Poisson multi-Bernoulli mixture trackers: continuity through random finite sets of trajectories

Karl Granström, Lennart Svensson
Yuxuan Xia
Dept. of Electrical Eng.
Chalmers Univ. of Tech.
Gothenburg, Sweden
`firstname.lastname@chalmers.se`

Jason Williams
National Security and ISR Division
Defence Science and Tech. Group
Edinburgh, Australia
`jason.williams@dst.defence.gov.au`

Ángel F. García-Fernández
Dept. of Electrical Eng. and Electronics
Univ. of Liverpool
Liverpool, UK
`angel.garcia-fernandez@liverpool.ac.uk`

Abstract—The Poisson multi-Bernoulli mixture (PMBM) is an unlabelled multi-target distribution for which the prediction and update are closed. It has a Poisson birth process, and new Bernoulli components are generated on each new measurement as a part of the Bayesian measurement update. The PMBM filter is similar to the multiple hypothesis tracker (MHT), but seemingly does not provide explicit continuity between time steps. This paper considers a recently developed formulation of the multi-target tracking problem as a random finite set (RFS) of trajectories, and derives two trajectory RFS filters, called PMBM trackers. The PMBM trackers efficiently estimate the set of trajectories, and share hypothesis structure with the PMBM filter. By showing that the prediction and update in the PMBM filter can be viewed as an efficient method for calculating the time marginals of the RFS of trajectories, continuity in the same sense as MHT is established for the PMBM filter.

Index Terms—data association, tracking, filtering, smoothing, trajectories, random finite sets

I. INTRODUCTION

Multi-target tracking (MTT) is a challenging problem due to the unknown correspondence of measurements and targets, referred to as data association. Each new measurement received could be the continuation of some previously detected target, the first detection of a new target, or a false alarm.

The multiple hypothesis tracker (MHT) introduced in [1] is a key method for addressing problems of this type. The algorithm maintains a series of global hypotheses for each possible measurement-target correspondence as measurements are received, along with a conditional state distribution for each target under each hypothesis. Each measurement could be hypothesised to represent the first detection of a new target, where the number of newly detected targets was given a Poisson distribution in order to provide a Bayesian prior.

In [2], the model was made rigorous through random finite sequences, under the assumption that the number of targets present is constant but unknown, with a Poisson prior. Target state sequences were formed under each global hypothesis, and the Poisson distribution of targets remaining to be detected provided a Bayes prior for events involving newly detected targets. Hypotheses were constructed as being data-to-data, since no a priori data was assumed on target identity. This formulation was revisited in [3], and compared to random finite sets, and finite point processes.

The framework of random finite sets (RFSs) was developed to provide a systematic methodology for dealing with problems involving an uncertain number of targets [4]. In [5], the Bayes filter for Poisson birth models was derived using RFSs, obtaining a result somewhat similar to [2], involving a Poisson distribution representing targets which are hypothesised to exist but remain to be detected, and a multi-Bernoulli mixture (MBM) representing targets that have been detected at some stage. Adopting the common modelling assumptions employed in RFS (see, e.g., [6]), target appearance and disappearance was modelled. The resulting, so called, PMBM filter, did not formally establish track continuity, although a hypothesis structure similar to MHT has been observed [7], [8].

MHT was extended in [9], [10] to address a time-varying number of targets, incorporating birth of targets which are not immediately detected, necessitating equivalence classes of indistinguishable hypotheses. In comparison, in [5] the hypotheses are purely data-to-data assignments, the hypothesis space explicitly includes targets that remain to be detected, and target death subsequent to the final detection.

The labelled RFS formulation of [11] addresses the lack of continuity of unlabelled RFSs by incorporating a label (uniqueness of which is ensured through the model) into the target state; continuity is maintained by connecting estimates with the same label at different times. PGFL expressions for the joint target-measurement density over a time interval were given in [12], including modelling of target birth, target death, and target spawning using branching and immigration processes. A solution was proposed using saddle point approximation, however, a practical implementation was not provided in [12]. In [13], a framework for performing inference on problems involving both distinguishable (previously detected) and indistinguishable (never detected) targets, was developed and applied to space situation awareness. Targets were characterized by sequences of measurements, under the non-standard assumption that the targets are *"immediately detected upon entering the surveillance scene"* [13, A.5].

A formulation of the target tracking problem as an unlabelled RFS of trajectories was provided recently in [14], [15]. Within this set of trajectories framework, the goal of Bayesian MTT is to compute the posterior density over the

set of trajectories. One approach to do this is to let the multi-target state be the set of trajectories, and compute the posterior density recursively. Assuming standard point target motion and measurement models, in this paper we formulate two different set of trajectory problems: one in which the set of current (i.e., remaining) trajectories is tracked, and one in which the set of all trajectories up to the current time is tracked. For the two problem formulations, we assume multi-target densities of the PMBM form and derive the PMBM predictions and the PMBM update. This results in two different PMBM set of trajectories filters: we call these tracking algorithms *PMBM trackers*, to distinguish them from the PMBM filter [5], which is for sets of target states.

The prediction and update of PMBM trackers are the main contribution in the paper. In addition to this, we discuss the PMBM trackers in relation to the PMBM filter [5], the MHT [1], [2], and the δ-GLMB filter [11]. Interestingly, the relations between the PMBM trackers and the PMBM filter show that the PMBM filter implicitly contains trajectory information. We also show how the PMBM trackers can be implemented for linear Gaussian models, and present results from a simulation study where we compare the results to the δ-GLMB filter [11].

Importantly, our paper shows that it is possible to compute the density over the set of trajectories even for complicated problems; in [10, Sec. 2] computing the set of trajectories density was thought to be *"extremely challenging... even for small problems"*. Further, we compute the density without requiring the assumption that targets are detected with probability one upon entering the scene.

II. PROBLEM FORMULATION

To clearly differentiate between a target state at a single time and a sequence of target states, we let *target* denote the state at some time, and we let *trajectory* denote a sequence of states. Thus, *"set of targets"* and *"target RFS"* refers to formulations involving RFSs of target states at a single time (e.g., the unlabelled RFS formulation in [5]), and *"set of trajectories"* and *"trajectory RFS"* refers to formulations involving a RFS of trajectories, or sequences of states.

Let x_k denote a target state at time k, and let z_k denote a measurement at time k. We utilise the standard multi-target dynamics model, defined in Assumption 1, and the standard point target measurement model, defined in Assumption 2.

Assumption 1. *The multiple target state evolves according to the following time dynamics process: targets arrive at each time according to a non-homogeneous Poisson Point Process (PPP) with birth intensity $\lambda^{\mathrm{b}}(x_k)$, independent of existing targets; targets depart according to iid Markovian processes; the survival probability in state x_k is $P^{\mathrm{S}}(x_k)$; target motion follows iid Markovian processes; the single-target transition density is $\pi^x(x_k|x_{k-1})$.*

Assumption 2. *The multiple target measurement process is as follows: each target may give rise to at most one measurement; probability of detection in state x_k is $P^{\mathrm{D}}(x_k)$; each measurement is the result of at most one target; false alarm measurements arrive according to a non-homogeneous PPP with intensity $\lambda^{\mathrm{FA}}(z_k)$, independent of targets and target-related measurements; each target-derived measurement is independent of all other targets and measurements conditioned on its corresponding target; the single target measurement likelihood is $f(z_k|x_k)$.*

There are many ways in which a Multiple Target Tracking (MTT) problem can be formulated. In this paper, we focus on the following two variants:

1) *The set of current trajectories:* the objective is to estimate the trajectories of the targets who are present in the surveillance area at the current time.
2) *The set of all trajectories:* the objective is to estimate the trajectories of all targets that have passed through the surveillance area at some point between time 0 and the current time, i.e., both the targets who are present in the surveillance area at the current time, and the targets that have left the surveillance area (but were in the surveillance area at at least one previous time).

