20th International Conference on Information Fusion
Xi'an, China - July 10-13, 2017

Performance Evaluation of Multi-Bernoulli Conjugate Priors for Multi-Target Filtering

Yuxuan Xia†, Karl Granström†, Lennart Svensson†, Ángel F. García-Fernández‡
†Department of Signals and Systems, Chalmers University of Technology, Göteborg, Sweden
‡Department of Electrical Engineering and Automation, Aalto University, Espoo, Finland

***Abstract*—In this paper, we evaluate the performance of labelled and unlabelled multi-Bernoulli conjugate priors for multi-target filtering. Filters are compared in two different scenarios with performance assessed using the generalised optimal sub-pattern assignment (GOSPA) metric. The first scenario under consideration is tracking of well-spaced targets. The second scenario is more challenging and considers targets in close proximity, for which filters may suffer from coalescence. We analyse various aspects of the filters in these two scenarios. Though all filters have pros and cons, the Poisson multi-Bernoulli filters arguably provide the best overall performance concerning GOSPA and computational time.**

I. Introduction

Multiple target tracking (MTT) involves the processing of sets of measurements obtained from multiple targets in order to estimate their current states. Solving the MTT problem is mainly complicated by the unknown correspondence between targets and measurements, known as data association. Popular solutions to MTT are the joint probabilistic data association (JPDA) filter [1], multiple hypothesis tracker (MHT) [2] and algorithms based on random finite sets (RFS) [3]. Developments using RFS have yielded a variety of tracking methods that avoid handling the data association uncertainty explicitly, such as probability hypothesis density (PHD) filter [4], [5], cardinalised PHD (CPHD) filter [6], [7], multiple target multi-Bernoulli (MeMBer) filter [3], and cardinality-balanced MeMBer (CB-MeMBer) filter [8].

A significant trend in RFS-based MTT is the development of conjugate distributions in Bayesian probability theory, which means that the posterior distribution has the same functional form as the prior. MTT algorithms based on conjugate priors for labelled RFS [9]–[12] and unlabelled RFS [13]–[15] can provide accurate approximations to the exact posterior distributions, and appealing performance for MTT has been demonstrated. Filters based on labelled multi-Bernoulli (MB) conjugate prior using MB birth model include Delta generalised labelled multi-Bernoulli (δ-GLMB) [9], [10] and its approximation labelled multi-Bernoulli (LMB) [12]. Filters based on unlabelled MB conjugate prior using Poisson birth model include Poisson multi-Bernoulli mixture (PMBM) [13], [14] and its variational approximation Poisson multi-Bernoulli (PMB) [15].

These algorithms are all important contributions to the MTT literature, but it is not yet clear how they compare to each other regarding computational time and performance. The main contribution of this paper is the performance evaluation of MB conjugate priors for multi-target filtering using the generalised optimal sub-pattern assignment (GOSPA) metric [16]. Compared with the unnormalised OSPA metric [17], the GOSPA metric allows for further breaking down the cardinality mismatch into errors due to missed and false targets. In this work, we assume that each target gives rise to at most one measurement and that targets do not have any shared measurements. Simulation results demonstrate that PMBM is a more computational efficient filter structure than δ-GLMB since it tends to yield better estimation performance than δ-GLMB with less computational time in a given test scenario. The advantage of a Poisson birth model over an MB birth model in MTT is also demonstrated in specific examples. In addition, the unlabelled RFS filter shows the ability to resolve the coalescence phenomenon, and its performance under this scenario can be further improved by applying the variational approximation method [15].

The paper is organised as follows. Background on RFS and Bayesian filtering is provided in Section II. Section III and IV summarise the labelled and unlabelled MB conjugate priors with their corresponding filter implementations. Section V discusses different methods that can be used to reduce the computational cost. Simulation results are presented in Section VI, and conclusions are drawn in Section VII.

TABLE I: Notations

- Single target states are represented by lower-case letter, e.g., x; multi-target states are represented by upper-case letters, e.g., X; labelled states are represented by bold letters, e.g., $\boldsymbol{x}, \boldsymbol{X}$; spaces are represented by blackboard bold letters, e.g., $\mathbb{X}$.
- $|X|$: set cardinality, i.e., number of elements in set X.
- Kronecker delta function
$$\delta_Y(X) = \begin{cases} 1 & \text{if } Y = X, \\ 0 & \text{otherwise.} \end{cases}$$
- $h^X = \prod_{x \in X} h(x)$.
- Π_N: set of permutation functions on $I_N \triangleq \{1, ..., N\}$
$$\Pi_N = \{\pi : I_N \to I_N | i \neq j \Rightarrow \pi(i) \neq \pi(j)\}.$$
- $\uplus$: disjoint set union, i.e., $Y \uplus U = X$ means that $Y \cup U = X$ and $Y \cap U = \emptyset$.
- I_m: identity matrix of size $m \times m$.

II. Background

In RFS-based methods [3], target states and observations are represented in the form of finite sets. The system state at time k is modelled as a set X_k, and the set of measurements obtained

978-0-9964-5270-0©2017 ISIF

at time k is denoted as Z_k, including clutter and target-generated measurements with unknown origin. The sequence of all the measurement sets received so far up to time k is denoted as Z^k. More notations are given in Table I.

