



CHALMERS

Chalmers Publication Library

Approximate Multi-Hypothesis Multi-Bernoulli Multi-Object Filtering Made Multi-Easy

This document has been downloaded from Chalmers Publication Library (CPL). It is the author's version of a work that was accepted for publication in:

IEEE Transactions on Signal Processing (ISSN: 1053-587X)

Citation for the published paper:

Granström, K. ; Willett, P. ; Bar-Shalom, Y. (2016) "Approximate Multi-Hypothesis Multi-Bernoulli Multi-Object Filtering Made Multi-Easy". IEEE Transactions on Signal Processing, vol. 64(7), pp. 1784-1797.

<http://dx.doi.org/10.1109/tsp.2015.2500884>

Downloaded from: <http://publications.lib.chalmers.se/publication/234762>

Notice: Changes introduced as a result of publishing processes such as copy-editing and formatting may not be reflected in this document. For a definitive version of this work, please refer to the published source. Please note that access to the published version might require a subscription.

Chalmers Publication Library (CPL) offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all types of publications: articles, dissertations, licentiate theses, masters theses, conference papers, reports etc. Since 2006 it is the official tool for Chalmers official publication statistics. To ensure that Chalmers research results are disseminated as widely as possible, an Open Access Policy has been adopted. The CPL service is administrated and maintained by Chalmers Library.

(article starts on next page)

Approximate Multi-Hypothesis Multi-Bernoulli Multi-Object Filtering Made Multi-Easy

Karl Granström, *Member, IEEE*, Peter Willett, *Fellow, IEEE*, and Yaakov Bar-Shalom, *Fellow, IEEE*

Abstract—In multiple target tracking (MTT) it becomes necessary to use a multi-hypothesis approach if the trajectories of two or more targets cross. However, multi-hypothesis approaches, e.g. the Multiple Hypothesis Tracker (MHT) or the emerging Generalized Labelled Multi-Bernoulli (GLMB) filter, are computationally demanding. In this paper we propose a simple multi-Bernoulli (MB) filter and a post processing method, which together deliver a multi-hypothesis tracking estimate at a computational cost that is only slightly larger than the cost of a single-hypothesis tracking filter even for many targets. The proposed MB filter is shown to be similar to the labeled MB filter, itself an approximation of the multi-hypothesis GLMB filter. In a simulation study with multiple targets and several trajectory crossings the proposed filter is shown to be capable of correctly estimating the multi-hypothesis output. The filter is also tasked with presenting to an operator a principled perspective on a scene with many feasible track switches.

I. INTRODUCTION

MTT is the processing of sets of measurements obtained from multiple sources in order to maintain estimates of targets' current states. The task is complicated by the fact that – in addition to noise, missed detections and clutter – the number of targets is unknown and time-varying. 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) [11], [12]. The MHT type approaches involve propagating target track hypotheses in time and calculating their likelihoods, the JPDA type approaches blend data association probabilities on a scan-by-scan basis, 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 (pdf) of the multiobject state in time. Because of the computational complexity of the data association problem it is generally considered infeasible to implement and use a multi-object filter without approximating the data association problem in some way. Computationally feasible approximate filters include the Probability Hypothesis Density (PHD) filters [13], the Cardinalized PHD (CPHD) filters [14], and the multi-Bernoulli filters [17], [18], [22]–[24], [26].

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Karl Granström is with the Department of Signals and Systems, Chalmers University of Technology, Gothenburg, Sweden. E-mail: karl.granstrom@chalmers.se.

Yaakov Bar-Shalom and Peter Willett are with the Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT, USA. E-mail: {ybs,willett}@engr.uconn.edu.

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.

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 of the PHD filter is the high variance of its cardinality estimate, this being a consequence of the underlying Poisson assumption. The CPHD filters are known to have better cardinality estimates, but remain susceptible to a “spooky effect” [8], [12], a phenomenon manifested by PHD mass shifted from undetected targets to detected targets that are far enough away that they ought to be statistically insulated. The multi-Bernoulli filters [17], [18], [22], [24], [26] approximate the multiobject density with a multi-Bernoulli distribution, which is then propagated and updated in time. Multi-Bernoulli filters estimate, for each target, its location and probability of existence; and they are known to be capable of matching the CPHD filters' cardinality performance without being susceptible to a “spooky effect” [8].

Ultimately the desired output from an MTT algorithm is a set of estimated trajectories (tracks), where a trajectory is defined as the sequence of states from the time the target appears to the time it disappears. Both the MHT and JPDA type algorithms estimate trajectories. In their most basic forms, none of the PHD [13], CPHD [14] nor multi-Bernoulli filters [23] formally estimates target trajectories – only a set of target state estimates is supplied at each time step – however, target trajectories can be obtained, e.g. using labeling schemes [16]. With the introduction of labeled RFS it has become possible to obtain trajectory estimates without the need for post-processing, leading to the Generalized Labeled Multi-Bernoulli (GLMB) filter [22], and its computationally efficient approximation the Labeled Multi-Bernoulli (LMB) filter [18]. An analysis of the approximation error in the GLMB filter is given in [25].

Some recent work [10] derived an interesting improvement to the PHD that integrates classical multi-target data association models very easily. Although [10] is attractive due to its PHD simplicity, a relationship to the better-performing MB approach can be suggested. In this paper we show that their combination has both simplicity and performance: we call it the Multi-Object Particle Multi-Bernoulli (MOP-MB) filter. Not only is it an excellent integrated target tracker, and not only can it track in dense situations (we show an example with 150 targets), but it turns out to be nicely adaptable to low-complexity recovery of multi-target hypotheses when targets cross and their identities become confused. This we call the Approximate Multi-Hypothesis Multi-Bernoulli (AMHMB) filter.

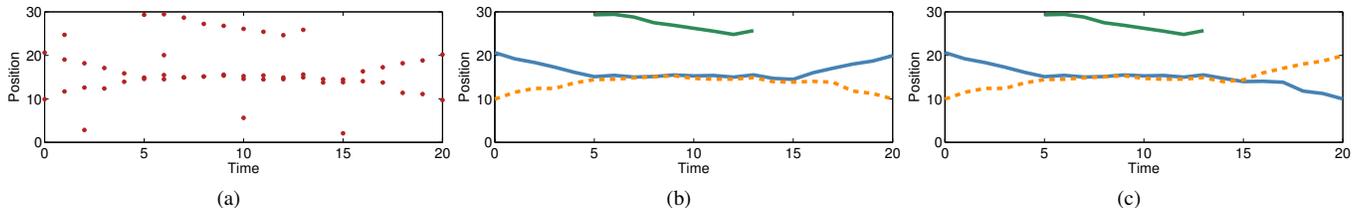


Fig. 1. Notional figure of trajectory crossing. Using linear Gaussian motion and measurement models, filtering the sequence of measurements in (a) has many solution hypotheses, however the two trajectory set hypotheses shown in (b) and (c) represent the most likely solutions. Because the measurements do not contain any information regarding target identity both hypotheses are equally likely.

TABLE I
LIST OF ACRONYMS

AMHMB	Approximate Multi Hypothesis MB
CPHD	Cardinalized PHD
EM	Expectation Maximization
GLMB	Generalized Labelled MB
JPDA	Joint Probabilistic Data Association
KL	Kullback Leibler
LMB	Labelled MB
MB	Multi Bernoulli
MHT	Multi Hypothesis Tracking
MOP	Multi Object Particle
MTT	Multi Target Tracking
OSPA	Optimal Sub-Pattern Assignment
PHD	Probability Hypothesis Density
RFS	Random Finite Set
SJPDA	Set JPDA
VMB	Variational MB

A list of acronyms is given in Table I. The paper¹ is organized as follows. In the next Section II we motivate by explaining the need for a joint multi-target multi-hypothesis tracking approach, as opposed to (for example) the PHD. The paper’s two main contributions (MOP-MB and AMHMB) are mapped in Section III. In Section IV we give background on multi-Bernoulli filters. In Section V we present the proposed multiple target tracking filter. Simulation results are presented in Section VII. The paper is concluded in Section VIII.

II. WHY A JOINT MULTI-TARGET REPRESENTATION?

Consider the measurements in Figure 1a, where the measurement provenances (target or clutter?) are assumed unknown. Due to measurement origin uncertainty the number of possible scenarios that could have caused this measurement sequence grows exponentially with time. One solution to this indeed is to consider all possible data association hypotheses, however this would come at the price of ever-increasing computation.

Under some realistic and non-restrictive assumptions, the two hypotheses in Figures 1b and 1c can reasonably be seen as the most likely solutions to the MTT problem given the measurements in Figure 1a. Note that, because the measurements do not provide information about target identity, the two hypotheses are (approximately) equally probable. In this paper we designate these as multi-object hypotheses.

An ideal solution to the MTT problem would output a set of trajectories that are “closest” to the true set of trajectories.

For example, the MHT is often implemented to output the most likely hypothesis. Indeed, if the targets are well separated in state space, the single most probable solution hypothesis will likely yield an accurate estimate. However, when target trajectories cross, a single multi-object hypothesis is insufficient. For example, in Figure 1 the true targets are located in the same position from time $k = 5$ to $k = 15$, after which they separate. When the targets separate starting at time $k = 16$ it cannot be inferred from the past measurements which target is going up (increasing position) and which is going down. In other words, both hypotheses are equally likely, and picking out one of them, based on an estimated hypothesis probability that numerically is perhaps only slightly larger than the second most probable hypothesis, does not correctly represent this ambiguity. Therefore a multi-hypothesis approach is necessary to correctly represent the scenario.

Because of the possibility of trajectory crossings, we argue that a solution to the MTT problem should output a set of the most probable multi-object hypotheses, along with probabilities of each hypothesis being true. Depending on the underlying scenario and the sequence of measurements, this output may, or may not, contain more than one hypothesis.

A. Non-Joint Trackers

The JPDA, PHD, CPHD, and LMB filters are all examples of MTT filters with a single multi-object hypothesis. The filters can be implemented with multi-modal representations of the single target densities (e.g. using a Gaussian mixture or a particle representation), which yields an output that is multi-hypothesis on the single target level. However, it is not multi-hypothesis on the multi-object level, because, given multiple single target densities, there is no information regarding which single target modes can coexist, and which are mutually exclusive. To obtain a multi-hypothesis multi-object output from multi-modal target densities post-processing is needed.