Note that both problem formulations are equally valid; which one is relevant depends on the application. Indeed, in some applications, a solution with only current states (no trajectories) is the objective.

III. RANDOM FINITE SETS OF TRAJECTORIES

Compact representations of the multiple trajectory density for the standard point target models, cf. Assumptions 1 and 2, are obtained using sets of trajectories, developed in the RFS formalism in [15]. Expressions for integration, Bayes prediction and Bayes update may be found there. In this section, we first review the trajectory state representation, and then present densities and probability generating functionals for sets of trajectories.

A. Single trajectory

Let $\mathcal{X}$ represent the base state space, e.g., $\mathcal{X} = \mathbb{R}^4$, representing position and velocity in two dimensions. We use the trajectory state model proposed in [14], [15], in which the trajectory state is a tuple

$$X = (\beta, \varepsilon, x_{\beta:\varepsilon}) \tag{1}$$

where β is the discrete time of the trajectory birth, i.e., the time the trajectory begins; ε is the discrete time of the trajectory's most recent state, i.e., the time the trajectory ends. If k is the current time, $\varepsilon < k$ means that the trajectory ended at time ε, and $\varepsilon = k$ means that the trajectory is ongoing; $x_{\beta:\varepsilon}$ is, given β and ε, the sequence of states

$$x_\beta, x_{\beta+1}, \ldots, x_{\varepsilon-1}, x_\varepsilon, \tag{2}$$

where $x_k \in \mathcal{X}$ for all $k \in \{\beta, \ldots, \varepsilon\}$. This gives a trajectory of length $\ell = \varepsilon - \beta + 1$ time steps. The trajectory state space at time k is [15]

$$\mathcal{T}_k = \uplus_{(\beta,\varepsilon) \in I_k} \{\beta\} \times \{\varepsilon\} \times \mathcal{X}^{\varepsilon-\beta+1}, \tag{3}$$

where $\uplus$ denotes disjoint set union, $I_k = \{(\beta, \varepsilon) : 0 \leq \beta \leq \varepsilon \leq k\}$ and $\mathcal{X}^\ell$ denotes ℓ Cartesian products of $\mathcal{X}$. The trajectory state density factorises as follows[1]

$$p_{k|k'}(X) = p_{k|k'}(x_{\beta:\varepsilon}|\beta, \varepsilon)P_{k|k'}(\beta, \varepsilon), \tag{4}$$

where, if $\varepsilon < \beta$, then $P_{k|k'}(\beta, \varepsilon)$ is zero. Integration is performed as follows [15],

$$\int p(X)\mathrm{d}X = \iiint p(x_{\beta:\varepsilon}|\beta, \varepsilon)P(\beta, \varepsilon)\mathrm{d}x_{\beta:\varepsilon}\mathrm{d}\varepsilon\mathrm{d}\beta \tag{5a}$$

$$= \sum_{\beta,\varepsilon} \left[\int \cdots \int p(x_\beta, \ldots, x_\varepsilon|\beta, \varepsilon)\mathrm{d}x_\beta \ldots \mathrm{d}x_\varepsilon \right] P(\beta, \varepsilon). \tag{5b}$$

B. Multiple trajectories

We now move to the case of a *set* of trajectories. Analogously to a set of targets, a set of trajectories is denoted as $\mathbf{X}_k \in \mathcal{F}(\mathcal{T}_k)$, where $\mathcal{F}(\mathcal{T}_k)$ is the set of all finite subsets of $\mathcal{T}_k$. Let $g(\mathbf{X}_k)$ be a function on a set of trajectories $\mathbf{X}_k$. Integration over sets of trajectories is defined as regular set integration:

$$\int g(\mathbf{X}_k)\delta\mathbf{X}_k \triangleq g(\emptyset) +$$

$$\sum_{n=1}^\infty \frac{1}{n!} \int \cdots \int g(\{X_k^1, \ldots, X_k^n\})\mathrm{d}X_k^1 \cdots \mathrm{d}X_k^n. \tag{6}$$

The set of trajectories density function $f_\mathbf{X}(\mathbf{X})$ is defined analogously to the set of targets density function. The probability generating functional (PGFL), see, e.g., [4], is a useful tool for manipulating and understanding RFS densities. The PGFL for a trajectory RFS density is defined like the PGFL for a target RFS density: using a test function $h(X) : \mathcal{T} \to \mathbb{R}$, the PGFL is

$$G_\mathbf{X}[h] \triangleq \int h^\mathbf{X} f_\mathbf{X}(\mathbf{X})\delta\mathbf{X} \tag{7}$$

$$\triangleq f_\mathbf{X}(\emptyset) + \sum_{n=1}^\infty \frac{1}{n!} \int \cdots \int \left[\prod_{i=1}^n h(X^i) \right]$$

$$\times f_\mathbf{X}(\{X^1, \ldots, X^n\})\mathrm{d}X^1 \cdots \mathrm{d}X^n, \tag{8}$$

where $h^\mathbf{X}$ is set power, defined as $h^\mathbf{X} = 1$ if $\mathbf{X} = \emptyset$ and $h^\mathbf{X} = \prod_{X \in \mathbf{X}} h(X)$ if $\mathbf{X} \neq \emptyset$.

A trajectory PPP is analogous to a target PPP, and has set density and PGFL

$$f^{\mathrm{PPP}}(\mathbf{X}) = e^{-\int \lambda(X)\mathrm{d}X} \lambda^\mathbf{X} \tag{9a}$$

$$G^{\mathrm{PPP}}[h] = \exp\{\langle \lambda; h - 1 \rangle\}, \tag{9b}$$

where $\langle f; g \rangle = \int f(X)g(X)\mathrm{d}X$. The trajectory PPP intensity $\lambda(\cdot)$ is defined on the trajectory state space $\mathcal{T}_k$, i.e., realisations of the PPP are trajectories with a birth time, a time of the most recent state, and a state sequence.

A trajectory Bernoulli process is analogous to a target Bernoulli process, and has set density and PGFL

$$f^{\mathrm{ber}}(\mathbf{X}) = \begin{cases} 1 - r, & \mathbf{X} = \emptyset \\ rf(X), & \mathbf{X} = \{X\} \\ 0, & \text{otherwise} \end{cases} \tag{10a}$$

$$G^{\mathrm{ber}}[h] = 1 - r + r\langle f; h \rangle \tag{10b}$$

Here, $f(\cdot)$ is a trajectory state density (4), and r is the Bernoulli probability of existence. Together, r and $f(\cdot)$ can be used to find the probability that the target trajectory existed at a specific time, or find the probability that the target state was in a certain area at a certain time. Trajectory MB RFS and trajectory MBM RFS are both defined analogously to target MB RFS and target MBM RFS: a trajectory MB is the disjoint union of a multiple trajectory Bernoulli RFS; trajectory MBM RFS is an RFS whose density is a mixture of trajectory MB densities.

