PREPRINT: Annual Conf. Inf. Fusion, Salamanca, Spain, July 2014

Multi-target Tracking with Merged Measurements Using Labelled Random Finite Sets

Michael Beard*†, Ba Tuong Vo†, and Ba-Ngu Vo†
*Maritime Division, Defence Science and Technology Organisation, HMAS Stirling, Rockingham, WA, Australia
†Department of Electrical and Computer Engineering, Curtin University, Bentley, WA, Australia
michael.beard@dsto.defence.gov.au, {ba-tuong.vo,ba-ngu.vo}@curtin.edu.au

***Abstract*—In real world multi-target tracking problems, the presence of merged measurements is a frequently occurring phenomenon, however, the vast majority of tracking algorithms in the literature assume that each target generates independent measurements. Allowing for the possibility of measurement merging increases the computational complexity of the multi-target tracking problem, and limited computing power has been a major factor in the dominance of algorithms that assume independent measurements. In the presence of merged measurements, these algorithms suffer from performance degradation, usually due to premature track termination. In this paper, we develop a principled Bayesian solution to this problem based on the theory of random finite sets (RFS), and a tractable implementation based on the recently proposed generalised labelled multi-Bernoulli (GLMB) filter. The performance of the proposed technique is demonstrated by simulation of a multi-target bearings-only tracking scenario, where measurements become merged due to finite resolution effects.**

***Index Terms*—Multi-target tracking, merged measurements, unresolved targets, random finite sets**

I. Introduction

Most traditional multi-target tracking algorithms assume that the sensor generates an independent measurement for each target. In some cases, this assumption is reasonable since it leads to computationally efficient algorithms that are scalable to large numbers of targets. However, most real world sensors violate this assumption when the target measurements draw close to each another. In these cases, sensors often generate fewer measurements than there are targets present. If the tracker assumes independent measurements, it will usually conclude that some targets have disappeared, when in fact their measurements have been merged. For example, passive sonar sensors can have beams of significant width, and multiple targets falling within the same beam may produce a merged measurement. Another example is in computer vision, where detection algorithms may produce merged measurements for objects that appear close together in an image.

It is fundamentally important to investigate algorithms that can deal with this problem, because real world sensors often produce merged observations, and trackers that assume independent measurements for tractability reasons perform poorly when this occurs. Traditional classes of multi-target tracking algorithms have been adapted to handle merged measurements, including Joint Probabilistic Data Association (JPDA) [1], [2], Multiple Model JPDA [3], [4], Integrated PDA [5], Multiple Hypothesis Tracking (MHT) [6], and Probabilistic MHT [7]. Although these algorithms are based on Bayesian techniques, their relationship to the full multi-object density has not been clearly established.

Recently, the concept of random finite sets (RFS) has received a great deal of attention, as it provides a principled framework for deriving algorithms for estimating the states for an unknown and time-varying number of objects, based on noisy measurements with false alarms and missed detections. One of the first algorithms to be derived using this framework was the Probability Hypothesis Density (PHD) filter, and in [8] Mahler proposed that the PHD filter could be applied to unresolved targets by modelling the multi-object state as an RFS of point clusters. This approach has some limitations, which shall be discussed in Section IV.

One critisism of the RFS framework has been that it yields algorithms that do not maintain target labels over time (performing multi-object filtering as opposed to tracking). It was shown in [9] that the RFS framework does admit target labelling, and a computationally feasible multi-target tracker for the standard sensor model was derived, known as the generalised labelled multi-Bernoulli (GLMB) filter.

In this paper, we generalise the GLMB filter to a sensor model that includes merged measurements, thereby making it suitable for application to a wider variety of real world problems. We have achieved this by deriving a multi-target likelihood function that takes into account the possible merging of target generated measurements, and a tractable solution to the multi-target posterior, by considering feasible partitions of the target set, and the feasible assignments of measurements to groups within these partitions. An advantage of our approach is that it can be parallelised to enable potential real-time implementation. To our knowledge, this is the first time that RFS principles have been used to derive and implement a full multi-object tracker that accommodates merged measurements.

The paper is organised as follows. Section II provides some background on labelled random finite sets and their application to tracking with standard sensor models. In Section III we formulate a sensor model that includes measurement merging. In Section IV we use labelled random finite sets to develop a filter for this sensor model, and Section V presents an approximation technique to aid in its tractable implementation. Section VI contains simulation results for a bearings-only multi-target tracking scenario, and finally, some concluding remarks are given in Section VII.

II. Background: RFS-based Tracking with Standard Sensor Models

Random finite sets provide a mathematically rigorous framework for developing Bayesian multi-object estimation algorithms. One of the first algorithms to be proposed based on this was the probability hypothesis density (PHD) filter, which achieves a tractable approximation to the full multi-object Bayes recursion by propagating only the first moment of the multi-object density. This was followed by the cardinalised PHD (CPHD) filter, which propagates the probability distribution of the number of targets, in addition to the first moment of the density. Another type of algorithm based on random finite sets involves approximating the density as a multi-Bernoulli RFS, known as multi-object multi-Bernoulli (MeMBer) filtering. A common feature of the PHD and multi-Bernoulli approaches is that they do not require explicit data association. However, a significant drawback is that they do not inherently produce target tracks, instead providing a set of unlabelled point estimates at each time step. A recently proposed technique for addressing this problem, which maintains the mathematical rigor of the RFS framework, is the concept of labelled random finite sets [9]. This technique involves assigning a distinct label to each element of the target set, so that the history of each object's trajectory can be naturally identified.

In this work, we are interested in algorithms that can produce continuous tracks, so we restrict our attention to methods based on labelled random finite sets. In [9], an algorithm was proposed for solving the standard multi-object tracking problem, based on a type of labelled RFS called 'generalised labelled multi-Bernoulli' (GLMB). We now review the main points of this technique, and in the next section we propose a generalisation which will enable it to handle merged measurements.

We begin by introducing some notation and definitions relating to labelled random finite sets. The multi-object exponential of a real valued function h raised to a set X is defined as

$$[h(\cdot)]^X = \prod_{x\in X} h(x) \tag{1}$$

where $h^{\emptyset} = 1$, and the elements of X may be of any type such as scalars, vectors, or sets, provided that the function $h(\cdot)$ takes an argument of that type. The generalised Kronecker delta function is defined as

$$\delta_Y(X) = \begin{cases} 1, & \text{if } Y = X \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

where again, X and Y may be of any type, such as scalars, vectors, or sets.