Furthermore, a multi-modal single target density representation, e.g. a Gaussian mixture, raises the question of how a point estimate is obtained:

- If there are multiple distinct modes with equal weights, which one should be chosen? Taking the weighted mean of distinct modes with equal weights will give an estimate that is clearly not correct, as it is not close to any of the modes. This is similar to the track coalescence problem [5], [7].
- Assume a point estimate is taken as the mean of the most probable mode. If there are two distinct modes with equal

¹A very preliminary version of this work is in [9].

weights, and two measurements each near a mode, the measurement update will give a slightly larger weight to the mode whose mean is closer to the corresponding measurement. If in the next time step the other mode is more probable, then the point estimate output will be “flickering” between the two modes. This is similar to the track switching problem, see [6].

Thus, from an output perspective, multi-modal single target density representations are sub-optimal. It must be noted though, that single target densities modeled by Gaussian mixtures or particle filters have been shown to be very useful, e.g. when the motion and/or measurement models are non-linear. In this case, to avoid the problems outlined above, care should be taken to ensure that the densities are essentially uni-modal.² Further study of target trajectory display can be found in, e.g., [6].

B. Joint Trackers

The MHT and GLMB are examples of filters that maintain multiple multi-object hypotheses, but those filters also have a computational cost (as measured by, e.g., average cycle time) that is considerably higher than that of the JPDA, PHD, CPHD and LMB filters. The drawbacks of an MTT filter with multiple multi-object hypotheses can be summarized in two points:

- High computational complexity: Slightly simplified, handling several multi-object hypotheses is akin to running several single-hypothesis tracking filters in parallel, one for each multi-object hypothesis. In case there are multiple trajectory crossings, the hypothesis space increases quickly, with a corresponding increase in computational complexity.
- Duplicate estimates: if multiple multi-object hypotheses are estimated there will be duplicate estimates. Consider the example in Figure 1: after the true targets separate starting at time 16, the blue and the orange estimates will quickly converge to the same values. In other words, assuming Gaussian single target densities, after convergence the means and covariances of the blue estimate in Figure 1b are equal to the means and covariances of the orange estimate in Figure 1c. The result is that the same information (means and covariances) is computed and stored twice, and if there are multiple trajectory crossings, there will be a lot of duplicated information.

The goal of this paper is to devise a way to obtain a multi-hypothesis output without using a computationally expensive multi-hypothesis filter. This is achieved by using a simple single-hypothesis multi-Bernoulli filter, and post-processing to produce a multi-hypothesis output. The reconstruction of the multi-hypothesis output is based on estimating the probability of trajectory crossings.

III. CONTRIBUTIONS

A. The MOP-MB filter

The first contribution of this paper is a filter with an update inspired by the work in [10]. Given a particle approximation

²If a Gaussian mixture is used, the sum is uni-modal. If a particle filter is used, the particles form a single cluster.

of the predicted PHD intensity, in [10] the following update was proposed:

- 1) Obtain multi-object particles (MOPs) with states that are randomly sampled (with replacement, and with varying cardinality) from the predicted PHD.
- 2) Compute the weight of each MOP by evaluating the multi-object measurement pdf and normalizing.
- 3) Obtain a particle approximation of the posterior PHD intensity from the updated MOPs.

The above procedure does improve the filter performance compared to a PHD filter with the standard update (see e.g. [13]), however the drawbacks of the Poisson assumption for the cardinality are retained. In this paper we relax the Poisson cardinality assumption and instead assume that an MB approximation of the predicted multi-object density is given. The MOP idea is used in the update of the parameters of the predicted MB density, because the MOPs are a great way of handling the unknown cardinality, i.e. the fact that at each time step it is not known how many new targets appear, how many extant targets persist, nor how many extant targets disappear. The update basically follows the same three steps as above, details are given in Section V.

In an early version of this work, see [9], JPDA association probabilities were used to simplify the data association in step 2 above; in this paper a single most probable association event is used, computed using the auction algorithm [3]. The posterior MB density is predicted using the MB prediction proposed in [18], [22], and the proposed filter is called Multi-Object Particle Multi-Bernoulli (MOP-MB) filter.

The MOP-MB has obvious relation to [10], and also shares its similarities to the MB work in [18], [22], [26]. The Generalized Labeled Multi-Bernoulli (GLMB) filter [22] predicts and updates a labeled MB mixture density, representing the target densities and different hypotheses for target existence and data association. The GLMB density is closed under the multi-object prediction and update [22], the GLMB filter formally produces tracks, and in [12] the filter is claimed to be the first fully Bayesian approach to MTT. The disadvantage of the filter is the complexity with an exponential growth in the number of components in the MB mixture, both in the prediction and the update.

The high computational cost of the GLMB filter can be alleviated by approximating the GLMB density with a labeled MB density, resulting in the Labeled Multi-Bernoulli (LMB) filter [18]. The LMB filter predicts a labeled MB density, uses it to approximate a GLMB density which is updated using the GLMB update [22], and then approximates the updated GLMB density with an LMB density that matches the first moment. A simulation study with a large number of targets and high clutter rate showed that the LMB filter has performance that matches the GLMB filter’s performance at lower computational cost, see [18].

The variational multi-Bernoulli (VMB) filter [26] approximates the multi-object posterior by finding the MB distribution that minimizes the set Kullback-Leibler (KL) divergence (defined in [11]). A solution to the minimization problem is found using variational approximation, specifically the expectation

maximization (EM) algorithm. Essentially, the variational approximation marginalizes the association uncertainty out of the update. In [26] the VMB filter is shown to be related to the set JPDA (SJPDA) filter [21], with the added benefit of handling an unknown number of targets (SJPDA assumes known number of targets). The VMB filter is also shown to approximate the minimum mean OSPA estimator [26].

Just like the LMB and VMB, the presented MOP-MB filter approximates the multi-object posterior with a multi-Bernoulli distribution. The construction of multi-object particles is similar to how, in the LMB filter, the LMB density is used to approximate the GLMB density. Exhaustive enumeration of the possible GLMB density components is used in the LMB filter, however the authors note that, for very large target numbers, sampling may be a better alternative [18]. Equivalently, exhaustive enumeration could be used in the MOP-MB filter. Thus the MOP-MB filter is similar to the LMB filter. It is also a “simpler” filter in the following sense:

- Sampling typically gives fewer MOPs than there are hypotheses in the GLMB after exhaustive enumeration.
- The GLMB update, used in both the LMB and GLMB filters, is suggested to be implemented using only the set of M most probable association events, e.g. computed using Murty’s algorithm [15]. The MOP-MB filter considers the single most probable association event, e.g. computed using the auction algorithm [3]. How many association events are considered is a trade-off between complexity and accuracy.
- The MOP-MB filter uses only a single Gaussian for each estimate, whereas the LMB filter uses a possibly multi-modal Gaussian mixture (or a particle filter, in severely non-linear scenarios).

The LMB approximation preserves the first moment of the posterior density, and as noted above the VMB minimizes the KL divergence. An important topic for future work is to investigate the relationship between the MOP-MB’s approximate posterior density and the actual posterior density.

B. The AMHMB Filter

Note that while the LMB, VMB and MOP-MB filters all consider multiple hypotheses for the target estimates to a varying degree, none of these filters is multi-hypothesis on the multi-object level in the sense that it can deliver an output like the one indicated in Figure 1. The reason for this lies in the approximation of the full multi-object density with a multi-Bernoulli density, which interestingly is also the key to the tractable complexity.

The second contribution in this paper is an approximative approach to multi-hypothesis multi-Bernoulli tracking that allows the MTT system to output an operationally-useful multi-hypothesis tracking solution at a computational cost that is nearly as low as a single hypothesis filter. The approximate multi-hypothesis method is based on estimating the probability that the true states corresponding to two estimates are equal, and then building multi-object identity switch hypotheses using these probabilities. Combined with the MOP-MB filter this results in the Approximate Multi-Hypothesis Multi-Bernoulli (AMHMB) filter.

IV. BACKGROUND: MULTI-BERNOULLI RANDOM FINITE SETS FOR MULTIPLE TARGET TRACKING

Some notations are given in Table II. Let \mathbf{x}_k^i denote the state of the i th target at time step k , and let the target set be denoted

$$\mathbf{X}_k = \{\mathbf{x}_k^i\}_{i=1}^{N_k^x} \quad (1)$$

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

$$\mathbf{Z}_k = \{\mathbf{z}_k^j\}_{j=1}^{N_k^z} \quad (2)$$

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 the measurement origin is assumed unknown. Note that the sets above are without order and the set indexing is arbitrary; the particular choices $i = 1, \dots, N_k^x$ and $j = 1, \dots, N_k^z$ are only used for notational simplicity and convenience.

The posterior multi-object 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}_κ from $\kappa = 0$ up to, and including, $\kappa = k - 1$. The predicted multi-object distribution is given by the Chapman-Kolmogorov equation

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

where $f(\mathbf{X}_k|\mathbf{X}_{k-1})$ is the multi-object transition density. Multi-object prediction involves modeling the time evolution of surviving targets (targets that remain in the surveillance area), target death (targets that do not remain), and target birth (new targets that appear in the surveillance area). The targets are assumed to evolve over time independently, with single target transition density $p_{k,k-1}(\mathbf{x}_k|\mathbf{x}_{k-1})$. The integral in (3) is a set integral, defined as

$$\int f(\mathbf{X})\delta\mathbf{X} = \sum_{n=0}^{\infty} \frac{1}{n!} \int f(\{\mathbf{x}^1, \dots, \mathbf{x}^n\}) d(\mathbf{x}^1, \dots, \mathbf{x}^n) \quad (4)$$

The posterior multi-object distribution at time t_k is given by the Bayes update

$$f(\mathbf{X}_k|\mathbf{Z}^k) = \frac{f(\mathbf{Z}_k|\mathbf{X}_k)f(\mathbf{X}_k|\mathbf{Z}^{k-1})}{\int f(\mathbf{Z}_k|\mathbf{X}_k)f(\mathbf{X}_k|\mathbf{Z}^{k-1})\delta\mathbf{X}_k} \quad (5)$$

where the integral in the denominator is a set integral. The multi-object measurement set density $f(\mathbf{Z}_k|\mathbf{X}_k)$ involves modeling target detection, measurement noise, and clutter measurements.

