

An Implementation of the Poisson Multi-Bernoulli Mixture Trajectory Filter via Dual Decomposition

Yuxuan Xia[†], Karl Granström[†], Lennart Svensson[†], Ángel F. García-Fernández[‡]

[†]Department of Electrical Engineering, Chalmers University of Technology, Göteborg, Sweden

[‡]Department of Electrical Engineering and Electronics, University of Liverpool, Liverpool, United Kingdom

Abstract—This paper proposes an efficient implementation of the Poisson multi-Bernoulli mixture (PMBM) trajectory filter. The proposed implementation performs track-oriented N-scan pruning to limit complexity, and uses dual decomposition to solve the involved multi-frame assignment problem. In contrast to the existing PMBM filter for sets of targets, the PMBM trajectory filter is based on sets of trajectories which ensures that track continuity is formally maintained. The resulting filter is an efficient and scalable approximation to a Bayes optimal multi-target tracking algorithm, and its performance is compared, in a simulation study, to the PMBM target filter, and the delta generalized labelled multi-Bernoulli filter, in terms of state/trajectory estimation error and computational time.

Index Terms—Bayesian estimation, multiple target tracking, random finite sets, set of trajectories, data association, multi-frame assignment.

I. INTRODUCTION

Multiple target tracking (MTT) refers to the problem of jointly estimating the number of targets and their trajectories from noisy sensor measurements [1]. The three most popular approaches to solving the MTT are the joint probabilistic data association filter [2], the multiple hypothesis tracker (MHT) [3], [4], [5] and multi-target filters based on random finite sets (RFS) [6].

During the last decade, RFS-based MTT algorithms have received a great deal of attention. In the early stage, RFS-based MTT algorithms were developed based on moment approximations of posterior multi-target densities, including the Probability Hypothesis Density (PHD) filter [7] and the Cardinalised PHD (CPHD) filter [8]. In recent years, a significant trend in RFS-based MTT is the development of conjugate multi-target distributions [9], [10], meaning that the posterior multi-target distribution has the same functional form as the prior. Comparisons have shown that PHD and CPHD filters have worse tracking performance, compared to filters based on MTT conjugate priors [11], [12].

We proceed to describe two MTT conjugate priors in the literature for the standard point target measurement model, which assumes that each target gives rise to at most one measurement and that targets do not have any shared measurement. The first conjugate prior, presented for unlabelled RFSs in [10], consists of the disjoint union of a Poisson point process (PPP) describing targets that have not yet been detected, and a multi-Bernoulli mixture (MBM) considering all the data association hypotheses. The second conjugate prior is the generalized labelled multi-Bernoulli (GLMB) density

[9] presented for labelled RFS, in which labels are appended to target states with the objective of the forming of target trajectories.

The Poisson multi-Bernoulli mixture (PMBM) filter [12] and the δ -GLMB filter [11] are two computationally tractable approaches, respectively, based on the PMBM and GLMB conjugate priors. The PMBM filter has computational advantages in relation to the δ -GLMB filter due to the structure of the hypotheses [12]. This was further analyzed in the performance evaluation study for point target filtering in [13]. However, unlike the δ -GLMB filter, the PMBM filter on its own does not provide information to estimate trajectories.

A Bayesian formulation of the MTT problem, in terms of unlabelled RFS of trajectories, was recently provided in [14]. The PMBM filtering recursions based on this framework have been derived in [15], which enables us to leverage on the benefits with the PMBM recursions while also obtaining track continuity. As a sequel to [15], the aim of this paper is to complement its theoretical contributions with practical algorithms. In particular, we present an efficient implementation of the PMBM trajectory filter that estimates the trajectories of the targets who are present in the surveillance area at the current time.

In this paper, we propose an approximation to the exact PMBM trajectory filter, by using track-oriented N-scan pruning [5] and dual decomposition [16] to solve the multi-frame assignment problem. The proposed algorithm therefore shares some of the key properties of certain MHT algorithms [5], [16], but is derived using RFS and birth/death models. Similar work has been done in [17], where MHT algorithms considering multi-frame data association based on random finite set/sequence formalism are derived but with the assumption that the number of targets is unknown but constant over time.

Compared to the δ -GLMB filter, our algorithm is unlabelled, and, given a global hypothesis, we obtain a distribution over the set of trajectories based on which it is convenient to construct trajectory estimates. The performance of the PMBM trajectory filter and the existing PMBM target filter [12], along with the δ -GLMB filter with joint prediction and update step [18], are compared in a simulated scenario with coalescence.

The paper is organized as follows. In Section II, we introduce the background on RFS modelling and the PMBM conjugate prior on the set of trajectories. In Section III, we present our proposed implementation. Simulation results are presented in Section IV, and conclusions are drawn in

II. BACKGROUND

In this section, we first outline the modelling assumptions utilized in this work. Next, we give a brief introduction to RFS of trajectories. Then, we give a summary to the PMBM conjugate prior on the set of trajectories.

A. Modelling assumptions

In the traditional RFS of targets formulation, target states and measurements are represented in the form of finite sets [6]. Let x_k denote a single target state at discrete time step k , and let z_k^j denote the j th measurement at discrete time step k . The set of measurements obtained at time step k is denoted as Z_k , including clutter and target-generated measurements with unknown origin.

In this paper, the state variable is a RFS of trajectories, as suggested in [14]. Single trajectories are parameterized using a tuple $X = (\beta, \epsilon, \mathbf{x}_{\beta:\epsilon})$ [15], [19], 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 (If k is the current time, $\epsilon < k$ means that the trajectory ends at time ϵ , and $\epsilon = k$ means that the trajectory is ongoing); $\mathbf{x}_{\beta:\epsilon}$ is, given β and ϵ , the sequence of states $(x_\beta, x_{\beta+1}, \dots, x_{\epsilon-1}, x_\epsilon)$.

