

The Cardinalized Probability Hypothesis Density Filter for Linear Gaussian Multi-Target Models

Ba-Tuong Vo*, Ba-Ngu Vo†, Antonio Cantoni*

*Western Australian Telecommunications Research Institute

The University of Western Australia, Crawley, WA 6009, Australia

†Department of Electrical and Electronic Engineering

The University of Melbourne, Parkville, VIC 3052, Australia

Email: vob@watri.org.au, bv@ee.unimelb.edu.au, cantoni@watri.org.au

Abstract—The probability hypothesis density (PHD) recursion propagates the posterior intensity of the random finite set of targets in time. The cardinalized PHD (CPHD) recursion is a generalization of the PHD recursion, which jointly propagates the posterior intensity and the posterior cardinality distribution. The incorporation of cardinality information naturally improves the accuracy and stability of state estimates. In general, the CPHD recursions are computationally intractable. This paper proposes a closed-form solution to the CPHD recursions under linear Gaussian assumptions on the target dynamics and birth process. Based on this solution, an effective multi-target tracking algorithm is developed. Extensions to non-linear models are also given using linearization and unscented transform techniques. The proposed CPHD implementations not only sidestep the need to perform data association found in traditional methods, but also dramatically improve the accuracy of individual state estimates as well as the variance of the estimated number of targets when compared to the standard PHD filter.

I. INTRODUCTION

The objective of multi-target filtering is to jointly estimate the time-varying number of targets and their states from observation sets in the presence of data association uncertainty, detection uncertainty, and noise. The problem of associating measurements to targets is the biggest challenge in multi-target filtering, and is the subject of numerous works [1], [2]. Mahler’s Finite set statistics (FISST) is a recent framework for multi-target filtering that avoids explicit associations between measurements and targets [3], [4]. Using random finite sets (RFSs) to model the collections of targets and observations led to an elegant multi-target generalization of the (single-target) Bayes filter. Moreover, FISST-based multi-target filters such as *multi-target Bayes filter*, the *Probability Hypothesis Density (PHD) filter* [3], [4] have attracted considerable interest.

The PHD filter is an approximation developed to alleviate the intractability in the multi-target Bayes filter [3]. This approximation propagates the PHD or intensity of the RFS of states in time, and has the distinct advantage that it operates only on the single-target state space and completely avoids any data association computations. Contrary to the belief that the PHD filter is intractable [5] and that no closed form recursions exist [6], a full sequential Monte Carlo (SMC) implementation was proposed in [7], and a closed form solution to the PHD recursion was derived for linear Gaussian multi-target models in [8], [9]. Moreover, this solution can be extended to

accommodate non-linear models, thereby providing a computationally efficient class of multi-target filtering algorithms.

The PHD filter assumes a Poisson distribution in the number of targets. Hence, it produces unreliable estimates of the number of targets when the number of targets is high. The PHD filter is herein referred to as the Poisson PHD filter to be precise. In [10], [11], Mahler relaxed the Poisson assumption on the number of targets and derived the cardinalized PHD (CPHD) filter, which jointly propagates the intensity function and the entire probability distribution of the number of targets. The CPHD filter is more complex than the PHD filter as the propagation equations for the posterior cardinality and intensity are coupled. Jointly propagating the cardinality distribution and the intensity promises better performance than the PHD filter. However, at present no closed form solutions are available for the CPHD recursions [10], [11].

This paper proposes an analytic solution to the CPHD recursions for the class of linear Gaussian multi-target models. In particular, it is shown that if the initial intensity is a Gaussian mixture, then so are all subsequent posterior intensities. Furthermore, closed form recursions for the posterior cardinality distribution and the posterior intensity in terms of the weights, means and covariances of its constituent components are derived. The resulting multi-target filter, herein referred to as the Gaussian mixture CPHD filter, has a computational complexity that is linear in the number of targets and cubic in the number of measurements. Extensions to non-linear models are also proposed using linearization and unscented transform techniques. The Gaussian mixture representation of the intensity also allows state estimates to be extracted much more efficiently and reliably than in particle-based techniques.

Our proposed multi-target filter is a generalization of the Gaussian mixture PHD filter described in [9]. Although both filters propagate Gaussian mixture intensities analytically in time, there are two key differences. Firstly, the intensity propagation equation in the CPHD filter is much more complex than that of the PHD filter. Secondly, the CPHD filter additionally propagates the posterior cardinality distribution which is coupled to the propagation of the posterior intensity. Indeed, the Gaussian mixture CPHD recursions reduce to the Gaussian mixture PHD recursions if the cardinality distributions of the posterior and predicted RFSs are Poisson.