Definition 1. A labelled RFS $\boldsymbol{X}$ with state space $\mathbb{X}$ and discrete label space $\mathbb{L}$, is an RFS on $\mathbb{X}\times\mathbb{L}$, such that the labels within each realisation are always distinct. That is, if $\mathcal{L}_{\boldsymbol{X}}$ is the set of unique labels in $\boldsymbol{X}$, and we define the distinct label indicator function as

$$\Delta(\boldsymbol{X}) = \begin{cases} 1, & \text{if } |\mathcal{L}_{\boldsymbol{X}}| = |\boldsymbol{X}| \\ 0, & \text{if } |\mathcal{L}_{\boldsymbol{X}}| \neq |\boldsymbol{X}| \end{cases} \tag{3}$$

then a labelled RFS $\boldsymbol{X}$ always satisfies $\Delta(\boldsymbol{X}) = 1$.

Definition 2. A generalised labelled multi-Bernoulli (GLMB) RFS is a labelled RFS with state space $\mathbb{X}$ and discrete label space $\mathbb{L}$, which satisfies the probability distribution

$$\boldsymbol{\pi}(\boldsymbol{X}) = \Delta(\boldsymbol{X}) \sum_{c\in\mathbb{C}} w^{(c)}(\mathcal{L}_{\boldsymbol{X}}) \left[p^{(c)}(\cdot)\right]^{\boldsymbol{X}} \tag{4}$$

where $\mathbb{C}$ is an arbitrary index set, and $w^{(c)}(\mathcal{L}_{\boldsymbol{X}})$ and $p^{(c)}(x,l)$ satisfy

$$\sum_{c\in\mathbb{C}} w^{(c)}(L) = 1 \tag{5}$$

$$\int_{x\in\mathbb{X}} p^{(c)}(x,l)\,dx = 1. \tag{6}$$

A. Multi-object Transition Kernel

Let $\boldsymbol{X}$ be the labeled RFS of objects at the current time with label space $\mathbb{L}$. A particular object $(x,l)\in\boldsymbol{X}$ has probability $p_S(x,l)$ of surviving to the next time with state (x_+,l_+) and probability density $f(x_+|x,l)\,\delta_l(l_+)$, and probability $q_S(x,l) = 1 - p_S(x,l)$ of being terminated. Thus, the set $\boldsymbol{S}$ of surviving objects at the next time is distributed according to

$$\boldsymbol{f}_S(\boldsymbol{S}|\boldsymbol{X}) = \Delta(\boldsymbol{S})\,\Delta(\boldsymbol{X})\,1_{\mathcal{L}_{\boldsymbol{X}}}(\mathcal{L}_S)\,[\Phi(\boldsymbol{S};\cdot)]^{\boldsymbol{X}} \tag{7}$$

where

$$\Phi(\boldsymbol{S};x,l) = \sum_{(x_+,l_+)\in\boldsymbol{S}} \delta_l(l_+)\,p_S(x,l)\,f(x_+|x,l) + \left(1 - 1_{\mathcal{L}_{\boldsymbol{S}}}(l)\right) q_S(x,l). \tag{8}$$

where $f(x_+|x,l)$ is the single target transition kernel. Now let $\boldsymbol{B}$ be the labelled RFS of new born objects with label space $\mathbb{B}$, where $\mathbb{L}\cap\mathbb{B} = \emptyset$. Since the births must have distinct labels, and assuming that their states are independent, $\boldsymbol{B}$ is distributed according to

$$\boldsymbol{f}_B(\boldsymbol{B}) = \Delta(\boldsymbol{B})\,w(\mathcal{L}_{\boldsymbol{B}})\,[p_B(\cdot)]^{\boldsymbol{B}} \tag{9}$$

where $p_B(\cdot)$ is the single target birth density. The overall multi-object state at the next time step is the union of surviving and new born objects, i.e. $\boldsymbol{X}_+ = \boldsymbol{S}\cup\boldsymbol{B}$. The label spaces $\mathbb{L}$ and $\mathbb{B}$ are disjoint, and the states of new born objects are independent of surviving objects, hence $\boldsymbol{S}$ and $\boldsymbol{B}$ are independent. The multi-object transition kernel is thus defined by

$$\boldsymbol{f}(\boldsymbol{X}_+|\boldsymbol{X}) = \sum_{\boldsymbol{S}\subseteq\boldsymbol{X}_+} \boldsymbol{f}_S(\boldsymbol{S}|\boldsymbol{X})\,\boldsymbol{f}_B(\boldsymbol{X}_+ - \boldsymbol{S}). \tag{10}$$

It is straightforward to show that the summation in (10) contains only one non-zero term, which occurs when $\boldsymbol{S} = \boldsymbol{X}_+\cap(\mathbb{X}\times\mathbb{L})$. Hence the transition kernel reduces to

$$\boldsymbol{f}(\boldsymbol{X}_+|\boldsymbol{X}) = \boldsymbol{f}_S(X_+\cap(\mathbb{X}\times\mathbb{L})\,|\boldsymbol{X})\,\boldsymbol{f}_B(\boldsymbol{X}_+ - \mathbb{X}\times\mathbb{L}). \tag{11}$$

It was shown in [9] that a GLMB density of the form (4) is closed under the Chapman-Kolmogorov prediction equation with the transition kernel defined in (11).

B. Standard Multi-object Observation Model

Let $\boldsymbol{X}$ be the labelled RFS of objects that exist at the observation time. A particular object $(x,l) \in \boldsymbol{X}$ has probability $p_D(x,l)$ of generating a detection z with likelihood $g(z|x)$, and probability $q_D(x,l) = 1 - p_D(x,l)$ of being misdetected. Let $\boldsymbol{D}$ be the set of target detections. Assuming the elements of $\boldsymbol{D}$ are conditionally independent, $\boldsymbol{D}$ is a multi-Bernoulli RFS distributed according to

$$\pi_D(\boldsymbol{D}|\boldsymbol{X}) = \{(p_D(x,l), g(\cdot|x)); (x,l) \in \boldsymbol{X}\}(\boldsymbol{D}). \quad (12)$$

Let $\boldsymbol{K}$ be the set of clutter observations, which are independent of the target detections. We model $\boldsymbol{K}$ as a Poisson RFS with intensity $\kappa(\cdot)$, hence $\boldsymbol{K}$ is distributed according to

$$\pi_K(\boldsymbol{K}) = e^{-\langle \kappa, 1\rangle} \kappa^{\boldsymbol{K}}. \quad (13)$$