IV. PMBM TRACKERS

As in [5], we hypothesise a multi-target conjugate prior[2] of the following Poisson Multi-Bernoulli Mixture (PMBM) form, and observe that the form is preserved through prediction and update. The PMBM PGFL is

$$G_{k|k'}[h] = \exp\left\{\left\langle \lambda_{k|k'}^\mathrm{u}; h(X_k) - 1 \right\rangle\right\} \tag{11}$$

$$\times \sum_{a \in \mathcal{A}^{k|k'}} w_{k|k'}^a \prod_{i \in \mathbb{T}_{k|k'}} \left(1 - r_{k|k'}^{i,a^i} + r_{k|k'}^{i,a^i} \left\langle f_{k|k'}^{i,a^i}; h \right\rangle \right).$$

This PGFL form states that the set of trajectories $\mathbf{X}_k$ is an independent union of a PPP with intensity $\lambda_{k|k'}^\mathrm{u}$, and a multi-Bernoulli mixture (MBM) with Bernoulli parameters $r_{k|k'}^{i,a^i}$ and $f_{k|k'}^{i,a^i}(\cdot)$.

The structure of the trajectory PMBM (11) is the same as the structure of the target PMBM from [5]. The PPP represents trajectories that are hypothesised to exist, but have never been detected, i.e., no measurement has been associated to them. We present *"track oriented"* (TO) PMBM trackers, where a track is initiated for each measurement. In the MBM in (11), $\mathbb{T}_{k|k'}$ is a track table with $n_{k|k'}$ tracks, $a \in \mathcal{A}^{k|k'}$ is a possible global data association hypothesis, and for each global hypothesis a and each track $i \in \mathbb{T}_{k|k'}$, a^i indicates which track hypothesis is used in the global hypothesis. For each track, there are $n_{k|k'}^i$ single trajectory hypotheses. The weight of the global hypothesis a is $w_{k|k'}^a \propto \prod_{i \in \mathbb{T}_{k|k'}} w_{k|k'}^{i,a^i}$, where $w_{k|k'}^{i,a^i}$ is the weight of single trajectory hypothesis a^i from track i. The set of all measurement indices up to time k is denoted $\mathcal{M}^k$, and $\mathcal{M}^k(i, a^i)$ is the history of measurements that are hypothesised to belong to hypothesis a^i from track i. Due to page length constraints, further elaboration is omitted; please refer to [5] for details and discussion.

Depending on the problem formulation, the multi-target state $\mathbf{X}_k$ that we are interested in is different. For the set

---

[1]"$p_{k|k'}$" is the density of X_k given measurements up to and including time $k' \leq k$.

[2]*Multi-target conjugate prior* was defined in [11] as meaning that "...*if we start with the proposed conjugate initial prior, then all subsequent predicted and posterior distributions have the same form as the initial prior.*"

of current trajectories, $\mathbf{X}_k$ is the set of trajectories for which $0 \leq \beta \leq \varepsilon = k$. For the set of all trajectories, $\mathbf{X}_k$ is the set of trajectories for which $0 \leq \beta \leq \varepsilon \leq k$. A PMBM process is defined by the parameters of the density/PGFL,

$$\lambda_{k|k'}^{\mathrm{u}}, \left\{ w_{k|k'}^{i,a^i}, r_{k|k'}^{i,a^i}, f_{k|k'}^{i,a^i} \right\} \tag{12}$$

In the following subsections we will show how the parameters are predicted and updated, in order to track either the set of current trajectories, or the set of all trajectories.

A. Prediction step

In this section we describe the time evolution of the set of trajectories. A standard PPP birth model is used, i.e., target birth at time step k is modelled by a PPP, with intensity

$$\lambda_k^{\mathrm{B}}(X) = \lambda_k^{\mathrm{B},x}(x_{\beta:\varepsilon}|\beta, \varepsilon)\Delta_k(\varepsilon)\Delta_k(\beta), \tag{13a}$$

$$\lambda_k^{\mathrm{B},x}(x_{k:k}|k, k) = \lambda_k^{\mathrm{b}}(x_k), \tag{13b}$$

where $\Delta(\cdot)$ denotes Kronecker's delta function. Note that it is possible to have alternative birth models, such as MB birth, or MBM birth. For those cases, the spatial densities of the Bernoulli birth components would be of the same form as in (13), i.e., for birth at time k we have $\beta = \varepsilon = k$.

For the transition of existing targets, there is one alternative for each of the two problem formulations. For the set of current trajectories, $P^{\mathrm{S}}(\cdot)$ is used in a way that is typical for tracking a set of targets, see, e.g., [5], [11]. For the set of all trajectories, $P^{\mathrm{S}}(\cdot)$ is used as in [15]. The trajectory state dependent probability of survival at time k is defined as

$$P_k^{\mathrm{S}}(X) = P^{\mathrm{S}}(x_\varepsilon)\Delta_k(\varepsilon). \tag{14}$$

1) Transition model for the set of current trajectories: The Bernoulli RFS transition density without birth is

$$f_{k|k-1}^{\mathrm{c}}(\mathbf{X}|\mathbf{X}') = \tag{15a}$$
$$\begin{cases} 1, & \mathbf{X}' = \emptyset, \mathbf{X} = \emptyset \\ 1 - P_{k-1}^{\mathrm{S}}(X'), & \text{if} \quad \mathbf{X}' = \{X'\}, \mathbf{X} = \emptyset \\ P_{k-1}^{\mathrm{S}}(X')\pi^{\mathrm{c}}(X|X'), & \mathbf{X}' = \{X'\}, \mathbf{X} = \{X\} \\ 0, & \text{otherwise} \end{cases}$$

$$\pi^{\mathrm{c}}(X|X') = \pi^{\mathrm{c},x}(x_{\beta:\varepsilon}|\beta, \varepsilon, X')\Delta_{\varepsilon'+1}(\varepsilon)\Delta_{\beta'}(\beta), \tag{15b}$$

$$\pi^{\mathrm{c},x}(x_{\beta:\varepsilon}|\beta, \varepsilon, X') = \pi^x(x_\varepsilon|x_{\varepsilon'}')\delta_{x_{\beta':\varepsilon'}'}(x_{\beta:\varepsilon-1}), \tag{15c}$$

where $\delta(\cdot)$ denotes Dirac's delta function. In this model, $P^{\mathrm{S}}(\cdot)$ is used as follows. If the target disappears, or "dies" (case $\mathbf{X}' = \{X'\}, \mathbf{X} = \emptyset$ in (15a)), then the entire trajectory will no longer be a member of the set of current trajectories. If the target survives, then the trajectory is extended by one time step. The PGFL of the multiple object transition density is

$$G_{k|k-1}^{\mathrm{c}}[h|\mathbf{X}] = e^{\langle \lambda_k^{\mathrm{B}};h-1 \rangle} \left(1 - P_{k-1}^{\mathrm{S}} + P_{k-1}^{\mathrm{S}} \langle \pi^{\mathrm{c}};h \rangle \right)^{\mathbf{X}} \tag{16}$$

The resulting prediction step is given in the theorem below.