II. THE CARDINALIZED PHD FILTER

The cardinalized PHD (CPHD) filter was developed in [10], [11] to address the practical limitations of the Poisson PHD filter. For complete details on the derivation, the reader is referred to [10], [11]. In essence, the strategy behind the CPHD filter is to propagate not only the first moment of the RFS of targets, but also the entire probability distribution of the number of the targets. The first order moment of a random finite set X on \mathcal{X} , also known as the PHD or intensity function, is a non-negative function v on \mathcal{X} with the property that for any closed subset $S \subseteq \mathcal{X}$

$$\mathbb{E}[|X \cap S|] = \int_S v(x) dx$$

where $|X|$ denotes the number of elements of X . In other words, for a given point x , the intensity $v(x)$ is the density of expected number of targets per unit volume at x . An RFS whose elements are i.i.d according to $v/\int v(x)dx$, but has arbitrary cardinality distribution is called an *i.i.d cluster process* [12] or a *generalized Poisson RFS* [10], [11].

The CPHD filter rests on the following assumptions:

- Each target evolves and generates measurements independently of one another
- The birth RFS and the surviving RFSs are independent of each other
- The clutter RFS is generalized Poisson and independent of the measurement RFSs.

The above assumptions are the same as those in the Poisson PHD filter, except that in this case there is no spawning and clutter is generalized Poisson.

Before proceeding to the CPHD recursions, some notation is in order. We denote by C_j^ℓ the *binomial coefficient* $\frac{\ell!}{j!(\ell-j)!}$, P_j^n the *permutation coefficient* $\frac{n!}{(n-j)!}$, $\langle \cdot, \cdot \rangle$ the *inner product* defined between two real valued functions¹ α and β by

$$\langle \alpha, \beta \rangle = \int \alpha(x)\beta(x) dx$$

(or $\sum_{\ell=0}^{\infty} \alpha(\ell)\beta(\ell)$ when α and β are real sequences), and $e_j(\cdot)$ the *elementary symmetric function* of order j defined for a finite set Z of real numbers by

$$e_j(Z) = \sum_{S \subseteq Z, |S|=j} \prod_{\zeta \in S} \zeta$$

with $e_0(Z) = 1$ by convention.

The CPHD filter proceeds as follows. Let $v_{k|k-1}$ and $p_{k|k-1}$ denote the intensity and cardinality distribution associated with the predicted multi-target state. Let v_k and p_k denote the intensity and cardinality distribution associated with the posterior multi-target state. Consider the multi-target Bayes filter and assume that the predicted and posterior multi-target RFSs at each time step are generalized Poisson. Then, the multi-target posterior intensity and posterior cardinality distribution can be propagated in time with the *CPHD prediction* and *CPHD update* [10], [11].

¹when α is constant $\langle \alpha, \beta \rangle = \alpha \int \beta(x) dx$

The CPHD prediction is given by (1)-(3)

$$p_{k|k-1}(n) = \sum_{j=0}^n p_{\Gamma,k}(n-j) \Pi_{k|k-1}[v_{k-1}, p_{k-1}](j), \quad (1)$$

$$v_{k|k-1}(x) = \int p_{S,k}(\zeta) f_{k|k-1}(x|\zeta) v_{k-1}(\zeta) d\zeta + \gamma_k(x), \quad (2)$$

where

$$\Pi_{k|k-1}[v, p](j) = \sum_{\ell=j}^{\infty} C_j^\ell \frac{\langle p_{S,k}, v \rangle^j \langle 1 - p_{S,k}, v \rangle^{\ell-j}}{\langle 1, v \rangle^\ell} p(\ell), \quad (3)$$

$$\begin{aligned} f_{k|k-1}(\cdot|\zeta) &= \text{single target transition density} \\ &\quad \text{at time } k \text{ given previous state } \zeta, \\ p_{S,k}(\zeta) &= \text{probability of target existence} \\ &\quad \text{at time } k \text{ given previous state } \zeta, \\ \gamma_k(\cdot) &= \text{intensity of spontaneous births,} \\ p_{\Gamma,k}(\cdot) &= \text{cardinality distribution of births.} \end{aligned}$$

The CPHD update is given by (4)-(6)

$$p_k(n) = \frac{\Upsilon_k^0[v_{k|k-1}; Z_k](n) p_{k|k-1}(n)}{\langle \Upsilon_k^0[v_{k|k-1}; Z_k], p_{k|k-1} \rangle}, \quad (4)$$

$$\begin{aligned} v_k(x) &= (1 - p_{D,k}) \frac{\langle \Upsilon_k^1[v_{k|k-1}; Z_k], p_{k|k-1} \rangle}{\langle \Upsilon_k^0[v_{k|k-1}; Z_k], p_{k|k-1} \rangle} v_{k|k-1}(x) \\ &\quad + \sum_{z \in Z_k} \psi_{k,z}(x) \frac{\langle \Upsilon_k^1[v_{k|k-1}; Z_k \setminus \{z\}], p_{k|k-1} \rangle}{\langle \Upsilon_k^0[v_{k|k-1}; Z_k], p_{k|k-1} \rangle} v_{k|k-1}(x), \quad (5) \end{aligned}$$

