_0^{(i)}, P_0^{(i)}). \quad (11)$$

These are distributed across the state space where each Gaussian $\mathcal{N}(x; m_0^{(i)}, P_0^{(i)})$ has mean state vector $m_0^{(i)}$, covariance $P_0^{(i)}$ and weight $w_0^{(i)}$. A unique identifier, or tag, is assigned to each Gaussian to form the set

$$T_0 = \{T_0^{(1)}, \dots, T_0^{(J_0)}\}. \quad (12)$$

Set $k = 1$.

Step 1: Prediction

The predicted intensity to time k is the Gaussian mixture,

$$v_{k|k-1}(x) = v_{S,k|k-1}(x) + \gamma_k(x), \quad (13)$$

where $\gamma_k(x)$ is defined in equation (5) and

$$v_{S,k|k-1}(x) = p_{S,k} \sum_{j=1}^{J_{k-1}} w_{k-1}^{(j)} \mathcal{N}(x; m_{S,k|k-1}^{(j)}, P_{S,k|k-1}^{(j)}), \quad (14)$$

$$m_{S,k|k-1}^{(j)} = F_{k-1} m_{k-1}^{(j)}, \quad (15)$$

$$P_{S,k|k-1}^{(j)} = Q_{k-1} + F_{k-1} P_{k-1}^{(j)} F_{k-1}^T, \quad (16)$$

Concatenate the set of tags from the previous time step with new tags from the Gaussians introduced for the spontaneous birth model,

$$T_{k|k-1} = T_k \cup \{T_{\gamma_k}^{(1)}, \dots, T_{\gamma_k}^{(J_{\gamma,k})}\}. \quad (17)$$

Step 2: Update

The posterior intensity at time k is given by Gaussian Mixture,

$$v_k(x) = (1 - p_{D,k}) v_{k|k-1}(x) + \sum_{z \in Z_k} v_{D,k}(x; z) \quad (18)$$

where

$$v_{D,k}(x; z) = \sum_{j=1}^{J_{k|k-1}} w_k^{(j)}(z) \mathcal{N}(x; m_{k|k}^{(j)}(z), P_{k|k}^{(j)}), \quad (19)$$

$$w_k^{(j)}(z) = \frac{p_{D,k} w_{k|k-1}^{(j)} q_k^{(j)}(z)}{\kappa_k(z) + p_{D,k} \sum_{\ell=1}^{J_{k|k-1}} w_{k|k-1}^{(\ell)} q_k^{(\ell)}(z)}, \quad (20)$$

$$q_k^{(j)}(z) = \mathcal{N}(z; H_k m_{k|k-1}^{(j)}, R_k + H_k P_{k|k-1}^{(j)} H_k^T), \quad (21)$$

$$m_{k|k}^{(j)}(z) = m_{k|k-1}^{(j)} + K_k^{(j)}(z - H_k m_{k|k-1}^{(j)}), \quad (22)$$

$$P_{k|k}^{(j)} = [I - K_k^{(j)} H_k] P_{k|k-1}^{(j)}, \quad (23)$$

$$K_k^{(j)} = P_{k|k-1}^{(j)} H_k^T (H_k P_{k|k-1}^{(j)} H_k^T + R_k)^{-1}. \quad (24)$$

There are $(1 + |Z_k|) J_{k|k-1}$ Gaussian components, $(1 + |Z_k|)$ for each prediction term. Assign the same tag as the related prediction component (the tag is no longer unique, but this will be adjusted in the following stages) to form the set,

$$T_{k,u} = T_{k|k-1}^{u_{k|k-1}} \cup T_{k|k-1}^{z_1} \cup \dots \cup T_{k|k-1}^{z_{|Z_k|}}. \quad (25)$$

