18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

PHD Filter with Approximate Multiobject Density Measurement Update

Karl Granström, Peter Willett, Yaakov Bar-Shalom
Department of Electrical and Computer Engineering, University of Connecticut, Storrs, Connecticut 06269
Email: karl@engr.uconn.edu, willett@engr.uconn.edu, ybs@engr.uconn.edu

***Abstract*—The PHD filter is a popular approach to the multiple target tracking problem, however, it suffers from the Poisson assumption which yields a cardinality estimate with too high variance. In recent work Le and Kaplan proposed to improve the performance of the PHD filter using a particle approximation of the predicted multiobject density and updating it using the multiobject measurement pdf. Following the work by Le and Kaplan, in this paper we use the predicted PHD to construct a particle approximation of the predicted multiobject density. Using joint probabilistic data association, the multiobject particle approximation can then be updated using the multiobject measurement likelihood, resulting in a particle approximation of the posterior multiobject density. The posterior multiobject particles are then used to approximate the posterior PHD, which is subsequently predicted using the standard PHD prediction. The proposed filter is implemented using a Gaussian mixture approximation of the PHD intensity, and a simulation study shows a significant performance improvement compared to using the standard PHD measurement update, especially in terms of the cardinality estimate.**

I. Introduction

Multiple target tracking is defined as the processing of sets of measurements obtained from multiple sources in order to maintain estimates of the targets' current states. The task is complicated by the fact that the number of targets is unknown and time-varying, there is detection uncertainty, and the measurements are affected by noise and are submerged in spurious clutter. As a result, at each time step only a subset of the measurements are generated by targets.

Broadly speaking there are three different approaches to multiple target tracking: Multiple Hypothesis Tracking (MHT) [4], Joint Probabilistic Data Association (JPDA) [1], and Random Finite Sets (RFS) [7], [8]. The MHT type approaches involve propagating target trajectory hypotheses in time and calculating their likelihoods, the JPDA type approaches use data association probabilities for the latest set of measurements, and the RFS type approaches rely on modeling the targets and the measurements as random sets.

The multiobject Bayes filter is an RFS type filter that propagates and updates the density of the multiobject state in time. Because of its high computational complexity it is generally considered infeasible to implement and use. Computationally feasible approximations include the Probability Hypothesis Density (PHD) filters [9], the Cardinalized PHD (CPHD) filters [10], and the multi-Bernoulli filters [18].

The PHD filters recursively estimate the first order moment of the multiobject state, called the PHD intensity, under an assumed Poisson distribution for the cardinality. The CPHD filters recursively estimate the PHD and also a truncated cardinality distribution. A known drawback is the PHD filters' high variance for the cardinality estimate, resulting from the Poisson assumption – the CPHD filters are known to have better cardinality estimates. Both the PHD and CPHD filters are susceptible to the "spooky effect" [5], [8] (PHD mass shifted from undetected targets to statistically noninteracting and possibly remote detected targets). The multi-Bernoulli filters approximate the multiobject density with a multi-Bernoulli distribution, which is then propagated and updated in time. Multi-Bernoulli filters estimate a for each target a probability of existence, in addition to the target densities, and are known to be capable of matching the CPHD filter's cardinality estimate without suffering from the "spooky effect".

In their most basic form, none of the PHD, CPHD, or multi-Bernoulli filters, formally estimate target trajectories – only point estimates are supplied at each time step – however, target trajectories can be obtained, e.g. using labeling schemes for PHD and CPHD filters [11] or labeled RFSs, leading to Labeled Multi-Bernoulli (LMB) filters [12], [17]. In contrast, both the MHT and JPDA type algorithms estimate trajectories.

Le and Kaplan [6] proposed to improve PHD filter performance by randomly sampling the predicted PHD intensity to construct a particle approximation of the predicted multiobject density. It is then possible to use the multiobject measurement set likelihood to update the multiobject density. The posterior multiobject density is then used to compute the posterior PHD intensity, allowing for the use of standard PHD prediction. The work [6] relied on a particle approximation of the PHD intensity, retained the Poisson model for the cardinality, and did not include a formal model of new target birth. In this work we extend their ideas by approximating the PHD intensity by a Gaussian mixture (GM), relaxing the Poisson assumption for the cardinality and instead interpreting the GM weights as probabilities of existence, and by including a formal model of new target birth.

The proposed multiple target tracking filter combines elements from RFS tracking and JPDA tracking to attain an improved update for the PHD filter, and thereby improve the performance of the multiple target tracking. Specifically, a particle approximation of the predicted multiobject density is constructed by randomly selecting Gaussian components from the the predicted PHD intensity. Using JPDA association probabilities, each multiobject particle (MOP) is updated using the multiobject measurement pdf. This results in a particle approximation of the posterior multiobject density, from which the posterior PHD intensity is computed. Lastly, the posterior PHD is predicted using the standard PHD prediction.

Results from a simulation study show that by using the proposed MOP measurement update instead of the standard PHD measurement update, significant performance improvements are obtained. Especially the estimate of cardinality has much lower variance, and the proposed filter is therefore less sensitive to missed detections. In addition, the PHD filters' known cardinality bias for detection probabilities lower than 1 is alleviated.

The paper is organized as follows. In the next section we review the multiple target tracking problem, and in Section III we present the proposed multiple target tracking filter. Simulation resutls are presented in Section IV, and a discussion of the proposed filter is given in Section V. The paper is concluded in Section VI.

II. Multiple target tracking

Let $\mathbf{x}_k^i$ denote the state of the ith target at time step k. There are N_k^x such targets, and the target set is denoted

$$\mathbf{X}_k = \left\{\mathbf{x}_k^i\right\}_{i=1}^{N_k^x} \tag{1}$$

The target set cardinality $|\mathbf{X}_k| = N_k^x$ is a time-varying discrete random variable, and each target state $\mathbf{x}_k^i$ is a random variable. The set of measurements obtained at time step k is denoted

$$\mathbf{Z}_k = \left\{\mathbf{z}_k^j\right\}_{j=1}^{N_k^z} \tag{2}$$

where $N_k^z = |\mathbf{Z}_k|$ is the cardinality of the measurement set at time k. There are two types of measurements: clutter measurements and target originated measurements, and measurement origin is assumed unknown. Note that the sets above are without order and the set indexing is arbitrary; the particular choices $i = 1, \ldots, N_k^x$ and $j = 1, \ldots, N_k^z$ are only used for notational simplicity and convenience.