where

$$\begin{aligned} \Upsilon_k^u[v; Z](n) &= \sum_{j=0}^{\min(|Z|, n)} (|Z| - j)! p_{K,k}(|Z| - j) P_{j+u}^n \\ &\quad \times \frac{\langle 1 - p_{D,k}, v \rangle^{n-(j+u)}}{\langle 1, v \rangle^n} e_j(\Xi_k(v, Z)), \quad (6) \end{aligned}$$

$$\begin{aligned} \psi_{k,z}(x) &= \frac{\langle 1, \kappa_k \rangle}{\kappa_k(z)} g_k(z|x) p_{D,k}(x), \\ \Xi_k(v, Z) &= \{ \langle v, \psi_{k,z} \rangle : z \in Z \}, \\ g_k(\cdot|x) &= \text{single target measurement likelihood} \\ &\quad \text{at time } k \text{ given current state } x, \\ p_{D,k}(x) &= \text{probability of target detection} \\ &\quad \text{at time } k \text{ given current state } x, \\ \kappa_k(\cdot) &= \text{intensity of clutter,} \\ p_{K,k}(\cdot) &= \text{cardinality distribution of clutter.} \end{aligned}$$

As previously mentioned, the main weakness of the Poisson PHD filter is that information on the number of targets is given by a single parameter. The CPHD filter addresses this weakness by propagating the entire posterior cardinality distribution in addition to the posterior intensity. The CPHD filter is thus first order in the multi-target state but higher order in the target number. Notice that although the CPHD filter is more complex than the Poisson PHD filter, the CPHD filter still avoids data association computations and still operates exclusively on the single target state space \mathcal{X} . The CPHD recursions are also intractable since they similarly encounter the ‘curse of dimensionality’. In the next section, we propose a closed form Gaussian mixture solution.

III. CLOSED FORM CPHD RECURSION FOR LINEAR GAUSSIAN MODELS

In this section, closed form expressions for the CPHD recursions (1)-(3) and (4)-(6) are derived for the special class of *linear Gaussian multi-target models*. The model assumptions are first given in Section III-A and closed form expressions for the CPHD recursions are derived in III-B. An efficient method for performing state extraction is described in Section III-C. Implementation issues are considered in Section IV.

A. Assumptions

The class of linear Gaussian multi-target models consists of standard linear Gaussian assumptions for the transition and observation models of individual targets, as well as certain assumptions on the birth, death and detection of targets:

- Each target follows a linear Gaussian dynamical model i.e.

$$f_{k|k-1}(x|\zeta) = \mathcal{N}(x; F_{k-1}\zeta, Q_{k-1}), \quad (7)$$

$$g_k(z|x) = \mathcal{N}(z; H_k x, R_k), \quad (8)$$

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, i.e.

$$p_{S,k}(x) = p_{S,k} \quad (9)$$

$$p_{D,k}(x) = p_{D,k} \quad (10)$$

- The intensity of the birth RFS is a Gaussian mixture of the form

$$\gamma_k(x) = \sum_{i=1}^{J_{\gamma,k}} w_{\gamma,k}^{(i)} \mathcal{N}(x; m_{\gamma,k}^{(i)}, P_{\gamma,k}^{(i)}), \quad (11)$$

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

B. CPHD Closed Form Recursions

For the linear Gaussian multi-target model, the following two propositions present a closed-form solution to the CPHD recursion (1)-(3) and (4)-(6). These propositions show that if the prior intensity at any given time is a Gaussian mixture, then so are the corresponding predicted and updated intensities. These propositions also show how the posterior intensity (in the form of its Gaussian components) and the cardinality distribution are analytically propagated in time.

Proposition 1 *Suppose at time $k-1$ that the posterior intensity v_{k-1} and posterior cardinality p_{k-1} are given, and that the posterior intensity is a Gaussian mixture of the form*

$$v_{k-1}(x) = \sum_{i=1}^{J_{k-1}} w_{k-1}^{(i)} \mathcal{N}(x; m_{k-1}^{(i)}, P_{k-1}^{(i)}). \quad (12)$$

Then, the predicted intensity at time k is also a Gaussian mixture, and the CPHD prediction simplifies to

$$p_{k|k-1}(n) = \sum_{j=0}^n p_{\Gamma,k}(n-j) \sum_{\ell=j}^{\infty} C_j^\ell p_{k-1}(\ell) p_{S,k}^j (1-p_{S,k})^{\ell-j} \quad (13)$$

$$v_{k|k-1}(x) = v_{S,k|k-1}(x) + \gamma_k(x), \quad (14)$$

where $\gamma_k(x)$ is given in (11),

$$v_{S,k|k-1}(x) = p_{S,k} \sum_{j=1}^{J_{k-1}} w_{k-1}^{(j)} \mathcal{N}(x; m_{S,k|k-1}^{(j)}, P_{S,k|k-1}^{(j)}), \quad (15)$$