Step 3: Pruning

In the pruning stage, the components with low weights are eliminated. The intensity function v_k can be represented by the set of weights, means and covariances of the Gaussian components, $v_k = \{w_k^{(i)}, m_k^{(i)}, P_k^{(i)}\}_{i=1}^{J_k}$. Without loss of generality, let the weights $w_k^{(1)}, \dots, w_k^{(N_P)}$ be those which are below the truncation threshold τ , and let

$$\bar{v}_k := \frac{\sum_{l=1}^{J_k} w_k^{(l)}}{\sum_{j=N_P+1}^{J_k} w_k^{(j)}} \sum_{i=N_P+1}^{J_k} w_k^{(i)} \mathcal{N}(x; m_k^{(i)}, P_k^{(i)}). \quad (26)$$

Now define new weights,

$$\bar{w}_k^{(i)} = w_k^{(i)} \frac{\sum_{l=1}^{J_k} w_k^{(l)}}{\sum_{j=N_P+1}^{J_k} w_k^{(j)}}, \quad (27)$$

so that the intensity \bar{v}_k is given by the set $\bar{v}_k = \{\bar{w}_k^{(i)}, m_k^{(i)}, P_k^{(i)}\}_{i=N_P+1}^{J_k}$.

Step 4: Merging

Merge the Gaussian components of the distance between their means falls within a merging threshold U . Given $\{w_k^{(i)}, m_k^{(i)}, P_k^{(i)}\}_{i=N_P+1}^{J_k}$, a merging threshold U , and a maximum allowable number of Gaussian terms J_{max} the procedure for merging the components is as follows:

Set $\ell = 0$, and $I = \{i = 1, \dots, J_k | w_k^{(i)} > \tau\}$.

Repeat

$$\ell := \ell + 1. \quad (28)$$

$$j := \arg \max_{i \in I} w_k^{(i)}.$$

$$L := \{i \in I | (m_k^{(i)} - m_k^{(j)})^T (P_k^{(i)})^{-1} (m_k^{(i)} - m_k^{(j)}) \leq U\}.$$

$$\tilde{w}_k^{(\ell)} = \sum_{i \in L} w_k^{(i)}.$$

$$\tilde{x}_k^{(\ell)} = \frac{1}{\tilde{w}_k^{(\ell)}} \sum_{i \in L} w_k^{(i)} x_k^{(i)}.$$

$$\tilde{P}_k^{(\ell)} = \frac{1}{\tilde{w}_k^{(\ell)}} \sum_{i \in L} w_k^{(i)} (P_k^{(i)} + (\tilde{m}_k^{(\ell)} - m_k^{(i)})(\tilde{m}_k^{(\ell)} - m_k^{(i)})^T).$$

$$I := I \setminus L.$$

Until $I = \emptyset$.

If $\ell > J_{max}$ then replace $\{\tilde{w}_k^{(i)}, \tilde{m}_k^{(i)}, \tilde{P}_k^{(i)}\}_{i=1}^{\ell}$ by those of the J_{max} Gaussians with largest weights.

Output $\{\tilde{w}_k^{(i)}, \tilde{m}_k^{(i)}, \tilde{P}_k^{(i)}\}_{i=1}^{\ell}$ as the merged Gaussian components. If two or more components still have the same tag $T_k^{(i)}$, then give this to the one with the largest weight $\tilde{w}_k^{(i)}$ and reassign new tags to the other components to form $\{\tilde{w}_k^{(i)}, \tilde{m}_k^{(i)}, \tilde{P}_k^{(i)}, \tilde{T}_k^{(i)}\}_{i=1}^{\ell}$ and set $(w_k, m_k, P_k, T_k) := (\tilde{w}_k, \tilde{m}_k, \tilde{P}_k, \tilde{T}_k)$.