We utilize 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 PPP with intensity $\lambda^b(x_k)$, independently of any existing targets.*
- *Targets survive with a probability $p^s(x_k)$, and targets depart the surveillance area according to i.i.d. Markov processes with probability $1 - p^s(x_k)$.*
- *Target motion follows i.i.d. Markov processes with transition density $\pi(x_k|x_{k-1})$.*

Assumption 2. *The multiple target measurement process is as follows:*

- *Each target can generate at most one measurement per time scan with detection probability $p^d(x_k)$.*
- *The clutter measurements are modelled a PPP with Poisson rate $\lambda^c(z_k)$ and spatial distribution $c(z_k)$, and the clutter PPP intensity is $\kappa(z_k)$. Clutter measurements are independently of any targets and target-generated measurements.*
- *Each target-generated measurement is only conditioned on its corresponding target. The single target measurement likelihood is $\ell(z_k|x_k)$.*

Expressions for recursively predicting and updating the multi-trajectory filtering densities based on the above modelling assumptions can be found in [14]. Following [15], we recursively express these densities as PMBMs.

B. Densities of sets of trajectories

The basic building blocks in a PMBM density are the PPP and the Bernoulli process. Let \mathbf{X} denote the set of trajectories, and let $f(\mathbf{X})$ denote the corresponding density function. A trajectory PPP is analogous to a target PPP, and has set density

$$f^{\text{PPP}}(\mathbf{X}) = e^{-\langle \lambda^u, 1 \rangle} \prod_{X \in \mathbf{X}} \lambda(X), \quad (1)$$

where $\langle a, b \rangle$ denotes $\int a(x)b(x)dx$. The trajectory PPP intensity $\lambda(\cdot)$ is defined on the trajectory state space, i.e., realizations of the PPP are trajectories with a birth time, a time of the most recent state, and a trajectory.

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

$$f^{\text{ber}}(\mathbf{X}) = \begin{cases} 1 - r & \mathbf{X} = \emptyset \\ r \cdot f(X) & \mathbf{X} = \{X\} \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where $f(X)$ is a trajectory state density, 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 multi-Bernoulli (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 multiple trajectory Bernoulli RFS; trajectory MBM RFS is an RFS whose density is a mixture of trajectory MB densities.

C. PMBM trajectory filter

1) *PMBM conjugate prior on the set of trajectories:* The PMBM conjugate prior was developed for point targets in [10] and extended to trajectories in [15]. In this context, conjugacy means that the family of PMBM densities is closed under the prediction and update steps, though the number of terms in the mixture grows rapidly with time. In the trajectory PMBM form, the trajectory set \mathbf{X} is a union of two disjoint sets \mathbf{X}^u and \mathbf{X}^d , i.e., $\mathbf{X} = \mathbf{X}^u \uplus \mathbf{X}^d$. The trajectories in the set \mathbf{X}^u are hypothesised to exist but to be undetected, and \mathbf{X}^u is presented by a PPP. The set of detected trajectories \mathbf{X}^d is presented by an MBM. The PMBM density on the set of trajectories can be expressed as

$$f(\mathbf{X}) = \sum_{\mathbf{X}^u \uplus \mathbf{X}^d = \mathbf{X}} f^{\text{PPP}}(\mathbf{X}^u) \sum_{j \in \mathbb{J}} W^j f^j(\mathbf{X}^d), \quad (3a)$$

$$f^{\text{PPP}}(\mathbf{X}^u) = e^{-\langle \lambda^u, 1 \rangle} \prod_{X \in \mathbf{X}^u} \lambda^u(X), \quad (3b)$$

$$f^j(\mathbf{X}^d) = \sum_{\uplus_{i \in \mathbb{J}} \mathbf{X}^i = \mathbf{X}^d} \prod_{i \in \mathbb{J}} f^{j,i}(\mathbf{X}^i), \quad (3c)$$

where $f^{j,i}(\mathbf{X}^i)$ denotes the i th Bernoulli component in the j th MB $f^j(\mathbf{X}^d)$ of the MBM, and W^j is the weight of the j th MB.

2) *Structure of the hypotheses:* The structure of the hypotheses in the PMBM trajectory filter is the same as the structure of the hypotheses in the PMBM target filter [10]. In the PMBM trajectory filter, each MB corresponds to a unique *global hypothesis* for the detected trajectories. A global hypothesis is a partitioning of all measurements received so far into subsets, where each subset is hypothesised to correspond to a *single trajectory hypothesis*. A *single trajectory hypothesis* is a subset of measurements, at most one from each time scan, that are hypothesised to correspond to the same target. A *track* is a collection of single trajectory hypotheses, representing different possibilities of measurement sequences corresponding to the trajectory. We initiate one new track per measurement, and all single trajectory hypotheses initiated by the same measurement belong to the same track. Each global hypothesis must incorporate one single trajectory hypothesis from each track.

Let m_k be the number of measurements at time $k \in \{1, \dots, \tau\}$ and $j \in \mathbb{M}_k = \{1, \dots, m_k\}$ be an index to each measurement with value $z_k^j \in Z_k$. Let \mathcal{M}_k denote the set of all measurement indices up to and including time step k ; the elements of \mathcal{M}_k are of the form (k', j) , where $j \in \{1, \dots, m_{k'}\}$ is an index of a measurement in scan $k' \leq k$. Let the number of tracks at time k be $n_{k|k}$ and $i \in \mathbb{X}_{k|k'} = \{1, \dots, n_{k|k'}\}$ be an index to each track, where $k' = k - 1$ corresponds to the prediction step and $k' = k$ corresponds to the update step. A global hypothesis at time k can be represented as $A_{k|k'} = (\mathbf{a}_{k|k'}^1, \dots, \mathbf{a}_{k|k'}^{n_{k|k'}})$, where $\mathbf{a}_{k|k'}^i$ denotes the single trajectory hypothesis utilized for the i th track. The set of single trajectory hypotheses for the i th track at time k is denoted as $\mathcal{A}_{k|k'}^i$. The set of measurement indices under single trajectory hypothesis $\mathbf{a}_{k|k'}^i$ is represented as $\mathcal{M}(\mathbf{a}_{k|k'}^i) \subseteq \mathcal{M}_k$.