$$m_{S,k|k-1}^{(j)} = F_{k-1} m_{k-1}^{(j)}, \quad (16)$$

$$P_{S,k|k-1}^{(j)} = Q_{k-1} + F_{k-1} P_{k-1}^{(j)} F_{k-1}^T. \quad (17)$$

Proposition 2 *Suppose at time $k-1$ that the predicted intensity $v_{k|k-1}$ and predicted cardinality $p_{k|k-1}$ are given, and that the predicted intensity is a Gaussian mixture of the form*

$$v_{k|k-1}(x) = \sum_{i=1}^{J_{k|k-1}} w_{k|k-1}^{(i)} \mathcal{N}(x; m_{k|k-1}^{(i)}, P_{k|k-1}^{(i)}). \quad (18)$$

Then, the posterior intensity at time k is also a Gaussian mixture, and the CPHD update simplifies to

$$p_k(n) = \frac{\Psi_k^0[w_{k|k-1}, Z_k](n) p_{k|k-1}(n)}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} \quad (19)$$

$$v_k(x) = (1-p_{D,k}) \frac{\langle \Psi_k^1[w_{k|k-1}, Z_k], p_{k|k-1} \rangle}{\langle \Psi_k^0[w_{k|k-1}, Z_k], p_{k|k-1} \rangle} v_{k|k-1}(x) + \sum_{z \in Z_k} \sum_{j=1}^{J_{k|k-1}} w_k^{(j)}(z) \mathcal{N}(x; m_k^{(j)}(z), P_k^{(j)}) \quad (20)$$

where

$$\Psi_k^u[w, Z](n) = \sum_{j=0}^{\min(|Z|, n)} (|Z| - j)! p_{K,k}(|Z| - j) P_{j+u}^n \times \frac{(1-p_{D,k})^{n-(j+u)}}{\langle 1, w \rangle^{j+u}} e_j(\Xi_k(w, Z)) \quad (21)$$

$$\Xi_k(w, Z) = \left\{ \frac{\langle 1, \kappa_k \rangle}{\kappa_k(z)} p_{D,k} w^T q_k(z) : z \in Z \right\}, \quad (22)$$

$$w_{k|k-1} = [w_{k|k-1}^{(1)}, \dots, w_{k|k-1}^{(J_{k|k-1})}]^T, \quad (23)$$

$$q_k(z) = [q_k^{(1)}(z), \dots, q_k^{(J_{k|k-1})}(z)]^T, \quad (24)$$

$$q_k^{(j)}(z) = \mathcal{N}(z; H_k m_{k|k-1}^{(j)}, R_k + H_k P_{k|k-1}^{(j)} H_k^T), \quad (25)$$

$$w_k^{(j)}(z) = p_{D,k} w_{k|k-1}^{(j)} q_k^{(j)}(z) \frac{\langle \Psi_k^1[W_{k|k-1}, Z_k \setminus \{z\}], p_{k|k-1} \rangle}{\langle \Psi_k^0[W_{k|k-1}, Z_k], p_{k|k-1} \rangle}, \quad (26)$$

$$m_k^{(j)}(z) = m_{k|k-1}^{(j)} + K_k^{(j)}(z - H_k m_{k|k-1}^{(j)}), \quad (27)$$

$$P_k^{(j)} = [I - K_k^{(j)} H_k] P_{k|k-1}^{(j)}, \quad (28)$$

$$K_k^{(j)} = P_{k|k-1}^{(j)} H_k^T (H_k P_{k|k-1}^{(j)} H_k^T + R_k)^{-1}. \quad (29)$$

Propositions 1 and 2 can be established by applying Lemmas 1 and 2 in [9] to the CPHD recursions. Following [9], Proposition 1 is established by applying Lemma 1 to evaluate the integral $\int f_{k|k-1}(x|\zeta) v_{k-1}(\zeta) d\zeta$. Similarly, Proposition 2 is established by applying Lemma 2 to $g_k(x|z) v_{k|k-1}(x)$ to convert this product to a Gaussian mixture.

It follows by induction from Propositions 1 and 2 that if the initial intensity v_0 is a Gaussian mixture (including the case where $v_0 = 0$), then all subsequent predicted intensities $v_{k|k-1}$ and posterior intensities v_k , are also Gaussian mixtures. Proposition 1 provides closed-form expressions for computing the means, covariances and weights of $v_{k|k-1}$ from those of v_{k-1} , and also for computing the distribution $p_{k|k-1}$ from p_{k-1} . Proposition 2 then provides closed-form expressions for computing the means, covariances and weights of v_k from those of $v_{k|k-1}$, and also for computing the distribution p_k from $p_{k|k-1}$, when a new set of measurements arrives. Propositions 1 and 2 are, respectively, the prediction and update steps of the CPHD recursion for linear Gaussian multi-target models, herein referred to as the *Gaussian mixture CPHD recursion*.