A labelled RFS can be formed from its unlabelled version by incorporating labels into target states such that each state $x \in \mathbb{X}$ is augmented with a unique discrete label $l \in \mathbb{L}$. Let $\mathcal{L}(\boldsymbol{X})$ be the set of unique labels in $\boldsymbol{X}$, then a finite subset $\boldsymbol{X}$ on space $\mathbb{X}\times\mathbb{L}$ has distinct labels if and only if $\delta_{|\boldsymbol{X}|}(|\mathcal{L}(\boldsymbol{X})|) = 1$ [10].

A non-homogeneous Poisson point process (PPP) with intensity function $\lambda(x)$ has RFS density

$$f^{ppp}(X) = \exp\Big(-\int \lambda(x)dx\Big)\cdot\lambda^X, \tag{1}$$

where $|X|$ is Poisson distributed, and elements $x \in X$ are independently and identically distributed (i.i.d.). A Poisson process is often used to model clutter and target birth in unlabelled RFS filters.

A Bernoulli process with probability of existence r and existence-conditioned probability density function (PDF) $f(x)$ has RFS density

$$f(X) = \begin{cases} 1-r & X = \emptyset \\ r\cdot f(x) & X = \{x\} \\ 0 & \text{otherwise}, \end{cases} \tag{2}$$

where $|X|$ is Bernoulli distributed with parameter r. A labelled Bernoulli RFS $\boldsymbol{X}$ is a Bernoulli RFS X augmented with label l corresponding to the non-empty Bernoulli component x, i.e., $\boldsymbol{X} = \{(x,l)\}$ (or $\boldsymbol{X} = \emptyset$). A Bernoulli process can capture the target uncertainty of both existence and state, and it is also used in labelled RFS filters to model target birth.

Multiple targets can be naturally represented through an MB RFS. An MB RFS X is a union of independent Bernoulli RFSs X_i, i.e., $X = \bigcup_{i=1}^{N} X_i$,

$$f^{mb}(X) = \sum_{\uplus_i^N X_i = X}\prod_{i=1}^{N} f_i(X_i). \tag{3}$$

An MB RFS can either be labelled or unlabelled.

A. Bayesian Filtering Recursion

An RFS density can be expressed using a cardinality distribution $c(n)$, and joint conditional state distributions $f_n(x_1,...,x_n|n)$, yielding [13]

$$f(\{x_1,...,x_n\}) = c(n)\sum_{\pi\in\Pi_n} f_n(x_{\pi(1)},...,x_{\pi(n)}|n), \tag{4}$$

where $\pi(i)$ is the i-th component of vector π. A labelled RFS and its unlabelled version have the same cardinality distribution [10].

The multi-target distribution at time k, given all measurement sets up to and including time k' is denoted as $f_{k|k'}(X_k|Z^{k'})$, and $f_k(Z_k|X_k)$ is the multi-target measurement likelihood. The multi-target Bayes filter propagates the target set PDF $f_{k|k-1}(X_k|Z^{k-1})$ in time using the Bayes update [3, p. 484]

$$f_{k|k}(X_k|Z^k) \propto f_k(Z_k|X_k)f_{k|k-1}(X_k|Z^{k-1}), \tag{5}$$

and the Chapman-Kolmogorov prediction [3, p. 484]

$$f_{k+1|k}(X_{k+1}|Z^k) = \int f_{k+1|k}(X_{k+1}|X_k)f_{k|k}(X_k|Z^k)\delta X_k, \tag{6}$$

where the set integral is defined as [3, p. 361]:

$$\int f(X)\delta X \triangleq f(\emptyset) + \sum_{n=1}^{\infty}\frac{1}{n!}\int f(\{x_1,...,x_n\})d(x_1,\cdots,x_n). \tag{7}$$

B. Standard Point Target Transition and Measurement Model

The dynamic point target transition model utilised to solve the RFS-based MTT problem is based on the following assumptions

- *Targets arrive according to a PPP or an MB process, independently of existing targets.*
- *At each time step, targets remain with a probability of survival $P^s(x)$. Targets depart according to i.i.d. Markov processes with probability $1 - P^s(x)$.*
- *Targets' next states only depend on their current states. Target motion follows i.i.d. Markov processes with transition density $f_{k|k-1}(x_k|x_{k-1})$.*

The multi-target transition kernel can be written as the convolution of transition density for survival targets $f_S^{mb}(\cdot|X)$ and transition density for new born targets $f_B(\cdot)$

$$f(X_{k+1}|X_k) = \sum_{S_k\uplus B_k = X_{k+1}} f_S^{mb}(S_k|X_k)f_B(B_k), \tag{8}$$

where S_k is the set of survival targets, and B_k is the set of new born targets.

In the measurement model, the following assumptions are made.

- *Each target-generated measurement is only conditioned on its corresponding target. The single target measurement likelihood is $f_k(z_k|x_k)$.*
- *At each time step, existing targets may or may not be detected. The detection probability is $P^d(x)$.*
- *The sensor may receive measurements not originating from any target, known as false alarm or clutter. At each time step, the clutter arrives according to a PPP with intensity $\lambda^{fa}(z)$, independently of targets and target-generated measurements.*

The multi-target likelihood can be presented in a convolution form

$$f_k(Z_k|X_k) = \sum_{T_k\uplus C_k = Z_k} f_k^{mb}(T_k|X_k)f_k^{ppp}(C_k), \tag{9}$$

where T_k is the set of target-generated measurements, and C_k is the set of clutter measurements.