In the PMBM trajectory filter, each global hypothesis should explain the association of each measurement received so far. In addition, every measurement should be associated to one and only one (pre-existing or new) track in each global hypothesis. In other words, the single trajectory hypotheses included in a given global hypothesis cannot have any shared measurement. Under these constraints, the set of global hypotheses at time scan k can be expressed as [10]

$$\begin{aligned} \mathfrak{A}_{k|k'} &= \left\{ (\mathbf{a}_{k|k'}^1, \dots, \mathbf{a}_{k|k'}^{n_{k|k'}}) \mid \mathbf{a}_{k|k'}^i \in \mathcal{A}_{k|k'}^i, \bigcup_{i \in \mathbb{X}_{k|k'}} \mathcal{M}(\mathbf{a}_{k|k'}^i) \right. \\ &= \left. \mathcal{M}_k, \mathcal{M}(\mathbf{a}_{k|k'}^i) \cap \mathcal{M}(\mathbf{a}_{k|k'}^j) = \emptyset \forall i \neq j \right\}. \quad (4) \end{aligned}$$

3) *Filtering recursion:* The form of the PMBM conjugate prior on the sets of trajectories is maintained by prediction and update steps. Given the sequences of measurements up to time step k , our objective is to calculate the PMBM multi-trajectory filtering density at time k recursively. In the prediction step, the MBs describing pre-existing tracks and the PPP describing undetected trajectories are predicted independently. By using a PPP birth model, the density of new born trajectories can be easily incorporated into the predicted density of undetected trajectories [10]. In the update step, the PPP and the MBs

are updated independently. The PPP intensity of undetected trajectories is thinned by the missed detection probability, i.e., $1 - p^d(x_k)$. According to the point target PMBM modelling assumptions, each measurement creates a new track; thus, the number of tracks after updating becomes $n_{k|k} = n_{k|k-1} + m_k$.

In what follows, we only present part of the update equations that are important to explain our proposed implementation due to page limits. We refer the reader to [15] for more details.

Let $w_{k|k'}^{\mathbf{a}^i}$ denote the weight of single trajectory hypothesis $\mathbf{a}_{k|k'}^i$, and let the corresponding Bernoulli density denote as $f_{k|k'}^{\mathbf{a}^i}(\mathbf{X})$. For pre-existing tracks ($i \in \{1, \dots, n_{k|k-1}\}$), each single trajectory hypothesis creates $1 + m_k$ new single trajectory hypotheses, one for missed detection,

$$\mathcal{M}_k(\mathbf{a}^i) = \mathcal{M}_{k-1}(\mathbf{a}^i), \quad (5a)$$

$$w_{k|k}^{\mathbf{a}^i} = w_{k|k-1}^{\mathbf{a}^i} \mathcal{L}_k^{\emptyset, \mathbf{a}^i}, \quad (5b)$$

$$\mathcal{L}_k^{\emptyset, \mathbf{a}^i} = (1 - r_{k|k-1}^{\mathbf{a}^i} + r_{k|k-1}^{\mathbf{a}^i} \langle f_{k|k-1}^{\mathbf{a}^i}, 1 - p^d \rangle), \quad (5c)$$

where $\mathcal{L}_k^{\emptyset, \mathbf{a}^i}$ is the likelihood that the potential trajectory represented by \mathbf{a}^i is missed detected, and the others for measurement updates ($j \in \mathbb{M}_k$),

$$\mathcal{M}_k(\mathbf{a}^i) = \mathcal{M}_{k-1}(\mathbf{a}^i) \cup \{(k, j)\}, \quad (6a)$$

$$w_{k|k}^{\mathbf{a}^i} = w_{k|k-1}^{\mathbf{a}^i} \mathcal{L}_k^{j, \mathbf{a}^i}, \quad (6b)$$

$$\mathcal{L}_k^{j, \mathbf{a}^i} = r_{k|k-1}^{\mathbf{a}^i} \langle f_{k|k-1}^{\mathbf{a}^i}, \ell(z_k^j | \cdot) p^d \rangle, \quad (6c)$$

where $\mathcal{L}_k^{j, \mathbf{a}^i}$ is the likelihood that single trajectory hypothesis \mathbf{a}^i is updated by the j th measurement from time k .

Each new track ($i \in \{n_{k|k-1} + 1, \dots, n_{k|k-1} + m_k\}$) contains two single trajectory hypotheses. The first one corresponds to a Bernoulli density with zero existence probability, and it covers the case that the new measurement is associated with one of the pre-existing tracks,

$$\mathcal{M}_k(\mathbf{a}^i) = \emptyset, \quad (7a)$$

$$w_{k|k}^{\mathbf{a}^i} = 1. \quad (7b)$$

The second one results from updating the PPP of undetected trajectories with a new measurement, which can be either a false alarm or the first detection of an undetected trajectory ($j \in \mathbb{M}_k$),

$$\mathcal{M}_k(\mathbf{a}^i) = \{(k, j)\}, \quad (8a)$$

$$w_{k|k}^{\mathbf{a}^i} = \mathcal{L}_k^{u, j, \mathbf{a}^i} = \kappa(z_k^j) + \langle D_k^u, \ell(z_k^j | \cdot) p^d \rangle, \quad (8b)$$

where $\mathcal{L}_k^{u, j, \mathbf{a}^i}$ is the likelihood that the j th measurement from time k is assigned to an undetected trajectory.

III. IMPLEMENTATION OF THE PMBM TRAJECTORY FILTER VIA MULTI-FRAME ASSIGNMENT

In this section, we present an efficient implementation of the PMBM trajectory filter that estimates the trajectories of the targets who are present in the surveillance area at the current time. The proposed implementation uses track-oriented N -scan pruning to reduce the number of global hypotheses, and

dual decomposition to obtain the most likely global hypothesis by solving a multi-frame assignment problem. In addition, we discuss the connections and differences between the existing PMBM target filter and the proposed PMBM trajectory filter.

A. Global hypothesis reduction

In the update step of the PMBM trajectory filter, each possible data association will create an updated global hypothesis so that the number of MBs in the PMBM posterior density (3) will increase exponentially over time. Hence, finding a feasible approach to reduce the number of global hypotheses after the update step is essential for designing a computationally tractable PMBM trajectory filter.

In the PMBM trajectory filter, the posterior over the set of detected trajectories is an MBM, i.e., a weighted mixture of global hypotheses. Typically, the estimates of trajectories are extracted from the most likely global hypothesis. Conditioning on the most likely global hypothesis, we make use of the track-oriented N -scan pruning [5], a conventional hypothesis reduction technique used in MHT, to prune global hypotheses with negligible weights.

Given the most likely global hypothesis A^* at current time scan τ , we trace the single trajectory hypotheses included in A^* back to their local hypotheses at scan $\tau - N$. The assumption behind the N -scan pruning method is that the data association ambiguity is resolved before scan $\tau - N$ [5]. In other words, global hypotheses which do not have the same local hypotheses at scan $\tau - N - 1$ as A^* are assumed to have negligible weights; these global hypotheses can then be pruned. In what follows, we show that the most likely global hypothesis A^* can be obtained as the solution of a multi-frame assignment problem.