The posterior multiobject distribution at time step $k-1$ is $f(\mathbf{X}_{k-1}|\mathbf{Z}^{k-1})$ where $\mathbf{Z}^{k-1}$ denotes all measurement sets $\mathbf{Z}_\ell$ from $\ell = 0$ up to, and including, $\ell = k-1$. The predicted multiobject distribution is given by the Chapman-Kolmogorov equation

$$f(\mathbf{X}_k|\mathbf{Z}^{k-1}) = \int f(\mathbf{X}_k|\mathbf{X}_{k-1}) f(\mathbf{X}_{k-1}|\mathbf{Z}^{k-1}) \delta \mathbf{X}_{k-1} \tag{3}$$

where $f(\mathbf{X}_k|\mathbf{X}_{k-1})$ is the multiobject transition density. Multiobject prediction involves modeling time evolution of surviving targets (targets that remain in the surveillance area), target death (targets that disappear from the surveillance area), and target birth (new targets that appear in the surveillance area).

The posterior multiobject distribution at time t_k is given by the Bayes update

$$f(\mathbf{X}_k|\mathbf{Z}^k) = \frac{f(\mathbf{Z}_k|\mathbf{X}_k) f(\mathbf{X}_k|\mathbf{Z}^{k-1})}{\int f(\mathbf{Z}_k|\mathbf{Y}) f(\mathbf{Y}|\mathbf{Z}^{k-1}) \delta \mathbf{Y}} \tag{4}$$

The multiobject measurement set density $f(\mathbf{Z}_k|\mathbf{X}_k)$ involves modeling target detection, measurement noise, and clutter measurements.

A typical assumption in multiple target tracking is that the targets are independent, see e.g. [2]. In this case the multiobject density can be assumed to be that of an i.i.d. cluster process,

$$f(\mathbf{X}_k|\mathbf{Z}^k) = P_{k|k}(n) \prod_{i=1}^{n} p_{k|k}(\mathbf{x}_k^i|\mathbf{Z}^k) \tag{5a}$$

Under the assumption that the targets and the clutter processes generate measurements independent of each other, the measurement set is the union of a set of clutter measurements $\mathbf{C}_k$ and sets of target generated measurements $\mathbf{W}_k\left(\mathbf{x}_k^i\right)$

$$\mathbf{Z}_k = \mathbf{C}_k \cup \left[\bigcup_{i=1}^{N_k^x} \mathbf{W}_k\left(\mathbf{x}_k^i\right)\right] \tag{6}$$

The clutter measurements are typically modeled as a Poisson process, meaning that the number of clutter measurements is Poisson distributed in number, and each clutter measurement is distributed with pdf $g_c(\mathbf{z})$, often assumed uniform. The clutter pdf is

$$\kappa(\mathbf{C}_k) = \frac{e^{-\lambda_c}}{N_k^c!} \prod_{i=1}^{N_k^c} \lambda_c g_c\left(\mathbf{z}_k^j\right) \tag{7}$$

where λ_c is the Poisson rate.

This paper is restricted to consideration of so called point targets, meaning that the ith target measurement set $\mathbf{W}_k\left(\mathbf{x}_k^i\right)$ is empty ($=\emptyset$) with probability $1 - p_D(\mathbf{x}_k^i)$, and with probability $p_D(\mathbf{x}_k^i)$ the set contains a single measurement $\mathbf{z}_k$ originating from $\mathbf{x}_k^i$, distributed according to the pdf $g_x(\mathbf{z}_k|\mathbf{x}_k^i)$.

Under the assumption of Poisson clutter and independent point target measurements the measurement set pdf is [7]

$$\begin{aligned} f(\mathbf{Z}_k|\mathbf{X}_k) = {} & e^{-\lambda_c}\left[\prod_{j=1}^{N_k^z} \lambda_c g_c\left(\mathbf{z}_k^j\right)\right]\left[\prod_{i=1}^{N_k^x}\left(1 - p_D\left(\mathbf{x}_k^i\right)\right)\right] \\ & \times \sum_{\theta\in\Theta} \prod_{i:\sigma_i>0} \frac{p_D\left(\mathbf{x}_k^i\right)}{1 - p_D\left(\mathbf{x}_k^i\right)} \frac{g_x\left(\mathbf{z}_k^{\sigma_i}|\mathbf{x}_k^i\right)}{\lambda_c g_c\left(\mathbf{z}_k^{\sigma_i}\right)} \end{aligned} \tag{8a}$$

$$\begin{aligned} = {} & \sum_{\theta\in\Theta} e^{-\lambda_c}\left[\prod_{j:\nexists\sigma_i=j} \lambda_c g_c\left(\mathbf{z}_k^j\right)\right]\left[\prod_{i:\sigma_i=0}\left(1 - p_D\left(\mathbf{x}_k^i\right)\right)\right] \\ & \times \left[\prod_{i:\sigma_i>0} p_D\left(\mathbf{x}_k^i\right) g_x\left(\mathbf{z}_k^{\sigma_i}|\mathbf{x}_k^i\right)\right] \end{aligned} \tag{8b}$$

Here $\theta = \{\sigma_i\}$, defined as in [7], is a set of associations σ_i, where $\sigma_i = 0$ if target $\mathbf{x}_k^i$ is not associated to any measurement, and $\sigma_i = j$ if target $\mathbf{x}_k^i$ is associated to measurement $\mathbf{z}_k^j$. The set of all associations θ is denoted Θ.

As mentioned above, the exact multiobject Bayes filter, with prediction (3) and update (4), is computationally intractable;

feasible approximations are, e.g., PHD, CPHD and multi-Bernoulli filters.