Step 5: Target State Estimation

The target states are determined by taking the Gaussians with weights above a given threshold and those which have previously defined to be a target, i.e. the set of live tracks from time k is

$$\hat{T}_k = \{T_k^{(i)} : w_k^{(i)} > 0.5\} \quad (29)$$

and the set of estimates is

$$\hat{X}_k = \{m_k^{(i)} : T_k^{(i)} \in \hat{T}_j, j = 1, \dots, k\}. \quad (30)$$

3.3 Determining Target Tracks

The above procedure allows the determination of the trajectories of the Gaussian components in the mixture by keeping the means associated with each identifying tag. In the original formulation of the GM PHD filter, estimates of the target states were taken at each stage of the algorithm by choosing the components with the maximum weights. In the version here, we have temporal continuity which enables us to keep track of targets when their weights fall below the desired threshold. In addition, the trajectory of the targets in the past can be determined by looking at the previous trajectory of the Gaussian after the weight is above a given threshold. Once the weight falls below another threshold, the Gaussian component is deleted indicating that it does not contribute significantly to the intensity function and so the target is not likely to still exist. Note that if $p_{D,k} < 1$, the component is not deleted when a measurement is not received for a target, so that we can continue to track even with missed detections. If the space requirements for this do not allow all of the Gaussians to be kept in memory, tracks could be deleted if the weight was not above a threshold for a specified

number of updates. This procedure is significantly better than the estimate-to-track association used in the particle implementation of the filter, which only considered estimates in the last frame and relied on the prediction instead of the updated Gaussian. This shows that the GM PHD filter has the inherent ability to track multiple targets with track continuity which shall be demonstrated in the simulations.

The probability of survival $p_{S,k}$ is adjusted for the expected lengths of the target tracks. When this is too low, target tracks are lost more often and when it is too high, the tracks continue for longer after the target has died. In the SMC version of the PHD filter, it was reported that when the probability of detection $p_{D,k}$ is low that targets are prematurely destroyed [14]. This could have been due to the particle mass being used to determine the number of targets and clustering to determine the state estimates. Similar problems were not encountered here since the weights of the Gaussians were used to determine the target states and these Gaussians were assumed to represent targets until the weight of the Gaussian fell below the pruning threshold.

4 Simulations

Simulated examples have been created to test the performance of the GM PHD Filter Multiple Target Tracker and the results of these are compared against the track-oriented Multiple Hypothesis Tracker [15] with a batch of 10 frames where the log-likelihood ratio was used to rank tracks and the best global hypothesis was selected for data outputs.

4.1 Example 1

In this example, a two-dimensional scenario with an unknown and time varying number of targets has been simulated in clutter over the region $[-1000, 1000] \times [-1000, 1000]$. The state $x_k = [p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k}]^T$, of each target consists of position $(p_{x,k}, p_{y,k})$ and velocity $(\dot{p}_{x,k}, \dot{p}_{y,k})$, while the measurement is a noisy version of the position.

Each target has survival probability $p_{S,k} = 0.9$, detection probability $p_{D,k} = 0.99$ and follows the linear Gaussian dynamics (8),

$$x_k = \begin{pmatrix} 1 & T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T \\ 0 & 0 & 0 & 1 \end{pmatrix} x_{k-1} + \begin{pmatrix} T^2/2 & 0 \\ T & 0 \\ 0 & T^2/2 \\ 0 & T \end{pmatrix} \sigma_v^2, \quad (31)$$

and observation model:

$$z_k = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} x_k + \sigma_\epsilon^2. \quad (32)$$

We assume no spawning, and that the spontaneous birth intensity is Poisson with four Gaussian terms distributed across the surveillance region,

$$\gamma_k(x) = \sum_{i=1}^4 0.1 \mathcal{N}(x; m_{\gamma,i}, P_\gamma).$$

Note that this does not need to sum to one but reflects the expected number of spontaneously appearing targets at time k .

The detected measurements are immersed in clutter that can be modelled as a Poisson RFS K_k with intensity

$$\kappa_k(z) = \lambda_k V u(z), \quad (33)$$

where $u(\cdot)$ is the uniform density over the surveillance region, $V = 4 \times 10^6 m^2$ is the area of the surveillance region, and $\lambda_k = 5 \times 10^{-6} m^{-2}$ is the average number of clutter returns per unit area which relates to 20 clutter measurements per scan.