III. Labelled Multi-Bernoulli Conjugate Prior

This section presents a general description of δ-GLMB and its efficient approximation LMB. The reader is referred to [9]–[12] for detailed analytic derivations and mathematical expressions.

A. Delta generalised Labelled Multi-Bernoulli

A δ-GLMB RFS is a labelled RFS with state space $\mathbb{X}$ and label space $\mathbb{L}$ distributed according to [10]

$$f(\boldsymbol{X}) = \delta_{|\boldsymbol{X}|}(|\mathcal{L}(\boldsymbol{X})|) \sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi} w^{(I,\xi)}\delta_I(\mathcal{L}(\boldsymbol{X}))(f^{(\xi)})^{\boldsymbol{X}}, \tag{10}$$

where $I \in \mathcal{F}(\mathbb{L})$ is the set of target labels, and $\xi \in \Xi$ represents a history of association maps. Each pair (I,ξ) is a hypothesis with probability $w^{(I,\xi)}$, which satisfies

$$\sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi} w^{(I,\xi)} = 1. \tag{11}$$

If the prior distribution is a δ-GLMB of the form (10), then under the multi-target likelihood function (9), the posterior and predicted distribution are both δ-GLMB [10]. According to this conjugacy, the δ-GLMB filter can recursively propagate a δ-GLMB RFS density in time via the Bayes update and prediction equation (5) and (6).

A labelled MB model is used for target birth, and we denote the label space for new born targets as $\mathbb{B}$. The likelihood for a birth hypothesis yields [9]

$$w_B(L) = [1 - r(\cdot)]^{(\mathbb{B}-L)}[r(\cdot)]^L, \tag{12}$$

where $r(l)$ denotes the existence probability of a target with label $l \in I$. For the surviving targets, their labels are kept from last time step, while the labels of birth targets are newly assigned. In the δ-GLMB prediction step, each component in the prior generates a new set of predicted components. In the δ-GLMB update step, each component in the predicted density generates a set of updated components for each possible measurement-to-label mapping. Instead of computing the filtering density in two separate steps, a fast implementation has been proposed in [11], in which the prediction and update steps are combined into a single step by formulating a new mapping between the components of the current and previous filtering density. Compared with the original mapping used in update step, the associations for non-survival and unconfirmed birth targets are included in this new mapping.

B. Labelled Multi-Bernoulli

The LMB filter [12] is based on approximating the posterior and predicted densities (5) and (6) as labelled MB RFS densities. In contrast to δ-GLMB, the number of components maintained in the recursive Bayesian filtering only grows linearly with time.

The multi-target prediction in the LMB filter is identical to the prediction for MeMBer [8] with target labels interpreted as component indices. The predicted existence probability and density distribution are re-weighted by the survival probability and transition density. In the update step, the posterior multi-target density is approximated by exact marginalisation; thus the LMB filter propagates only one component. The predicted LMB is first represented as δ-GLMB, then a full δ-GLMB update is applied directly before the posterior collapsing back to a matching LMB approximation.

IV. Unlabelled Multi-Bernoulli Conjugate Prior

This section gives a general introduction to PMBM and its variational approximation PMB. We refer readers to [13]–[15] for further derivation and mathematical details.

A. Poisson Multi-Bernoulli Mixture

The PMBM conjugate prior for point target MTT was developed in [13]. It is a linear combination of independent PPP and multi-Bernoulli mixture (MBM) components with the following form

$$f(X) = \sum_{Y\uplus W=X} f^{ppp}(Y) f^{mbm}(W), \tag{13}$$

where Y stands for targets that have not yet been detected, and W stands for targets that have been detected at least once. For targets that have already been detected, their distributions can be described as an MBM of the form [13]:

$$f^{mbm}(X) = \sum_{a=\{h_1,\ldots,h_N\}\in\mathcal{A}} w_a \sum_{\uplus_i^N X_i = X} \prod_{i=1}^{N} f_{h_i}(X_i), \tag{14}$$

where each of the MB components $a \in \mathcal{A}$ corresponds to a particular data association with weight w_a satisfying

$$\sum_{a=\{h_1,\ldots,h_N\}\in\mathcal{A}} w_a = 1. \tag{15}$$

Each data association hypothesis is made up of single target hypotheses $\{h_1, ..., h_N\}$ on each target [13]. One single target hypothesis can incorporate events, including that the target never existed, that the target existed before and that the target continues to exist, represented via a Bernoulli process. Missed detections may occur on some proportions of targets that are hypothesised to be born. A target that is hypothesised to exist but has never been detected is treated as an unknown target [13]. The distribution of unknown targets is represented by a PPP.