TABLE II
NOTATIONS

-
- \mathbb{R}^n is the set of real column vectors of length n .
 - \mathbb{S}_+^d is the set of positive semi-definite $d \times d$ matrices.
 - $|\mathbf{Y}|$ is the cardinality of the set \mathbf{Y} .
 - $\mathcal{N}(\mathbf{x}; \mathbf{m}, P)$ denotes a multi-variate Gaussian pdf over the vector $\mathbf{x} \in \mathbb{R}^{n_x}$ with mean vector $\mathbf{m} \in \mathbb{R}^{n_x}$, and covariance matrix $P \in \mathbb{S}_+^{n_x}$.
-

A Bernoulli RFS \mathbf{X}^j is a type of RFS that is empty with probability $1 - w^j$ or, with probability w^j , contains a single element with distribution $p^j(\mathbf{x})$. The cardinality is Bernoulli distributed with parameter w^j and the pdf of \mathbf{X}^j is

$$f(\mathbf{X}^j) = \begin{cases} 1 - w^j & \mathbf{X}^j = \emptyset \\ w^j \cdot p^j(\mathbf{x}) & \mathbf{X}^j = \{\mathbf{x}\} \\ 0 & |\mathbf{X}^j| \geq 2 \end{cases} \quad (6)$$

A typical assumption in multiple target tracking is that the targets are independent, see e.g. [2]. A multi-Bernoulli (MB) RFS \mathbf{X} is the union of a fixed number M of independent Bernoulli RFSs \mathbf{X}^j ,

$$\mathbf{X} = \bigcup_{j=1}^M \mathbf{X}^j \quad (7)$$

and is defined by the set of existence probabilities $\{w^j\}$ and distributions $\{p^j(\cdot)\}$. Here M is the maximum number of targets that the MB RFS can represent. A realization with $N \leq M$ targets has probability density

$$f(\{\mathbf{x}^1, \dots, \mathbf{x}^i, \dots, \mathbf{x}^N\}) = \sum_{\{i_j\}} \prod_{j:i_j=0} (1 - w^j) \prod_{j:i_j>0} w^j p^j(\mathbf{x}^{i_j}) \quad (8)$$

Here $\{i_j\}$ is a set of M indicators that specify which Bernoulli RFSs are empty ($i_j = 0$) and which are not ($i_j > 0$), where $i_{j_1} \neq 0, i_{j_2} \neq 0 \Rightarrow i_{j_1} \neq i_{j_2}$ for $j_1 \neq j_2$. The summation is over all indicator sets such that there are N targets. The MB probability density is abbreviated as

$$f(\mathbf{X}) = \{(w^j, p^j)\}_{j=1}^M \quad (9)$$

The MB cardinality distribution follows from (8),

$$P_{\mathbf{X}}(N) = \sum_{\{i_j\}} \prod_{j:i_j=0} (1 - w^j) \prod_{j:i_j>0} w^j \quad (10)$$

Analogous to mixtures of distributions, e.g. Gaussian mixtures, we can define MB mixtures [12], [22], [26], i.e. convex combinations of multi-Bernoulli distributions

$$f(\mathbf{X}) = \sum_{\ell} \mathcal{W}^{\ell} \prod_{\mathbf{x} \in \mathbf{X}} p^{\ell}(\mathbf{x}); \quad \sum_{\ell} \mathcal{W}^{\ell} = 1 \quad (11)$$

In a target tracking scenario the weights can, e.g., correspond to different existence and data association histories for the target estimates.

The measurement set is the union of a set of clutter measurements \mathbf{C}_k and sets of target generated measurements $\mathbf{W}_k(\mathbf{x}_k^i)$

$$\mathbf{Z}_k = \mathbf{C}_k \cup \left[\bigcup_{i=1}^{N_k^x} \mathbf{W}_k(\mathbf{x}_k^i) \right] \quad (12)$$

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

$$\kappa(\mathbf{C}_k) = e^{-\lambda_c} \prod_{i=1}^{N_k^c} \lambda_c g_c(\mathbf{z}_k^i) \quad (13)$$

where $\lambda_c g_c(\mathbf{z}_k^j)$ is the Poisson process intensity.

This paper is restricted to consideration of so called point targets, meaning that the i th target measurement set $\mathbf{W}_k(\mathbf{x}_k^i)$ is a Bernoulli RFS that 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 [11]

$$f(\mathbf{Z}_k|\mathbf{X}_k) = e^{-\lambda_c} \left[\prod_{j=1}^{N_k^c} \lambda_c g_c(\mathbf{z}_k^j) \right] \left[\prod_{i=1}^{N_k^x} (1 - p_D(\mathbf{x}_k^i)) \right] \times \sum_{\theta \in \Theta} \prod_{i:\sigma_i>0} \frac{p_D(\mathbf{x}_k^i)}{1 - p_D(\mathbf{x}_k^i)} \frac{g_x(\mathbf{z}_k^{\sigma_i}|\mathbf{x}_k^i)}{\lambda_c g_c(\mathbf{z}_k^{\sigma_i})} \quad (14a)$$

$$= \sum_{\theta \in \Theta} e^{-\lambda_c} \left[\prod_{j:\# \sigma_i=j} \lambda_c g_c(\mathbf{z}_k^j) \right] \left[\prod_{i:\sigma_i=0} (1 - p_D(\mathbf{x}_k^i)) \right] \times \left[\prod_{i:\sigma_i>0} p_D(\mathbf{x}_k^i) g_x(\mathbf{z}_k^{\sigma_i}|\mathbf{x}_k^i) \right] \quad (14b)$$

Here $\theta = \{\sigma_i\}$, defined as in [11], 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 .

V. MB FILTER WITH MULTI-OBJECT PARTICLE UPDATE

In this section we present the proposed multi-object particle probability hypothesis density (MOP-MB) filter. We assume Gaussian single target densities, Gaussian transition density, and Gaussian measurement pdf,

$$p_{k|k}(\mathbf{x}_k|\mathbf{Z}^k) = \mathcal{N}(\mathbf{x}_k; m_{k|k}, P_{k|k}) \quad (15a)$$

$$p_{k,k-1}(\mathbf{x}_k|\mathbf{x}_{k-1}) = \mathcal{N}(\mathbf{x}_k; f_{k,k-1}(\mathbf{x}_{k-1}), Q_k) \quad (15b)$$

$$g_x(\mathbf{z}_k|\mathbf{x}_k) = \mathcal{N}(\mathbf{z}_k; h_k(\mathbf{x}_k), R_k) \quad (15c)$$

The birth process is assumed unknown and an adaptive birth process is used, for some previous work see, e.g., [18], [19]. 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 region.

For the probability of detection and probability of survival the following is assumed to hold

$$p_D(\mathbf{x}) \mathcal{N}(\mathbf{x}; m, P) \approx p_D(m) \quad (16)$$

$$p_S(\mathbf{x}) \mathcal{N}(\mathbf{x}; m, P) \approx p_S(m) \quad (17)$$

The above is trivially exact when $p_D(\cdot)$ and $p_S(\cdot)$ are constants. Further, it is true if $p_D(\cdot)$ and $p_S(\cdot)$ are sufficiently smooth functions in the uncertainty area defined by $\mathcal{N}(\mathbf{x}; m, P)$. If this were to not be true, a non-linear approach, e.g. a particle filter, could be utilized.

A. Update

The update is an adaptation of the PHD update in [10] and has three main steps:

- 1) The predicted MB density is used to create a particle approximation of the predicted multi-object density $f(\mathbf{X}_k|\mathbf{Z}^{k-1})$.
- 2) Update each multi-object particle using the multi-object update (5), using an approximation for the data association computed using the auction algorithm.³
- 3) Use the posterior multi-object particles to approximate the posterior MB density.

To mitigate computational cost, measurement gating should be used to define groups of estimates and measurements, such that given the gating decisions the groups are statistically independent. The update is then performed for each group. The details of the proposed filter are below.

1) *Approximate predicted multi-object density:* Given a predicted Gaussian MB density

$$\left\{ \left(w_{k|k-1}^j, m_{k|k-1}^j, P_{k|k-1}^j \right) \right\}_j \quad (18)$$

we approximate the predicted multi-object density by M multi-object particles $\mathbf{X}_{k|k-1}^\ell$

$$f(\mathbf{X}_k|\mathbf{Z}^{k-1}) \approx \sum_{\ell=1}^M \mathcal{W}_{k|k-1}^\ell \phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^\ell) \quad (19)$$

where, for the ℓ th particle, $\phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^\ell) = 0$ if $|\mathbf{X}_k| \neq |\mathbf{X}_{k|k-1}^\ell|$ and

$$\phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^\ell) = \prod_{i \in I_\ell} \mathcal{N}(\mathbf{x}_i; m_{k|k-1}^i, P_{k|k-1}^i) \quad (20)$$

if $|\mathbf{X}_k| = |\mathbf{X}_{k|k-1}^\ell|$. Expressed in words, the meaning of the set I_ℓ is that in the ℓ th particle the j th predicted Gaussian component is included with probability $w_{k|k-1}^j$. For each particle, the cardinality is the cardinality of the set I_ℓ , meaning that within each multi-object particle the included Gaussians are interpreted as representing targets that do exist, i.e. target existence is here represented by the random sampling u_ℓ^i .

Comparing to (11) we see that (19) is an MB mixture where the probabilities of existence are either zero or one. Note that, strictly speaking, (19) is a Gaussian sum approximation and not a particle approximation, because (20) defines a Gaussian distribution. However, for the sake of brevity and simplicity, in the remainder of the paper we will follow [10] and use the terminology ‘‘multi-object particle.’’

The sets I_ℓ that define the particles $\mathbf{X}_{k|k-1}^\ell$ are constructed using random sampling. Specifically, the set I_ℓ is defined as $I_\ell = \{i | u_\ell^i \leq w_{k|k-1}^i\}$, where u_ℓ^i are randomly sampled from the uniform distribution $\mathcal{U}(0, 1)$. The MOPs have equal weights, $\mathcal{W}_{k|k-1}^\ell = M^{-1}$, $\forall \ell$. When sampling the M particles are not necessarily unique, in which case computations should be saved by only considering only unique particles.

2) *Update:* Given the particle approximation of the predicted multi-object density (19), the posterior multi-object density is given by the Bayes update (5),

$$f(\mathbf{X}_k|\mathbf{Z}^k) = \frac{\sum_{\ell=1}^M f(\mathbf{Z}_k|\mathbf{X}_k) \phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^\ell)}{\sum_{\ell=1}^M \int f(\mathbf{Z}_k|\mathbf{X}_k) \phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^\ell) \delta \mathbf{X}_k} \quad (21)$$

³JPDA association probabilities [2], cheap JPDA [7] or Murty’s algorithm [15] can also be used.