B. Data association modelling and problem formulation

The posterior global hypothesis weight $W_{k|k}^A$ is proportional to the product of different weights of the single trajectory hypotheses $w^{\mathbf{a}^i|k}$, one from each track [10]:

$$W_{k|k}^A \propto \prod_{i \in \mathbb{X}_{k|k}} w^{\mathbf{a}^i|k}, \quad (9)$$

where the proportionality denotes that normalization is required to ensure that $\sum_{A_{k|k} \in \mathfrak{A}_{k|k}} W_{k|k}^A = 1$. Omitting time indices and introducing the notation $c^A = -\log(W^A)$ and $c(\mathbf{a}^i) = -\log(w^{\mathbf{a}^i})$, yields

$$c^A = \sum_{i \in \mathbb{X}} c(\mathbf{a}^i) + C, \quad (10)$$

where C is the logarithm of the normalization constant in (9). The most likely global hypothesis is the collection of single trajectory hypotheses that minimizes the total cost, i.e.,

$$\arg \min_{A^* = (\mathbf{a}^1, \dots, \mathbf{a}^n) \in \mathfrak{A}} \sum_{i \in \mathbb{X}} c(\mathbf{a}^i). \quad (11)$$

The minimization problem (11) can be further posed as a multi-frame assignment problem by decomposing the con-

straint $A = (\mathbf{a}^1, \dots, \mathbf{a}^n) \in \mathfrak{A}$ into a set of smaller constraints, as first stated in [16, Section III], in the form of

$$\arg \min_{\rho \in \bigcup_{k=0}^{\tau} \mathcal{P}^k} \sum_{i \in \mathbb{X}} \sum_{\mathbf{a}^i \in \mathcal{A}^i} c(\mathbf{a}^i) \rho(\mathbf{a}^i), \quad (12)$$

with the constraints sets denoted as

$$\mathcal{P}^0 = \left\{ \rho \mid \sum_{\mathbf{a}^i \in \mathcal{A}^i} \rho(\mathbf{a}^i) = 1, \forall i \in \mathbb{X} \right\}, \quad (13)$$

$$\mathcal{P}^k = \left\{ \rho \mid \sum_{i \in \mathbb{X}} \sum_{\substack{\mathbf{a}^i \in \mathcal{A}^i: \\ (k,j) \in \mathcal{M}(\mathbf{a}^i)}} \rho(\mathbf{a}^i) = 1, \forall j \in \mathbb{M}_k \right\}, \quad (14)$$

where $k = 1, \dots, \tau$, $\rho(\mathbf{a}^i) \in \{0, 1\}$ is a binary indicator variable, indicating whether single trajectory hypothesis \mathbf{a}^i in the i th track is included in a global hypothesis or not, and $\rho = \{\rho(\mathbf{a}^i) \in \{0, 1\} \mid \mathbf{a}^i \in \mathcal{A}^i \forall i \in \mathbb{X}\}$ is the set of all binary indicator variables. The first constraint (13) enforces that each global hypothesis should include one and only one single trajectory hypothesis from each track. The set of τ constraints (14) enforce that each measurement from each scan should be associated to one and only one track.

C. Multi-frame assignment via dual decomposition

The multi-dimensional assignment problem (12) is NP-hard for two or more scans of measurements. An effective solution to this is to use Lagrangian relaxation; this technique has been widely used to solve the multi-scan data association problem in the MHT framework [20], [21]. In this work, we focus on the dual decomposition formulation [22], i.e., a special case of Lagrangian relaxation, which has been successfully applied to solve the multi-frame assignment problem [16].

1) *Decomposition of Lagrangian dual*: We follow similar implementation steps as in [16]. The original (primal) problem (12) is separated into τ subproblems, one for each scan of measurements, and for each subproblem a binary variable ρ^k is used. The constraint that ρ^k are equal for each time scan k is enforced through Lagrange multipliers that are incorporated into the subproblems to act as penalty weights. The k th subproblem can be written as [16]

$$\min_{\rho^k \in \{\mathcal{P}^0, \mathcal{P}^k\}} \sum_{i \in \mathbb{X}} \sum_{\mathbf{a}^i \in \mathcal{A}^i} \left(\frac{c(\mathbf{a}^i)}{\tau} + \delta^k(\mathbf{a}^i) \right) \rho^k(\mathbf{a}^i) \triangleq \min_{\rho^k \in \{\mathcal{P}^0, \mathcal{P}^k\}} \mathcal{S}(\rho^k, \delta^k), \quad (15)$$

where the binary indicator variables and Lagrange multipliers used for the k th subproblem are denoted, respectively, by

$$\rho^k = \{\rho^k(\mathbf{a}^i) \in \{0, 1\} \mid \mathbf{a}^i \in \mathcal{A}^i \forall i \in \mathbb{X}\}, \quad (16a)$$

$$\delta^k = \{\delta^k(\mathbf{a}^i) \mid \mathbf{a}^i \in \mathcal{A}^i \forall i \in \mathbb{X}\}. \quad (16b)$$

The Lagrange multipliers $\delta^k(\mathbf{a}^i)$ can be any real number, but with the constraint that, for each single trajectory hypothesis, they must add up to zero over different subproblems [22]. Thus, the set of Lagrange multipliers has the form

$$\Lambda = \{\delta^k \mid \sum_k \delta^k(\mathbf{a}^i) = 0, \forall \mathbf{a}^i \in \mathcal{A}^i \forall i \in \mathbb{X}\}. \quad (17)$$

2) *Subproblem solving*: After eliminating all the constraints sets except two, i.e., \mathcal{P}^0 and \mathcal{P}^k , we obtain a 2-D assignment problem. For the k th assignment problem, the objective is to associate each measurement received at time scan $k \leq \tau$, i.e., $j \in \mathbb{M}_k$, either to a pre-existing track or a new track at current time scan τ , i.e., $i \in \mathbb{X}_\tau$, such that the total assignment cost is minimized. Problems of this type can be solved in polynomial time using a modified auction algorithm [1, Chapter VII].