The overall multi-object observation is the union of target detections and clutter observations, i.e. $Z = \boldsymbol{D} \cup \boldsymbol{K}$. Since $\boldsymbol{D}$ and $\boldsymbol{K}$ are independent, the multi-object likelihood is defined by

$$g(Z|\boldsymbol{X}) = \sum_{\boldsymbol{D} \subseteq Z} \pi_D(\boldsymbol{D}|\boldsymbol{X}) \pi_K(Z - \boldsymbol{D}). \quad (14)$$

As demonstrated in [10], this can be equivalently expressed as

$$g(Z|\boldsymbol{X}) = e^{-\langle \kappa, 1\rangle} \kappa^Z \sum_{\theta \in \Theta} [\psi_Z(\cdot; \theta)]^{\boldsymbol{X}} \quad (15)$$

where Θ is the set of all one-to-one mappings of labels in $\boldsymbol{X}$ to measurement indices in Z,

$$\Theta = \{\theta : \mathcal{L}_{\boldsymbol{X}} \to \{0 : |Z|\}\} \quad (16)$$

such that $[\theta(i) = \theta(j) > 0] \Rightarrow [i = j]$, and $\psi_Z(\cdot; \theta)$ is given by

$$\psi_Z(x,l;\theta) = \begin{cases} \frac{p_D(x,l) g(z_{\theta(l)}|x,l)}{\kappa(z_{\theta(l)})}, & \theta(l) > 0 \\ q_D(x,l), & \theta(l) = 0 \end{cases} \quad (17)$$

It was demonstrated in [9] that a GLMB density of the form (4) is closed under the Bayes update with likelihood function defined by (15).

III. A Sensor Model for Merged Measurements

It is often the case that when targets appear close together in the measurement space, it may not be possible for a sensor to generate separate measurements for each target. Instead, a group of closely spaced targets may generate a single measurement consisting of contributions from all targets in the group. This is in contrast to the standard sensor model, which assumes that all targets generate measurements independently, regardless of their position relative to each other. In what follows, we describe the details of a sensor model in which merged measurements may occur.

Let $\mathbb{X}$ be the single-object state space, and $\mathbb{M}$ the sensor measurement space. At time index k, consider a single sensor with known control input U_k, and let $X_k = \{\boldsymbol{x}_{k,1}, \boldsymbol{x}_{k,2}, ..., \boldsymbol{x}_{k,M}\}$ represent the unknown multi-target state, in which $\boldsymbol{x}_k^i \in \mathbb{X}$ for $i = 1..M$. Let $h(\boldsymbol{x}_k^t, U_k)$ be the standard measurement function for target t, such that

$$h(\boldsymbol{x}_{k,t}, U_k) : \mathbb{X} \to \mathbb{M}. \quad (18)$$

Let $C = \{c_1, c_2, ...c_{N_C}\}$ be a set of disjoint cells covering $\mathbb{M}$,

$$\bigcup_{i=1}^{N_C} c_i = \mathbb{M} \quad (19)$$

$$c_i \cap c_j = \emptyset, \quad \forall i \neq j \quad (20)$$

and let $T_{k,i}$ be the set of targets whose true state falls in cell i at time k,

$$T_{k,i} = \{j : j \in 1..|X_k|, h(\boldsymbol{x}_{k,j}, U_k) \in c_i\}. \quad (21)$$

Then cell i produces the following measurement

$$z_{k,i} = \begin{cases} \frac{1}{|T_{k,i}|} \sum_{j \in T_i} h(\boldsymbol{x}_{k,j}, U_k), & |T_{k,i}| > 0 \\ \emptyset, & |T_{k,i}| = 0 \end{cases} \quad (22)$$

with probability $p_D(T_{k,i})$, and $z_{k,i} = \emptyset$ with probability $q_D(T_{k,i}) = 1 - p_D(T_{k,i})$. The set of target generated measurements at time k is thus

$$\Theta_k = \bigcup_{i=1}^{N_C} z_{k,i}. \quad (23)$$

In addition, at each time k, the sensor generates a Poisson RFS of false measurements denoted K_k, the intensity of which is known at all times. The overall set of measurements can thus be expressed as

$$Z_k = \Theta_k \cup K_k \quad (24)$$

IV. RFS-based Tracking with Merged Measurements

A. Multi-object Likelihood with Merged Measurements

In [10], Mahler proposed a technique to handle unresolved targets using the RFS framework. It was based on modelling the multi-target state as a set of point clusters, where each cluster has a location in the state space, and a number determining how many targets are effectively co-located at that point. The resulting likelihood function was defined in terms of a sum over partitions of the measurement set, making it computationally very demanding. A PHD filter based on this model was proposed in [8], however, a working implementation has not been developed.

The point cluster model is intuitively appealing in cases where measurements may only merge when the targets are close together in the state space. However, this implicitly assumes that merging depends on the target states only, and the sensor control input has no influence. There are cases in which this assumption is too restrictive, such as bearings-only tracking, and visual tracking. In these cases, targets that cannot be resolved along the line of sight may be separated by a considerable distance in the state space. The sensor's position clearly has an impact on whether the measurements become merged, but this cannot be accounted for using point clusters.

To handle these cases, we need a likelihood function that considers that the targets may give rise to merged measurements, not only when they are nearby in the state space, but more generally when they are nearby in the measurement space. To achieve this, we consider partitions of the target set, where each group of targets within a partition represents a merged group. This leads to the following form for the likelihood of observing a measurement set Z given a target set $\boldsymbol{X}$:

$$g(Z|\boldsymbol{X}) = \frac{1}{|\mathcal{P}(\boldsymbol{X})|} \sum_{\mathcal{U}\in\mathcal{P}(\boldsymbol{X})} \mathring{g}(Z|\mathcal{U}) \tag{25}$$

where $\mathcal{P}(\boldsymbol{X})$ is the set of all partitions of $\boldsymbol{X}$, and $\mathring{g}(Z|\mathcal{U})$ is the measurement likelihood conditioned upon the targets being observed according to partition $\mathcal{U}$. The latter is obtained by generalising the standard likelihood in (15) to target groups as follows

$$\mathring{g}(Z|\mathcal{U}) = e^{-\langle\kappa,1\rangle}\kappa^Z \sum_{\theta\in\Theta_{\mathcal{U}}} \left[\check{\psi}_Z(\cdot;\theta)\right]^{\mathcal{U}} \tag{26}$$