C. Multi-target State Extraction

Given the posterior intensity v_k and posterior cardinality p_k at each time step k , estimation of individual target states is generally non-trivial since the form of the intensity function is not known. When particle approximations are used, clustering techniques are additionally required to partition the particle population into an estimated number of clusters, each of which corresponds to state estimates. However, this approach is not only computationally expensive but also extremely sensitive to the estimate of the number of targets, as demonstrated in [7] using the Poisson PHD filter. In contrast, performing state extraction with the Gaussian mixture CPHD filter is very straightforward since the means of the mixture components are the local maxima of the posterior intensity. Thus, all that is required is to first estimate the number of targets, and then to extract as state estimates the relevant number of components from the posterior intensity having the highest weights. The number of targets can be estimated using for example an MAP or EAP estimator on the posterior cardinality distribution.

IV. IMPLEMENTATION ISSUES

Closer examination of the Gaussian mixture CPHD recursions reveals a number of numerical issues which need to be addressed in order to design an efficient and practical filter. Each of these issues is considered in turn below.

A. Computing Cardinality Distributions

It is generally not possible to represent and propagate the cardinality distribution in full since it may have an infinite tail. However, if it is known in advance that there is a maximum possible number of targets, the distribution can be accurately represented with a finite number of terms $\{p_k(n)\}_{n=0}^{N_{\max}}$. This is a reasonable approximation when N_{\max} is significantly greater than the number of targets on the scene at any time. Moreover, it can be seen from the CPHD recursions that if the initial cardinality distribution has a finite support, then the equations for computing the mixtures weights and cardinality distributions reduce to finite sums.

B. Computing Elementary Symmetric Functions

Attempting to evaluate the elementary symmetric functions directly from the definition is clearly infeasible. Fortunately, there is a much quicker way to perform the evaluation using a basic result from algebra and combinatorics theory known as the Newton-Girard formulas or equivalently Vieta's Theorem. The result can be found in [13] and is summarized as follows. Let $\rho_1, \rho_2, \dots, \rho_M$ be distinct roots of the polynomial $\alpha_M x^M + \alpha_{M-1} x^{M-1} + \dots + \alpha_1 x + \alpha_0$. Then, the elementary symmetric functions $e_j(\cdot)$ for orders $j = 0, \dots, M$ are given by $e_j(\rho_1, \rho_2, \dots, \rho_M) = (-1)^j \alpha_{M-j} / \alpha_M$. The elementary symmetric functions can thus be evaluated by expanding out the polynomial with roots given by the arguments to the symmetric function. This can be implemented with a variety of methods, for example an appropriate recursion or convolution.

C. Managing Mixture Components

Similar to Gaussian mixture PHD filter [9], the number of Gaussian components required to represent the posterior increases without bound. Fortunately, to mitigate this problem the 'pruning' procedure described in [9] is also directly applicable for the CPHD filter. The basic idea is to discard components with negligible weights and merge components that are close together.

V. EXTENSION TO NON-LINEAR MODELS

Analogous to the approach in [9] for extending the Gaussian mixture PHD filter to non-linear models, two non-linear extensions of the Gaussian mixture CPHD filter are proposed. In essence, the original assumptions of the CPHD filter are retained, but the assumptions on the form of the single target dynamical model given by the transition equation $f_{k|k-1}(\cdot|\cdot)$ and measurement equation $g_k(\cdot|\cdot)$ are relaxed to the non-linear functions in the state and noise variables. Then, based on taking local linearizations of $f_{k|k-1}$ and g_k , the extended Kalman CPHD (EK-CPHD) filter is proposed. Also, based on applying the unscented transform to propagate $f_{k|k-1}$ and g_k , the unscented Kalman CPHD (UK-CPHD) filter is proposed. Note that in comparison to a particle implementation, the EK-CPHD and UK-CPHD filters are much less computationally expensive, and that state estimates can still be extracted very easily as a result of the underlying Gaussian implementation.

VI. SIMULATIONS

In the following examples, a time varying number of targets is observed in clutter over a two dimensional region. The first example illustrates in detail a linear Gaussian multi-target scenario whereas the second illustrates a non-linear scenario with the EK-CPHD and UK-CPHD approximations discussed in Section V. In both examples, the filter calculates the cardinality distribution to $N_{\max} = 100$ terms and pruning is performed at each time step using a weight threshold of $T = 10^{-5}$, a merging threshold of $U = 4m$, and a maximum of $J_{\max} = 100$ Gaussian components.