For a track that is created after time scan k , no measurement from time scan k should be assigned to it; therefore, the measurement-to-track assignment cost is infinity. For a track that existed before and up to time scan k , i.e., $i \in \mathbb{X}_k$, if measurement \mathbf{z}_k^j has never been associated to this track, let the measurement-to-track assignment cost be infinity; if otherwise, let the cost first be the minimum cost of the single trajectory hypotheses in this track that were updated by \mathbf{z}_k^j [1, Chapter VII, Equation (7.24)], i.e.,

$$\min \sum_{\substack{\mathbf{a}^i \in \mathcal{A}^i: \\ (k,j) \in \mathcal{M}(\mathbf{a}^i)}} \left(\frac{c(\mathbf{a}^i)}{\tau} + \delta^k(\mathbf{a}^i) \right). \quad (18)$$

In order to keep the cost of a hypothesis that does not assign a measurement to a track the same for a pre-existing and a new track, the cost (18) should then be subtracted by the minimum cost of hypotheses that this track is not updated by any of the measurements at time scan k , i.e.,

$$\min \sum_{\substack{\mathbf{a}^i \in \mathcal{A}^i: \\ (k,j) \notin \mathcal{M}(\mathbf{a}^i), \forall j \in \mathbb{M}_k}} \left(\frac{c(\mathbf{a}^i)}{\tau} + \delta^k(\mathbf{a}^i) \right). \quad (19)$$

Note that, in the context of Lagrangian relaxation, the costs of single trajectory hypotheses refer to the costs that are penalized by the Lagrangian multipliers.

After solving this 2-D assignment problem, we can obtain the associations for each measurement at time scan k . For tracks not being associated to any measurements at time scan k , if the track is created before and up to time scan k , i.e., $i \in \mathbb{X}_k$, the single trajectory hypothesis

$$\arg \min_{\mathbf{a}^i} \sum_{\substack{\mathbf{a}^i \in \mathcal{A}^i: \\ (k,j) \notin \mathcal{M}(\mathbf{a}^i), \forall j \in \mathbb{M}_k}} \left(\frac{c(\mathbf{a}^i)}{\tau} + \delta^k(\mathbf{a}^i) \right) \quad (20)$$

is included in the most likely global hypothesis; if otherwise, i.e., $i \in \mathbb{X}_\tau \setminus \mathbb{X}_k$, we can simply choose the single trajectory hypothesis

$$\arg \min_{\mathbf{a}^i} \sum_{\mathbf{a}^i \in \mathcal{A}^i} \left(\frac{c(\mathbf{a}^i)}{\tau} + \delta^k(\mathbf{a}^i) \right) \quad (21)$$

to be included in the most likely global hypothesis.

3) *Subgradient updates*: The objective of Lagrange relaxation is to find the tightest lower bound of the summation of the cost of each subproblem (15). The dual problem can be expressed as [16]

$$\max_{\{\delta^k\} \in \Lambda} \left(\sum_{k=1}^{\tau} \min_{\rho^k \in \{\mathcal{P}^0, \mathcal{P}^k\}} \mathcal{S}(\rho^k, \delta^k) \right), \quad (22)$$

where the maximum can be found using subgradient methods [23]. The Lagrange multipliers $\{\delta^k\}$ are updated using

$$\delta^k(\mathbf{a}^i) = \delta^k(\mathbf{a}^i) + \alpha_t \cdot g^k(\mathbf{a}^i), \quad (23)$$

where $g^k(\mathbf{a}^i)$ is the projected subgradient that can be calculated as

$$g^k(\mathbf{a}^i) = \rho^k(\mathbf{a}^i) - \frac{1}{\tau} \sum_{k'=1}^{\tau} \rho^{k'}(\mathbf{a}^i), \quad (24)$$

and α_t is the step size at iteration t . There are many rules to set the step size, see [22]. In this work, we choose to use the same setting as in [16], which has the form

$$\alpha_t = \frac{\text{BESTPRIMAL}_t - \text{DUAL}_t}{\|\{\mathbf{g}^k\}\|^2}, \quad (25)$$

where BESTPRIMAL_t is the best (minimum) feasible primal cost so far obtained, and DUAL_t is the dual cost calculated at iteration t from (22). The optimal solution is assumed to be attained when the relative gap between the primal cost and the dual cost is less than a specified threshold ϵ [22], where

$$\text{gap} = \frac{\text{BESTPRIMAL}_t - \text{DUAL}_t}{\text{BESTPRIMAL}_t}. \quad (26)$$

Each subproblem solution will, in general, be infeasible with respect to the primal problem (12); nevertheless, subproblem solutions will usually be nearly feasible since large constraints violations got penalized [22]. Hence, feasible solutions ρ can be obtained by correcting the minor conflicting binary elements on which subproblem solutions ρ^k disagree. For tracks for which we have not yet selected which single trajectory hypothesis to included in the most likely global hypothesis, we use the branch and bound technique to reconstruct the best feasible solution at each iteration of the Lagrange relaxation. Note that there are many other ways to recover a feasible primal solution from subproblem solutions, see [22].

D. Discussion

As discussed in [10], the structure of the data association hypotheses in the PMBM density is similar to that of the track-oriented MHT. Benefiting from the separability of the global hypothesis weight (9), in this work, we make use of the same approximation technique in the track-oriented MHT to find the most likely global hypothesis by solving a multi-frame assignment problem. It should be noted that, the approximated filtering density describing the set of detected trajectories in our proposed implementation, strictly speaking, is not an MBM but a set of Bernoulli components. The objective of solving the multi-frame assignment problem is to know which Bernoulli components are included in the MB with the highest weight. At each update step, Bernoulli components forming the highest weight MB, together with the other Bernoulli components that survive from N -scan pruning, are carried over to the next filtering recursion. Because the data association ambiguity is assumed to be resolved before time $\tau - N$, obtaining the most likely global hypothesis at time τ , which explains the origin of each measurement from scan $\tau - N$ to

current scan τ , requires the solution of a $N + 2$ dimensional assignment problem [5].