where $\Theta_{\mathcal{U}}$ is defined as the set of all one-to-one mappings of target groups in $\mathcal{U}$ to measurement indices in Z,

$$\Theta_{\mathcal{U}} = \begin{cases} \{\theta : \mathcal{U} \to \{0 : |Z|\}\} & \text{if } \mathcal{U} \in \mathcal{P}(\boldsymbol{X}) \\ \varnothing & \text{if } \mathcal{U} \notin \mathcal{P}(\boldsymbol{X}) \end{cases} \tag{27}$$

where $[\theta(\boldsymbol{X}) = \theta(\boldsymbol{Y}) > 0] \Rightarrow [\boldsymbol{X} = \boldsymbol{Y}]$, and $\check{\psi}_Z(\boldsymbol{Y};\theta)$ is a 'group likelihood' defined by

$$\check{\psi}_Z(\boldsymbol{Y};\theta) = \begin{cases} \frac{\check{p}_D(\boldsymbol{Y})\check{g}\left(z_{\theta(\boldsymbol{Y})}|\boldsymbol{Y}\right)}{\kappa\left(z_{\theta(\boldsymbol{Y})}\right)}, & \theta(\boldsymbol{Y}) > 0 \\ \check{q}_D(\boldsymbol{Y}), & \theta(\boldsymbol{Y}) = 0 \end{cases} \tag{28}$$

where $\check{p}_D(\boldsymbol{Y})$ is the detection probability for group $\boldsymbol{Y}$, $\check{q}_D(\boldsymbol{Y}) = 1 - \check{p}_D(\boldsymbol{Y})$ is the misdetection probability for group $\boldsymbol{Y}$, and $\check{g}\left(z_{\theta(\boldsymbol{Y})}|\boldsymbol{Y}\right)$ is the likelihood of measurement $z_{\theta(\boldsymbol{Y})}$ given group $\boldsymbol{Y}$. Note that in (26), the exponent $\mathcal{U}$ is a set of target sets, and the base is a real valued function whose argument is a target set. Substituting (26) into (25) yields the likelihood function

$$g(Z|\boldsymbol{X}) = \frac{e^{-\langle\kappa,1\rangle}\kappa^Z}{|\mathcal{P}(\boldsymbol{X})|} \sum_{\substack{\mathcal{U}\in\mathcal{P}(\boldsymbol{X})\\ \theta\in\Theta_{\mathcal{U}}}} \left[\check{\psi}_Z(\cdot;\theta)\right]^{\mathcal{U}}. \tag{29}$$

B. General Form for the Tracker

We now define a general form for the multi-object density, and demonstrate that it is closed under both the multi-object Champan-Kolmogorov prediction with transition kernel (11), and the Bayes update with likelihood function (29).

Definition 3. A labelled RFS mixture density on state space $\mathbb{X}$ and discrete label space $\mathbb{L}$, is a density of the form

$$\pi(\boldsymbol{X}) = \Delta(\boldsymbol{X}) \sum_{c\in\mathbb{C}} w^{(c)}(\mathcal{L}_{\boldsymbol{X}})\, p^{(c)}(\boldsymbol{X}) \tag{30}$$

where

$$\sum_{c\in\mathbb{C}} \sum_{L\subseteq\mathbb{L}} w^{(c)}(L) = 1, \tag{31}$$

$p^{(c)}(\boldsymbol{X})$ is symmetric in the elements of $\boldsymbol{X}$, and $p^{(c)}((\cdot,l_1),\dots(\cdot,l_n))$ is a joint pdf in $\mathbb{X}^n$.

Proposition 4. *If the multi-object prior is a labelled RFS mixture density with probability distribution of the form (30), then the predicted multi-object density under the transition kernel (11) is also a labelled RFS mixture density with probability distribution given by*

$$\pi_+(\boldsymbol{X}_+) = \Delta(\boldsymbol{X}_+) \sum_{c\in\mathbb{C}} \sum_{L\subseteq\mathbb{L}} w_{+,L}^{(c)}(\mathcal{L}_{\boldsymbol{X}_+})\, p_{+,L}^{(c)}(\boldsymbol{X}_+) \tag{32}$$

where

$$w_{+,L}^{(c)}(J) = 1_L(J\cap\mathbb{L})\, w_B(J-\mathbb{L})\, w^{(c)}(L)\, \eta_S^{(c)}(L) \tag{33}$$

$$p_{+,L}^{(c)}(Y) = [p_B(\cdot)]^{Y-\mathbb{X}\times\mathbb{L}}\, p_{S,L}^{(c)}(Y\cap\mathbb{X}\times\mathbb{L}) \tag{34}$$

$$p_{S,L}^{(c)}(Y) = \frac{\int p_L^{(c)}\left(x_1,\dots,x_{|L|}\right) \prod_{i=1}^{|L|} \Phi(Y;x_i,l_i)\, dx_1\dots dx_{|L|}}{\eta_S^{(c)}(L)} \tag{35}$$

$$\eta_S^{(c)}(L) = \int\int p_L^{(c)}\left(x_1,\dots,x_{|L|}\right) \prod_{i=1}^{|L|} \Phi(Y;x_i,l_i)\, dx_1\dots dx_{|L|}\delta Y \tag{36}$$

Proposition 5. *If the prior is a labelled RFS mixture density with probability distribution of the form (30), then the posterior multi-object density under the likelihood function (29) is also a labelled RFS mixture density with probability distribution given by*

$$\pi(\boldsymbol{X}|Z) = \Delta(\boldsymbol{X}) \sum_{c\in\mathbb{C}} \sum_{\substack{\mathcal{U}\in\mathcal{P}(\boldsymbol{X})\\ \theta\in\Theta_{\mathcal{U}}}} w_Z^{(c,\mathcal{U},\theta)}(\mathcal{L}_{\boldsymbol{X}})\, p^{(c,\mathcal{U},\theta)}(\boldsymbol{X}|Z) \tag{37}$$

where

$$w_Z^{(c,\mathcal{U},\theta)}(L) = \frac{\frac{1}{|\mathcal{P}(L)|} w^{(c)}(L)\, \eta_Z^{(c,\mathcal{U},\theta)}(L)}{\sum_{c\in\mathbb{C}} \sum_{J\subseteq\mathbb{L}} \sum_{\substack{\mathcal{U}\in\mathcal{P}(J)\\ \theta\in\Theta_{\mathcal{U}}}} \frac{1}{|\mathcal{P}(J)|} w^{(c)}(J)\, \eta_Z^{(c,\mathcal{U},\theta)}(J)} \tag{38}$$

$$p^{(c,\mathcal{U},\theta)}(\boldsymbol{X}|Z) = \frac{\left[\check{\psi}_Z(\cdot;\theta)\right]^{\mathcal{U}} p^{(c)}(\boldsymbol{X})}{\eta_Z^{(c,\mathcal{U},\theta)}(\mathcal{L}_{\boldsymbol{X}})} \tag{39}$$

$$\eta_Z^{(c,\mathcal{U},\theta)}(L) = \int \left[\check{\psi}_Z(\cdot;\theta)\right]^{\mathcal{U}} p^{(c)}\left(x_1,\dots x_{|L|}\right) dx_1\dots dx_{|L|} \tag{40}$$