In the prediction step, the MB describing pre-existing tracks and the PPP describing unknown targets are predicted individually. By having a Poisson birth model, PPP for new born targets can be easily incorporated into the predicted PPP. In the update step, PPP and MBM are updated independently. Two single target hypotheses are created for each measurement, and then the PPP intensity is updated by the miss detection probability. The first single target hypothesis covers the case that the measurement is associated with a previous target so that this hypothesis has zero existence probability and the corresponding PDF of the Bernoulli distribution has no effect. The second single target hypothesis covers the case that the measurement corresponds to a false alarm or a new target. For targets surviving from previous time steps, new single target

hypotheses are included from missed detections or updates of previous hypotheses using one of the new measurements.

B. Poisson Multi-Bernoulli

A PMB is a union of a PPP describing unknown targets and an MB process describing already detected targets. In the PMB recursion, the PMB density is preserved in prediction step, whereas the MB component becomes an MBM due to data association. A variational approximation method was presented in [15] to obtain the best-fitting MB $g(X)$ that minimises the Kullback-Leibler (KL) divergence from the MBM distribution $f(X)$. In [15], it is shown that this optimisation problem can be solved approximately as:

$$\min_{q(h,j)\in\mathcal{M}} -\sum_{j=1}^{N}\int\left(\sum_{h\in\mathcal{H}} q(h,j)f_h(X)\right)\cdot\log g_j(X)\delta X, \quad (16)$$

where $q(h,j)$ is the probability that a Bernoulli component in $f(X)$ that is utilising hypothesis h is assigned to the j-th Bernoulli component of $g(X)$, and $\mathcal{M}$ is an approximated polytope needed for tractability

$$\mathcal{M}=\left\{q(h,j)\geq 0\middle|\sum_{h\in\mathcal{H}} q(h,j)=p_j\ \forall\ j\in\{1,...,N\},\ \sum_{j=1}^{N} q(h,j)=p_h\ \forall\ h\in\mathcal{H}\right\}. \quad (17)$$

The algorithm is initialized with

$$p_j(h)=\sum_{a=(h_1,...,h_N)\in\mathcal{A}|h_j=h} w_a, \quad (18)$$

which can either be calculated using the best data association hypotheses and their corresponding weights obtained from Murty's algorithm [18] or be approximated based on loopy belief propagation (LBP) [19].

V. Approximations for Computational Tractability

In previous sections, we have summarised the δ-GLMB, LMB, PMBM and PMB. The common bottleneck for all the filters is the large number of hypotheses generated in the update step. Since it is not tractable to compute all the possible components, efficient truncation methods should be implemented by only maintaining components with most significant weights[1].

If hypothesis weights are generated in non-decreasing order, the M-best mappings (i.e. highest weights) can be selected without computing the weights of all possible mappings exhaustively. The M-best mappings can be found, e.g., by solving the ranked assignment problem using Murty's algorithm. Implementation details, e.g., how to construct the cost function for an association map, can be found in [10], [11] and [14]. In the implementation of δ-GLMB with separate prediction and update steps, additional approximation is needed to truncate the independent surviving and birth hypotheses separately. The hypotheses with highest weights can be determined by solving the K-shortest path problem [20].

To further reduce the computational complexity of data association, gating can be implemented in all the proposed filters to remove unlikely target-to-measurement pairs before the update step. Moreover, clustering can be used in LMB, PMBM and PMB to partition well-spaced targets and the corresponding measurements falling into their gates into independent groups, which allows for parallel computing.

For unlabelled RFS filters, in each MB component, the recycling method of [21] can be applied to Bernoulli components with low existence probability. The recycled components are approximated as being Poisson and are incorporated into the PPP representing unknown targets for generating possible new targets in subsequent steps. Thus, the number of single target hypotheses in each data association hypothesis is reduced while maintaining the significant information of already detected targets. In addition, recycling can help recover performance loss due to the pruning of Bernoulli components with small existence probability [19]. However, having a PPP for each MB in PMBM would result in multiple PPP updates, which may instead increase the computational cost. To address this problem, we first find the best-fitting MB of the MBM to be recycled, and then the approximated MB is further approximated as being PPP and recycled.

VI. Simulations and Results

In this section we show simulation results that compare six different filters: 1) δ-GLMB with separate prediction and update steps [9], 2) δ-GLMB with joint prediction and update steps [11], 3) LMB [12], 4) PMBM [14], 5) PMB [15] using Murty's algorithm, and 6) PMB using LBP [19]. All the codes are written in MATLAB, except Murty's algorithm, which is written in C++ [2]. The benefit of recycling for PMBM is also studied. The estimation performances of the filters are evaluated in two different scenarios. In the first scenario, targets are well-separated most of the time. The second scenario is more complicated, involving targets that get in close proximity first and then separate.

A. State Space Models

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

$$F_k=I_2\otimes\begin{bmatrix}1 & T\\ 0 & 1\end{bmatrix},\quad Q_k=\sigma_v^2 I_2\otimes\begin{bmatrix}T^4/4 & T^3/2\\ T^3/2 & T^2\end{bmatrix}$$

where $T=1s$ is the sampling period, and σ_v is the standard deviation of motion noise. The linear-Gaussian measurement

[1]For PMB using LBP, marginal association probabilities are used to merge hypotheses.

[2]For the implementations of δ-GLMB with separate prediction and update steps and LMB, we use the codes that Profs Ba-Ngu Vo and Ba-Tuong Vo share online. The authors thank them for providing the codes.