Theorem 1. *Assume that the distribution from the previous time step $G_{k-1|k-1}[h]$ is the PGFL of the form given in (11), and that the transition model is of the kind (16). Then the*

predicted distribution for the next step $G_{k|k-1}[h]$ is of the form (11), with:

$$\lambda_{k|k-1}^{\mathrm{u}}(X_k) = \lambda_k^{\mathrm{B}}(X_k) + \left\langle \lambda_{k-1|k-1}^{\mathrm{u}}; \pi^{\mathrm{c}}P_{k-1}^{\mathrm{S}} \right\rangle \tag{17a}$$

$$n_{k|k-1} = n_{k-1|k-1} \tag{17b}$$

$$h_{k|k-1}^i = h_{k-1|k-1}^i \; \forall \, i \tag{17c}$$

$$w_{k|k-1}^{i,a^i} = w_{k-1|k-1}^{i,a^i} \; \forall \, i, a^i \tag{17d}$$

$$r_{k|k-1}^{i,a^i} = r_{k-1|k-1}^{i,a^i} \left\langle f_{k-1|k-1}^{i,a^i}; P_{k-1}^{\mathrm{S}} \right\rangle, \; \forall \, i, a^i \tag{17e}$$

$$f_{k|k-1}^{i,a^i} = \frac{\left\langle f_{k-1|k-1}^{i,a^i}; \pi^{\mathrm{c}}P_{k-1}^{\mathrm{S}} \right\rangle}{\left\langle f_{k-1|k-1}^{i,a^i}; P_{k-1}^{\mathrm{S}} \right\rangle}, \; \forall \, i, a^i \tag{17f}$$

2) Transition model for the set of all trajectories: The Bernoulli RFS transition density without birth is

$$f_{k|k-1}^{\mathrm{a}}(\mathbf{X}|\mathbf{X}') = \tag{18a}$$
$$\begin{cases} 1 & \mathbf{X}' = \emptyset, \mathbf{X} = \emptyset \\ \pi^{\mathrm{a}}(X|X') & \text{if} \quad \mathbf{X}' = \{X'\}, \mathbf{X} = \{X\} \\ 0 & \text{otherwise} \end{cases}$$

$$\pi^{\mathrm{a}}(X|X') = \pi^{\mathrm{a},x}(x_{\beta:\varepsilon}|\beta, \varepsilon, X')\pi^\varepsilon(\varepsilon|\beta, X')\Delta_{\beta'}(\beta), \tag{18b}$$

$$\pi^\varepsilon(\varepsilon|\beta, X') = \begin{cases} 1, & \varepsilon = \varepsilon' < k-1 \\ 1 - P_{k-1}^{\mathrm{S}}(X'), & \varepsilon = \varepsilon' = k-1 \\ P_{k-1}^{\mathrm{S}}(X'), & \varepsilon = \varepsilon' + 1 = k \\ 0, & \text{otherwise} \end{cases}, \tag{18c}$$

$$\pi^{\mathrm{a},x}(x_{\beta:\varepsilon}|\beta, \varepsilon, X') = \tag{18d}$$
$$\begin{cases} \delta_{x_{\beta':\varepsilon'}'}(x_{\beta:\varepsilon}), & \varepsilon = \varepsilon' \\ \pi^x(x_\varepsilon|x_{\varepsilon'}')\delta_{x_{\beta':\varepsilon'}'}(x_{\beta:\varepsilon-1}), & \varepsilon = \varepsilon' + 1 \end{cases}.$$

In this model, the interpretation of the probability of survival is that it governs whether or not the trajectory ends, or if it extends by one more time step. However, importantly, regardless of whether or not the trajectory ends, the trajectory remains in the set of all trajectories.

In this case, the PGFL of the multiple object transition density with PPP birth is

$$G_{k|k-1}^{\mathrm{a}}[h|\mathbf{X}] = e^{\langle \lambda_k^{\mathrm{B}};h-1 \rangle} \langle \pi^{\mathrm{a}};h \rangle^{\mathbf{X}} \tag{19}$$

Note the difference to (16), which comes from the way in which the probability of survival is used. In this case, the target always remains in the set of targets, even if its trajectory has ended. The prediction step is presented in the theorem below.

Theorem 2. *Assume that the distribution from the previous time step $G_{k-1|k-1}[h]$ is of the form given in (11), and that the transition model is of the kind (19). Then the predicted distribution for the next step $G_{k|k-1}[h]$ is of the form (11),*

with:

$$\lambda^{\mathrm{u}}_{k|k-1}(X_k) = \lambda^{\mathrm{B}}_k(X_k) + \left\langle \lambda^{\mathrm{u}}_{k-1|k-1}; \pi^{\mathrm{a}} \right\rangle \tag{20a}$$

$$n_{k|k-1} = n_{k-1|k-1} \tag{20b}$$

$$h^i_{k|k-1} = h^i_{k-1|k-1} \; \forall \, i \tag{20c}$$

$$w^{i,a^i}_{k|k-1} = w^{i,a^i}_{k-1|k-1} \; \forall \, i, a^i \tag{20d}$$

$$r^{i,a^i}_{k|k-1} = r^{i,a^i}_{k-1|k-1}, \; \forall \, i, a^i \tag{20e}$$

$$f^{i,a^i}_{k|k-1} = \left\langle f^{i,a^i}_{k-1|k-1}; \pi^{\mathrm{a}} \right\rangle, \; \forall \, i, a^i \tag{20f}$$

B. Update step

The target measurement model of assumption 2 is extended by defining a Bernoulli measurement density as follows:

$$\varphi_k(\mathbf{W}_k|\mathbf{X}) = \tag{21a}$$

$$\begin{cases} 1, & \mathbf{X} = \emptyset, \mathbf{W}_k = \emptyset \\ 1 - P^{\mathrm{D}}_k(X), & \mathbf{X} = \{X\}, \mathbf{W}_k = \emptyset \\ P^{\mathrm{D}}_k(X)\varphi(z_k|X), & \mathbf{X} = \{X\}, \mathbf{W}_k = \{z_k\} \\ 0, & \text{otherwise} \end{cases}$$

$$P^{\mathrm{D}}_k(X) = P^{\mathrm{D}}(x_\varepsilon)\Delta_k(\varepsilon), \tag{21b}$$

$$\varphi(z|X) = f(z|x_\varepsilon). \tag{21c}$$

The clutter is modelled as a PPP with intensity $\lambda^{\mathrm{FA}}_k(z)$, and it follows that the measurement PGFL is

$$G_k[g|\mathbf{X}_k] = e^{\left\langle \lambda^{\mathrm{FA}}_k; g-1 \right\rangle} \left(1 - P^{\mathrm{D}} + P^{\mathrm{D}}\left\langle \varphi; g \right\rangle\right)^{\mathbf{X}_k}. \tag{22}$$

We see that $G_k[g|\mathbf{X}_k]$ is a trajectory equivalent to the standard point target PGFL, and thus the trajectory measurement update is analogous to the target measurement update in [5].

Theorem 3. *Assume that the predicted distribution $G_{k|k-1}[h]$ is of the form given in (11), and that the measurement model is of the kind (22). Then, the updated distribution $G_{k|k}[h]$ (updated with the measurement set $Z_k = \{z^1_k, \dots, z^{m_k}_k\}$) is of the same form, with $n_{k|k} = n_{k|k-1} + m_k$,*

$$\lambda^{\mathrm{u}}_{k|k}(X_k) = \left(1 - P^{\mathrm{D}}(X_k)\right) \lambda^{\mathrm{u}}_{k|k-1}(X_k), \tag{23}$$

$$\mathcal{M}^k = \mathcal{M}^{k-1} \cup \left\{(k,j)|j \in \{1, \dots, m_k\}\right\}. \tag{24}$$