Proofs of propositions 4 and 5 are omitted due to space limitations. Although an exact implementation may be possible, it would quickly become intractable because there is no clear way of truncating the sum over the space of measurement-to-group associations $\Theta_{\mathcal{U}}$. The reason is that each component in the overall density (30) consists of a joint density encompassing all targets. Thus there may be dependencies between target groups, and standard ranked assignment techniques cannot be applied. In the next section we propose an approximation that avoids this problem, allowing us to tractably implement a multi-target Bayes filter for the merged measurement likelihood. The family of labeled RFS mixture densities can be regarded as a conjugate prior for the merged measurement likelihood function, however, we refrain from using this term because this family covers almost every labeled multi-target density in practice, as well as being generally numerically intractable.

V. Approximate Implementation

The exact form of the tracker, as specified in propositions 4 and 5, is intractable due to the presence of joint densities in the prior. To derive an algorithm that can be implemented in practice, we make the assumption that the targets are a-priori independent, thus reducing the prior from a labelled RFS mixture density, to a generalised labelled multi-Bernoulli RFS. After applying the merged measurement likelihood to this prior, we obtain a posterior density which is no longer in the form of a GLMB, i.e. the GLMB is not a conjugate prior with respect to the proposed likelihood. However, by performing marginalisation on the joint densities in the posterior, we can obtain an approximate posterior in GLMB form, such that a recursive filter may be implemented (as shown in subsection V-B). We refer to this as the GLMB-M filter.

A. Prediction

The first step is to compute the density of surviving targets, which is done by iterating through the components of the previous density, generating a set of highly weighted predicted components as we go. The predicted component weight is determined by the product of the survival/termination probabilities of its constituent targets, thus we can list the components in order of weight using a k-shortest paths algorithm. To do this, we generate the following $n \times 2$ matrix, where n is the number of targets in the current component, and the negative logarithm is used to transform the problem from maximum-product to minimum-sum form

$$C = -\log \begin{pmatrix} p_S\left(X^{(n,i,1)}\right) & 1 - p_S\left(X^{(n,i,1)}\right) \\ \vdots & \vdots \\ p_S\left(X^{(n,i,n)}\right) & 1 - p_S\left(X^{(n,i,n)}\right) \end{pmatrix}. \quad (41)$$

This matrix is used to create a directed graph, where each element is a node, each node is connected to both nodes in the following row, and the cost associated with each edge is the value of the node that the edge points to. K-shortest paths [11] is then applied in order to list the cheapest paths from top to bottom. The rows in which column 1 was visited correspond to survivals, and rows in which column 2 was visited correspond to deaths. A new component is formed from the surviving targets, with their pdfs propagated to the current time using the single target transition kernel. The new component weight is the product of the previous weight and e^{-c}, where c is the cost of the path used to generate that component.

In addition to predicting the density of surviving targets, we need to generate a density for newborn targets. In the general case, targets may appear anywhere within the sensor's detection region. It is therefore important that the birth model not restrict the locations where new tracks can be initiated. A general measurement driven birth model would allow any number of measurements on any scan to initiate new tracks. However, this approach is very computationally demanding, as it requires enumerating all subsets of the measurements on each scan. Instead, we make the following three restrictions to reduce computation to a more managable level:

1) a maximum of one new target can appear on each scan,
2) measurements that are likely to have originated from existing targets cannot originite from newborn targets,
3) new targets are always detected on their first two scans.

Applying these can drastically reduce the number of terms in the birth density, thereby significantly cutting the computational requirements. Once the survival and birth densities have been computed, their product is taken, which yields the overall predicted multi-target density.

B. Update

For a prior density in the form a GLMB as in (4), applying the likelihood function (29) yields a posterior which is a labelled RFS mixture density with probability distribution given by

$$\pi\left(\boldsymbol{X}|Z\right) = \Delta\left(\boldsymbol{X}\right) \sum_{c\in\mathbb{C}} \sum_{\substack{\mathcal{U}\in\mathcal{P}(\boldsymbol{X}) \\ \theta\in\Theta_{\mathcal{U}}}} w_Z^{(c,\mathcal{U},\theta)}\left(\mathcal{L}_{\boldsymbol{X}}\right) \left[p^{(c,\mathcal{U},\theta)}\left(\cdot|Z\right)\right]^{\mathcal{U}(\boldsymbol{X})} \quad (42)$$

where

$$w_Z^{(c,\mathcal{U},\theta)}\left(L\right) = \frac{\frac{1}{|\mathcal{P}(L)|} w^{(c)}\left(L\right) \left[\eta_Z^{(c,\mathcal{U},\theta)}\left(\cdot\right)\right]^{\mathcal{U}(L)}}{\sum_{c\in\mathbb{C}} \sum_{J\subseteq\mathbb{L}} \sum_{\substack{\mathcal{U}\in\mathcal{P}(J) \\ \theta\in\Theta_{\mathcal{U}}}} \frac{1}{|\mathcal{P}(J)|} w^{(c)}\left(J\right) \left[\eta_Z^{(c,\mathcal{U},\theta)}\right]^{\mathcal{U}(J)}} \quad (43)$$

$$p^{(c,\mathcal{U},\theta)}\left(\boldsymbol{Y}|Z\right) = \frac{\tilde{p}^{(c,\mathcal{U})}\left(\boldsymbol{Y}\right) \tilde{\psi}_Z\left(\boldsymbol{Y};\theta\right)}{\eta_Z^{(c,\mathcal{U},\theta)}\left(\mathcal{L}_{\boldsymbol{Y}}\right)} \quad (44)$$

$$\eta_Z^{(c,\mathcal{U},\theta)}\left(L\right) = \left\langle \tilde{p}^{(c,\mathcal{U})}\left(\cdot,L\right), \tilde{\psi}_Z\left(\cdot,L;\theta\right) \right\rangle \quad (45)$$

$$\tilde{\psi}_Z\left(\boldsymbol{Y};\theta\right) = \begin{cases} \tilde{q}_D\left(\boldsymbol{Y}\right), & \theta\left(\boldsymbol{Y}\right) = 0 \\ \tilde{g}\left(z_{\theta(\boldsymbol{Y})}|\boldsymbol{Y}\right) \tilde{p}_D\left(\boldsymbol{Y}\right), & \theta\left(\boldsymbol{Y}\right) > 0 \end{cases} \quad (46)$$