A. Example 1

The surveillance region is the square $[-1000, 1000] \times [-1000, 1000]$ (units are in m). For clarity, a 5 target scenario is considered here. Targets appear from various birth locations at various times and also occasionally disappear. The target state variable is a vector of planar position and velocity $x_k = [p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k}]^T$. The single-target transition model is linear Gaussian specified by

$$F_k = \begin{bmatrix} I_2 & \Delta I_2 \\ 0_2 & I_2 \end{bmatrix}, \quad Q_k = \sigma_v^2 \begin{bmatrix} \frac{\Delta^4}{4} I_2 & \frac{\Delta^3}{2} I_2 \\ \frac{\Delta^3}{2} I_2 & \Delta^2 I_2 \end{bmatrix},$$

where I_n and 0_n denote the $n \times n$ identity and zero matrices, $\Delta = 1s$ is the sampling period, and $\sigma_v = 5(m/s^2)$ is the standard deviation of the process noise. The probability of target survival $p_{S,k} = 0.99$. The birth process is Poisson with intensity $\gamma_k(x) = 0.03\mathcal{N}(x; m_\gamma^{(1)}, P_\gamma) + 0.02\mathcal{N}(x; m_\gamma^{(2)}, P_\gamma)$ where $m_\gamma^{(1)} = [250, -10, -250, 0]^T$, $m_\gamma^{(2)} = [-500, 10, -500, 0]^T$, $P_\gamma = \text{diag}([100, 100, 25, 25]^T)$. The probability of target detection $p_{D,k} = 0.98$. The single-target measurement model is also linear Gaussian with

$$H_k = [I_2 \quad 0_2], R_k = \sigma_\varepsilon^2 I_2,$$

where $\sigma_\varepsilon = 10m$, is the standard deviation of the measurement noise. Clutter is assumed uniform Poisson with intensity $\kappa_k(z) = \lambda_c V u(z)$, where $u(\cdot)$ is a uniform probability density over the surveillance region, $V = 4 \times 10^6 m^2$ is the ‘volume’ of the surveillance region, and $\lambda_c = 12.5 \times 10^{-6} m^{-2}$ is the clutter intensity (giving an average of 50 clutter returns per scan).

Fig. 1 plots the x and y components of the true trajectories with cluttered measurements against time. The Gaussian mixture CPHD filter’s position estimates are shown in Fig. 2 for the x and y components against time. It can be seen that the filter correctly identifies all target births from the two locations and tracks them accordingly, and correctly identifies two target deaths at times $k = 60$ and $k = 70$. Also note that the filter has no difficulty resolving the two targets which cross paths at time $k = 48$.

To evaluate the performance for the current scenario, 1000 Monte Carlo runs are performed on the same track data but with randomly generated target and clutter measurements. In Fig. 3(a) the Wasserstein miss distance [14] between the estimated and true multi-target states is shown, and in Fig. 3(b) the mean and standard deviation of the estimated cardinality distribution is shown along with the true number of targets. It can be seen that the filter estimates the time varying number of targets accurately and that the variance of the estimate is reasonable (and much improved from the Poisson PHD filter). Note that the peaks in Fig. 3(a) coincide with time instants where there is change in the true number of targets Fig. 3(b), and also that the peaks settle to smaller values very quickly. This can be expected since the filter makes an error when the number of target suddenly changes, and subsequently requires several time steps to validate new or deleted tracks using measurement data.

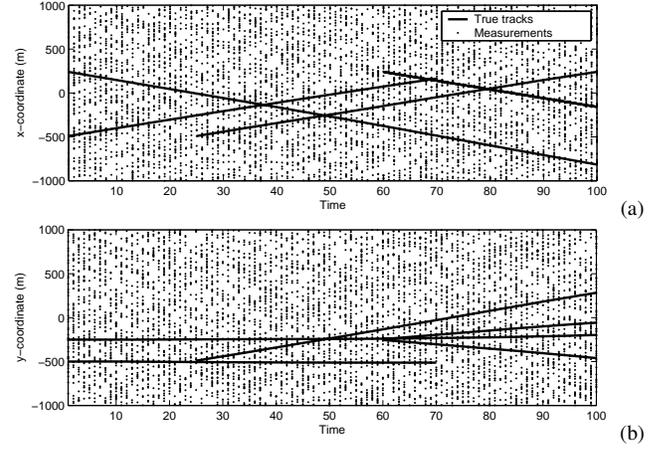


Fig. 1. Measurements and true target positions.

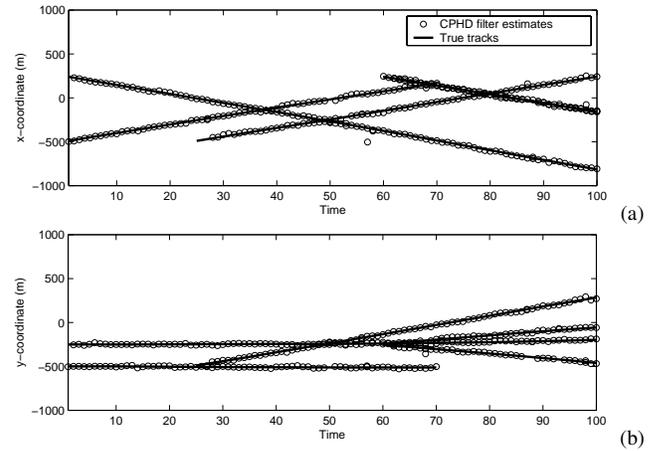


Fig. 2. Position estimates of the Gaussian mixture CPHD filter.