Using the measurement set pdf (14) and the Kalman filter update, or one of its non-linear variants, for each multi-object particle we have

$$f(\mathbf{Z}_k|\mathbf{X}_k) \phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^\ell) = 0 \quad (22a)$$

if $|\mathbf{X}_k| \neq |\mathbf{X}_{k|k-1}^\ell|$, and when $|\mathbf{X}_k| = |\mathbf{X}_{k|k-1}^\ell|$ we have

$$\begin{aligned} & f(\mathbf{Z}_k|\mathbf{X}_k) \phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^\ell) \\ &= \sum_{\theta \in \Theta} e^{-\lambda_c} \left[\prod_{j: \# \sigma_i = j} \frac{\lambda_c}{V} \right] \left[\prod_{i \in I_\ell: \sigma_i = 0} \left(1 - p_D(m_{k|k-1}^i) \right) \right] \\ & \quad \times \left[\prod_{i \in I_\ell: \sigma_i > 0} p_D(m_{k|k-1}^i) \mathcal{N}(\mathbf{z}_k^{\sigma_i}; m_{k|k-1}^i, P_{k|k-1}^i) \right] \\ & \quad \times \left[\prod_{i \in I_\ell} \mathcal{N}(\mathbf{x}_i; m_{k|k-1}^i, P_{k|k-1}^i) \right] \end{aligned} \quad (22b)$$

$$= \sum_{\theta \in \Theta} \mathcal{L}_{k|k-1}^{\ell, \theta} \prod_{i \in I_\ell} \mathcal{N}(\mathbf{x}_i; m_{k|k-1}^{i, \sigma_i}, P_{k|k-1}^{i, \sigma_i}) \quad (22c)$$

$$= \sum_{\theta \in \Theta} \mathcal{L}_{k|k-1}^{\ell, \theta} \phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^{\ell, \theta}) \quad (22d)$$

Note that $m_{k|k-1}^{i, \sigma_i} = m_{k|k-1}^i$ and $P_{k|k-1}^{i, \sigma_i} = P_{k|k-1}^i$ for $\sigma_i = 0$. The likelihoods of the detections, given the MOP and the association event, are

$$\begin{aligned} \mathcal{L}_{k|k-1}^{\ell, \theta} &= e^{-\lambda_c} \left(\frac{\lambda_c}{V} \right)^{N_{FA}(\theta)} \left[\prod_{i \in I_\ell: \sigma_i = 0} \left(1 - p_D(m_{k|k-1}^i) \right) \right] \\ & \quad \times \left[\prod_{i \in I_\ell: \sigma_i > 0} p_D(m_{k|k-1}^i) \right] \left[\prod_{i \in I_\ell: \sigma_i > 0} \mathcal{N}(\mathbf{z}_k^{\sigma_i}; \hat{\mathbf{z}}_k^i, S_k^i) \right] \end{aligned} \quad (23)$$

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

$$S_k^i = H_k P_{k|k-1}^i H_k^T + R_k \quad (24a)$$

$$K_k^i = P_{k|k-1}^i H_k (S_k^i)^{-1} \quad (24b)$$

$$\hat{\mathbf{z}}_k^i = H_k m_{k|k-1}^i \quad (24c)$$

$$m_{k|k-1}^{i, \sigma_i} = m_{k|k-1}^i + K_k^i (\mathbf{z}_k^{\sigma_i} - \hat{\mathbf{z}}_k^i) \quad (24d)$$

$$P_{k|k-1}^{i, \sigma_i} = P_{k|k-1}^i - K_k^i S_k^i (K_k^i)^T \quad (24e)$$

where H_k is the gradient of $h_k(\mathbf{x})$ evaluated at $\mathbf{x} = m_{k|k-1}^i$.

Note that (22) 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 mitigate computational complexity data association is used. The JPDA algorithm was used in [9], however this solution is highly susceptible to track coalescence, and it also suffers from high computational complexity when there are many targets.

Instead, for each MOP, the auction algorithm [3] is used to compute a single most probable association event $\hat{\theta}(\ell)$. Under this approximation the Bayes normalization constant $f(\mathbf{Z}_k|\mathbf{Z}^{k-1})$ is zero for $|\mathbf{X}_k| \neq |\mathbf{X}_{k|k-1}^\ell|$, and for $|\mathbf{X}_k| =$

$|\mathbf{X}_{k|k-1}^\ell|$ it becomes

$$\int f(\mathbf{Z}_k|\mathbf{X}_k)\phi_{\mathbf{X}_k}(\mathbf{X}_{k|k-1}^\ell)\delta\mathbf{X}_k = \mathcal{L}_{k|k-1}^{\ell,\hat{\theta}(\ell)} \quad (25)$$

We thus have a multi-object particle approximation of the posterior multi-object density

$$f(\mathbf{X}_k|\mathbf{Z}^k) = \frac{\sum_{\ell=1}^M \mathcal{L}_{k|k-1}^{\ell,\hat{\theta}(\ell)} \phi_{\mathbf{X}_k}(\mathbf{X}_{k|k}^{\ell,\hat{\theta}(\ell)})}{\sum_{\ell=1}^M \mathcal{L}_{k|k-1}^{\ell,\hat{\theta}(\ell)}} \quad (26a)$$

$$\triangleq \sum_{\ell=1}^M \mathcal{W}_{k|k}^\ell \phi_{\mathbf{X}_k}(\mathbf{X}_{k|k}^{\ell,\hat{\theta}(\ell)}) \quad (26b)$$

3) *Approximate posterior multi-Bernoulli density*: For each gating group there will be multiple different MOPs created, each with a unique combination of estimates. A predicted estimate may be included in multiple MOPs, and it follows that there may be multiple updated estimates that correspond to the same predicted estimate. As noted in the discussion in Section II we want each estimate to be uni-modal; hence merging of distributional modes is required.

A posterior multi-Bernoulli density with uni-modal target estimates

$$\left\{ \left(w_{k|k}^j, m_{k|k}^j, P_{k|k}^j \right) \right\}_j \quad (27)$$

is obtained as follows,

$$w_{k|k}^j = \sum_{\ell:j \in I_\ell} \mathcal{W}_{k|k}^\ell \quad (28)$$

$$m_{k|k}^j = \frac{1}{w_{k|k}^j} \sum_{\ell:j \in I_\ell} \mathcal{W}_{k|k}^\ell m_{k|k}^{j,\hat{\sigma}_j(\ell)} \quad (29)$$

$$M_{k|k}^{j,\hat{\sigma}_j(\ell)} = \left(m_{k|k}^{j,\hat{\sigma}_j(\ell)} - m_{k|k}^j \right) \left(m_{k|k}^{j,\hat{\sigma}_j(\ell)} - m_{k|k}^j \right)^\top \quad (30)$$

$$P_{k|k}^j = \frac{1}{w_{k|k}^j} \sum_{\ell:j \in I_\ell} \mathcal{W}_{k|k}^\ell \left(P_{k|k}^{j,\hat{\sigma}_j(\ell)} + M_{k|k}^{j,\hat{\sigma}_j(\ell)} \right) \quad (31)$$

In other words, for each predicted estimate w, m, P an updated probability of existence is found by taking the sum of updated MOP weights for the MOPs that the estimate was included in. A single updated Gaussian density is found by merging the updated Gaussian densities from the MOPs that the estimate was included in. Note that this merging minimizes the Kullback-Leibler divergence between the Gaussian mixture and the single Gaussian (see [26]).

B. Prediction

The MB prediction used here follows the MB prediction used in, e.g., the LMB filter [18]. Given a posterior Gaussian MB density (27), the Gaussian MB density representing a surviving target is

$$\left\{ \left(w_{k+1|k}^j, m_{k+1|k}^j, P_{k+1|k}^j \right) \right\}_j \quad (32)$$

where

$$w_{k+1|k}^j = ps \left(m_{k|k}^j \right) w_{k|k}^j \quad (33)$$

$$m_{k+1|k}^j = f_{k+1,k} \left(m_{k|k}^j \right) \quad (34)$$

$$P_{k+1|k}^j = F_{k|k}^j P_{k|k}^j \left(F_{k|k}^j \right)^\top + Q_k \quad (35)$$

and $F_{k|k}^j$ is the gradient of $f_{k+1,k}(\mathbf{x})$ evaluated at $\mathbf{x} = m_{k|k}^j$. The Gaussian MB density representing new targets (due to birth, or spawning) is

$$\left\{ \left(w_{k+1}^b, \beta \left(\mathbf{z}_k^j \right), P_{k+1}^b \right) \right\}_{j:\bar{z}_{\hat{\sigma}_i(\ell)}=j, \forall i,\ell} \quad (36)$$

where the weights w_{k+1}^b and the covariances P_{k+1}^b are user defined parameters that are equal for all birth components. In other words, each measurement from the previous time step that was not associated to any estimate in any multi-object particle is taken as a potential new target. The function $\beta(\mathbf{z})$ returns a Gaussian mean vector where the position is given by the detection \mathbf{z} and the velocity is initiated by an all-zero vector. Note that this handles both new target birth and target spawning. In the simulation study presented below it was sufficient to set the weight equal for all birth targets, however the weight can be set adaptively, see e.g. [18], [19].

The predicted Gaussian MB density at time step $k+1$, cf (18), is the union of the surviving targets MB density (32) and the new targets MB density (36).

C. Pruning, confirmation and extraction

To reduce the complexity after the update, multi-Bernoulli components with a probability of existence $w_{k|k}^j$ lower than a threshold τ_P are pruned (i.e. removed). A target estimate is considered confirmed if the probability of existence is larger than a threshold τ_C for at least one time step. An estimate of the target set is taken as the set of estimates that have been confirmed and have probabilities of existence larger than a threshold τ_E .

VI. APPROXIMATE MULTI HYPOTHESIS APPROACH

In this section we apply the MOP MB toward a simple approach to the trajectory crossing problem that allows us to compute a multiple hypothesis output at a computational cost that is only marginally larger than the cost of a single hypothesis filter⁴. The approximate multi hypothesis procedure presented below works on the confirmed target estimates, i.e. it is a post-processing procedure. Combined with the MOP-MB filter it gives the approximate multi-hypothesis multi-Bernoulli (AMHMB) filter.

A. Probability of identity switch

If the true trajectories cross at some point we say that the corresponding estimates may switch their identities, where switching identity should be understood as meaning confusion as to, for example, which target went left and which went right.

An identity switch implies that the true states were (approximately) equal at some prior time, excluding that the states are (approximately) equal at the current time. In the example in Figure 1, up until time step 15 there is no confusion as to which target is located where, i.e. there is no reasonable identity switch. However, from time step 16 there is confusion as to which target went up and which went down.

⁴This section presents one nice application of the MOP MB filter. Note that we will show the MOP MB filter is a fine integrated target tracker taken all on its own.