The Gaussian mixture PHD filter, with pruning parameters elimination threshold $T = 10^{-5}$, merging threshold $U = 4$, and maximum number of Gaussian terms $J_{max} = 200$.

Figure 1 shows the simulated scenario with the true target trajectories and an average of 20 clutter points per scan. Figures 2 and 3 give the results of the PHD and MHT filters respectively on a set of measurements over 100 time-steps. The dots show the true target locations and the lines show the estimated trajectories. It can be seen that the GM PHD Filter Tracker has very few false tracks, can pick up a track very quickly, does not drop the tracks while the target still exists and eliminates tracks shortly after the target leaves the surveillance region. The MHT filter gives more false trajectories and has more occasions where the track is eliminated despite the target still existing.

Five hundred sets of measurements for these target trajectories have been generated to compare the two algorithms. The Wasserstein multi-target miss distance has been used to compare the accuracy of the estimates and also the expected absolute error in the estimated number of targets.

4.1.1 Wasserstein Distance

The Wasserstein distance from theoretical statistics was adopted as a means of defining a metric for multi-target distances which penalises when the estimate of the number of targets is incorrect [16]. When the number of targets is estimated correctly, the Wasserstein distance is the same as the Hausdorff distance but the Hausdorff does not penalise for incorrectly estimating the number of targets. This metric has been used for assessing the performance of the PHD filter [2] [10].

Let X_k be the RFS of target states at time k and \hat{X}_k be the RFS of estimated target states. The L^P Wasserstein distance between the two sets is defined as follows:

$$d_P^W(X_k, \hat{X}_k) = \inf_C \left(\sum_{x_i \in X_k} \sum_{\hat{x}_j \in \hat{X}_k} C_{ij} d(x_i, \hat{x}_j)^P \right)^{\frac{1}{P}}. \quad (34)$$

where C represents an $|X_k| \times |\hat{X}_k|$ matrix $\{C_{ij}\}$ such

that:

$$\forall i = 1 \dots |X_k|, j = 1 \dots |\hat{X}_k|: \quad (35)$$

$$\sum_{i=1}^{|X_k|} C_{ij} = \frac{1}{|\hat{X}_k|}, \quad \sum_{i=1}^{|\hat{X}_k|} C_{ij} = \frac{1}{|X_k|}, \quad C_{ij} \geq 0.$$

The L^∞ Wasserstein distance is:

$$d_\infty^W(X_k, \hat{X}_k) = \inf_C \max_{x_i \in X_k, \hat{x}_j \in \hat{X}_k} \tilde{C}_{ij} d(x_i, \hat{x}_j), \quad (36)$$

where $\tilde{C}_{ij} = 1$ if $C_{ij} > 0$ and $\tilde{C}_{ij} = 0$ if $C_{ij} = 0$.

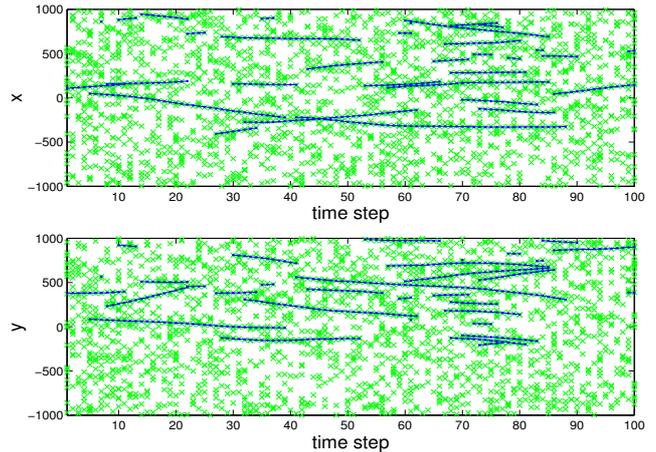


Figure 1: True target positions (lines) and measurements (crosses).