In contrast, the existing PMBM target filter in [10] operates by propagating a number of global hypotheses with explicitly calculated weights over time. The number of global hypotheses grows exponentially over time due to data association; thus it is intractable to enumerate all the possible global hypotheses. In order to have a tractable solution, each predicted global hypothesis only generates a set of most likely updated global hypotheses by approximately selecting the M most likely data associations, using, e.g., Murty’s algorithm [24] or Gibbs sampling [18]. Further, updated global hypotheses with negligible normalized weights are assumed to be unlikely to generate any high-weight global hypothesis in subsequent time steps; thus they can be pruned. Only maintaining global hypotheses with non-negligible weights means that we prune single target hypotheses that are not included in any of the remaining global hypotheses. One can relate this pruning procedure to the N-scan pruning used in the PMBM trajectory filter, in which the same intuition applies: single trajectory hypotheses are pruned if their updates are assumed to be unlikely to be included in the most likely global hypothesis in subsequent time steps.

The computational complexity of both filters can be further reduced by limiting the number of single target/trajectory hypotheses. For the PMBM target filter, we can prune single target hypothesis with small enough Bernoulli existence probability. A more efficient approximation method is called recycling [25], in which Bernoulli components are first approximated as being Poisson and then added to the PPP density representing undetected targets. The benefits of applying recycling in the PMBM target filter has been demonstrated in [13]. As for the PMBM trajectory filter, pruning single trajectory hypotheses will sometimes harm the solvability of the multi-frame assignment problem, since the problem is formulated using the measurement assignment information contained in single trajectory hypotheses. Instead, we can choose single trajectory hypotheses $\mathbf{a}^i \in \mathcal{A}^i, \forall i \in \mathbb{X}$ with small Bernoulli existence probability $r^{\mathbf{a}^i}$ at time step k to be updated only by missed detection at time step $k + 1$.

IV. SIMULATIONS AND RESULTS

In this section, we show simulation results that compare the proposed PMBM trajectory filter with the PMBM target filter [12], and the δ -GLMB filter with joint prediction and update steps [18]. For the δ -GLMB filter and the PMBM target filter, the M most likely data associations are approximately found using the Gibbs sampling solution proposed in [18]. In the implementation, all the codes are written in MATLAB, except the Gibbs sampling and the auction algorithm, which are written in C++.

A. State space model

A two-dimensional Cartesian coordinate system is used to define measurement and target kinematic parameters. The kinematic target state is a vector of position and velocity $x_k = [p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k}]^T$. A single measurement is a




Fig. 1. Scenario of simulations. There are six targets born at times $\{1,11,21,31,41,51\}$ and dead at times $\{61,71,81,91,101,101\}$ respectively. Targets move from left to right, and they are in close proximity around the mid-point.

vector of position $z_k = [z_{x,k}, z_{y,k}]^T$. Targets follow a linear Gaussian constant velocity model $f_{k|k-1}(x_k|x_{k-1}) = \mathcal{N}(x_k; F_k x_{k-1}, Q_k)$, with parameters

$$F_k = I_2 \otimes \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix}, \quad Q_k = 0.002I_2 \otimes \begin{bmatrix} T^3/3 & T^2/2 \\ T^2/2 & T \end{bmatrix},$$

where \otimes is the Kronecker product, I_m is an identity matrix of size $m \times m$, and $T = 1$. The linear Gaussian measurement likelihood model has density $g(z_k|x_k) = \mathcal{N}(z_k; H_k x_k, R_k)$, with parameters $H_k = I_2 \otimes [1, 0]$ and $R_k = I_2$.

B. Parameter setting

We consider 101 time steps and the scenario in Figure 1. For each trajectory, we initiate the midpoint from a Gaussian with mean $[0, 0, 0, 0]^T$ and covariance matrix $10^{-6}I_4$, and the rest of the trajectory is generated by running forward and backwards dynamics. This scenario is challenging due to the high number of targets in close proximity and the fact that one target is born when and where the other five are in close proximity. In the simulation, we consider cases with constant target survival probability $p^s = 0.99$, constant detection probability $p^d \in \{0.7, 0.9\}$, and Poisson clutter uniform in the region of interest with rate $\lambda^c \in \{10, 30\}$.

For the PMBM target/trajectory filter, targets are born according to a PPP with intensity 0.05 and Gaussian density with mean $[0, 0, 0, 0]^T$ and covariance matrix $\text{diag}([100^2, 1, 100^2, 1])$, which covers the region of interest. For the PPP part, components are pruned with weight smaller than 10^{-4} . For the PMBM target filter, Bernoulli components with existence probability less than 0.1 are recycled. For the δ -GLMB filter, targets are born according to a Bernoulli process with existence probability 0.05 and the same Gaussian density as the PPP birth model. For the PMBM target filter and the δ -GLMB filter, for a global hypothesis with weight \mathcal{W} , we use Gibbs sampling to select at most $\lceil \mathcal{W} \cdot N_h \rceil$ updated global hypotheses with the highest weights. In addition, global hypotheses are truncated to only contain those that correspond to 99.99% of the likelihood. For all the filters, gating is performed to remove unlikely target-to-measurement association.

An implementation of the PMBM trajectory filter propagating the whole trajectory (full state sequence) using information Gaussian representation has been provided in [15]. A more computational efficient implementation is to consider only the target states of the last L steps [26]. In the simulation, we compare three filters: the δ -GLMB filter, which provides filtered estimates; a PMBM trajectory filter that provides filtered estimates ($L = 1$); and a PMBM trajectory filter that provides smoothed estimates (L is equal trajectory length). Smoothed estimates can be easily obtained using Rauch-Tung-Striebel smoother, since the measurement assignment information is contained in single trajectory hypotheses. As for the δ -GLMB filter, a closed form for forward-backward smoothing was developed in [27] but the implementation has not yet been proposed.

C. Performance evaluation

Given a multi-target posterior density, several state estimators are available [11]. In this work, we choose to extract the target states by finding the maximum a posteriori cardinality estimate, following the methods suggested in [11], [12]. Filtering performance is assessed using the generalized optimal subpattern (GOSPA) metric [28] ($\alpha = 2, c = 20, p = 1$), which allows for decomposing the estimation error into location error, missed detection error and false detection error.

The δ -GLMB filter provides estimates of a target trajectory by estimating the current state at each time, and connecting estimates from different times with the same label. For the PMBM trajectory filter, we can choose to either provide the trajectory estimation directly from single trajectory hypothesis density [15] or connect current state estimates which belong to the same track. In the simulation, the latter one is used. For assessing the estimation performance on the sets of trajectories for the δ -GLMB filter and the PMBM trajectory filter, we use the metric proposed in [29] ($c = 20, p = 1, \gamma = 2$), which allows for decomposing the estimation error into location error, missed detection error, false detection error and error due to track switches.