The scope of this paper is limited to considering PHD filters. The PHD intensity $D_{k|k}(\mathbf{x})$, defined on the single target state $\mathbf{x}$, is the first-order moment of the multiobject density $f(\mathbf{X}_k|\mathbf{Z}^k)$. Let $D_{k-1|k-1}(\mathbf{x})$ be the posterior PHD intensity at time t_k. Omitting target spawning, the predicted PHD intensity is

$$D_{k|k-1}(\mathbf{x}) = D_k^b(\mathbf{x}) + \int p_S(\mathbf{y}) p_{k,k-1}(\mathbf{x}|\mathbf{y}) D_{k-1|k-1} \mathrm{d}\mathbf{y} \tag{9}$$

where $p_S(\cdot)$ is the probability of survival, and $p_{k,k-1}(\cdot|\cdot)$ is the single target transition density. The updated PHD intensity is [7]

$$\begin{aligned} D_{k|k}(\mathbf{x}) &\approx (1 - p_D(\mathbf{x})) D_{k|k-1}(\mathbf{x}) \\ &+ \sum_{\mathbf{z} \in \mathbf{Z}} \frac{p_D(\mathbf{x}) g_x(\mathbf{z}|\mathbf{x}) D_{k|k-1}(\mathbf{x})}{\lambda_c g_c(\mathbf{z}) + \int p_D(\mathbf{y}) g_x(\mathbf{z}|\mathbf{y}) D_{k|k-1}(\mathbf{y}) \mathrm{d}\mathbf{y}} \end{aligned} \tag{10}$$

As mentioned above, a known drawback of the PHD filter is that the cardinality distribution is iteratively approximated with a Poisson distribution, which has equal mean and variance. In practical applications an oversensitive cardinality estimate is manifested, e.g., when there are missed detections – often a missed detection results in a lost target estimate. Another drawback is that the PHD filter has a positive bias on the cardinality estimate, where the bias increases as the probablity of detection decreases.

In this paper we improve upon the standard PHD filter by doing an approximate multiobject measurement update instead of the standard measurement update (10). The new measurement update is based on using the predicted PHD intensity to generate a particle representation of the predicted multiobject density $f(\mathbf{X}_k|\mathbf{Z}^{k-1})$, and then updating each particle using the JPDA filter measurement update and evaluating the multiobject set likelihood (8).

III. Proposed MOP-PHD filter

In this section we present the proposed multi-object particle probability hypothesis density (MOP-PHD) filter. Under the assumption of Gaussian single target densities, linear Gaussian motion and measurement models, and Gaussian mixture birth PHD intensity,

$$p_{k|k}(\mathbf{x}_k) = \mathcal{N}\left(\mathbf{x}_k\,;\, m_{k|k}, P_{k|k}\right) \tag{11a}$$

$$p_{k,k-1}(\mathbf{x}_k|\mathbf{x}_{k-1}) = \mathcal{N}\left(\mathbf{x}_k\,;\, F_k \mathbf{x}_{k-1}, Q_k\right) \tag{11b}$$

$$g_x(\mathbf{z}_k|\mathbf{x}_k) = \mathcal{N}\left(\mathbf{z}_k\,;\, H_k \mathbf{x}_k, R_k\right) \tag{11c}$$

$$D_k^b(\mathbf{x}) = \sum_{j=1}^{J_k^b} w_k^{b,j} \mathcal{N}\left(\mathbf{x}\,;\, m_k^{b,j}, P_k^{b,j}\right) \tag{11d}$$

the PHD intensity can be approximated by a mixture of Gaussian densities [16]

$$D_{k|k}(\mathbf{x}) = \sum_{j=1}^{J_{k|k}} w_{k|k}^j \mathcal{N}\left(\mathbf{x}\,;\, m_{k|k}^j, P_{k|k}^j\right) \tag{12}$$

The birth PHD is assumed known, however, the work can be extended to unknown birth distribution, e.g. using a uniform birth PHD [3] or an adaptive birth process [13]. We assume that the clutter is uniformly distributed in the surveillance area, $g_c(\mathbf{z}) = 1/V$ where V is the volume of the surveillance area. Further, we assume that the probability of detection and probability of survival are constant, $p_D(\mathbf{x}) = p_D$ and $p_S(\mathbf{x}) = p_S$.

In the proposed multiple target filter the standard prediction (9) is used. In the measurement update step, instead of using the standard PHD correction (10) we use a measurement update that is inspired by ideas presented in [6]. The measurement update has three main steps:

1) Sample the predicted PHD to create a particle approximation of the predicted multiobject density $f(\mathbf{X}_k|\mathbf{Z}^{k-1})$.
2) Obtain a particle approximation of the posterior multiobject density $f(\mathbf{X}_k|\mathbf{Z}^k)$ by using JPDA association probabilities to approximate the multiobject measurement update (4) for each particle.
3) Compute the posterior PHD intensity from the particle approximation of the posterior multiobject density.

The details of the proposed filter are presented below.

Given a GM approximation of the predicted PHD intensity

$$D_{k|k-1}(\mathbf{x}) = \sum_{j=1}^{J_{k|k-1}} w_{k|k-1}^j \mathcal{N}\left(\mathbf{x}\,;\, m_{k|k-1}^j, P_{k|k-1}^j\right) \tag{13}$$

we approximate the predicted multiobject density by P randomly sampled multiobject particles $\mathbf{X}_{k|k-1}^p$

$$f(\mathbf{X}_k|\mathbf{Z}^{k-1}) \approx \sum_{p=1}^{P} P^{-1} \delta(\mathbf{X}_k, \mathbf{X}_{k|k-1}^p) \tag{14}$$

where, for the pth particle, $\delta(\mathbf{X}_k, \mathbf{X}_{k|k-1}^p) = 0$ if $|\mathbf{X}_k| \neq |\mathbf{X}_{k|k-1}^p|$ and

$$\delta(\mathbf{X}_k, \mathbf{X}_{k|k-1}^p) = \prod_{i \in I_p} \mathcal{N}\left(\mathbf{x}_i\,;\, m_{k|k-1}^i, P_{k|k-1}^i\right) \tag{15}$$

if $|\mathbf{X}_k| = |\mathbf{X}_{k|k-1}^p|$. Note that, strictly speaking, (14) is a Gaussian sum approximation and not a particle approximation, because (15) defines a Gaussian distribution. However, for the sake of brevity and simplicity, in the remainder of the paper will use the terminology multiobject particle.

In (15) the set I_p is defined as $I_p = \{i | u_p^i \leq w_{k|k-1}^i\}$, where u_p^i are randomly sampled from the uniform distribution $\mathcal{U}(0,1)$. Expressed in words, the meaning of the set I_p is that in the pth particle the ith predicted Gaussian component is included with probability $w_{k|k-1}^i$. For each particle, the cardinality is the cardinality of the set I_p, meaning that within each multiobject particle the included Gaussians are interpreted as representing targets that do exist, i.e. target birth and death is here represented by the random sampling u_p^i.