The derivation is omitted due to space limitations. In order to bring the posterior back to the required GLMB form, we marginalise the joint densities within each component of the labelled RFS mixture density as follows

$$p^{(c,\mathcal{U},\theta)}\left(\boldsymbol{Y}|Z\right) \approx \left[p_m^{(c,\mathcal{U},\theta)}\left(\cdot|Z\right)\right]^{\boldsymbol{Y}} \quad (47)$$

$$p_m^{(c,\mathcal{U},\theta)}(x,l|Z) = \int p^{(c,\mathcal{U},\theta)}\left(\boldsymbol{Y}|Z\right) \delta\left(\boldsymbol{Y} - \{(x,l)\}\right) \quad (48)$$

This allows us to make the following approximation, since the set of joint densities representing the groups in partition $\mathcal{U}(\boldsymbol{X})$ reduces to a set of independent densities representing the targets in $\boldsymbol{X}$,

$$\left[p^{(c,\mathcal{U},\theta)}(\cdot|Z)\right]^{\mathcal{U}(\boldsymbol{X})} \approx \left[p_m^{(c,\mathcal{U},\theta)}(\cdot|Z)\right]^{\boldsymbol{X}}. \tag{49}$$

This leads to the following GLMB approximation to the posterior multi-object density

$$\begin{aligned}&\pi(\boldsymbol{X}|Z) \qquad\qquad (50)\\ &\approx \Delta(\boldsymbol{X}) \sum_{c\in\mathbb{C}} \sum_{\substack{\mathcal{U}\in\mathcal{P}(\boldsymbol{X})\\ \theta\in\Theta_{\mathcal{U}}}} w_Z^{(c,\mathcal{U},\theta)}(\mathcal{L}_{\boldsymbol{X}}) \left[p_m^{(c,\mathcal{U},\theta)}(\cdot|Z)\right]^{\boldsymbol{X}}.\end{aligned}$$

The procedure for approximating (50) involves generating a set of posterior components for each prior component, by enumerating the feasible partitions and measurement-to-group assignments. Note that this can be easily parallelised, since there is no dependency between prior components. The maximum overall number of components (denoted N_{max}), which is set as a parameter, is split up among all possible cardinalities, such that each cardinality $c \in 1..M$ is allowed to generate N_c posterior components. We have chosen to do this based on a Poisson distribution centered around the mean prior cardinality, which ensures a reasonable spread of components, allowing for target births and deaths. We then split up each N_c among the prior components of cardinality c, such that component (c,k) is allocated a number $N_{c,k}$, proportional to its prior weight. Since it is usually impossible to do this exactly for integer numbers of components, we use a randomised weighted allocation algorithm.

A brute force approach to computing (50) would list all partitions of the target set, and generate assignments of the measurements to groups within each partition. However, since we are limited in the number of terms we can compute, we must find a way to efficiently disregard insignificant components. Since targets often appear in clusters, many partitions may be very unlikely, and therfore do not warrant inclusion in the posterior. For each component, we compute clusters of the target set (i.e. groups which have reasonable probability of being unresolved), and enumerate the partitions for each cluster. These are then combined across clusters by taking the Cartesian product, yielding the feasible set of global partitions.

We now evaluate a score function, denoted $\gamma(\cdot)$, for each global partition, which is used to allocate the number of measurement-to-group associations to generate. The definition of $\gamma(\cdot)$ will be problem-dependent, as it needs to capture the relative likelihood that a set of targets will be observed according to different groupings. For the purposes of this study we assume a relatively simple model, such that on average, the measurements will be merged if their separation is less than W_R. We may argue that a partition is more likely if the targets within each group are clustered within distance W_R of each other, and the distances between groups is greater than W_R. We therefore define the score function as

$$\begin{aligned}\gamma(X,\mathcal{U}) = &\min\left(\frac{D_{min}\left(\{\bar{Y}; Y\in\mathcal{U}(\boldsymbol{X})\}\right)}{W_R}, 1\right)^2 \qquad (51)\\ &\times \prod_{Y\in\mathcal{U}(\boldsymbol{X})} C\left(\max\left(\frac{D_{max}(Y)}{W_R}, 1\right)^{-1}, \frac{1}{2}\right)\end{aligned}$$

where $\bar{Y}$ is the mean of the set Y, $D_{min}(Y)$ and $D_{max}(Y)$ are the minimum and maximum pairwise distances between points in Y, and C is a cutoff function defined as $C(x,y) = x$ when $x \geq y$, and $C(x,y) = 0$ when $x < y$. The first term ensures that the distance between the means of any pair of groups in $\mathcal{U}$ must be greater than W_R, and the likelihood increases as the square of the minimum distance between groups. The second term enforces that any group which is spread over a distance of more than $2W_R$ cannot be merged, and that any group which is spread over a distance of less than W_R has maximum likelihood of being merged.

Having computed the partition scores using (51), the number of allocated components $N_{c,k}$ is then divided among them proportionally by carrying out a randomised weighted allocation. The number allocated to partition $\mathcal{U}$ of the k-th component of cardinality c is denoted $N_{c,k,\mathcal{U}}$. We then proceed to compute the terms of the posterior by generating ranked measurement-to-group assignments for each partition using Murty's algorithm [12]. The cost matrix for a set of targets $\boldsymbol{X}$ and a particular partion $\mathcal{U} \in \mathcal{P}(\boldsymbol{X})$, is of the form $C = [D; M]$, where D is a $|\mathcal{U}| \times |Z|$ matrix with elements

$$D_{i,j} = -\log\left(\frac{p_D(\mathcal{U}^{(i)})\breve{g}\left(z^{(j)}|\mathcal{U}^{(i)}\right)}{\kappa\left(z^{(j)}\right)}\right) \tag{52}$$

and M is a $|\mathcal{U}| \times |\mathcal{U}|$ matrix with diagonal elements $M_{i,i} = -\log\left(q_D(\mathcal{U}^{(i)}\right)$, and all off-diagonal elements set to ∞. In the above, $\mathcal{U}^{(i)}$ denotes the i-th group in partition $\mathcal{U}$. For each assignment generated by Murty, we create a new posterior component, consisting of one updated joint density for each group of targets according to (44)-(46). Finally, we approximate the posterior as a GLMB by marginalising the joint densities for each target according to (47)-(49)

To extract the track estimates from the posterior, we find the maximum a-posteriori cardinality estimate, and then the highest weighted component with that cardinality. The track labels within the selected component are compared to a list of previously confirmed tracks. Tracks that have already been confirmed are updated with a new state extracted from the corresponding single target pdf, and those that have not are used to initiate new tracks.