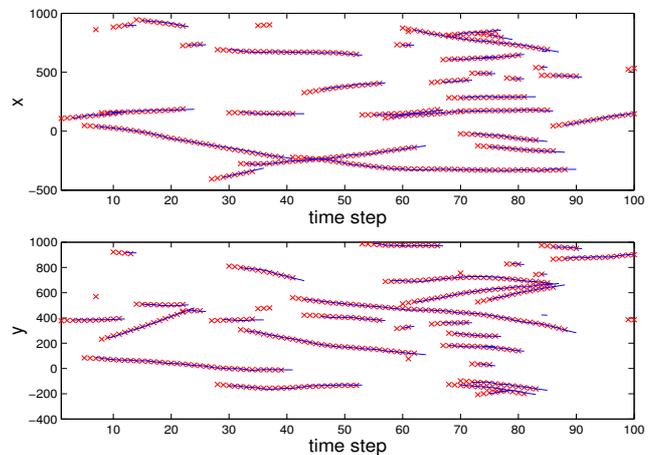


Figure 2: PHD estimated target trajectories (lines) and true positions (crosses).

When the estimated number of targets is incorrect, the Wasserstein distance puts all the weight on the outliers. Figure 4 shows the results of the mean Wasserstein distance over the 500 measurement sets for each time step. The spikes in the result for the PHD filter usually indicate that either a new target has entered the scene but has not yet been detected or has died and has not been eliminated. The PHD filter is lower than the MHT filter for most of the iterations since it can more accurately estimate the correct number of targets.

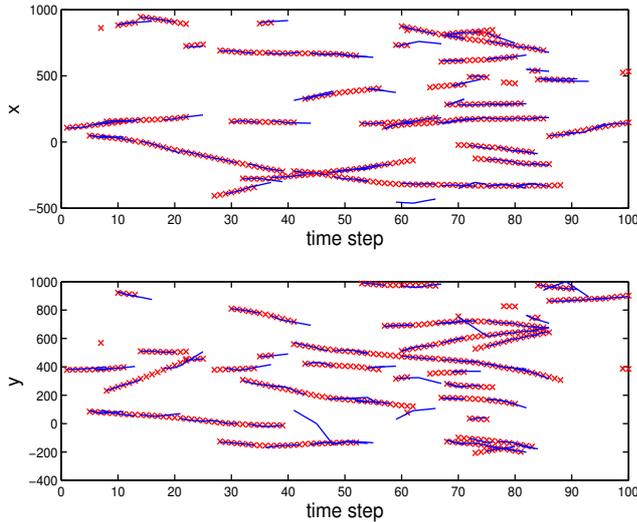


Figure 3: MHT estimated target trajectories (lines) and true positions (crosses).

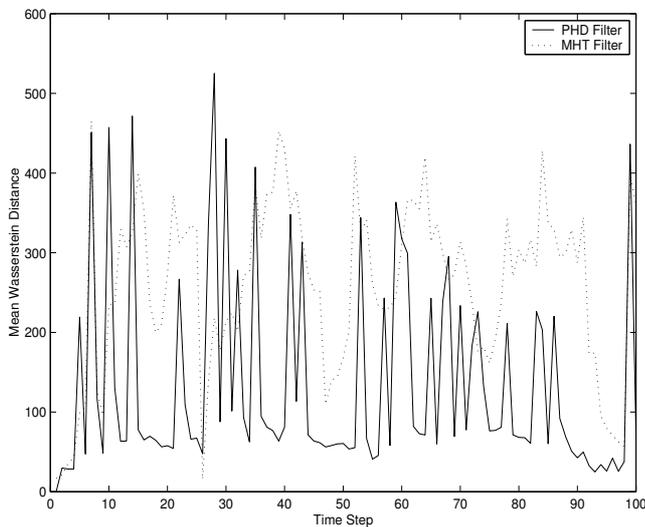


Figure 4: Mean Wasserstein Distance.