Note that the GM-PHD weights $w_{k|k-1}^j$ are not probabilities of existence, instead the weights are related to the cardinality estimation. Specifically, the sum of weights $\sum_{j=1}^{J_{k|k-1}} w_{k|k-1}^j$

is the estimated number of targets that are predicted to be in the surveillance area. However, as will be demonstrated in the results section, interpreting the weights as probabilities of target existence actually gives results that are rather accurate.

Given the particle approximation of the predicted multiobject density (14), the posterior multiobject density is given by the Bayes update (4),

$$f(\mathbf{X}_k|\mathbf{Z}^k) = \frac{\sum_{p=1}^{P} f(\mathbf{Z}_k|\mathbf{X}_k)\delta(\mathbf{X}_k, \mathbf{X}^p_{k|k-1})}{\sum_{p=1}^{P}\int f(\mathbf{Z}_k|\mathbf{X}_k)\delta(\mathbf{X}_k, \mathbf{X}^p_{k|k-1})\delta\mathbf{X}_k} \tag{16}$$

Using the measurement set likelihood (8) and the Kalman filter measurement update, for each multiobject particle we have

$$\begin{aligned} f(\mathbf{Z}_k|\mathbf{X}_k)&\delta(\mathbf{X}_k, \mathbf{X}^p_{k|k-1}) \\ &= \sum_{\theta\in\Theta} e^{-\lambda_c}\left[\prod_{j:\nexists\sigma_i=j}\frac{\lambda_c}{V}\right]\left[\prod_{i\in I_p:\sigma_i=0}(1-p_D)\right] \\ &\quad\times\left[\prod_{i\in I_p:\sigma_i>0} p_D\mathcal{N}\left(\mathbf{z}_k^{\sigma_i}\,;\,H_k\mathbf{x}_k^i, R_k\right)\right] \\ &\quad\times\left[\prod_{i\in I_p}\mathcal{N}\left(\mathbf{x}_k^i\,;\,m^i_{k|k-1}, P^i_{k|k-1}\right)\right] &\text{(17a)} \\ &= \sum_{\theta\in\Theta}\mathcal{L}^{p,\theta}_{k|k-1}\prod_{i\in I_p}\mathcal{N}\left(\mathbf{x}_i\,;\,m^{i,\sigma_i}_{k|k}, P^{i,\sigma_i}_{k|k}\right) &\text{(17b)} \end{aligned}$$

when $|\mathbf{X}_k| = |\mathbf{X}^p_{k|k-1}|$, and $f(\mathbf{Z}_k|\mathbf{X}_k)\delta(\mathbf{X}_k, \mathbf{X}^p_{k|k-1}) = 0$ otherwise. The last equality follows from using the Kalman filter measurement update for each association, note that $m^{i,\sigma_i}_{k|k} = m^i_{k|k-1}$ and $P^{i,\sigma_i}_{k|k} = P^i_{k|k-1}$ for $\sigma_i = 0$. The likelihoods are

$$\begin{aligned} \mathcal{L}^{p,\theta}_{k|k-1} =& e^{-\lambda_c}\left(\frac{\lambda_c}{V}\right)^{N_{FA}(\theta)}(1-p_D)^{N_{MD}(\theta)}p_D^{N_D(\theta)} \\ &\times\left[\prod_{i\in I_p:\sigma_i>0}\mathcal{N}\left(\mathbf{z}_k^{\sigma_i}\,;\,\hat{\mathbf{z}}_k^i, S_k^i\right)\right] \end{aligned} \tag{18}$$

where $N_{FA}(\theta)$ is the number of measurements that are not associated to a target, $N_{MD}(\theta)$ is the number of targets that are not associated to any measurement, and $N_D(\theta)$ is the number of targets that are associated to a measurement.

Note that (17) includes a summation over Θ, the set of all possible measurement associations θ. Except for very simple scenarios with few targets and high signal to noise ratio, this is computationally infeasible. To alleviate the computational complexity data association can be used, using any of the standard association algorithms, see e.g. [2]. Here we use the JPDA algorithm to compute association probabilities

$$\pi^{i,j}_{k|k-1} = \begin{cases} P(\mathbf{x}^i_{k|k-1}\leftrightarrow\mathbf{z}^j_k) & \forall i, j = 1,\ldots,N^z_k \\ P(\mathbf{x}^i_{k|k-1}\leftrightarrow\emptyset) & \forall i, j = 0\end{cases} \tag{19}$$

where $x \leftrightarrow z$ means "estimate x is associated to measurement z". Using the association probabilities we approximate (17) with

$$f(\mathbf{Z}_k|\mathbf{X}_k)\delta(\mathbf{X}_k, \mathbf{X}^p_{k|k-1}) \approx \delta(\mathbf{X}_k, \widetilde{\mathbf{X}}^p_{k|k}) \tag{20a}$$

$$\delta(\mathbf{X}_k, \widetilde{\mathbf{X}}^p_{k|k}) = \widetilde{\mathcal{L}}^p_{k|k-1}\prod_{i\in I_p}\mathcal{N}\left(\mathbf{x}_i\,;\,\widetilde{m}^i_{k|k}, \widetilde{P}^i_{k|k}\right) \tag{20b}$$

if $|\mathbf{X}_k| = |\widetilde{\mathbf{X}}^p_{k|k}|$ and zero otherwise. The natural logarithm of $\widetilde{\mathcal{L}}^p_{k|k-1}$ is computed as

$$\begin{aligned} \log(\widetilde{\mathcal{L}}^p) =& -\lambda_c + \log\left(\frac{\lambda_c}{V}\right)\left(N^z - \sum_i\sum_{j>0}\pi^{i,j}\right) \\ &+ \log(1-p_D)\sum_i\pi^{i,0} + \log(p_D)\sum_i\sum_{j>0}\pi^{i,j} \\ &+ \sum_i\sum_{j>0}\pi^{i,j}\log\left(\mathcal{N}\left(\mathbf{z}^j\,;\,\hat{\mathbf{z}}^i, S^i\right)\right) \end{aligned} \tag{21}$$

where we have omitted the time indexing for the sake of brevity. The measurement updated Gaussian distributions in (20) are given by the JPDA measurement update equations