To make the algorithm efficient, we prune the posterior density after each update so that insignificant components are removed. There are two conditions under which a component is retained. The first is if the component is weighted in the top N_{tot} overall, and its ratio to the maximum weighted component is less than a threshold R_{max}. Alternatively, a component is retained if it is in the top N_{cmin} for its cardinality.

VI. Simulation Results

We now demonstrate the performance of the merged measurement GLMB filter on a simulated multi-target passive sonar tracking scenario. The sensor generates noisy bearings-only measurements with false alarms and misdetections, and undergoes two manoeuvres to ensure that the target states become observable over the course of the simulation. The scenario involves multiple crossing targets, each following a white noise acceleration dynamic model. The target state space is defined in terms of 2D Cartesian position and velocity vectors

$$\boldsymbol{x} = \left[\begin{array}{cccc} x & y & \dot{x} & \dot{y} \end{array}\right]^T \tag{53}$$

and all targets follow the dynamic model

$$\boldsymbol{x}_{k+1} = \boldsymbol{F}\boldsymbol{x}_k + \boldsymbol{\Gamma}\boldsymbol{w}_k \tag{54}$$

$$\boldsymbol{F} = \left[\begin{array}{cc} 1 & T \\ 0 & 1 \end{array}\right] \otimes \boldsymbol{I}_2, \qquad \boldsymbol{\Gamma} = \left[\begin{array}{c} T^2/2 \\ T \end{array}\right] \otimes \boldsymbol{I}_2 \tag{55}$$

where T is the sensor sampling period, and $\boldsymbol{w}_k \sim \mathcal{N}(0, \boldsymbol{Q})$ is a 2×1 independent and identically distributed Gaussian process noise vector with $\boldsymbol{Q} = \sigma_w^2 \boldsymbol{I}_2$ where the standard deviation of the target acceleration is $\sigma_w = 10^{-3} m/s^2$.

The measurements are simulated using the scheme described in Section III, with single target measurement function

$$h(\boldsymbol{x}_t, \boldsymbol{x}_s) = \arctan\left(\frac{x_t - x_s}{y_t - y_s}\right). \tag{56}$$

For the purposes of this analysis, we model the pdf of each target using a single Gaussian, and the extended Kalman filter (EKF) is used to perform the measurement updates. It is clearly possible to use other types of non-linear filters such as the unscented Kalman filter or cubature Kalman filter, or to model the target pdfs using more accurate representations such as Gaussian mixtures or particle distributions. These techniques may yield some improvement, but this analysis is beyond the scope of the paper.

The scenario consists of 4 targets running parallel to each other, with geometry as shown in Figure 2. One target is present at the beginning, with another three arriving during the first 250 seconds, and three dying during the last 400 seconds. Between time 1000 and 1400, the bearings cross each other, and their measurements become merged. The sensor sampling period is 5 seconds, measurement noise has a standard deviation of 0.5 degrees, resolution cell width is 2.5 degrees, detection probability is 0.98, and clutter is uniformly distributed on the interval $[0, 2\pi]$ with a Poisson cardinality distribution with a mean of 40 points per scan. A single realisation of the measurements is shown in Figure 1.

The filter generates a maximum of $N_{max} = 1000$ components during the update, which is pruned back to $N_{tot} = 100$ (with $N_{cmin} = 10$, and $R_{max} = 10^{-10}$) before processing the next scan. Both filters were implemented in C++, and executed on an Intel Core i7 2600 processor. No attempt at parallelising the generation of components has been made.

Figure 2a shows the tracks from a single run when the standard GLMB filter is applied in the presence of merged measurements. When the measurements are merged, the filter drops three tracks to account for the 'missing' measurements. When the targets are resolved again, new tracks (with a default prior pdf) are initiated on those that were dropped. Since the prior is quite uninformative, and no further sensor manoeuvres are carried out to re-establish observability, the localisation performance is severely degraded. Figure 2b shows the track output from the GLMB-M filter for the same measurement data. This time, all four tracks are maintained through the crossing, resulting in significantly better localisation and cardinality estimation performance. Note that localisation is poor during the early stages of the scenario, since the targets are not observable until the sensor performs its first manoeuvre.

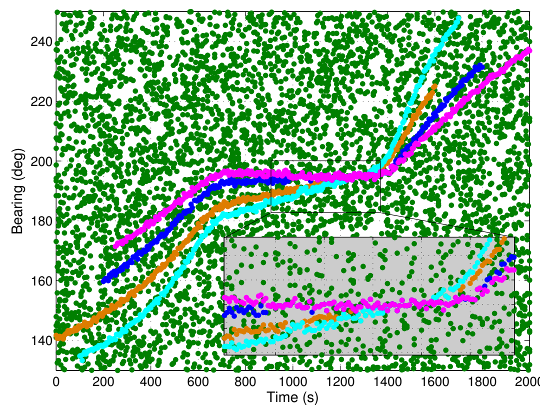

Figure 1. Simulated bearing measurements with merging. The merged measurements can be seen in the inset, which shows fewer than four target generated measurements, despite the fact that all four targets are still present.

Figure 3a shows the average optimal subpattern assignment (OSPA) distance [13] (with cutoff 1 km and order 2) between the true and estimated target sets for the two filters over 100 Monte Carlo runs. To avoid biassing the results, a different random offset is applied to the location of the cell boundaries on each run. The standard GLMB performs poorly since it cannot account for the merged measurements. Whenever two or more targets produce a merged measurement, tracks become prematurely terminated, which is clearly evident from the average cardinality estimates shown in Figure 3b. The GLMB-M performs well in this scenario, since it is able to reliably maintain tracks when the measurements become merged. By considering the possible groupings of the target set in the measurement likelihood calculation, the algorithm is capable of assigning a single measurement to a group of targets, thereby keeping all tracks in that group alive.