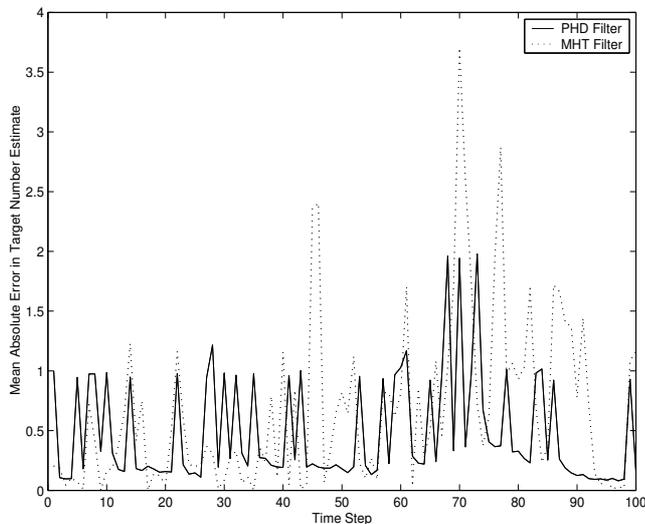


Figure 5: Absolute Error in Target Number Estimate

4.1.2 Error in Estimating the Number of Targets

The expected absolute error on the number of targets has been calculated for each of the algorithms,

$$E\{|\hat{X}_k| - |X_k|\}.$$

Note that standard performance measures such as the mean square distance error are not applicable to multi-target filters that jointly estimate number of targets and their states.

Figure 5 shows the absolute error in the estimation of the number of targets, averaged over 500 measurement sets. We can see that the estimate for the PHD filter is generally lower than for the MHT filter, the mean over all iterations is 0.4711 compared with 0.6692 for the MHT filter. The maximum error is also much less, 1.978 for the PHD filter compared with 3.707 for the MHT filter. The PHD filter can more reliably estimate the correct number of targets, it has fewer false tracks and can initiate the correct tracks more easily.

4.2 Example 2

In this example, we consider the theoretical constraints of the algorithm and illustrate this through a simulation. Consider a situation where we have two targets, then ideally this would be represented by two Gaussians,

$$v_k(x) = w_k^{(1)} \mathcal{N}(x; m_k^{(1)}, P_k) + w_k^{(2)} \mathcal{N}(x; m_k^{(2)}, P_k). \quad (37)$$

(For simplicity, it is assumed that the covariance matrix is the same for each Gaussian. This can be achieved through diagonalisation, since the covariance matrix is symmetric, nonnegative and semi-definite.)

Suppose that the targets cross, then $v_k(x)$ is unimodal with mean $(m_k^{(1)} + m_k^{(2)})/2$ when $(m_k^{(1)} - m_k^{(2)})^T P_k^{-1} (m_k^{(1)} - m_k^{(2)}) < 4$, see [17]. This means that the PHD will fail to distinguish between targets within this separation. Furthermore, these components could actually be merged into the same Gaussian if the means fall within the merging threshold, U . Thus, if the tracks of the targets are to be maintained when the targets are too close, alternative methods for data association need to be used. If the trajectories of the targets are known in the past, these could be used to separate the tracks after the targets have crossed.

A simulation of the above scenario has been created but with Gaussians from the spontaneous birth are included to ensure that if a track is lost, then it can be recaptured. Targets 1 and 2 are born at the same time but at two different locations. These two targets travel along straight lines and their tracks cross at $k = 53$ s, see figure 6 for the paths of the targets.

Two sets of measurements have been generated to show how the tracker behaves with crossing targets, figure 7 shows the crossing region with two outcomes. In the first outcome, the target trajectories are correctly estimated through the crossing point. In the second outcome however, whilst the estimates from

the PHD filter are not affected, the tracks follow the wrong trajectories after the crossing point. It is anticipated that this problem could be resolved by associating tracks from predictions before the Gaussians are in the merging region with estimates after the targets have crossed, similar to techniques used with multiple hypothesis tracking.