$$\begin{aligned} S_k^i =& H_kP^i_{k|k-1}H_k^{\mathrm{T}} + R_k &\text{(22a)} \\ K_k^i =& P^i_{k|k-1}H_k\left(S_k^i\right)^{-1} &\text{(22b)} \\ \hat{\mathbf{z}}_k^i =& H_km^i_{k|k-1} &\text{(22c)} \\ \hat{\mathbf{z}}_k^{i,Eq} =& \pi^{i,0}_{k|k-1}\hat{\mathbf{z}}_k^i + \sum_{j>0}\pi^{i,j}_{k|k-1}\mathbf{z}_k^j &\text{(22d)} \\ \widetilde{m}^i_{k|k} =& m^i_{k|k-1} + K_k^i\left(\hat{\mathbf{z}}_k^{i,Eq} - \hat{\mathbf{z}}_k^i\right) &\text{(22e)} \\ m_k^{i,j} =& m^i_{k|k-1} + K_k^i\left(\mathbf{z}_k^j - \hat{\mathbf{z}}_k^i\right) &\text{(22f)} \\ M_k^i =& (m^i_{k|k-1} - \widetilde{m}^i_{k|k})(m^i_{k|k-1} - \widetilde{m}^i_{k|k})^{\mathrm{T}} &\text{(22g)} \\ M_k^{i,j} =& (m_k^{i,j} - \widetilde{m}^i_{k|k})(m_k^{i,j} - \widetilde{m}^i_{k|k})^{\mathrm{T}} &\text{(22h)} \\ \widetilde{P}^i_{k|k} =& \pi^{i,0}\left(P^i_{k|k-1} + M_k^i\right) &\text{(22i)} \\ &+ \sum_{j>0}\pi^{i,j}\left(P_{k|k-1} - K_k^iS_k^i(K_k^i)^{\mathrm{T}} + M_k^{i,j}\right) \end{aligned}$$

Under the approximation (20), the Bayes normalization constant $f(\mathbf{Z}_k|\mathbf{Z}^{k-1})$ is zero for $|\mathbf{X}_k| \neq |\mathbf{X}^p_{k|k-1}|$, and for $|\mathbf{X}_k| = |\mathbf{X}^p_{k|k-1}|$ it becomes

$$\int f(\mathbf{Z}_k|\mathbf{X}_k)\delta(\mathbf{X}_k, \mathbf{X}^p_{k|k-1})\delta\mathbf{X}_k = \widetilde{\mathcal{L}}^p_{k|k-1} \tag{23}$$

We thus have a multiobject particle approximation of the posterior multiobject density

$$f(\mathbf{X}_k|\mathbf{Z}^k) = \frac{\sum_{p=1}^P\widetilde{\mathcal{L}}^p_{k|k}\delta(\mathbf{X}_k, \widetilde{\mathbf{X}}^p_{k|k})}{\sum_{p=1}^P\widetilde{\mathcal{L}}^p_{k|k}} \tag{24a}$$

$$= \sum_{p=1}^P\mathcal{W}^p_{k|k}\delta(\mathbf{X}_k, \widetilde{\mathbf{X}}^p_{k|k}) \tag{24b}$$

and the posterior PHD intensity is

$$D_{k|k}(\mathbf{x}) = \sum_{p=1}^{P} \mathcal{W}_{k|k}^{p} \sum_{i\in I_p} \mathcal{N}\left(\mathbf{x}\,;\, \widetilde{m}_{k|k}^{p,i}, \widetilde{P}_{k|k}^{p,i}\right) \tag{25}$$

After the measurement update there are $\sum_{p=1}^{P} |I_p|$ components in the Gaussian mixture approximation of the PHD intensity. However, many of the components will be identical and can thus be merged without approximation. Further, many components will have very low weight and can thus be discarded with very little approximation. In a mixture reduction algorithm components with weight lower than a threshold τ are pruned (i.e. removed from the mixture), followed by mixture merging where the components are merged such that the merged weight is less than one. The mixture reduction results in a PHD intensity

$$D_{k|k}(\mathbf{x}) = \sum_{i=1}^{J_{k|k}} w_{k|k}^{i} \mathcal{N}\left(\mathbf{x}\,;\, m_{k|k}^{i}, P_{k|k}^{i}\right) \tag{26}$$

Target estimates are extracted from the reduced posterior PHD intensity (26) by picking the expected value of the Gaussians whose weights are larger than a threshold T. In the simulation studies that we have performed, the reduced Gaussian mixture often has a single Gaussian component for each target estimate, where the weights are indicative of the probability that the estimate corresponds to an actual target.

The multiobject particles (MOP) can also be used to compute approximations of the predicted and posterior cardinality distributions. The probability of n targets in the surveillance area is given by the sum of multiobject particle weights for particles with cardinality n,

$$P_{k|k-1}(N_k^x = n) = \sum_{p:|I_p|=n} P^{-1} \tag{27a}$$

$$P_{k|k}(N_k^x = n) = \sum_{p:|I_p|=n} \mathcal{W}_{k|k}^{p} \tag{27b}$$

Note that the predicted particle weights in (27a) are all equal to P^{-1}, see (14). In comparison, the standard PHD filter approximates the cardinality distributions with a Poisson distribution that – when a Gaussian mixture is used to approximate the PHD intensity – has expected values $\sum_i w_{k|k-1}^{i}$ and $\sum_i w_{k|k}^{i}$. The approximations (27) typically have much lower variance than Poisson distributions – this is an attractive property since the high Poisson variance causes the standard PHD filter to be sensitive missed detections and clutter.

When the particle approximation of the posterior multi-object density is transformed into an approximation of the posterior PHD intensity, the only information in the posterior cardinality distribution that is formally retained is the expected value

$$\hat{N}_{k|k}^{x} = \sum_{n\geq 0} n P_{k|k}(N_k^x = n) \tag{28}$$

However, by the interpretation taken here – the component weights are probabilities of existence – an estimate of the cardinality distribution can be computed using the component weights $w_{k|k}^{i}$ in the reduced posterior PHD intensity (26).

Lastly, we note that through the use of the joint probabilistic data association target trajectories can easily be obtained by maintaining a history of data associations for each target estimate.