For tracks continuing from previous time steps ($i \in \{1, \dots, n_{k|k-1}\}$), a hypothesis is included for each combination of a hypothesis from a previous time and either a missed detection or an update using one of the m_k new measurements, such that the number of hypotheses becomes $h^i_{k|k} = h^i_{k|k-1}(1 + m_k)$. For missed detection hypotheses ($i \in \{1, \dots, n_{k|k-1}\}$, $a^i \in \{1, \dots, h_{k|k-1}\}$):

$$\mathcal{M}^k(i,a^i) = \mathcal{M}^{k-1}(i,a^i) \tag{25a}$$

$$w^{i,a^i}_{k|k} = w^{i,a^i}_{k|k-1} \left(1 - r^{i,a^i}_{k|k-1}\left\langle f^{i,a^i}_{k|k-1}; P^{\mathrm{D}} \right\rangle\right) \tag{25b}$$

$$r^{i,a^i}_{k|k} = \frac{r^{i,a^i}_{k|k-1}\left\langle f^{i,a^i}_{k|k-1}; 1 - P^{\mathrm{D}} \right\rangle}{1 - r^{i,a^i}_{k|k-1}\left\langle f^{i,a^i}_{k|k-1}; P^{\mathrm{D}} \right\rangle} \tag{25c}$$

$$f^{i,a^i}_{k|k}(X_k) = \frac{\left(1 - P^{\mathrm{D}}(X)\right) f^{i,a^i}_{k|k-1}(X)}{\left\langle f^{i,a^i}_{k|k-1}; 1 - P^{\mathrm{D}} \right\rangle} \tag{25d}$$

For hypotheses updating existing tracks ($i \in \{1, \dots, n_{k|k-1}\}$, $a^i = \tilde{a}^i + h^i_{k|k-1}j$, $\tilde{a}^i \in \{1, \dots, h^i_{k|k-1}\}$, $j \in \{1, \dots, m_k\}$, i.e., the previous hypothesis $\tilde{a}^i$, updated with measurement z^j_k).[3]

$$\mathcal{M}^k(i,a^i) = \mathcal{M}^{k-1}(i,\tilde{a}^i) \cup \{(t,j)\} \tag{26a}$$

$$w^{i,a^i}_{k|k} = w^{i,\tilde{a}^i}_{k|k-1} r^{i,\tilde{a}^i}_{k|k-1} \left\langle f^{i,\tilde{a}^i}_{k|k-1}; \varphi(z^j_k|\cdot)P^{\mathrm{D}} \right\rangle \tag{26b}$$

$$r^{i,a^i}_{k|k} = 1 \tag{26c}$$

$$f^{i,a^i}_{k|k}(X_k) = \frac{\varphi(z^j_k|X_k)P^{\mathrm{D}}_k(X_k)f^{i,\tilde{a}^i}_{k|k-1}(X_k)}{\left\langle f^{i,\tilde{a}^i}_{k|k-1}; \varphi(z^j_k|\cdot)P^{\mathrm{D}}_k \right\rangle} \tag{26d}$$

Finally, for new tracks, $i \in \{n_{k|k-1} + j\}$, $j \in \{1, \dots, m_k\}$ (i.e., the new track commencing on measurement z^j_k),

$$h^i_{k|k} = 2 \tag{27a}$$

$$\mathcal{M}^k(i,1) = \emptyset, \quad w^{i,1}_{k|k} = 1, \quad r^{i,1}_{k|k} = 0 \tag{27b}$$

$$\mathcal{M}^k(i,2) = \{(t,j)\} \tag{27c}$$

$$w^{i,2}_{k|k} = \lambda^{\mathrm{FA}}(z^j_k) + \left\langle \lambda^{\mathrm{u}}_{k|k-1}; \varphi(z^j_k|\cdot)P^{\mathrm{D}}_k \right\rangle \tag{27d}$$

$$r^{i,2}_{k|k} = \frac{\left\langle \lambda^{\mathrm{u}}_{k|k-1}; \varphi(z^j_k|\cdot)P^{\mathrm{D}}_k \right\rangle}{\lambda^{\mathrm{FA}}(z^j_k) + \left\langle \lambda^{\mathrm{u}}_{k|k-1}; \varphi(z^j_k|\cdot)P^{\mathrm{D}}_k \right\rangle} \tag{27e}$$

$$f^{i,2}_{k|k}(X_k) = \frac{\varphi(z^j_k|X_k)P^{\mathrm{D}}_k(X_k)\lambda^{\mathrm{u}}_{k|k-1}(X_k)}{\left\langle \lambda^{\mathrm{u}}_{k|k-1}; \varphi(z^j_k|\cdot)P^{\mathrm{D}}_k \right\rangle} \tag{27f}$$

C. Density/intensity representation and estimation

Note that the birth PPP intensity (13) is an un-normalized mixture density of the form

$$p(X) = \sum_j w^j p^j(x_{\beta:\varepsilon}|\beta, \varepsilon)\Delta_{e^j}(\varepsilon)\Delta_{b^j}(\beta), \tag{28}$$

where each mixture component is characterised by a weight w^j, a distinct birth time b^j, a distinct most recent time e^j where $b^j \le e^j$ for all j, and a state sequence density $p^j(\cdot)$. For the weights we have that $\sum_j w^j = 1$ if $p(\cdot)$ is a density, and $\sum_j w^j \ge 0$ if $p(\cdot)$ is an intensity function, e.g., a PPP intensity. Because of the assumption that the trajectory birth intensity is an un-normalized density of the form (28), it follows that the undetected intensity $\lambda^{\mathrm{u}}_{k|k'}(X_k)$, and all Bernoulli densities $f^{i,a^i}_{k|k'}(X_k)$, will be of the form (28). This type of state density facilitates simple representations for the state sequence $x_{\beta:\varepsilon}$, conditioned on β and ε.

Trajectory estimation, or trajectory extraction, is the process of obtaining estimates of the set of trajectories (or set of targets) from the multi-target density. A typical approach to MTT estimation is to select a certain global hypothesis and report estimates from it. In this paper, we use a PMBM filter estimator—e.g., any of the three discussed in [16, Sec. VI] is

---

[3] A hypothesis at the previous time with $r^{i,a^i}_{k|k-1} = 0$ need not be updated since the posterior weight in (26b) would be zero. For simplicity, the hypothesis numbering does not account for this exclusion.

valid—to identify a global hypothesis from which to extract estimates, and then extract the full trajectories directly. Note that alternative estimators are possible, representing different trade-offs between tracking accuracy and computational cost. Comparing estimators will be adressed in future work.

D. Structure of PMBM trackers

Two different PMBM trackers result from the theorems: a PMBM tracker for the set of current trajectories is given by the prediction in Theorem 1 and the update in Theorem 3; a PMBM tracker for the set of all trajectories is given by the prediction in Theorem 2 and the update in Theorem 3. Note that both PMBM trackers are TO. For each measurement, a potential new track is initiated, see (27). As in the PMBM filter [5], for each track there is a hypothesis tree, where each hypothesis corresponds to different data association sequences for the track. The predictions (17) and (20) preserve the number of tracks and the number of hypotheses[4], meaning that the predictions can be implemented without approximation. In the update, additional hypotheses are created, as indicated in (25) and (26).

The Bernoulli probabilities of existence r have different meanings in the two trackers: for the set of current trajectories problem formulation, r is the probability that the trajectory exists at the current time step k and has not ended yet; in the set of all trajectories problem formulation, r represents the probability that the trajectory existed at any time between 0 and the current time step k.