TABLE II: FILTER PARAMETERS

- Gating size in percentage is 0.999.
- δ-GLMB: capping threshold on the number of hypotheses is 3000; pruning threshold on the weight of hypothesis is 10^{-5}; requested number of components used in K-shortest path and Murty's algorithm is $\lceil 3000 \cdot w \rceil$, where w is the weight of hypothesis; requested number of birth components used in K-shortest path algorithm is 10.
- LMB: capping threshold on the number of targets is 100; pruning threshold is 10^{-4}; requested number of birth and survival components used in K-shortest path algorithm are 10 and 30 respectively; requested number of components used in Murty's algorithm is $\lceil 100 \cdot w \rceil$; the number of Gaussian components in each Gaussian mixture is limited to 10 and components are merged with Mahalanobis distance smaller than 1.
- PMBM: capping threshold on the number of hypotheses is 100; pruning threshold on the weight of hypotheses and weight of Gaussian components in PPP is 10^{-4}; requested number of components used in Murty's algorithm is $\lceil 100 \cdot w \rceil$ in the first testing scenario and $\lceil 300 \cdot w \rceil$ in the second; pruning threshold on target existence probability is 10^{-4} without recycling, and 10^{-1} with recycling.
- PMB: corresponding pruning threshold is set to the same as PMBM; requested number of components used in Murty's algorithm is 20; convergence threshold used in LBP is 10^{-6}.

likelihood model has density $g(z_k|x_k) = \mathcal{N}(z_k; H_k x_k, R_k)$ with parameters $H_k = I_2 \otimes \begin{bmatrix} 1 & 0 \end{bmatrix}$ and $R_k = \sigma_\epsilon^2 I_2$, where σ_ϵ is the standard deviation of measurement noise. The target survival probability P_k^s and the detection probability P_k^d are assumed to be constant throughout the simulation. The clutter follows a Poisson RFS with uniform density, giving an average of λ^{fa} per time step. We consider cases with $P_k^s = 0.99$, $P_k^d \in \{0.75, 0.98\}$, and $\lambda^{fa} \in \{10, 30\}$.

For unlabelled filters, the Poisson birth intensity is a Gaussian mixture $\lambda^b(x) = \sum_{n=1}^{n_b} \lambda^{b,n} \mathcal{N}(x; \mu^{b,n}, \Sigma^{b,n})$ with all the Gaussian components sharing the same weight $\lambda^{b,n} = 0.1$ and covariance. The PPP intensity of unknown targets is represented as $\lambda^u(x) = \sum_{n=1}^{n_u} \lambda^{u,n} \mathcal{N}(x; \mu^{u,n}, \Sigma^{u,n})$, with the initial value assumed to be $\lambda_{0|0}^u(\boldsymbol{x}) = \lambda^b(\boldsymbol{x})$. For labelled filters, the birth density is a labelled MB RFS with all the Bernoulli components sharing the same existence probability $r^{b,n} = 0.1$. The existence-conditioned PDF of each Bernoulli is Gaussian distributed. A labelled Poisson RFS can also be used to model target birth in δ-GLMB [10], but this is not a suitable option [14]. As a Poisson RFS has a cardinality distribution that expands from zero to infinity, it is not efficient to consider each possible birth hypothesis separately via a δ-GLMB filter. Parameter settings for different filters are stated in Table II.

B. Performance Evaluation

Given a multi-target posterior density, several state estimators are available [10]. In this work, we choose to extract the target states by finding the maximum a posteriori (MAP) cardinality estimate, following the methods suggested in [10] and [14].

Filtering performance is assessed using GOSPA metric. It was mentioned in [16] that the GOSPA metric with parameter $\alpha = 2$ can be rewritten using the 2D assignment set as

$$d_p^{(c,2)}(X,Y) \triangleq \left[\min_{\gamma \in \Gamma^{(|X|,|Y|)}} \left(\sum_{(i,j)\in\gamma} d(x_i, y_j)^p + \frac{c^p}{2}(|X| + |Y| - 2|\gamma|) \right) \right]^{\frac{1}{p}}, \quad (19)$$

where $\Gamma^{(|X|,|Y|)}$ is the set of all possible 2D assignment sets, c denotes the cut-off value at base distance and p determines the severity of penalising the outliers in the localisation component. GOSPA allows for decomposing the multi-target metric into localisation error $\sum_{(i,j)\in\gamma} d(x_i, y_j)^p$, missed detection error $c^p(|X| - |\gamma|)/2$, and false detection error $c^p(|Y| - |\gamma|)/2$, considering X as the ground truth and Y as the estimates. In the simulations, we choose $p = 1, c = 100$, and results are shown over 200 Monte Carlo trials.

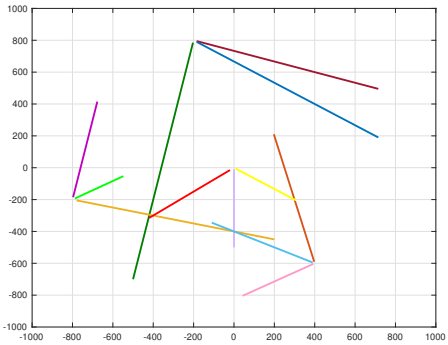

Fig. 1: Target trajectories on the region $[-1000, 1000]m \times [-1000, 1000]m$. Up to twelve targets travel along straight lines with constant speed in the duration time of $100s$ with three targets born at time 1 and 20, two targets born at time 40, 60 and 80 and two targets dead at time 50 and 70. Specifically, the three targets born at time 20 are from roughly the same position, and the two targets born at time 40 are from roughly the same position.