D. Results

The results are obtained over 100 Monte Carlo trials. The average localization error per target, missed and false detection error, measured using GOSPA, of different filters are listed in Table I. The performance comparison regarding computational time and GOSPA error among different filters is presented in Fig. 2. Overall, the PMBM target filter has the best filtering performance, followed by the PMBM trajectory filter, with the δ -GLMB filter in last.

Table II presents the trajectory estimation error¹ of the δ -GLMB filter and the PMBM trajectory filter. By only looking at the trajectory estimation results without smoothing, the PMBM trajectory filter has significantly smaller missed detection error, similar track switch error, and slightly larger localization and false detection error than the δ -GLMB filter.

¹We compared all the trajectories that had existed in the simulation to the ground truth (six true trajectories).

In the simulation, we found that, for the δ -GLMB, missed detection frequently occurs when targets become in close proximity, which leads to the result that the trajectories obtained using the δ -GLMB filter are usually shorter than those obtained by the PMBM filter. This explains the results since a shorter trajectory implies lower error due to track switches, localization mismatch and possible false detections.

V. CONCLUSION

This paper has proposed an efficient implementation of the PMBM trajectory filter. Compared with the existing PMBM target filter, the PMBM trajectory filter is able to estimate all target trajectories. Also, a simulation study shows that the PMBM trajectory filter has better performance than the δ -GLMB filter in terms of state/trajectory estimation error and computational time in a scenario with coalescence.

REFERENCES

- [1] Y. Bar-Shalom, "Multitarget-multisensor tracking: advanced applications," *Norwood, MA, Artech House, 1990, 391 p.*, 1990.
- [2] T. Fortmann, Y. Bar-Shalom, and M. Scheffe, "Sonar tracking of multiple targets using joint probabilistic data association," *IEEE J. Ocean. Eng.*, vol. 8, no. 3, pp. 173–184, 1983.
- [3] D. Reid, "An algorithm for tracking multiple targets," *IEEE Trans. Autom. Control*, vol. 24, no. 6, pp. 843–854, 1979.
- [4] T. Kuiren, "Issues in the design of practical multitarget tracking algorithms," *Multitarget-multisensor tracking: advanced applications*, pp. 43–87, 1990.
- [5] S. S. Blackman, "Multiple hypothesis tracking for multiple target tracking," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 19, no. 1, pp. 5–18, 2004.
- [6] R. P. Mahler, *Statistical multisource-multitarget information fusion*. Artech House, Inc., 2007.
- [7] —, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, 2003.
- [8] R. Mahler, "PHD filters of higher order in target number," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 43, no. 4, 2007.
- [9] 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.
- [10] 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, pp. 1664–1687, 2015.
- [11] B.-N. Vo, B.-T. Vo, and D. Phung, "Labeled random finite sets and the bayes multi-target tracking filter," *IEEE Trans. Signal Process.*, vol. 62, no. 24, pp. 6554–6567, 2014.
- [12] Á. F. García-Fernández, J. L. Williams, K. Granström, and L. Svensson, "Poisson multi-bernoulli mixture filter: direct derivation and implementation," *IEEE Trans. Aerosp. Electron. Syst.*, 2018, to be published.
- [13] Y. Xia, K. Granström, L. Svensson, and Á. F. García-Fernández, "Performance evaluation of multi-Bernoulli conjugate priors for multi-target filtering," in *Proc. 2017 Information Fusion Conf*, 2017, pp. 1–8.
- [14] Á. F. García-Fernández, L. Svensson, and M. R. Morelande, "Multiple target tracking based on sets of trajectories," *arXiv preprint arXiv:1605.08163*, 2016.
- [15] K. Granström, J. L. Williams, Á. F. García-Fernández, L. Svensson, and Y. Xia, "Poisson multi-Bernoulli mixture trackers: continuity through random finite sets of trajectories," submitted to 2018 Information Fusion Conference for publication.
- [16] R. A. Lau and J. L. Williams, "Multidimensional assignment by dual decomposition," in *Proc. 2011 ISSNIP Conf*, 2011, pp. 437–442.
- [17] S. Mori, C.-Y. Chong, and K.-C. Chang, "Three formalisms of multiple hypothesis tracking," in *Proc. 2016 Information Fusion Conf*, 2016, pp. 727–734.
- [18] B.-N. Vo, B.-T. Vo, and H. G. Hoang, "An efficient implementation of the generalized labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, vol. 65, no. 8, pp. 1975–1987, 2017.

TABLE I

SIMULATION RESULTS: THE AVERAGE STATE ESTIMATION ERROR PER TIME STEP. LEGEND: TE– TOTAL GOSPA ERROR; LE–AVERAGE LOCATION ERROR PER TARGET; MT–MISSED TARGETS; FT–FALSE TARGETS; N_h : THE MAXIMUM NUMBER OF UPDATED GLOBAL HYPOTHESES PER PREDICTED GLOBAL HYPOTHESIS; N: N-SCAN PRUNING; GLMB: δ -GLMB FILTER WITH JOINT PREDICTION AND UPDATE STEPS; PMMT: PMBM TARGET FILTER; PMDD: PMBM TRAJECTORY FILTER

Filter	$p^d = 0.9, \lambda^c = 10$				$p^d = 0.9, \lambda^c = 30$				$p^d = 0.7, \lambda^c = 10$				$p^d = 0.7, \lambda^c = 30$			
	TE	LE	MT	FT	TE	LE	MT	FT	TE	LE	MT	FT	TE	LE	MT	FT
GLMB($N_h = 300$)	7.36	0.81	4.40	0.34	10.73	0.77	8.12	0.36	11.33	0.86	8.27	0.61	16.65	0.82	14.21	0.49
GLMB($N_h = 200$)	7.97	0.80	5.08	0.34	11.89	0.76	9.39	0.35	12.12	0.86	9.16	0.57	18.50	0.79	16.36	0.42
GLMB($N_h = 100$)	9.04	0.79	6.21	0.38	15.47	0.74	13.35	0.28	13.27	0.83	10.37	0.62	21.86	0.76	20.16	0.31
PMMT($N_h = 200$)	6.07	0.79	2.75	0.52	6.33	0.80	3.02	0.54	8.79	0.88	5.19	0.78	9.75	0.87	6.19	0.82
PMMT($N_h = 100$)	6.40	0.76	3.12	0.53	6.37	0.79	3.06	0.55	8.85	0.88	5.24	0.79	9.80	0.88	6.24	0.80
PMMT($N_h = 50$)	6.71	0.74	3.45	0.54	6.48	0.78	3.19	0.54	8.97	0.86	5.38	0.78	9.87	0.87	6.32	0.81
PMDD($N = 5$)	5.94	0.83	2.63	0.42	6.60	0.82	3.35	0.45	9.18	0.90	5.57	0.78	10.65	0.89	7.20	0.78
PMDD($N = 4$)	5.98	0.83	2.69	0.40	6.63	0.82	3.41	0.44	9.26	0.90	5.67	0.76	10.91	0.89	7.52	0.73
PMDD($N = 3$)	6.06	0.83	2.81	0.39	6.80	0.82	3.62	0.42	9.46	0.91	5.91	0.72	11.51	0.89	8.28	0.64
PMDD($N = 2$)	6.50	0.82	3.41	0.34	7.44	0.81	4.39	0.38	10.54	0.91	7.25	0.61	13.06	0.89	10.10	0.53