The estimates at time step κ corresponding to two targets with states \mathbf{x}_κ^i and \mathbf{x}_κ^j are subject to a possible identity switch if there is at least one time step $k < \kappa$ for which the true states were

$$\mathbf{x}_k^i \cong \mathbf{x}_k^j \quad (37)$$

meaning, approximately, but not necessarily exactly, equal. In this case the sequence of measurements will not contain sufficient information to enable telling the two targets apart. Note that (approximate) equality of positions, on its own, is generally not sufficient to alert for a possible identity switch because if, e.g., the kinematics (velocity, turn rate, etc) are different enough the sequence of measurements will generally contain sufficient information to tell the two targets apart.

The difference between approximately equal and exactly equal is important, and the difference must be understood not only in terms of the true states but also in terms of the sensor's and the targets' characteristics. For example, if two targets are separated by 1 meter, and the measurement noise standard deviation is 10 meters, then there is relatively large chance of identity switch. However, if the measurement noise standard deviation is 0.1 meters, then there is relatively little chance of identity switch.

We denote by $i \stackrel{k}{\rightleftharpoons} j$ the event that at time step k the estimates of targets i and j underwent an identity switch. The probability of an identity switch is equal to the probability of approximately equal true states (37), and is here approximated by the probability of the states being exactly equal,

$$P(i \stackrel{k}{\rightleftharpoons} j) = P(\mathbf{x}_k^i \cong \mathbf{x}_k^j) \approx P(\mathbf{x}_k^i = \mathbf{x}_k^j) \quad (38)$$

The probabilities of identity switch between two targets will subsequently be used to construct multi-object hypotheses that include identity switches involving an arbitrary number of targets.

Denote the difference between the true states as

$$\mathbf{d}_k^{ij} = \mathbf{x}_k^i - \mathbf{x}_k^j \quad (39)$$

and consider the two hypotheses

$$H_0^{ij} : \mathbf{d}_k^{ij} = \mathbf{0} \quad \text{and} \quad H_1^{ij} : \mathbf{d}_k^{ij} \neq \mathbf{0} \quad (40)$$

Note that the probability of the null hypothesis $P(H_0^{ij})$ corresponds to the probability of identity switch (38). Let the estimates of \mathbf{x}_k^i and \mathbf{x}_k^j have respective expected values $m_{k|k}^i$ & $m_{k|k}^j$, and covariances $P_{k|k}^i$ & $P_{k|k}^j$; and cross-covariance $P_{k|k}^{ij}$. The estimate of the true difference (39) is

$$\hat{\mathbf{d}}_{k|k}^{ij} = m_{k|k}^i - m_{k|k}^j \quad (41)$$

The error in the difference between the target estimates

$$\tilde{\mathbf{d}}_{k|k}^{ij} = \mathbf{d}_k^{ij} - \hat{\mathbf{d}}_{k|k}^{ij} \quad (42)$$

is zero-mean and has covariance [2, Eqn. 9.3.3-3]

$$T_{k|k}^{ij} = P_{k|k}^i + P_{k|k}^j - P_{k|k}^{ij} - \left(P_{k|k}^{ij}\right)^\top \quad (43)$$

Most MTT algorithms assume that the targets are independent and subsequently do not estimate cross-covariances

$P_{k|k}^{ij}$. However, during a true trajectory crossing the target kinematics are equal,⁵ assuming equivalent motion models (15b). The longer the target kinematics are equal, the more significantly non-zero the crosscovariance will be. A recursion for $P_{k|k}^{ij}$, assuming Gaussian models, is given in [2]. In the remainder of this paper we will assume that the trajectory crossings are sufficiently short to allow the cross-covariances to be approximated by all-zero matrices. For some standard 2D motion models, e.g. constant velocity and constant acceleration, "sufficiently short" means less than 15 time steps.

Below we will outline two different approaches to estimating the probability of identity switch (38).

1) *Alternative 1:* Assuming the estimation errors to be Gaussian, we accept hypothesis H_0 if

$$\Delta_{k|k}^{ij} \triangleq \left(\hat{\mathbf{d}}_{k|k}^{ij}\right)^\top \left(T_{k|k}^{ij}\right)^{-1} \hat{\mathbf{d}}_{k|k}^{ij} \leq \Delta_\alpha \quad (44)$$

This test is called the "track-2-track-association test" [2, Sec. 9.2], and the threshold Δ_α is chosen such that

$$P(\Delta_{k|k}^{ij} > \Delta_\alpha | H_0) = \alpha \quad (45)$$

From the Gaussian assumption the threshold is the $1 - \alpha$ point of the chi-square distribution with n_x degrees of freedom [1],

$$\Delta_\alpha = \chi_{n_x}^2(1 - \alpha) \quad (46)$$

where $\chi_{n_x}^2(\cdot)$ denotes the inverse of the chi-squared cumulative distribution function (cdf). From this we derive the approximate probability of identity switch

$$P(i \stackrel{k}{\rightleftharpoons} j) = 1 - F_{n_x} \left(\Delta_{k|k}^{ij} \right) \quad (47)$$

where $F_n(\cdot)$ is the chi-squared cdf with n degrees of freedom.

2) *Alternative 2:* Let $P_{k|k-1}(H_0^{ij})$ be the predicted probability of H_0^{ij} at time step k given data up to time step $k - 1$. The updated probability follows from Bayes' theorem

$$P_{k|k}(H_0^{ij}) = \frac{p(\hat{\mathbf{d}}_{k|k}^{ij} | H_0^{ij}) P_{k|k-1}(H_0^{ij})}{p(\hat{\mathbf{d}}_{k|k}^{ij})} \quad (48a)$$

$$= \frac{p(\hat{\mathbf{d}}_{k|k}^{ij} | H_0^{ij}) P_{k|k-1}(H_0^{ij})}{\sum_{h=0}^1 p(\hat{\mathbf{d}}_{k|k}^{ij} | H_h^{ij}) P_{k|k-1}(H_h^{ij})} \quad (48b)$$

where $P_{k|k-1}(H_1^{ij}) = 1 - P_{k|k-1}(H_0^{ij})$. Here $\hat{\mathbf{d}}_{k|k}^{ij}$ is taken as a pseudo-measurement of \mathbf{d}_k^{ij} , with pseudo-measurement pdfs for the two hypotheses modeled as

$$p\left(\hat{\mathbf{d}}_{k|k}^{ij} | H_0^{ij}\right) = \mathcal{N}\left(\hat{\mathbf{d}}_{k|k}^{ij}; \mathbf{0}, T_{k|k}^{ij}\right) \quad (49a)$$

$$p\left(\hat{\mathbf{d}}_{k|k}^{ij} | H_1^{ij}\right) = \int_{\mathbf{d} \neq \mathbf{0}} \mathcal{N}\left(\hat{\mathbf{d}}_{k|k}^{ij}; \mathbf{d}_k^{ij}, T_{k|k}^{ij}\right) p(\mathbf{d}_k^{ij}) d\mathbf{d}_k^{ij} \\ \approx p\left(\hat{\mathbf{d}}_{k|k}^{ij}\right) \quad (49b)$$

It is assumed that if the true states are not equal (hypothesis H_1^{ij}) the true states are independently uniformly distributed in the state space. With knowledge of the sensor's surveillance area, the pdf $p(\mathbf{d}_k^{ij})$ can be derived under this assumption. For example, if the state space is one dimensional with minimum

⁵For at least a single time step.

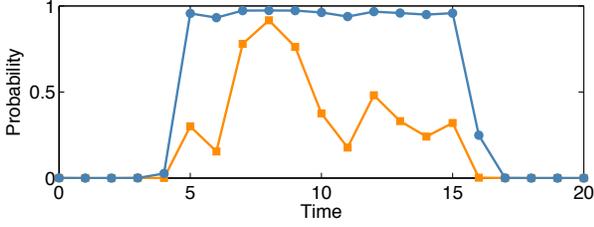


Fig. 2. Probabilities of trajectory crossing computed for the two crossing targets in Figure 1. Alternative 1 is in orange, alternative 2 in blue. From time step 5 to 15 two of the true targets are located in (approximately) the same position, and the corresponding estimates are thus susceptible to an identity switch. Alternative 2 estimates the probability of trajectory cross to just less than 1, while alternative 1 estimates a much lower probability of trajectory cross.

value d_1 and maximum value d_2 ($d_1 = 0$ and $d_2 = 30$ in Figure 1), then the state difference distribution is

$$p(\mathbf{d}) = \frac{1}{d_2 - d_1} - \frac{\text{sign}(\mathbf{d})\mathbf{d}}{(d_2 - d_1)^2} \quad (50)$$

The generalization to higher dimensions is straightforward.

Remark: The approximation in (49b) is accurate as long as the pdf $p(\mathbf{d}_k^{ij})$ is a sufficiently smooth function inside the uncertainty region defined by $T_{k|k}^{ij}$ (the 99% probability region). For the example in (50), this means that the uncertainty region is sufficiently smaller than $2(d_2 - d_1)$. \square

Remark: Note that because the target estimates at time steps k and $k - 1$ are dependent, it follows that the state difference pseudo-measurements are also dependent. \square

The probability of identity switch is approximated as

$$P(i \stackrel{k}{=} j) = P_{k|k}(H_0^{ij}) \quad (51)$$

The predicted probability $P_{k+1|k}(H_0^{ij})$ is computed as

$$P_{k+1|k}(H_0^{ij}) = P_{00}P_{k|k}(H_0^{ij}) + P_{10}P_{k|k}(H_1^{ij}) \quad (52)$$

where P_{00} and P_{10} are the probabilities that two targets with equal/unequal states will maneuver such that their states remain/become equal. $P_{00} = P_{10} = 0.5$ corresponds to there being no memory left from previous time steps.

3) *Comparison:* Probabilities $P(i \stackrel{k}{=} j)$ corresponding to the example in Figure 1 are shown in Figure 2. We see that alternative 2 (blue) gives better results than alternative 1 (orange). The two alternatives have been tested extensively in simulations, and empirically it was found that alternative 2 gives better performance.

B. Pairwise switch hypotheses

A pairwise switch hypothesis $H = (I, S)$ is defined by an index set I indicating the two estimates in the pair, and a sequence S of consecutive time steps for which it holds $P(i \stackrel{s}{=} j) > 0, \forall s \in S$. Note that for a particular pair I of estimates there may be several switch hypotheses for different time sequences, in which case the intersection of the time sequences are the empty set.

C. Merged switch hypotheses

After computing all pairwise switch hypotheses, the hypotheses are merged to form switch hypotheses involving an arbitrary number of targets.

Two hypotheses $H_1 = (I_1, S_1)$ and $H_2 = (I_2, S_2)$ are merged if $I_1 \cap I_2 \neq \emptyset$ and $|S_1 \cap S_2| > S_{\min}$, i.e. the hypotheses involve the same targets and overlap in time by at least S_{\min} time steps. The merged hypothesis resulting from merging H_1 and H_2 is $H = (I_1 \cup I_2, S_1 \cup S_2)$. Hypotheses are merged until there are no two hypotheses that fulfill the merging criterion.