C. Scenario with Well-spaced Targets

In this scenario, targets are born from four localised positions with ground truth shown in figure 1. The standard deviations of motion and measurement noises are $\sigma_v = 5m/s^2$ and $\sigma_\epsilon = 10m$, respectively. In the birth intensity, each Gaussian component has the same covariance $\Sigma^{b,n} = 100 \times I_4$ and the mean values are $\mu^{b,1} = [0, 0, 0, 0]^T$, $\mu^{b,2} = [400, -600, 0, 0]^T$, $\mu^{b,3} = [-200, 800, 0, 0]^T$ and $\mu^{b,4} = [-800, -200, 0, 0]^T$. The initial target positions are randomly sampled from these Gaussian densities.

Figure 2 presents the GOSPA errors and the average computation time required for a complete Monte Carlo simulation. The GOSPA errors were averaged over time, and are shown as box plots. In cases with high detection probability, the median of the GOSPA error due to localisation errors is very similar for all filters, while in cases of low detection probability, filters using variational approximations have slightly higher median value compared with the rest. It is clearly visible that PMBM exhibits a similar or better estimation performance than δ-GLMB while saving computational time notably. Errors

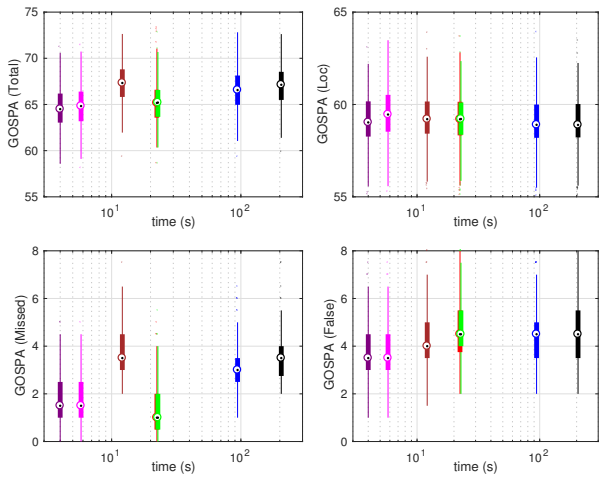

(a) $\lambda^{fa} = 10$, $P_k^d = 0.98$

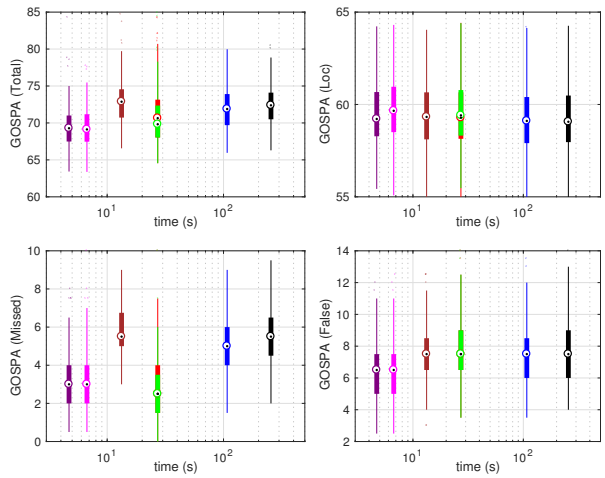

(b) $\lambda^{fa} = 30$, $P_k^d = 0.98$

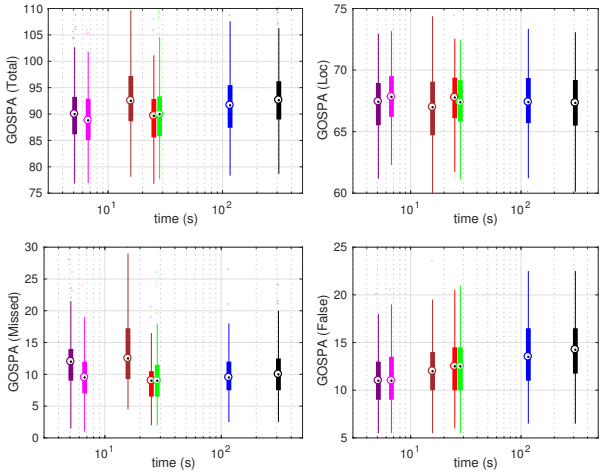

(c) $\lambda^{fa} = 10$, $P_k^d = 0.75$

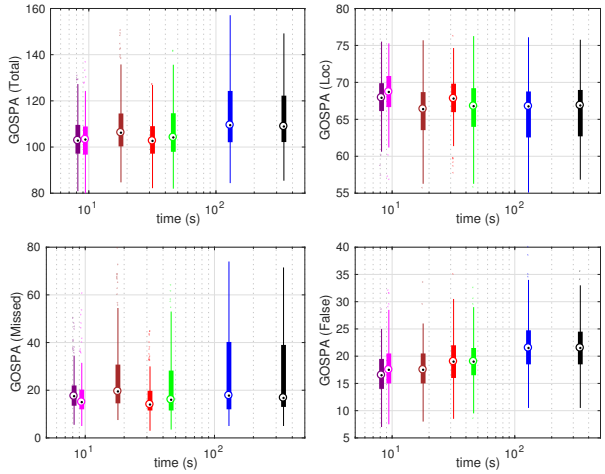

(d) $\lambda^{fa} = 30$, $P_k^d = 0.75$

Fig. 2: Filtering performance evaluation of PMB (LBP) with recycling (purple), PMB (Murty) with recycling (magenta), LMB (brown), PMBM with recycling (red), PMBM without recycling (green), δ-GLMB with joint prediction and update (blue) and δ-GLMB (black) using GOSPA in a scenario where targets are well-spaced. Results are presented in box plots. On each box, '$\odot$' indicates the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively. The end of 'Whisker' corresponds to approximately +/2.7 units of standard deviation.

due to missed and false targets can be reduced by increasing the requested number of components used in K-shortest path or Murty's algorithm, but this would further increase the computational burden.

The benefit of recycling is most clear when the probability of detection is low. In figure 2 we see that PMBM with recycling requires less computation time, and has lower GOSPA error, compared to PMBM without recycling. As for the two variants of PMB, LBP is a faster implementation than Murty's algorithm, and it tends to yield fewer false targets, but more missed targets.

Figure 3 shows the mean GOSPA error cost for missed targets against time. Here, we can see that labelled RFS filters experience higher peaks than unlabelled filters around initialization period, and at time 20 and 40, when there are two or more new targets born from roughly the same position.

D. Scenario with Coalescence

In the following simulation, filters are evaluated in a more challenging scenario involving six targets which are in close proximity at the mid-point of the simulation, achieved by initializing at the midpoint and running forward and backwards dynamically. When multiple targets are in close proximity, multiple estimates may be placed on the same target that leads to missed detection [13].