IV. Results

We compare the proposed approach to PHD measurement update to the standard PHD measurement update using simulated data. Both of the compared PHD filters use a Gaussian mixture to approximate the PHD.

A. Simulation setup and performance evaluation

Position measurements were generated with probability of detection $p_D = 0.75$ and noise covariance matrix $R_k = 10^2\mathbf{I}_2$. A nearly constant velocity motion model (white noise acceleration [1]) was used with acceleration noise standard deviation $\sigma_a = 2$ m/s^2. The probability of survival was set to 0.98. For mixture reduction we use $\tau = 10^{-5}$, for target extraction we use threshold $T = 0.5$. For the data assocation, gating probability $P_G = 0.99$ was used.

For performance evaluation we compare the estimated cardinality and the optimal subpattern assignment metric (OSPA) [15]. The OSPA is implemented using the Euclidean norm with cut-off parameter $c = 300$ and $p = 1$. For further details, please see [15]. For the cardinality, we compare two cardinality estimates. The first is the sum of weights, and the other is the number of extracted targets, i.e. the number of Gaussian components for which $w_{k|k}^{i} > T$. In addition to this, for a simple scenario we also compare the approximate MOP cardinality distributions (27) to the standard PHD filter's Poisson distribution.

B. Single measurement update example

In Figure 1 we show the benefits of the MOP-PHD filter, compared to the standard PHD filter, for a single measurement update step. The predicted PHD, Figure 1a, has two birth components and four predicted components, all four with high weights. There are seven detections, of which four are located close to PHD components. Intuitively we can expect the following:

- the weights corresponding to detected predicted estimates should increase to 1 (or just below), since these estimates already had a high weight and were detected;
- the weight corresponding to the not-detected predicted estimate should decrease conservatively, since there is fairly high probability that there was just a missed detection;
- the weight of the detected birth estimate should increase conservatively, since there is a chance that it was only a clutter detection;
- the weight of the not-detected birth estimate should decrease significantly.

Using the MOP measurement update, see Figure 1b, we see that all expectations are met. Using the standard measurement update, see Figure 1c, the weights of the detected predicted

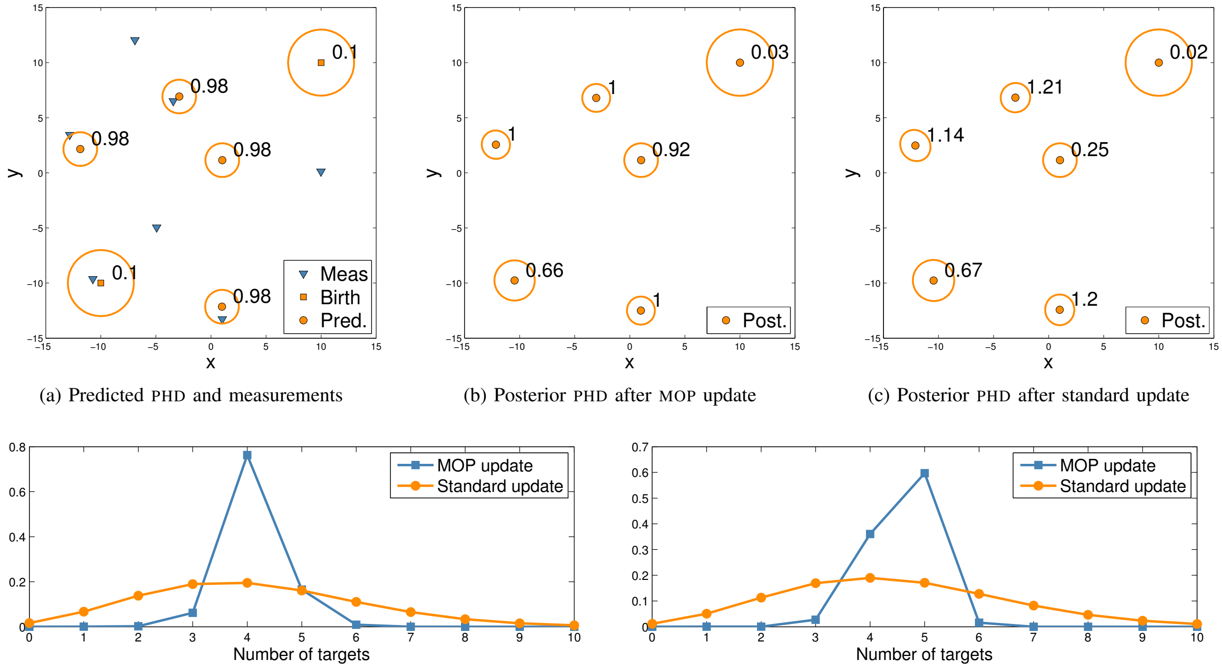

(a) Predicted PHD and measurements (b) Posterior PHD after MOP update (c) Posterior PHD after standard update

(d) Predicted cardinality distribution (e) Posterior cardinality distribution

Fig. 1. Comparison of MOP and standard update. (a): seven measurements (blue triangles) and a predicted PHD with two birth components (orange squares) and four predicted components (orange circles). The ellipses show the covariances $P^j_{k|k-1}$ and the numbers are the weights $w^j_{k|k-1}$. (b), (c): the PHD after MOP measurement update and standard measurement update, respectively. (d) and (e) show the predicted and posterior cardinality distributions, respectively.

estimates have a positive bias of around 0.2, and the weight of the not detected predicted estimate has dropped rather abruptly to 0.25. The standard measurement update can only match the MOP update in terms of the updated weights of the birth components.

The predicted and posterior cardinality distributions are shown in Figure 1d and Figure 1e. The expected values are almost equal in both cases, about 4.1 (predicted) and 4.6 (updated), however the variance is considerably lower using the MOP measurement update.

C. Multiple target scenario

A multiple target scenario was simulated with ten targets. A birth PHD with two components was used. For each target a birth time was randomly simulated, and the initial state was sampled from the birth PHD. A nearly constant velocity model was used to simulate the motion. The death times were randomly sampled such that all ten targets existed at the same time for at least a couple of time steps. Clutter was simulated with Poisson rate 100, i.e. on average 100 clutter measurements per time step. In each measurement update $P = 1000$ multiobject particles were created.