D. Probability of switch hypothesis

For a merged switch hypothesis, there are $N_I = |I|!$ possible switch outcomes J_1, \dots, J_{N_I} . A switch outcome J_ℓ is defined as a set of indices $j_\ell(i) \in J_\ell$, where $j_\ell(i)$ indicate the identity that estimate $i \in I$ switches to. The probability of a particular switch outcome $I \rightarrow J_\ell$ is

$$P(I \rightarrow J_\ell) = \frac{\max_{s \in S} \prod_{i \in I} P(i \stackrel{s}{=} j_\ell(i))}{\sum_{\ell=1}^{N_I} \max_{s \in S} \prod_{i \in I} P(i \stackrel{s}{=} j_\ell(i))} \quad (53)$$

For a hypothesis, a single switch time s for all switch outcomes is estimated as

$$s = \max_k \left\{ \arg \max_k \sum_{\ell=1}^{N_I} \prod_{i \in I} P(i \stackrel{k}{=} j_\ell(i)) \right\} \quad (54)$$

The outermost max-function ensures that if the maximum is attained at more than one time step, then the most recent time step is taken. Note that switch times can be defined individually for each outcome J_ℓ , however empirically we have found that defining a single switch time as in (54) is a better compromise between accuracy and complexity.

E. Overall scenario hypotheses

An overall scenario hypothesis \mathcal{H} is constructed by ordering the merged switch hypotheses chronologically, and enumerating all possible combinations of the merged hypotheses. In total there are $\prod_{I \in \mathcal{H}} |I|!$ possible overall hypotheses, and the probability of an overall scenario hypothesis is

$$P(\mathcal{H}) = \frac{\prod_I P(I \rightarrow J_{\ell(h)})}{\sum_{\mathcal{H}} \prod_I P(I \rightarrow J_{\ell(h)})} \quad (55)$$

The forming of trajectories among the overall hypotheses is illustrated with an example. Let there be three confirmed targets $\hat{\mathbf{x}}_{k|k}^i, \hat{\mathbf{x}}_{k|k}^j$, and $\hat{\mathbf{x}}_{k|k}^\ell$ with initiation times t_b^i, t_b^j, t_b^ℓ , and pruning times t_d^i, t_d^j, t_d^ℓ , respectively. Consider an overall hypothesis \mathcal{H} with two switches $i \stackrel{s_1}{=} j$ and $i \stackrel{s_2}{=} \ell$, $s_1 < s_2$. The trajectories in this hypothesis are formed as

$$\hat{\mathbf{x}}_{k|k}^{i, \mathcal{H}} = \begin{cases} \hat{\mathbf{x}}_{k|k}^i & t_b^i \leq k \leq s_1 \\ \hat{\mathbf{x}}_{k|k}^j & s_1 < k \leq t_d^j \end{cases} \quad (56)$$

$$\hat{\mathbf{x}}_{k|k}^{j, \mathcal{H}} = \begin{cases} \hat{\mathbf{x}}_{k|k}^j & t_b^j \leq k \leq s_1 \\ \hat{\mathbf{x}}_{k|k}^i & s_1 < k \leq s_2 \\ \hat{\mathbf{x}}_{k|k}^\ell & s_2 < k \leq t_d^\ell \end{cases} \quad (57)$$

$$\hat{\mathbf{x}}_{k|k}^{\ell, \mathcal{H}} = \begin{cases} \hat{\mathbf{x}}_{k|k}^\ell & t_b^\ell \leq k \leq s_2 \\ \hat{\mathbf{x}}_{k|k}^i & s_2 < k \leq t_d^i \end{cases} \quad (58)$$

F. Complexity

The approximate multi hypothesis approach requires probabilities of identity switch for each possible pair of confirmed estimates, which has complexity of order

$$\mathcal{O}(N_k^x(N_k^x - 1)) \quad (59)$$

VII. SIMULATION RESULTS

The MOP-MB filter and AMHMB filter are tested in different scenarios. The first has many targets and confirms that the MOP-MB filter can estimate the trajectories. This scenario is mainly included as a “proof-of-concept”: it is intended to show that the MOP-MB can handle many targets and high clutter rate.

As noted in e.g. [26], an increasing number of targets is not necessarily the most challenging scenario, since a large number of targets spread out in the surveillance area can be treated like a number of single target tracking problems. More challenging is to handle multiple targets that have identical states. The MOP-MB filter is tested in scenarios where multiple targets have identical states at the mid point in time. For these scenarios we compare the performances of the MOP-MB and LMB filters. A comparison between an early version of the MOP-MB filter and the PHD filter was included in [9]; the MOP-MB was shown to give superior performance, specifically the cardinality error is much lower. For a comparison between the LMB and GLMB filter please refer to [18].

Lastly, the AMHMB filter is tested in scenarios where there are multiple trajectory crossings at a number of different times. Here a trajectory crossing means that the target states are (approximately) equal at a point in time, i.e. the true target states “coalesce” at a point in time.

In all simulations a nearly constant velocity motion model (white noise acceleration [1]) was used with acceleration noise standard deviation $\sigma_a = 2 \text{ m/s}^2$, and a linear measurement model was simulated. In the update, for each gate group the maximum number of MOPs was set to 8192. However, in practice the number of MOPs is typically much smaller. The probability of survival was $p_S = 0.99$ and for the data association gating probability $P_G = 0.999$ was used. For pruning, confirmation and extraction the following thresholds were used: $\tau_P = 10^{-5}$, $\tau_C = 0.75$, and $\tau_E = 0.25$. Alternative 2 was used to estimate the switch probabilities, with $P_{00} = 0.90$ and $P_{10} = 0.10$ and $S_{\min} = 5$.

For performance evaluation we compare the estimated cardinality and the optimal subpattern assignment metric (OSPA) [20]. The OSPA is implemented using the Euclidean norm with cut-off parameter $c = 300$ and $p = 1$. For the cardinality, we show the number of extracted targets; the cardinality error is the number of extracted targets minus the true cardinality.

A. High target number

A scenario was generated with a total of 150 targets: 110 appear at the first time step, 20 appear at time step 20, and the last 20 appear at time step 40. The time of disappearance was either randomly sampled from the set $\{50, 60, \dots, 90, 100\}$, or set to the time when the target left the surveillance area $[-1000, 1000] \times [-1000, 1000]$. The initial

positions were uniformly sampled in the surveillance area. Measurements were simulated with probability of detection 0.75, measurement noise $R_k = 10^2 \mathbf{I}$ and clutter measurements with Poisson rate 100 per time step, distributed uniformly in the surveillance area. The scenarios were simulated 500 times, the mean OSPA and mean cardinality is given in Figure 3. The results show that there is some convergence time following target appearance/disappearance, but eventually the MOP-MB filter gives a correct estimate.

B. Target coalescence

Multiple target trajectories, each 100 time steps long, were generated by sampling the state at time step 50 from a Gaussian distribution with mean \mathbf{x}_{50} and covariance $\mathbf{I} \times 10^{-6}$, and then using forward and backward simulation to generate the true trajectories. The mean position was set as the origin, with Cartesian velocity $[20 \ 0]^T$. The scenario was simulated with $N = 4, 6$, and 8 targets, and for each target number two probabilities of detection were tested: 0.75 and 0.99. As in the simulation study in [26], the measurement noise covariance was set to $R_k = 1^2 \mathbf{I}$.

The results of 100 Monte Carlo simulations are shown in Figure 4, where we compare the performances of the MOP-MB filter and the LMB filter. The plots show time steps 25 to 75, to highlight the results when the targets are very close. For the MOP-MB filter we see that for the OSPA there is an increase, and for the cardinality there is a decrease, both around time step 50, where the targets are extremely close. The MOP-MB filter errors are larger when the target number is higher, which is expected. The higher the target number, the more possible association events there are. For spatially separated targets a single association hypothesis is often enough to obtain accurate results. In a highly ambiguous situation such as this, as the target number increases, only considering a single association event is a progressively worse approximation, leading to a worse MB approximation. As expected, *ceteris paribus*, the MOP-MB filter results are better for higher p_D .

By comparison, for the LMB filter we see that the cardinality error is low during the entire simulation scenario. The OSPA error is lower when the targets are still approaching each other (up until time step 50), due to the fact that the LMB considers more association hypotheses than the MOP-MB does. However, after the targets separate, the LMB filter’s estimated target densities become multi-modal, causing the target output to “flicker” between the multiple modes, which in turn causes the OSPA to grow significantly. This “flickering” behaviour was mentioned in Section II-A as one of the motivations behind constructing a target tracker that estimates uni-modal single target densities.

C. Multi-Hypothesis Results

The AMHMB filter has been tested in several scenarios with multiple targets and multiple trajectory crossings. Here we present results from two selected scenarios in which the true trajectories are generated such that the results are clear when visualized. Both scenarios were simulated with

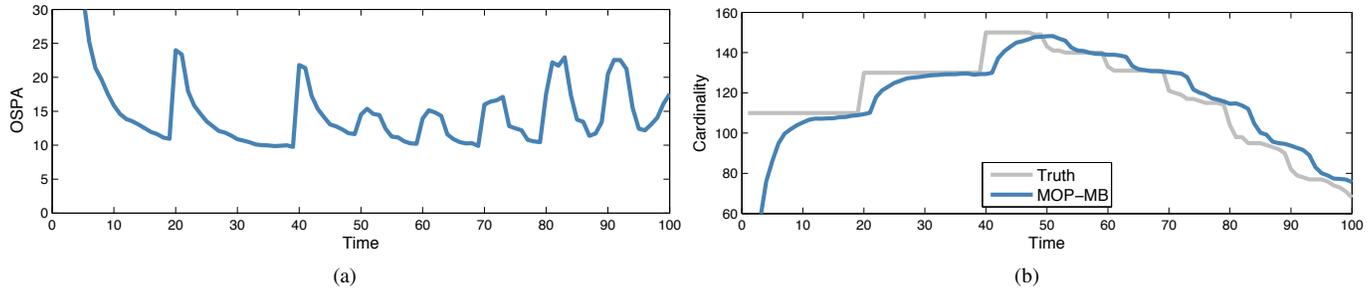


Fig. 3. Scenario with 150 targets. (a): OSPA (b): Cardinality, estimated by the number of confirmed tracks.