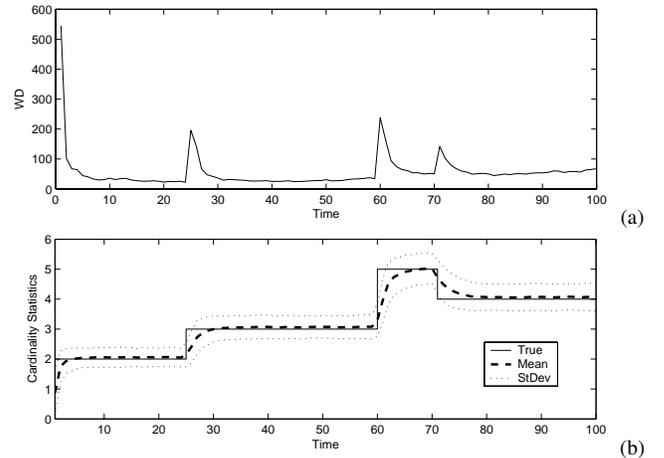


Fig. 3. (a) Wasserstein distance (b) Cardinality statistics

B. Example 2

In this non-linear example, a nearly constant turn model having varying turn rate together with bearing and range measurements is considered. The observation region is the

half disc of radius 2000m. Again, for clarity only 5 targets are considered in which there are various births and deaths throughout the simulation. The target state variable $x_k = [\tilde{x}_k^T, \omega_k]^T$ comprises the planar position and velocity $\tilde{x}_k^T = [p_{x,k}, \dot{p}_{x,k}, p_{y,k}, \dot{p}_{y,k}]^T$ as well as the turn rate ω_k . The state transition model is

$$\tilde{x}_k = F(\omega_{k-1})\tilde{x}_{k-1} + Gw_{k-1} \quad (30)$$

$$\omega_k = \omega_{k-1} + \Delta u_{k-1} \quad (31)$$

where

$$F(\omega) = \begin{bmatrix} 1 & \frac{\sin \omega \Delta}{\omega} & 0 & -\frac{1 - \cos \omega \Delta}{\omega} \\ 0 & \cos \omega \Delta & 0 & -\sin \omega \Delta \\ 0 & \frac{1 - \cos \omega \Delta}{\omega} & 1 & \frac{\sin \omega \Delta}{\omega} \\ 0 & \sin \omega \Delta & 0 & \cos \omega \Delta \end{bmatrix}, \quad G = \begin{bmatrix} \frac{\Delta^2}{2} & 0 \\ T & 0 \\ 0 & \frac{\Delta^2}{2} \\ 0 & \Delta \end{bmatrix},$$

$w_{k-1} \sim \mathcal{N}(\cdot; 0, \sigma_w^2 I)$, and $u_{k-1} \sim \mathcal{N}(\cdot; 0, \sigma_u^2 I)$ with $\Delta = 1s$, $\sigma_w = 15m/s^2$, and $\sigma_u = \pi/180rad/s$. The sensor observation is a noisy bearing and range vector given by

$$z_k = \begin{bmatrix} \arctan(p_{x,k}/p_{y,k}) \\ \sqrt{p_{x,k}^2 + p_{y,k}^2} \end{bmatrix} + \varepsilon_k \quad (32)$$

where $\varepsilon_k \sim \mathcal{N}(\cdot; 0, R_k)$, with $R_k = \text{diag}([\sigma_\theta^2, \sigma_r^2]^T)$, $\sigma_\theta = 2(\pi/180)rad$, and $\sigma_r = 10m$. The birth process is Poisson with intensity $\gamma_k(x) = 0.1\mathcal{N}(x; m_\gamma^{(1)}, P_\gamma) + 0.1\mathcal{N}(x; m_\gamma^{(2)}, P_\gamma)$ where $m_\gamma^{(1)} = [-1000, 0, -500, 0, 0]^T$, $m_\gamma^{(2)} = [1050, 0, 1070, 0, 0]^T$, and $P_\gamma = \text{diag}([50, 50, 50, 50, 6(\pi/180)]^T)$. The probability of target survival and detection are $p_{S,k} = 0.99$ and $p_{D,k} = 0.98$ respectively. Clutter is Poisson with intensity $\lambda_c = 3.2 \times 10^{-3} (radm)^{-1}$ over the region $[-\pi/2, \pi/2]rad \times [0, 2000]m$ (giving an average of 20 clutter points per scan).

Both the EK-CPHD and UK-CPHD filters are run on the same measurement data. The true trajectories and filter outputs for the UK-CPHD are shown in Fig. 4. The results for EK-CPHD are very similar and are not shown here. It can be seen that the EK and UK approximation to the Gaussian mixture CPHD are able to identify all target births and track the non-linear motion well. Notice also that the filters have no trouble resolving two targets which cross paths at time $k = 80$.