In Theorem 1, the predicted Bernoulli existence probability (17e) is the product of the posterior probability and $\langle P_k^S; f \rangle$. This results in a predicted probability $r \leq 1$, where equality only holds for the atypical choice $P_k^S(x_k) = 1$. The predicted Bernoulli density (17f) results in a density in which $\varepsilon = k$ with probability $= 1$ and the state sequence is augmented by one more time step. This follows directly from the set of *current* trajectories problem formulation: the filter does not maintain trajectories that ended before time k, because those are not current trajectories. Notice the important difference in Theorem 2: here the Bernoulli existence probability is constant, and $P_k^S(\cdot)$ enters in the integrals with π^a. In the set of *all* trajectories problem formulation, the existence probability represents existence from time 0 to time k, and this is unaffected by whether or not the trajectory ends at this time or not, or if it has already ended. The probabilities of different ε, and the corresponding state sequences, are captured in the Bernoulli density (20f).

In Theorem 3, the updated Bernoulli existence probability decreases following a missed detection (25c), unless the predicted probability is $= 1$. For the set of current trajectories, the predicted probability of existence (17e) is typically < 1, however, for the set of all trajectories the predicted probability of existence can be $= 1$. However, remember that trajectory death ($\varepsilon < k$) is modelled in the Bernoulli densities (25d) via the mixture components, cf. (28). For the detection updates

---

[4]The prediction (20f) results in additional mixture components (28).

(26) and (27), an implication of $P^D(\cdot)$ (21b) is that a detected trajectory must necessarily exist at time k. It follows that, for both problem formulations, in the updated state densities (26d) and (27f) the probability of $\varepsilon = k$ is $= 1$.

V. DISCUSSION

A. Relations to PMBM filter and MHT

Theorem 1 is a direct re-statement of the PMBM filter prediction, see [5, Thm. 1], to the trajectory RFS dynamics model. Further, Theorem 3 is a direct re-statement of the PMBM filter update, see [5, Thm. 2], to the trajectory RFS measurement model. Note that regardless of which problem formulation is considered, current trajectories or all trajectories, the update is the same. An important difference between the work presented here and the work in [5] is that the former is for trajectory sets, whereas the latter is for target sets.

The predictions and update of the PMBM trackers are TO, with one track initiated for each measurement. By comparing the track hypotheses in the two TO-PMBM trackers with the track hypotheses in the TO-PMBM filter [5], we see that the hypothesis structure is identical for the prediction and update steps. This can be seen by comparing the components of the distribution of the form (11), and by considering the marginalisation of the trajectories to leave only the current time step k. The result of the marginalisation of a trajectory Bernoulli density with probability of existence $r_{k|k'}$ and trajectory density $p_{k|k'}(X)$ is a target Bernoulli density, with probability of existence

$$r_k^m = r_{k|k'} \sum_{\beta \leq k} \sum_{\varepsilon = k} P_{k|k'}(\beta, \varepsilon), \tag{29a}$$

and, if $r_k^m > 0$, state density

$$p_k^m(x_k) = \frac{1}{r_k^m} \sum_{\beta \leq k} \sum_{\varepsilon = k} \int p_{k|k'}(x_{\beta:\varepsilon}|\beta, \varepsilon) \mathrm{d}x_{\beta:k-1} P_{k|k'}(\beta, \varepsilon). \tag{29b}$$

If $r^m = 0$, a state density need not be defined, because the target does not exist at time k.

Regarding the general relation between target densities and trajectory densities, a result was presented in [15, Thm. 11] showing that if a trajectory RFS distribution is marginalised to the final time to obtain a target RFS distribution, and prediction and update steps are performed on that target RFS distribution, then the result is equivalent to the marginalisation of the trajectory RFS distribution *after* prediction and update using the equivalent trajectory models.

Using the marginalisation (29), it can be shown that marginalising each predicted/updated trajectory hypothesis is the same as predicting/updating the marginalised trajectory hypothesis; the details are omitted due to page length constraints. From this, it follows that the marginalisation of the posterior PMBM tracker density is equivalent to the PMBM filter update of the marginalisation of the prior PMBM tracker density. The same type of relation can be shown for the prediction.

Since the system state in the PMBM filter [5] was formally an unlabelled set of target states, track continuity was not

provided formally, yet a hypothesis structure similar to MHT was observed, which was exploited to provide an implicit form of track continuity. The MHT and the PMBM filter were compared in detail in [7], [8], and conditions for equivalence were identified. It follows from the PMBM tracker/filter-relations established here that the PMBM filter [5] can be understood to operate implicitly on the trajectories: it marginalises the state sequence in each prediction, and it only ever explicitly outputs the current time step. This shows the implicit track continuity in the PMBM filter. Note also that the general similarities between the MHT and the PMBM filter carry over to the PMBM trackers.

B. Track initiation

The track initiation in the PMBM trackers is of the same type as in the PMBM filter [5], as well as in the TO-MHT: a trajectory is defined by the measurement that initiates it. In tracking algorithms based on labelled multi-Bernoulli densities, such as the δ-GLMB filter [11], a trajectory is defined by the labelled Bernoulli birth component.[5]

One disadvantage of labelled birth is that, by its intrinsic modelling properties, it does not permit targets to be indistinguishable up until first detection. This is a significant practical drawback in cases where targets are actually indistinguishable before being detected for the first time, e.g.: when the first detection occurs a long time after target birth; or when target births are iid, which is a common modelling assumption. In fact, in the majority of cases, it is an essential character of the tracking problem that targets are indistinguishable up until first detection, after which continuity is able to be maintained (for as long as sufficient information persists). Thus, we consider it to be an important advantage that in the PMBM tracker, the birth is modelled by a PPP (with intensity (13)).

Note that after the first detection, the targets are distinguishable. From a trajectory and a given data association hypothesis, we can infer at all times the location of the target. Thus, there is no need to label the target upon initialisation.

C. Estimation

The TO-MHT provides estimates of a target trajectory by estimating the current state at each time, and connecting estimates from different times. The δ-GLMB filter estimates trajectories in a similar fashion: the set of current target states is estimated at each time, and the estimates from different times are connected into trajectories using the labels. Both of these methods can be understood as answering the trajectory question *"for the chosen data association hypothesis, what is the current location of the target that was initiated by measurement z_β at time β (in the case of MHT), or was initiated by birth component l_β at time β (in the case of δ-GLMB)?"*

The estimation in the PMBM tracker (Section IV-C) can be understood as answering the trajectory question *"for the chosen data association hypothesis, what is the state sequence*

---

[5]Specifically, each Bernoulli birth component is uniquely labelled at birth, which is modelled through the prediction step; refer to [11] for details.

---

of the target that was initiated by measurement z_β at time β?" A benefit of this type of estimation is that instead of appending the sequence of estimates by one more state estimate, it estimates the full state sequence, i.e., smoothed estimates. This is possible, thanks to the fact that each global hypothesis contains full trajectory information.

Thus, the PMBM tracker produces full trajectory estimates upon receipt of each new set of measurements, whereas the PMBM filter, the MHT, and the δ-GLMB filter produce estimates of the latest time. The beginning of the trajectory will reveal the origin of the target, and thus correspondence between trajectories estimated at different times is established. Practically, these linkages can be made simply through hypothesis metadata, similar to MHT, and if all that is required is to estimate the current state at each time, the PMBM filter suffices, since it may be interpreted as implicitly operating on the full trajectory (cf. Section V-A).