The execution time for the two filters is shown in Figure 3c. The GLMB-M filter has a significanly higher peak execution time due to the enumeration of the feasible partitions for each component in the density. When targets are closely spaced, the number of partitions increases significantly, leading to increased computation. The execution time of the standard GLMB drops in the presence of merged measurements, but this is due to the filter incorrectly terminating some of its tracks, so fewer components are needed in the density.

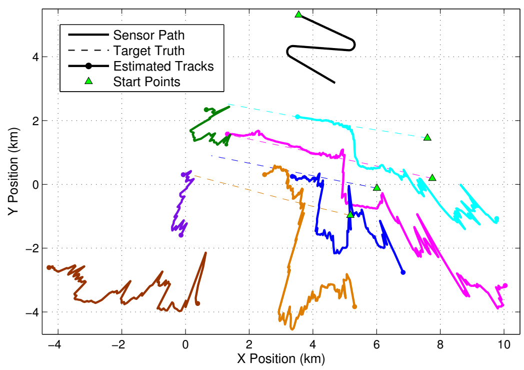

(a) GLMB filter with standard likelihood

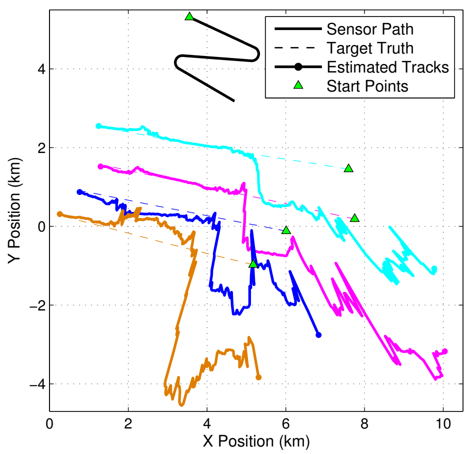

(b) GLMB-M filter

Figure 2. Single run with merged measurements

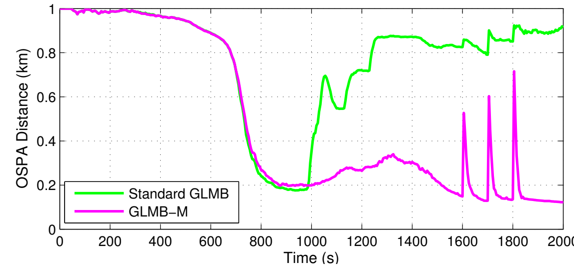

(a) Optimal sub-pattern assignment distance

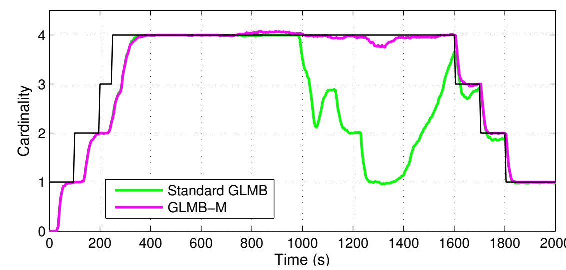

(b) Cardinality estimates

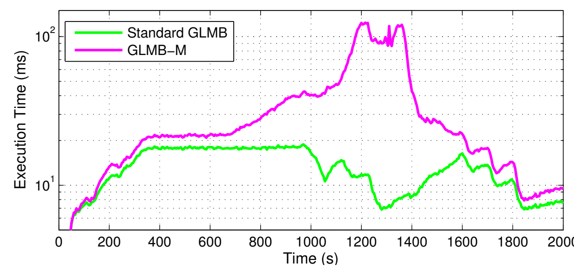

(c) Execution Time

Figure 3. Monte Carlo results (100 runs)

VII. Conclusion

We have proposed an algorithm for multi-object tracking with merged measurements, using the framework of labelled random finite sets. The algorithm is a generalisation of the GLMB filter, originally presented in [9] for the standard sensor model without measurement merging. The exact form of our proposed merged measurement tracker is intractable, so we have also proposed an approximation that allows for its practical implementation. Simulation of a bearings-only multi-target tracking scenario shows that the algorithm performs well, however, it remains computationally demanding compared to the standard GLMB filter. Future work in this area will involve investigating ways of reducing the computational load whilst maintaining acceptable tracking performance, and comparing the performance to some existing techniques for tracking in the presence of merged measurements.

References

[1] K.-C. Chang, Y. Bar-Shalom, "Joint probabilistic data association for multitarget tracking with possibly unresolved measurements and maneuvers", *IEEE Transactions on Automatic Control*, vol. 29, no. 7, pp. 585-594, July 1984.

[2] D. Svensson, M. Ulmke, and L. Danielsson, "Joint probabilistic data association filter for partially unresolved target groups", *Proc. 13th International Conference on Information Fusion*, Edinburgh, UK, July 2010.

[3] H. A. P. Blom, E. A. Bloem, "Bayesian tracking of two possibly unresolved maneuvering targets", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 2, pp. 612-627, April 2007.

[4] S. Jeong, J. K. Tugnait, "Tracking of two targets in clutter with possibly unresolved measurements", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 44, no. 2, pp. 748-765, April 2008.

[5] D. Musicki, M. R. Morelande, "Finite resolution multitarget tracking", *Proc. SPIE Signal and Data Processing of Small Targets 2005*, vol. 5913, September 2005.

[6] W. Koch, G. van Keuk, "Multiple hypothesis track maintenance with possibly unresolved measurements", *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 33, No. 3, pp. 883-892, July 1997.

[7] S. J. Davey, "Tracking possibly unresolved targets with PMHT", *Information, Decision and Control (IDC) 2007*, February 2007.

[8] R. Mahler, "PHD filters for nonstandard targets, II: Unresolved targets", *Proc. 12th International Conference on Information Fusion*, pp. 922-929, Seattle, Washington, USA, July 2009.

[9] 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, July 2013.

[10] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*, Artech House, 2007.

[11] J. Y. Yen, "Finding the K-shortest loopless paths in a network", *Management Science*, vol. 17, pp. 712-716, 1971.

[12] K. G. Murty, "An algorithm for ranking all the assignments in increasing order of cost", *Operations Research*, vol. 16, no. 3, pp. 682-678, 1968.

[13] D. Schuhmacher, B.-T. Vo, 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, August 2008.