Fig. 2. Performance comparison among the δ -GLMB filter (triangle), the PMBM target filter (circular) and the PMBM trajectory filter (rectangle): simulation time versus GOSPA error (in logarithm). Plots from left to right, respectively, correspond to scenario parameter setting: 1) $p^d = 0.9, \lambda^c = 10$, 2) $p^d = 0.9, \lambda^c = 30$, 3) $p^d = 0.7, \lambda^c = 10$ and 4) $p^d = 0.7, \lambda^c = 30$. For the same filter in each subplot, the scatter points from left to right marked with the same color, respectively, correspond to $N_h = \{100, 200, 300\}$ for the δ -GLMB filter, $N_h = \{50, 100, 200\}$ for the PMBM target filter, and $N = \{2, 3, 4, 5\}$ for the PMBM trajectory filter.

TABLE II

SIMULATION RESULTS: THE SUM OF TRAJECTORY ESTIMATION ERROR. LEGEND: TE–TOTAL ERROR; LE–LOCATION ERROR; ME–MISSED DETECTION; FE–FALSE DETECTION; TS–TRACK SWITCH ERROR; PMDDS: PMBM TRAJECTORY FILTER WITH SMOOTHING

Filter	$p^d = 0.9, \lambda^c = 10$					$p^d = 0.9, \lambda^c = 30$					$p^d = 0.7, \lambda^c = 10$					$p^d = 0.7, \lambda^c = 30$				
	TE	LE	ME	FE	TS	TE	LE	ME	FE	TS	TE	LE	ME	FE	TS	TE	LE	ME	FE	TS
GLMB($N_h = 300$)	857	348	440	34	35	1176	299	812	36	29	1242	327	826	60	29	1759	269	1420	42	28
GLMB($N_h = 200$)	918	340	508	35	35	1287	287	938	35	27	1319	317	915	56	31	1927	241	1628	34	24
GLMB($N_h = 100$)	1021	329	621	37	34	1631	247	1333	27	24	1434	304	1040	61	29	2324	195	2010	24	95
PMDD($N = 5$)	717	379	262	42	34	778	366	335	45	32	1032	370	556	77	29	1174	353	720	74	27
PMDD($N = 4$)	714	375	268	39	32	777	362	341	44	30	1039	370	566	76	27	1198	350	750	71	27
PMDD($N = 3$)	727	375	281	38	33	795	360	362	42	31	1057	371	589	70	27	1257	343	826	62	26
PMDD($N = 2$)	769	360	341	34	34	857	348	439	38	32	1163	354	722	58	29	1405	326	1005	48	26
PMDDS($N = 5$)	420	263	123	14	20	472	258	181	14	19	701	256	411	16	18	842	245	561	18	18
PMDDS($N = 4$)	416	261	127	9	19	471	256	185	12	18	720	255	430	17	18	883	241	609	16	17
PMDDS($N = 3$)	453	257	170	7	19	514	250	234	12	18	766	255	478	14	19	976	233	710	15	18
PMDDS($N = 2$)	529	245	259	6	19	618	241	348	11	18	926	241	654	11	20	1177	220	925	13	19

- [19] L. Svensson and M. Morelande, “Target tracking based on estimation of sets of trajectories,” in *Proc. 2014 Information Fusion Conf*, 2014, pp. 1–8.
- [20] A. B. Poore and A. J. Robertson III, “A new lagrangian relaxation based algorithm for a class of multidimensional assignment problems,” *Computational Optimization and Applications*, vol. 8, no. 2, pp. 129–150, 1997.
- [21] S. Deb, M. Yeddapanudi, K. Pattipati, and Y. Bar-Shalom, “A generalized SD assignment algorithm for multisensor-multitarget state estimation,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 33, no. 2, pp. 523–538, 1997.
- [22] N. Komodakis, N. Paragios, and G. Tziritas, “MRF energy minimization and beyond via dual decomposition,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 3, pp. 531–552, 2011.
- [23] B. Polyak, “Subgradient methods: A survey of soviet research,” in *Nonsmooth optimization: Proceedings of the IIASA workshop March*, 1978, pp. 5–30.
- [24] K. Murthy, “An algorithm for ranking all the assignments in order of increasing costs,” *Operations Research*, vol. 16, no. 3, pp. 682–687, 1968.
- [25] J. L. Williams, “Hybrid Poisson and multi-Bernoulli filters,” in *Proc. 2012 Information Fusion Conf*, 2012, pp. 1103–1110.
- [26] Á. F. García-Fernández and L. Svensson, “Trajectory probability hypothesis density filter,” *arXiv preprint arXiv:1605.07264*, 2016.
- [27] M. Beard, B. T. Vo, and B.-N. Vo, “Generalised labelled multi-Bernoulli forward-backward smoothing,” in *Proc. 2016 Information Fusion Conf*, 2016, pp. 688–694.
- [28] A. S. Rahmthullah, Á. F. García-Fernández, and L. Svensson, “Generalized optimal sub-pattern assignment metric,” in *Proceedings of International Conference on Information Fusion*, 2017, pp. 1–8.
- [29] —, “A metric on the space of finite sets of trajectories for evaluation of multi-target tracking algorithms,” *arXiv preprint arXiv:1605.01177*, 2016.