VI. LINEAR GAUSSIAN IMPLEMENTATION

Let the single target transition density and measurement likelihoods both be linear-Gaussian,

$$\pi^x(x|x') = \mathcal{N}\left(x;\, Fx', Q\right), f(z|x) = \mathcal{N}\left(z;\, Hx, R\right), \quad (30)$$

and let the probabilities of detection and of survival be constant, $P_k^D(\cdot) = P^D$ and $P_k^S(\cdot) = P^S$. Further, let the state sequence, conditioned on β and ε, be Gaussian distributed with mean m and covariance P. The state sequence density can be expressed in information form,

$$p_{k|k'}(x_{\beta:\varepsilon}|\beta, \varepsilon) = \frac{e^{-\frac{1}{2}y_{k|k'}^T Y_{k|k'}^{-1} y_{k|k'} - \frac{1}{2}x_{\beta:\varepsilon}^T Y_{k|k'} x_{\beta:\varepsilon} + y_{k|k'}^T x_{\beta:\varepsilon}}}{\sqrt{|2\pi Y_{k|k'}^{-1}|}}$$

(31)

with information vector $y_{k|k'} = P_{k|k'}^{-1} m_{k|k'}$ and information matrix $Y_{k|k'} = P_{k|k'}^{-1}$.

Let $y_{k|k}$ and $Y_{k|k}$ be the posterior information parameters. Given a linear Gaussian motion model (30), the predicted trajectory density has information vector and information matrix [17]

$$y_{k+1|k} = \begin{bmatrix} y_{k|k} \\ \mathbb{O}_{n_x \times 1} \end{bmatrix} \quad (32a)$$

$$Y_{k+1|k} = \begin{bmatrix} Y_{k|k}^{[\beta:\varepsilon-1,\beta:\varepsilon-1]} & Y_{k|k}^{[\beta:\varepsilon-1,\varepsilon]} & \mathbb{O}_{(\ell-1)n_x \times n_x} \\ Y_{k|k}^{[\varepsilon,\beta:\varepsilon-1]} & Y_{k|k}^{[\varepsilon,\varepsilon]} + F^T Q^{-1} F & -F^T Q^{-1} \\ \mathbb{O}_{n_x \times (\ell-1)n_x} & -Q^{-1}F & Q^{-1} \end{bmatrix}$$

(32b)

where n_x is the dimension of the single-time-step state x, $\mathbb{O}_{m \times n}$ is an m by n all-zero matrix, and $Y^{[a:b,c:d]}$ denotes the part the information matrix with rows for time steps a to b and columns for time steps c to d. Given a linear Gaussian

measurement model (30), the posterior trajectory density has information vector and information matrix [17]

$$y_{k+1|k+1} = y_{k+1|k} + \begin{bmatrix} \mathbb{O}_{(\ell-1)n_x \times 1} \\ H^T R^{-1} z \end{bmatrix}, \tag{33a}$$

$$Y_{k+1|k+1} = Y_{k+1|k} + \begin{bmatrix} \mathbb{O}_{(\ell-1)n_x \times (\ell-1)n_x} & \mathbb{O}_{(\ell-1)n_x \times n_x} \\ \mathbb{O}_{n_x \times (\ell-1)n_x} & H^T R^{-1} H \end{bmatrix}. \tag{33b}$$

Note that in the prediction (32), both the information vector and the information matrix are augmented to account for the additional time step that, following the prediction, is included by the trajectory. A key result is that the bottom left and top right corners of the predicted information matrix (32b) are exactly zero, and that the update (33b) only affects the part of the information matrix that is related to the current state. This means that the information matrix is exactly sparse, a direct consequence of the Markov property associated with $\pi^x(\cdot|\cdot)$. The number of non-zero elements in the information matrix increases linearly: a trajectory from β to ε includes $\ell = \varepsilon - \beta + 1$ time steps. In the information matrix this leads to $n_x^2 (3\ell - 2)$ non-zero elements. In comparison, with a Gaussian representation (with a mean vector and covariance matrix) the covariance is full and has $n_x^2 \ell^2$ non-zero elements. Alternatively, an approximate solution has to be used, in which the trajectory density is marginalised to only contain the most recent time steps.

The PMBM trackers were implemented using Bernoulli state densities and Poisson intensities of the form (28), with information form densities (31) for the state sequences. To limit the number of data associations in each update (cf. Theorem 3), analogously to the PMBM filter, see [16, Sec V.C.3], the K best global hypotheses are found using Murty's algorithm [18]. Given a Bernoulli state density, an estimate of the trajectory is obtained by selecting the most probable mixture component $j^\star = \arg\max_j w_{k|k'}^j$, and computing the expected value of the state sequence $\hat{x}_{\hat{\beta}:\hat{\varepsilon}} = (Y_{k|k'}^{j^\star})^{-1} y_{k|k'}^{j^\star}$, where $\hat{\beta} = b_{k|k'}^{j^\star}$ and $\hat{\varepsilon} = e_{k|k'}^{j^\star}$. Note that in theory, the weights (26b) and (27d), and the estimate $\hat{x}_{\hat{\beta}:\hat{\varepsilon}}$, involves the inverse of the information matrix $Y_{k|k'}^{j^\star}$. However, it is not necessary to compute the inverse in practice. Instead, multiplications with the inverse information matrix are solved efficiently as a sparse least squares problem. Further discussion about computationally efficient data association, as well as state recovery (both full and partial), can be found in, e.g., [17].

VII. RESULTS

In this section we present results from two simulated scenarios, one with a high number of targets and long trajectories, and one in which there is coalescence, i.e., the targets are all very close at one point. Both scenarios were generated with a 2D constant velocity motion model, see, e.g., [19, Sec. 3], with acceleration standard deviation σ_v, linear position measurements with covariance R The PMBM tracker for all targets is compared to the δ-GLMB filter [11], in both the birth density was assumed known.

A. Measuring performance

To evaluate tracking performance, the trajectory metric $d(\cdot, \cdot)$ [20] was used with location error cut-off $c = 100$, order $p = 1$, and switch cost $\gamma = 20$; the trajectory metric can be decomposed into a location cost, a missed target cost, a false target cost, and a switch cost. In the simulated scenarios, we apply the metric at each time step, and normalise it by the time step. This allows a comparison of how the metric evolves over time in the scenario, as opposed to only computing the metric at the final time step.

For scenarios with many targets, computing the metric becomes costly; in this case we have used the GOSPA metric [21] instead, which decomposes into location cost, missed target cost, and false target cost. The GOSPA metric is applied to the set of target states at each time step, and a trajectory metric is obtained by summing the GOSPA metric over time.

B. Scenario with many targets

A scenario with 200 time steps and 117 true trajectories was generated and then processed by the tracker for all trajectories. Average GOSPA over 100 Monte Carlo simulations, the location cost per target per time is 0.25, the number of missed targets per time step is 5×10^{-4}, and the number of false targets per time step is 0.33. Processing the full scenario took, on averge, 154 seconds, i.e., 0.77 seconds per time step.[6] This shows that the PMBM trackers are accurate in terms of tracking performance, with low location error and few cardinality errors (missed and false targets), and that they are computationally feasible even for a high number of long trajectories.

C. Scenario with coalescence

As noted in, e.g., [5], a large number of spatially separated targets is not necessarily the most difficult scenario. Therefore, we highlight a challenging scenario with three targets that start separated, come together at the mid point of the scenario, and then separate again. The parameters were $\sigma_v = 0.5$, $R = \text{diag}([100, 100])$, $P^S = 0.99$, $P^D = 0.98$ and $\lambda^{FA} = 2.5 \times 10^{-8}$ in the area $[-10^3, 10^3] \times [-10^3, 10^3]$. The birth density had a single component, centered in the surveillance area, with a covariance covering the whole surveillance area.