The scenario was simulated 250 times, the Monte Carlo results are shown in Figure 2. For the three performance measures we show Monte Carlo median, as well as uncertainty regions given by the 10th and 90th percentiles. The OSPA is significantly smaller using the MOP update, which is largely an effect of the higher robustness to missed detections. Using the standard measurement update, a missed detection typically leads to a cardinality error.

For the sum of weights, the positive bias and sensitivity to missed detections appear to average out for the standard PHD filter, and the median results are more or less identical for the two filters. Notice that the uncertainty region is much smaller using the MOP update.

In contrast to the sum of weights, the number of extracted targets show a significant difference between the two filters. Using the MOP update, the cardinality errors are typically due to the filter being slow at “killing” target estimates after the true targets have indeed disappeared from the surveillance area. In contrast, the standard PHD filter show considerable error when there are many targets.

The differences that can be observed here are a manifestation of the improved cardinality model: randomly sampling the existence of predicted estimates in a multiobject particle representation is much more accurate than modeling the cardinality distribution with a Poisson distribution.

D. Target trajectory estimation

A scenario with two targets was simulated, where the trajectories cross about halfway through the simulation. True trajectories and estimates are shown in Figure 3, where we have marked the time steps for which the same measurements fall inside both estimates’ data association gates. Using the

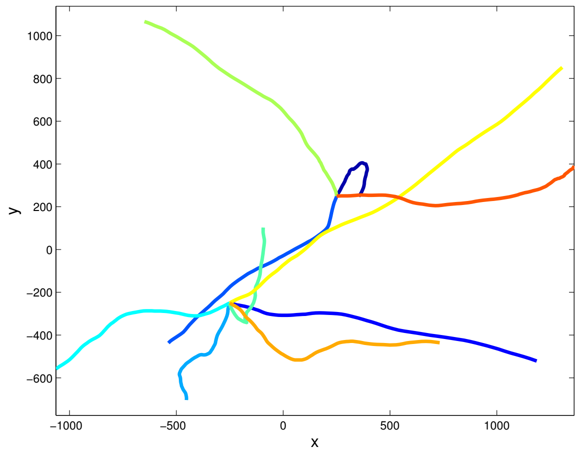

(a) True target tracks

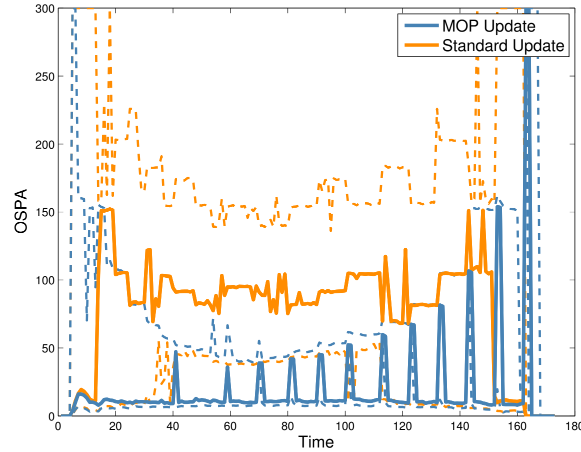

(b) OSPA metric

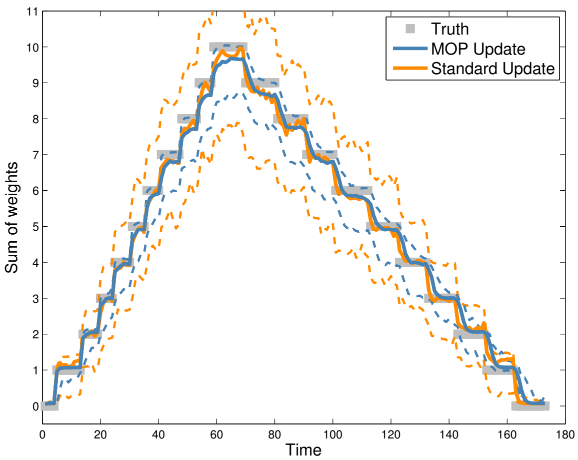

(c) Sum of weights

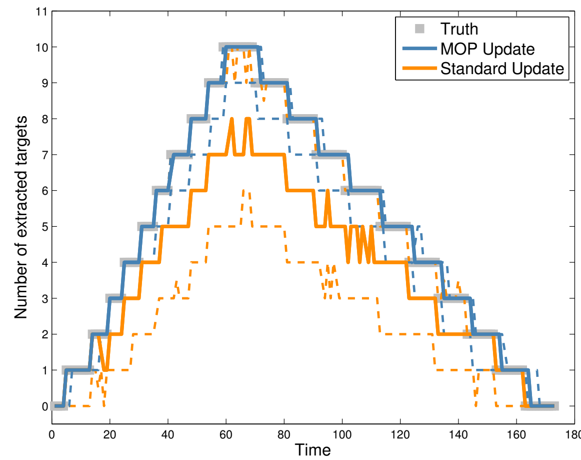

(d) Number of extracted targets

Fig. 2. Results from a scenario with ten targets. (a) shows the true target tracks. (b) shows the OSPA metric, (c) shows the sum of weights, and (d) shows the number of extracted targets. Solid lines are Monte Carlo median, dashed lines of the same color show the 10th and 90th percentiles. The results using the MOP update are clearly better than the results using the standard update.

JPDA measurement update the proposed filter can correctly estimate both trajectories.

V. Discussion

In this section we discuss the MOP-PHD filter in relation to other multiple target tracking filters, and we also discuss the computational complexity.

A Bernoulli RFS is a single target tracking filter that is defined by the parameters $(r, p(\mathbf{x}))$, representing the probability of existence and a spatial distribution, respectively. In multi-Bernoulli filters the set of targets is modeled as a union of Bernoulli RFSs, and the multi-Bernoulli RFS is defined by a parameter set $\{(r^i, p^i(\mathbf{x}))\}_i$. In the filters, the probabilities of existence and the spatial distributions are recursively updated. The interpretation of the PHD intensity weights as probabilities of existence makes the MOP-PHD filter rather similar to multi-Bernoulli filters. The multiobject particles $\delta(\mathbf{X}, \mathbf{X}^p)$ can actually be seen as special cases of multi-Bernoulli densities where cardinality and existence is assumed known. In other words, the cardinality distribution is $P(n) = 1$ if $n = |\mathbf{X}^p|$ and zero otherwise, and all Bernoulli existence probabilities are exactly one. In both the predicted and posterior particle approximation of the multiobject density, the Gaussian distributions are (typically) included in more than one particle. It then follows that the relative frequency of each Gaussian can be taken as an estimate of the probability of existence.