VII. CONCLUSION

This paper has proposed a Gaussian mixture implementation of the cardinalized Probability Hypothesis Density (CPHD) filter as a solution to the multi-target detection and estimation problem. We have shown that for the class of the linear Gaussian multi-target models, the CPHD recursions admit closed form solutions. In particular, closed form expressions for propagating the Gaussian mixture intensity, as well as for the cardinality distribution have been derived. Furthermore, efficient techniques for propagating the intensity and cardinality distribution have been given. Extension to non-linear models have been provided via linearization and unscented transform techniques. Simulations have verified that the proposed Gaussian mixture CPHD filter performs well and shows a dramatic reduction in the variance of the estimated target number when compared to the Poisson PHD filter.

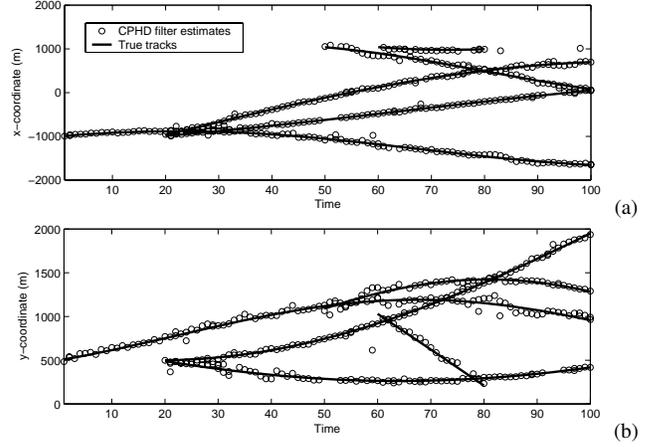


Fig. 4. Position estimates of the UK-CPHD filter.

ACKNOWLEDGMENTS

The authors would like thank Dr. Ronald Mahler for his valuable discussions and feedback. This work is also supported by the Australian Telecommunications Cooperative Research Centre (ATCRC) and Australian Research Council (ARC) Discovery Grant DP0345215.

REFERENCES

- [1] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Academic Press, San Diego, 1988.
- [2] S. Blackman, *Multiple Target Tracking with Radar Applications*. Artech House, Norwood, 1986.
- [3] R. Mahler, "Multi-target Bayes filtering via first-order multi-target moments," *IEEE Trans. AES*, vol. 39, no. 4, pp. 1152–1178, 2003.
- [4] I. Goodman, R. Mahler, and H. Nguyen, *Mathematics of Data Fusion*. Kluwer Academic Publishers, 1997.
- [5] G. Pulford, "Taxonomy of multiple target tracking methods," *IEE Proc. Radar Sonar Navig.*, vol. 152, no. 5, pp. 291–304, Oct. 2005.
- [6] L. Lin, Y. Bar-Shalom, and T. Kirubarajan, "Data association combined with the Probability Hypothesis Density filter for multitarget tracking," in *O. E. Drummond (ed.) Signal and Data Processing of Small Targets, Proc. SPIE*, vol. 5428, pp. 464–475, 2004.
- [7] B. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multi-target filtering with random finite sets," *IEEE Trans. Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1224–1245, 2005, also: <http://www.ee.unimelb.edu.au/staff/bv/publications.html>.
- [8] B. Vo and W.-K. Ma, "A closed-form solution to the Probability Hypothesis Density filter," in *Proc. Int'l Conf. on Information Fusion*, Philadelphia, 2005, also: <http://www.ee.unimelb.edu.au/staff/bv/publications.html>.
- [9] B. Vo and W.-K. Ma, "The Gaussian mixture Probability Hypothesis Density filter," *IEEE Trans. SP*, to appear 2006, also: <http://www.ee.unimelb.edu.au/staff/bv/publications.html>.
- [10] R. Mahler, "A theory of PHD filters of higher order in target number," in I. Kadar (ed.), *Signal Processing, Sensor Fusion, and Target Recognition XV, SPIE Defense & Security Symposium*, Orlando, April 17–22, 2006 (to appear).
- [11] R. Mahler, "PHD filters of higher order in target number," Submitted to *IEEE Trans. AES*, December 2005.
- [12] D. Daley and D. Vere-Jones, *An introduction to the theory of point processes*. Springer-Verlag, 1988.
- [13] P. Borwein and T. Erdélyi, "Newton's Identities" Section 1.1.E.2 in *Polynomials and Polynomial Inequalities*, Springer-Verlag, New York, 1995.
- [14] J. Hoffman and R. Mahler, "Multitarget miss distance via optimal assignment," *IEEE Trans. Sys., Man, and Cybernetics-Part A*, Vol. 34, No. 3, pp. 327–336, 2004.