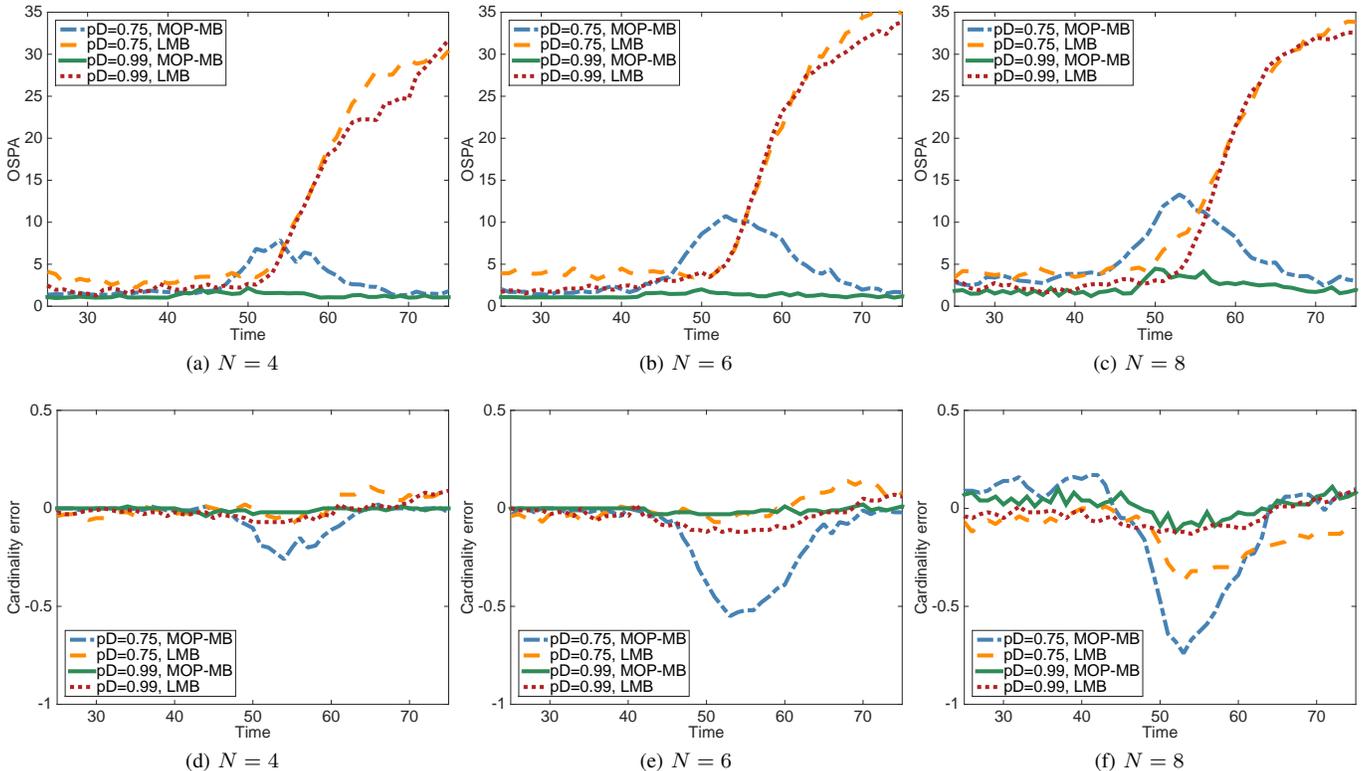


Fig. 4. Coalescence scenario with N targets. (a), (b), (c): OSPA; (d), (e), (f): Cardinality error. The N targets are closest at time step 50. The MOP-MB filter has larger cardinality error than the LMB filter when the targets are close, i.e. around time step 50. After the targets separate the LMB filter has larger OSPA error, because the output of the LMB filter “flickers” between the different single target modes.

probability of detection 0.75 and mean clutter number 200, uniformly distributed in the surveillance area $[-1500, 1500] \times [-1500, 1500]$. The measurement noise had $R_k = 1^2 \mathbf{I}$.

1) *Two targets*: Two targets have trajectories that cross at time steps 37 and 109, where each crossing lasts for a single time step (i.e. equal states for a single time step). Each crossing has 2 possible switching outcomes, and there are 4 different scenario outcomes. The true tracks and the results are shown in Figure 5. The AMHMB filter finds the trajectory crossings and correctly reconstructs all four multi-object hypotheses. The identity switches have estimated switch times 38 and 107, and estimated probabilities,

$$P(1 \stackrel{38}{=} 2) = 0.999997; \quad P(1 \stackrel{107}{=} 2) = 0.999998 \quad (60)$$

The resulting four multi-object hypotheses are equiprobable (0.25) to the 5th decimal.

2) *Eight targets*: Eight targets and three trajectory cross at time steps 75, 100 and 125, involving 2, 3, and 2 targets, respectively. For this scenario there are 24 different scenario outcomes. The true tracks are shown in Figure 6. To conserve space we do not show all 24 multi-object hypotheses, instead we show the different hypotheses for each of the eight targets, see Figure 7. The AMHMB filter correctly estimates the multiple trajectory hypotheses for the eight estimates. The identity switches have estimated switch times 74, 105 and 130, and the resulting 24 multi-object hypotheses are equiprobable (1/24) to the 5th decimal.

VIII. CONCLUSIONS AND FUTURE WORK

In this paper we presented a very simple implementation of a multi-object multi-Bernoulli filter, called the MOP-MB filter. The MOP-MB filter is similar to the LMB filter [18], however

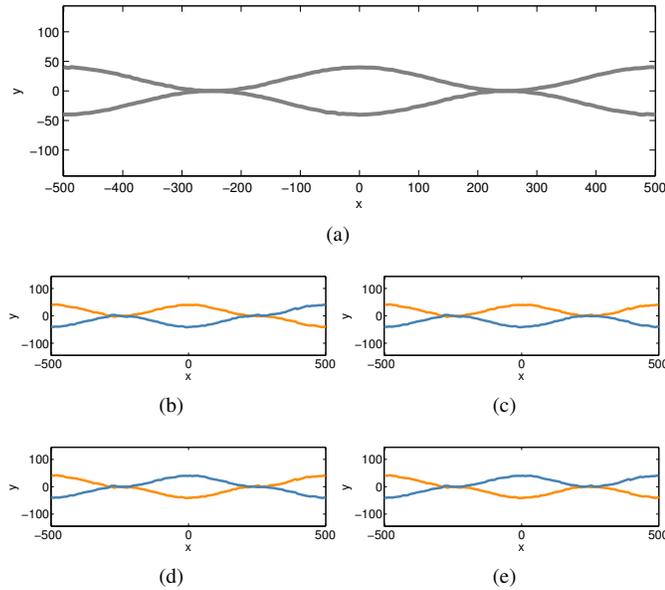


Fig. 5. Scenario with two targets. (a): true tracks. The trajectories cross twice, around $(x, y) = (-250, 0)$ and $(x, y) = (250, 0)$. (b)–(e): AMHMB output, the four hypotheses are equiprobable to the seventh decimal. (b) shows the result from the MOP-MB filter, the other three hypotheses are constructed using the approximate multi-hypothesis method.

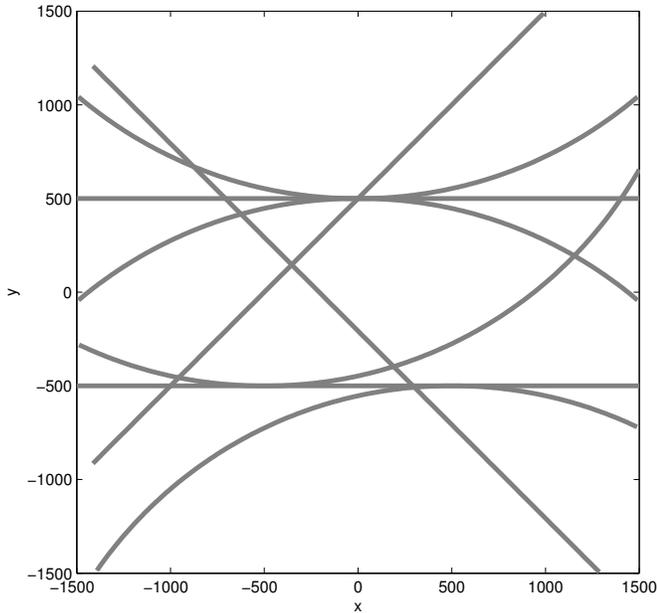


Fig. 6. True tracks for scenario with eight targets. Around $(x, y) = (0, 500)$ four targets have equal positions, but only three of them have equal velocity vectors. Around $(x, y) = (500, -500)$ and $(x, y) = (-500, -500)$ two targets have equal states. Around $(x, y) = (-350, 150)$ two targets have equal positions, but perpendicular velocity vectors. There are $3! \cdot 2! \cdot 2! = 24$ identity switch combinations.

it is also slightly simpler than the LMB filter. Additionally, we presented a postprocessing method that can reconstruct the multi-hypothesis picture of the tracking scenario using the output from a single hypothesis filter. Coupled with the MOP-MB filter this leads to the AMHMB filter.

The use of the postprocessing together with a very simple single hypothesis filter is deliberate. It shows that it is suffi-

cient to have a most simple single hypothesis filter, and then use the postprocessing technique on the output, allowing the multi-hypothesis output to be obtained at lower computational cost than the cost of a full multi-hypothesis approach such as the MHT [4] or the GLMB filter [22].

The results from the simulation study showed that the AMHMB filter can correctly approximate the multi-hypothesis multi-object output, however the MOP-MB filter's estimation error increases in scenarios where very many targets have equal states. The reason is the approximation of the full set of association events with a single association event compute using the auction algorithm. A topic for future work is to investigate the relationship between the MOP-MB's approximate posterior density and the actual posterior density.

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 W911NF-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] D. Bertsekas, "The auction algorithm: A distributed relaxation method for the assignment problem," *Annals of Operations Research*, vol. 14, no. 1, pp. 105–123, 1988.
- [4] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA, USA: Artech House, 1999.
- [5] H. Blom and E. Bloem, "Probabilistic data association avoiding track coalescence," *IEEE Transactions on Automatic Control*, vol. 45, no. 2, pp. 247–259, Feb. 2000.
- [6] D. Crouse, W. P., and Y. Bar-Shalom, "Developing a real-time track display that operators do not hate," *IEEE Transactions on Signal Processing*, vol. 59, no. 7, pp. 3441–3447, Jul. 2011.
- [7] R. Fitzgerald, "Development of practical PDA logic for multitarget tracking by microprocessor," in *Multitarget-Multisensor Tracking: Advanced Applications*, Y. Bar-Shalom, Ed. Artech House, 1990, pp. 1–23.
- [8] 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.
- [9] K. Granström, P. Willett, and Y. Bar-Shalom, "PHD filter with approximate multiobject density measurement update," in *Proceedings of the International Conference on Information Fusion*, Washington, DC, USA, Jul. 2015, pp. 1802–1809.
- [10] 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.
- [11] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2007.
- [12] —, *Advances in Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2014.
- [13] —, "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.
- [14] —, "PHD filters of higher order in target number," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 4, pp. 1523–1543, Oct. 2007.
- [15] K. Murty, "An algorithm for ranking all the assignments in order of increasing cost," *Operations Research*, vol. 16, no. 3, pp. 682–687, 1968.
- [16] 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.