The midpoint is initialized as $x_{50} = \mathcal{N}(x; 0, 10^{-6} \times I_4)$. One possible realisation of target trajectories is shown in figure 4. Standard deviation of motion and measurement noises are $\sigma_v = 0.2m/s^2$ and $\sigma_\epsilon = 1m$ respectively. For unlabelled filters, the Poisson birth intensity only contains a single Gaussian, $\lambda^b(x) = 0.1\mathcal{N}(x; 0, \Sigma^b)$, where $\Sigma^b = \text{diag}[100^2, 1, 100^2, 1]$ covers the position and velocity region of interest. For labelled filters, the birth intensity is a single

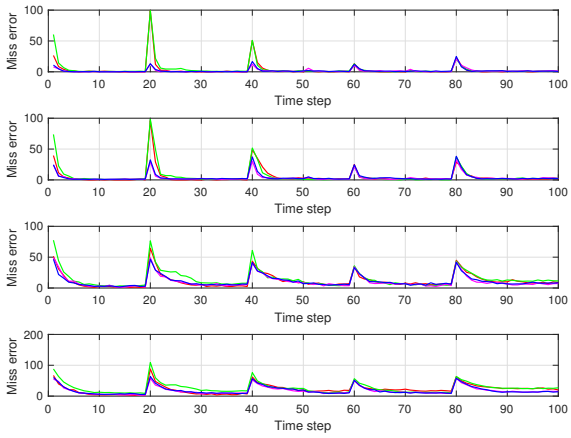

Fig. 3: Mean GOSPA error cost for missed targets of δ-GLMB with joint prediction and update (red), LMB (green), PMBM with recycling (magenta) and PMB using Murty's algorithm with recycling (blue). Testing cases from top to bottom are (a), (b), (c) and (d) respectively.

Bernoulli with $r^b = 0.1$ and $f^b(x) = \mathcal{N}(x; 0, \Sigma^b)$. To generate clear plots, for each type of filter we choose to evaluate only the one that yields fewer missed targets: PMB using Murty's algorithm with recycling, PMBM with recycling, δ-GLMB with joint prediction and update and LMB. The results of compared filters are shown in figure 5.

One significant difference between the unlabelled and labelled filters is that unlabelled filters initiate the tracker faster at initialization period. In cases with high detection probability, none of these filters suffers from the coalescence phenomenon. However, when detection probability decreases to $P_k^d = 0.75$, we can observe a substantial increase in missed targets for all the filters after targets become closely spaced. In this case, the labelled filters present fewer missed targets around midpoint that correspond to coalescence than unlabelled filters, though it is difficult for LMB to detect targets when the false alarm rate increases to $\lambda^{fa} = 30$. Compared with labelled filters, unlabelled filters exhibit good robustness to coalescence effects, and the performance is restored after targets become separated. As for the error due to false targets, we can observe that labelled filters suffer less from false detections than unlabelled filters. Regarding the localisation error, it is increased abnormally shortly after the midpoint for labelled filters, though δ-GLMB is an exception to this in cases of high detection probability.

Furthermore, the results demonstrate the success of PMB in resolving the coalescence phenomenon, which presents even less deterioration than PMBM. The advantage with PMB is that it merges tracks using a (variational) technique that reduces coalescence [15], [22]. In the simulation, we found that the data association hypotheses have multiple dominant weights when coalescence happens. In this case, information is lost if we only extract the best states from a single MB component in PMBM. However, approximating MBM as a single MB by minimising the KL divergence yields an MB RFS density which tends to be closer to the full RFS density when targets are in close proximity.