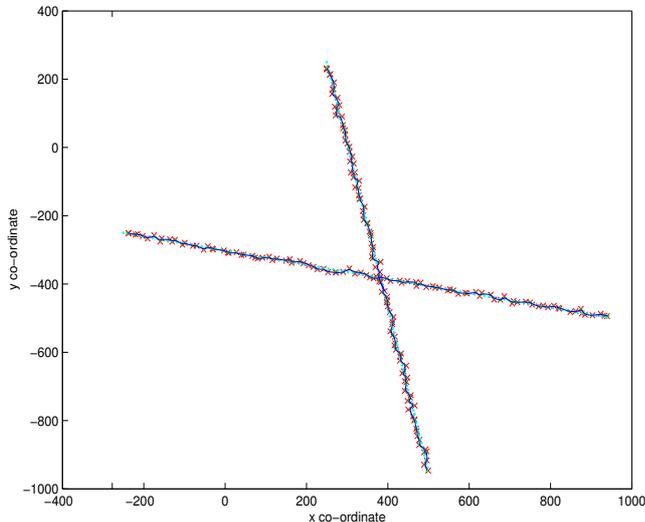


Figure 6: Example 2 - Targets Crossing.

5 Conclusions

An algorithm has been presented for tracking multiple targets in high clutter density which has the ability to estimate the number of targets, track the trajectories of the targets over time, operate with missed detections and give the trajectories of the targets in the past once a target has been identified. It has been shown to outperform the track-oriented Multiple Hypothesis Tracker in its ability to operate in clutter with fewer false tracks and can initiate and eliminate targets more accurately. The theoretical constraints of the proposed tracking algorithm have been discussed in the case of crossing targets. It is anticipated that the problem of retaining the correct target identity in this scenario can be resolved by considering the previous trajectories of targets.

References

- [1] B. Vo and W. K. Ma. The Gaussian Mixture Probability Hypothesis Density Filter. *IEEE Transactions on Signal Processing*, to appear, 2006.
- [2] B-N. Vo, S. Singh, and A. Doucet. Sequential Monte Carlo methods for Multi-target Filtering with Random Finite Sets. *IEEE Trans. Aerospace Elec. Systems*, 41, No.4:1224–1245, 2005.
- [3] D.E. Clark and J. Bell. Bayesian Multiple Target Tracking in Forward Scan Sonar Images Using the