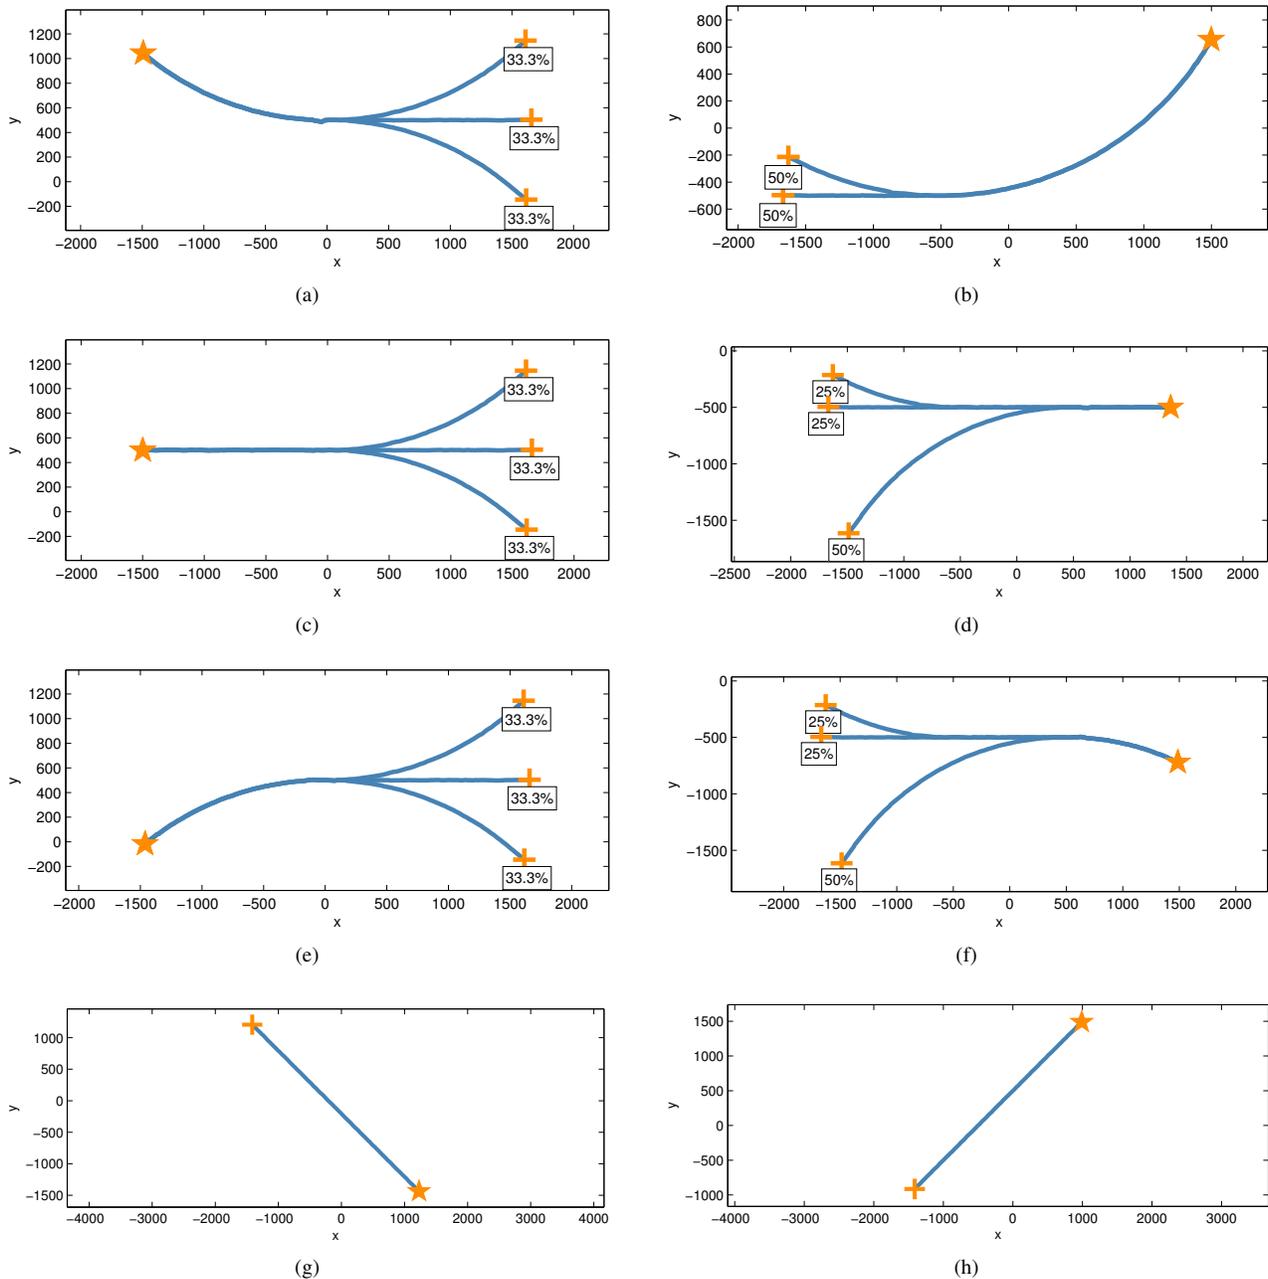


Fig. 7. Results from scenario with eight targets, true tracks in Figure 6. (a)–(h): AMH-MB output for each of the eight estimated targets. Stars mark birth locations, crosses mark death locations. For estimated tracks subject to identity switch, the probability of each death location is printed.

- [17] V. C. Ravindran, L. Svensson, L. Hammarstrand, and M. Morelande, "A cardinality preserving multitarget multi-Bernoulli RFS tracker," in *Proceedings of the International Conference on Information Fusion*, Singapore, Jul. 2012, pp. 832–839.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] L. Svensson, D. Svensson, M. Guerriero, and P. Willett, "Set JPDA filter for multitarget tracking," *IEEE Transactions on Signal Processing*, vol. 59, no. 10, pp. 4677–4691, Oct. 2011.
- [22] 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.
- [23] B.-T. Vo, B.-N. Vo, and A. Cantoni, "The cardinality balanced multi-target multi-bernoulli filter and its implementations," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 409–423, Feb. 2009.
- [24] 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.
- [25] B.-T. Vo, B.-N. Vo, and D. Phung, "Labeled random finite sets and the Bayes multi-target tracking filter," *IEEE Transactions on Signal Processing*, vol. 62, no. 24, pp. 6554–6567, Dec. 2014.
- [26] J. Williams, "An efficient, variational approximation of the best fitting multi-Bernoulli filter," *IEEE Transactions on Signal Processing*, vol. 63, no. 1, pp. 258–273, Jan. 2015.



Karl Granström (M'08) is a postdoctoral research fellow at the Department of Signals and Systems, Chalmers University of Technology, Gothenburg, Sweden. He received the MSc degree in Applied Physics and Electrical Engineering in May 2008, and the PhD degree in Automatic Control in November 2012, both from Linköping University, Sweden. He previously held postdoctoral positions at the Department of Electrical and Computer Engineering at University of Connecticut, USA, from September 2014 to August 2015, and at the Department of Electrical Engineering of Linköping University from December 2012 to August 2014. His research interests include estimation theory, multiple model estimation, sensor fusion and target tracking, especially for extended targets. He has received paper awards at the FUSION 2011 and FUSION 2012 conferences.



Peter Willett (F'03) received his B.A.Sc. (engineering science) from the University of Toronto in 1982, and his Ph.D. degree from Princeton University in 1986. He has been a faculty member at the University of Connecticut ever since, and since 1998 has been a Professor. His primary areas of research have been statistical signal processing, detection, machine learning, data fusion and tracking. He also has interests in and has published in the areas of change/abnormality detection, optical pattern recognition, communications and industrial/security condition monitoring. He is editor-in-chief of IEEE Signal Processing Letters. He was editor-in-chief for IEEE TRANSACTIONS ON AEROSPACE AND ELECTRONIC SYSTEMS (2006-2011), and was Vice President for Publications for AESS (2012-2014). He is a member of the IEEE AESS Board of Governors 2003-2009, 2011 to present. He was General Co-Chair (with Stefano Coraluppi) for the 2006 ISIF/IEEE Fusion Conference in Florence, Italy and for the 2008 ISIF/IEEE Fusion Conference in Cologne, Germany, Program Co-Chair (with Eugene Santos) for the 2003 IEEE Conference on Systems, Man & Cybernetics in Washington DC, and Program Co-Chair (with Pramod Varshney) for the 1999 Fusion Conference in Sunnyvale.



Yaakov Bar-Shalom received the B.S. and M.S. degrees from the Technion, Israel Institute of Technology, in 1963 and 1967 and the Ph.D. degree from Princeton University in 1970, all in electrical engineering. Currently he is Board of Trustees Distinguished Professor in the Dept. of Electrical and Computer Engineering and Marianne E. Klewin Professor in Engineering at the University of Connecticut. He is also Director of the ESP (Estimation and Signal Processing) Lab. His current research interests are in estimation theory, target tracking and

data fusion. He has published over 500 papers and book chapters in these areas and in stochastic adaptive control as well as 8 books including Estimation with Applications to Tracking and Navigation (Wiley, 2001), Tracking and Data Fusion (YBS Publishing, 2011). He has been elected Fellow of IEEE for "contributions to the theory of stochastic systems and of multitarget tracking". He served as President of the International Society of Information Fusion in 2000 and 2002 and Vice President for Publications in 2004-13. He is corecipient of the M. Barry Carlton Award for the best paper in the IEEE Transactions on Aerospace and Electronic Systems in 1995 and 2000 and recipient of the J. Mignona Data Fusion Award from the DoD JDL Data Fusion Group. He is a member of the Connecticut Academy of Science and Engineering. In 2008 he was awarded the IEEE Dennis J. Picard Medal for Radar Technologies and Applications, and in 2012 the Connecticut Medal of Technology. He has been listed by academic.research.microsoft (top authors in engineering) as #1 among the researchers in Aerospace Engineering based on the citations of his work.