The PMBM tracker for all trajectories is compared to the δ-GLMB filter, which provides estimates at each time using measurements up to and including that time. For trajectory estimates extracted at the each time step of the scenario, 100 Monte Carlo simulations gave the results shown in Figure 1. If the results are summed over time, the following is obtained,

	Metric	Loc	Miss	False	Switch
PMBM tracker	2593.4	2361.7	181.8	3.9	46.1
δ-GLMB filter	5546.6	4442.4	854.3	104.2	145.7

Example results are shown in Figure 2. The PMBM results is three valid trajectories; however, there are switches in the δ-GLMB results around $(x, y) = (-350, -50)$ and $(x, y) = (300, -50)$. The switches around $(x, y) = (-350, -50)$

---

[6]Matlab implementation on laptop with 3.1GHz processor, 16GB memory.
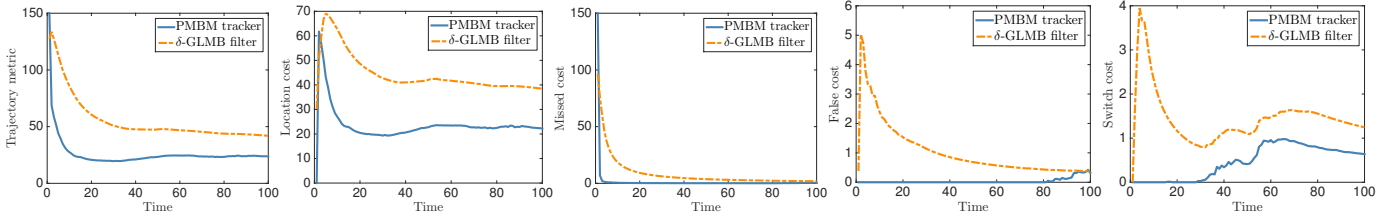
Fig. 1. Results from simulated scenario with coalescence; the lines show mean over 100 Monte Carlo runs.
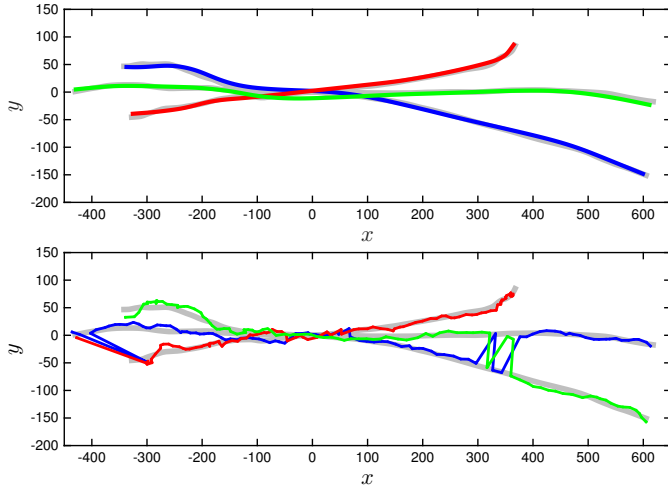

Fig. 2. Example results from scenario with coalescence. Information form PMBM results (top), Gaussian δ-GLMB results (bottom). Three true targets move left to right.

are due to label ambiguity upon birth: all targets are born at the same time and it is not known which birth label belongs to which measurement. The track switches around $(x, y) = (300, -50)$ are a result of ambiguity regarding which target goes where as they separate: following the separation, multiple data association sequences, i.e., global hypotheses, are approximately equally likely. Note that both the PMBM tracker and the δ-GLMB filter have target switches around the mid-point of the scenario; however, it is only the δ-GLMB filter that sometimes produces these kinds of unrealistic switches. An important difference between the δ-GLMB filter and the PMBM trackers is that, at any time, the PMBM trackers always provide a valid trajectory, i.e., not one that is flipping between different hypotheses at different times.

VIII. CONCLUSION

In this paper we have presented two PMBM trackers for the set of target trajectories, and we have established connections to the PMBM filter for the set of targets. The connections show the implicit track continuity in the PMBM filter.

REFERENCES

[1] D. B. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Automatic Control*, vol. AC-24, no. 6, pp. 843–854, December 1979.
[2] S. Mori, C.-Y. Chong, E. Tse, and R. Wishner, "Tracking and classifying multiple targets without a priori identification," *IEEE Trans. Automatic Control*, vol. 31, no. 5, pp. 401–409, May 1986.
[3] S. Mori, C.-Y. Chong, and K. C. Chang, "Three formalisms of multiple hypothesis tracking," in *Proc. 19th International Conference on Information Fusion*, July 2016.
[4] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA: Artech House, 2007.
[5] J. L. Williams, "Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA and association-based MeMBer," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 51, no. 3, July 2015.
[6] R. P. S. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003.
[7] E. F. Brekke and M. Chitre, "The multiple hypothesis tracker derived from finite set statistics," in *Proc of FUSION*, July 2017.
[8] E. Brekke and M. Chitre, "Relationship between finite set statistics and the multiple hypothesis tracker," *IEEE Trans. Aerosp. Electron. Syst.*, vol. PP, 02 2018.
[9] S. Coraluppi, C. Carthel, and M. Guerriero, "Hypothesis structure in enhanced multiple-hypothesis tracking," in *Asilomar Conference on Signals, Systems and Computers*, Nov 2013, pp. 1738–1742.
[10] S. Coraluppi and C. A. Carthel, "If a tree falls in the woods, it does make a sound: multiple-hypothesis tracking with undetected target births," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 50, no. 3, pp. 2379–2388, July 2014.
[11] B.-T. Vo and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Trans. Signal Process.*, vol. 61, no. 13, pp. 3460–3475, 2013.
[12] R. Streit, "Interval/smoothing filters for multiple object tracking via analytic combinatorics," in *Proc. 20th International Conference on Information Fusion*, July 2017.
[13] E. Delande, C. Frueh, J. Franco, J. Houssineau, and D. Clark, "Novel multi-object filtering approach for space situational awareness," *Journal of Guidance, Control, and Dynamics*, vol. 41, no. 1, pp. 59–73, January 2018.
[14] L. Svensson and M. Morelande, "Target tracking based on sets of trajectories," in *Proceedings of the International Conference on Information Fusion*, Salamanca, Spain, Jul. 2014.
[15] Á. F. García-Fernández, L. Svensson, and M. R. Morelande, "Multiple target tracking based on sets of trajectories," arXiv, e-print 1605.08163, May 2016. [Online]. Available: http://arxiv.org/abs/1605.08163
[16] A. F. Garcia-Fernandez, J. Williams, K. Granström, and L. Svensson, "Poisson multi-Bernoulli mixture filter: direct derivation and implementation," *IEEE Transactions on Aerospace and Electronic Systems*.
[17] R. M. Eustice, H. Singh, and J. J. Leonard, "Exactly sparse delayed-state filters for view-based SLAM," *IEEE Transactions on Robotics*, vol. 22, no. 6, pp. 1100–1114, 2006.
[18] K. Murty, "An algorithm for ranking all the assignments in order of increasing cost," *Operations Research*, vol. 16, no. 3, pp. 682–687, 1968.
[19] X. Rong-Li and V. Jilkov, "Survey of maneuvering target tracking: Part I. Dynamic models," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1333–1364, Oct. 2003.
[20] A. S. Rahmathullah, Á. F. García-Fernández, and L. Svensson, "A metric on the space of finite sets of trajectories for evaluation of multi-target tracking algorithms," *arXiv pre-print*, 2016. [Online]. Available: arxiv.org/abs/1605.01177
[21] A. S. Rahmathullah, A. F. García Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *Proceedings of International Conference on Information Fusion*, Jul. 2017, pp. 1–8.