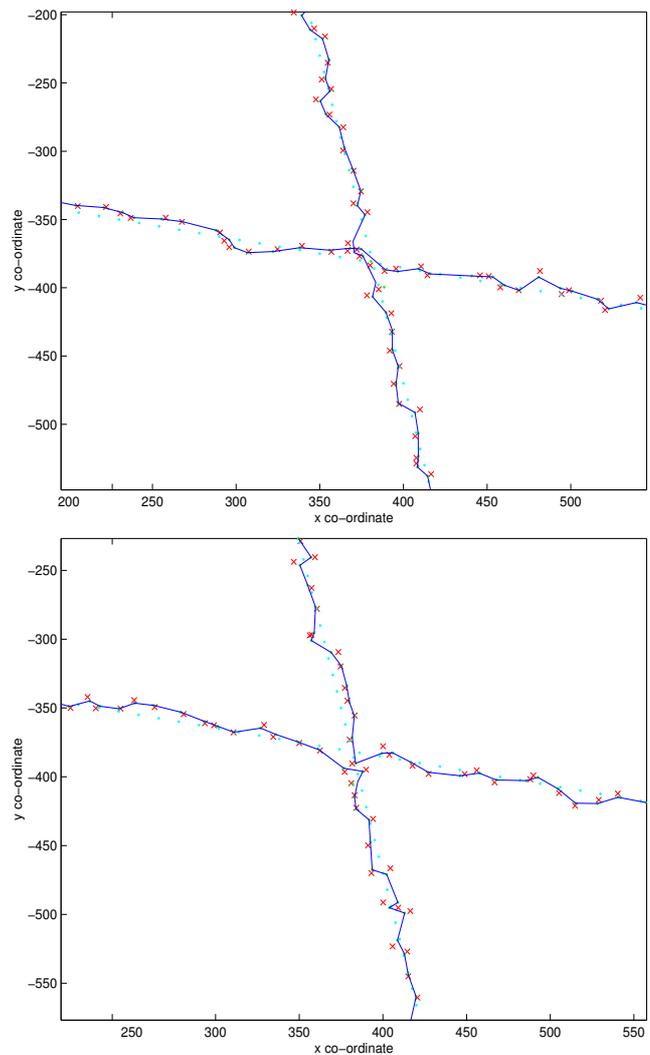


Figure 7: Example 2 - Two different results.

PHD Filter. *IEE Radar, Sonar and Navigation, Volume 152, Issue 5*, p. 327-334., 2005.

- [4] H. Sidenbladh and S.L. Wirkander. Tracking random sets of vehicles in terrain. *IEEE Workshop on Multi-Object Tracking, Madison, WI, USA*, 2003.
- [5] M. Tobias and A. D. Lanterman. Probability Hypothesis Density-based multi-target tracking with bistatic range and Doppler observations. *IEE Radar, Sonar and Navigation, Volume 152, Issue 3*, p. 195-205., 2005.
- [6] N. Ikoma, T. Uchino, and T. Maeda. Tracking of feature points in image sequence by SMC implementation of PHD filter. *ICE 2004 Annual Conference, 4-6 Aug, 2004. p 1696 - 1701 vol. 2*.
- [7] R. Mahler. Multitarget Bayes filtering via first-order multitarget moments. *IEEE Transactions on Aerospace and Electronic Systems*, 39, No.4:1152–1178, 2003.
- [8] D. E. Clark and J. Bell. Convergence Results for the Particle PHD Filter. *IEEE Transactions on Signal Processing*, to appear, 2006.

- [9] A. M. Johansen, S. S. Singh, A. Doucet, and B-N. Vo. Convergence of the SMC implementation of the PHD filter. *Technical Report CUED/F-INFENG/TR-517, University of Cambridge*, 2005.
- [10] K. Panta, B. Vo, S. Singh, and A. Doucet. Probability hypothesis density filter versus multiple hypothesis tracking. *Proceedings of SPIE – Volume 5429 Signal Processing, Sensor Fusion, and Target Recognition XIII, Ivan Kadar, Editor, August 2004*, pp. 284-295.
- [11] Lin Lin. Parameter estimation and data association for multitarget tracking. *PhD Thesis, The University of Connecticut*, 2004.
- [12] D. E. Clark and J. Bell. Data Association for the PHD Filter. *ISSNIP, Melbourne, Australia. 5th-8th December 2005*.
- [13] K. Panta, B. Vo, and S. Singh. Improved probability hypothesis density filter (PHD) for multitarget tracking. *Proceedings ICISIP, Bangalore 12th-15th December 2005*.
- [14] O. Erdinc, P. Willet, and Y. Bar-Shalom. Probability Hypothesis Density Filter for Multitarget Multisensor Tracking. *Proc. FUSION 2005*.
- [15] T. Kurien. Issues in the design of practical multi-target tracking algorithms. *Multi-target Multi-sensor Tracking: Advanced Applications*, pages 43–83, 1990.
- [16] J. Hoffman and R. Mahler. Multitarget miss distance via optimal assignment. *IEEE Trans. Sys., Man, and Cybernetics-Part A*, 34(3):327–336, 2004.
- [17] R. Mahler. Multi-target Bayes filtering via first-order multi-target moments. *IEEE Trans. AES*, 39(4):1152–1178, 2003.