To obtain target trajectory estimates labeled RFSs can be used in the multi-Bernoulli framework, see [12]. In the proposed filter target trajectory estimates are given by the use of the JPDA state update. An important topic for future work is further theoretical analysis of the proposed update, along with a thorough comparison to other multiple target tracking filters,

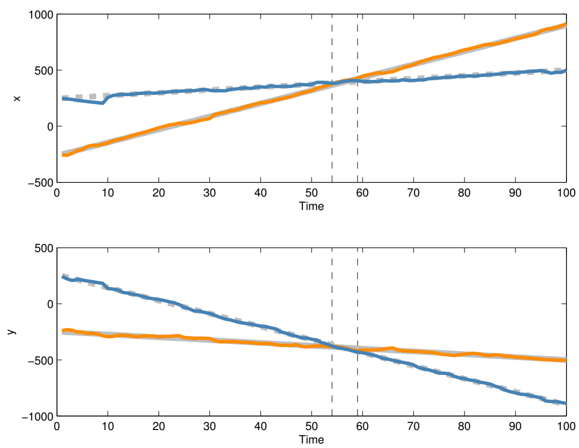

Fig. 3. Results from trajectory estimation, x-position (top) and y-position (bottom) shown over time. The true trajectories are shown in gray, the estimates are shown in blue and orange. The time steps inbetween which both estimates are associated to the same measurements are indicated by the black dashed lines.

e.g. multi-Bernoulli filters.

Implementation of the MOP-PHD filter is simple using the well known Gaussian mixture PHD prediction [16], sampling from a uniform distribution, and the JPDA update. For a predicted PHD with J components, the larger J is, the more multiobject particles are needed to obtain an accurate approximation of the predicted multiobject density. There are 2^J possible ways to construct unique multiobject particles, and for smaller J it is possible to exhaustively enumerate all possibilities, however for larger J this is infeasible. One alternative is to limit the possibilities by considering only the more probable ones, as indicated by the weights. Similar ideas are proposed for the measurement update in the labeled multi-Bernoulli (LMB) filter [12].

Another idea that can be utilized here is state space partitioning, i.e. the grouping of estimates and measurements into distinct subsets that are approximately statistically independent. The updates are then performed separately on each group. This idea has been used successfully to alleviate computational complexity in point target tracking, e.g. in the LMB filter [12], and in extended target tracking using a PHD filter [14].

VI. Conclusions and Future Work

In this paper we proposed a new update that can be used to estimate the PHD intensity. The new update has three main ingredients: particle approximation of the predicted multiobject density by sampling from the predicted PHD; update of each particle using the multiobject measurement pdf and JPDA filter association probabilities; and approximation of the posterior PHD using the measurement-updated particle approximation of the multiobject density. Simulation studies showed that the proposed filter significantly outperforms the standard PHD filter, especially in terms of the estimated cardinality.

Future work include more testing of the filter, using both simulated and real world data, as well as further theoretical analysis of the MOP-PHD filter. An analysis of the computational complexity in comparison to other filters will also be performed.

Acknowledgment

This research was supported by the Naval Postgraduate School, via ONR N00244-14-1-0033, and by ONR directly via N00014-13-1-0231; Y. Bar-Shalom is also supported by ARO W991NF-10-1-0369.

References

[1] Y. Bar-Shalom, X. Rong-Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York, NY, USA: John Wiley & Sons, 2001.

[2] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, 2011.

[3] M. Beard, B. Vo, B.-N. Vo, and S. Arulampalam, “A partially uniform target birth model for gaussian mixture phd/cphd filtering,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 4, pp. 2835–2844, Oct. 2013.

[4] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA, USA: Artech House, 1999.

[5] D. Fränken, M. Schmidt, and M. Ulmke, “'Spooky Action at a Distance‘ in the Cardinalized Probability Hypothesis Density Filter,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 4, pp. 1657–1664, Oct. 2009.

[6] Q. Le and L. Kaplan, “Probability Hypothesis Density-Based Multitarget Tracking for Proximity Sensor Networks,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 3, pp. 1476–1496, Jul. 2013.

[7] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2007.

[8] ——, *Advances in Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2014.

[9] ——, “Multitarget Bayes filtering via first-order multi target moments,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003.

[10] ——, “PHD filters of higher order in target number,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 4, pp. 1523–1543, Oct. 2007.

[11] K. Panta, D. Clark, and B.-N. Vo, “Data association and track management for the Gaussian mixture probability hypothesis density filter,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 3, pp. 1003–1016, Jul. 2009.

[12] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, “The Labeled Multi-Bernoulli Filter,” *IEEE Transactions on Signal Processing*, vol. 62, no. 12, pp. 3246–3260, Jul. 2014.

[13] B. Ristic, D. Clark, B.-N. Vo, and B.-T. Vo, “Adaptive target birth intensity for PHD and CPHD filters,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 48, no. 2, pp. 1656–1668, Apr. 2012.

[14] A. Scheel, K. Granström, D. Meissner, S. Reuter, and K. Dietmayer, “Tracking and Data Segmentation Using a GGIW Filter with Mixture Clustering,” in *Proceedings of the International Conference on Information Fusion*, Salamanca, Spain, Jul. 2014.

[15] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, “A consistent metric for performance evaluation of multi-object filters,” *IEEE Transactions on Signal Processing*, vol. 56, no. 8, pp. 3447–3457, Aug. 2008.

[16] B.-N. Vo and W.-K. Ma, “The Gaussian mixture probability hypothesis density filter,” *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4091–4104, Nov. 2006.

[17] B.-T. Vo and B.-N. Vo, “Labeled random finite sets and multi-object conjugate priors,” *IEEE Transactions on Signal Processing*, vol. 61, no. 13, pp. 3460–3475, Apr. 2013.

[18] B.-T. Vo, B.-N. Vo, R. Hoseinnezhad, and R. Mahler, “Robust multi-bernoulli filtering,” *IEEE Journal of Selected Topics in Signal Processing, Special Issue on Multi-target Tracking*, vol. 7, no. 3, pp. 399–409, Jun. 2013.