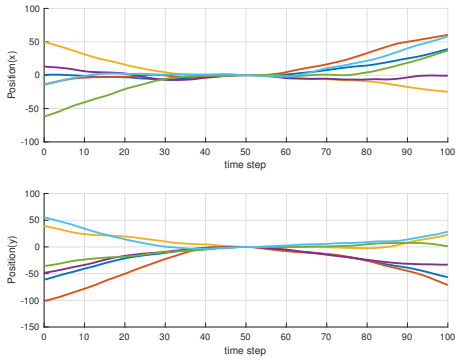

Fig. 4: One dimension of target trajectories in a single Monte Carlo trial on the region $[-100, 100]m \times [-100, 100]m$. Six targets are all born at time step 1 and exist throughout the duration time of $100s$.

VII. Conclusions and Future Work

In this paper, we have evaluated the δ-GLMB, LMB, PMBM and PMB filters in two different scenarios. Estimation error due to localisation, false targets and missed targets have been assessed using the GOSPA metric. Results from the scenario with well-spaced targets show that PMBM is a more efficient filter structure than δ-GLMB since it yields similar or better performance with less computational time. The Poisson birth model provides the unlabelled filter with the ability to detect multiple targets born from roughly the same position quickly. Also, approximating the distribution of unknown targets as being Poisson enables the use of recycling methods to improve estimation performance further. Under the simulation with coalescence, the labelled filters have fewer missed targets when targets are in close proximity and fewer false targets than unlabelled filters. The unlabelled filter can resolve the coalescence phenomenon, whereas δ-GLMB and LMB are affected at various degrees in cases with low detection probability. By applying variational approximation, PMB presents the best overall performance in this challenging scenario.

In the simulation study, filters are evaluated using the MAP cardinality estimator with fixed parameter settings. It would be interesting to evaluate and compare filters with different estimators and with different complexities by tuning their parameters, e.g., pruning threshold and requested number of components used in Murty's algorithm. A technique to solve the data association problem based on Gibbs sampling was proposed in [23], and it has been applied in δ-GLMB with joint prediction and update [11] to replace Murty's algorithm to truncate hypotheses. Thus, a potential direction for future work is to evaluate the performance and computational time of different filters using Gibbs sampling.

References

[1] 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.

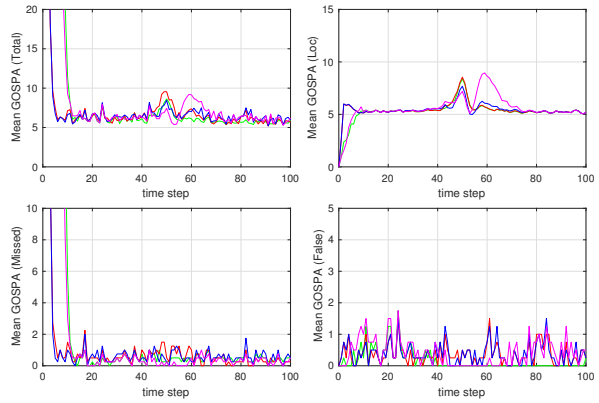

(a) $\lambda^{fa} = 10$, $P_k^d = 0.98$

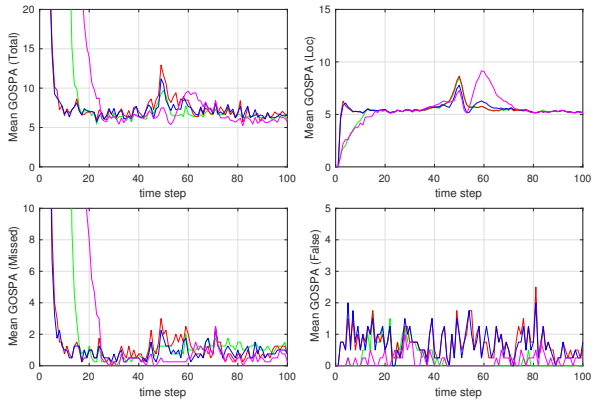

(b) $\lambda^{fa} = 30$, $P_k^d = 0.98$

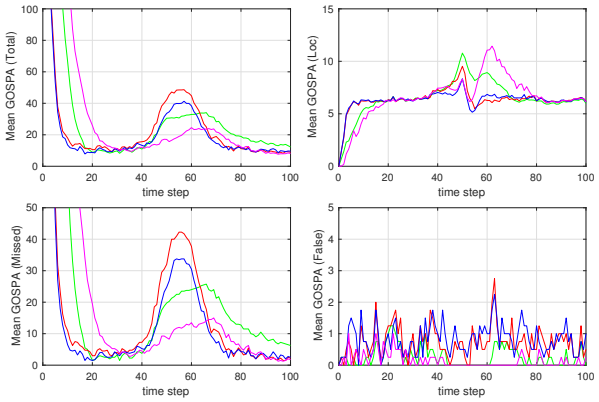

(c) $\lambda^{fa} = 10$, $P_k^d = 0.75$

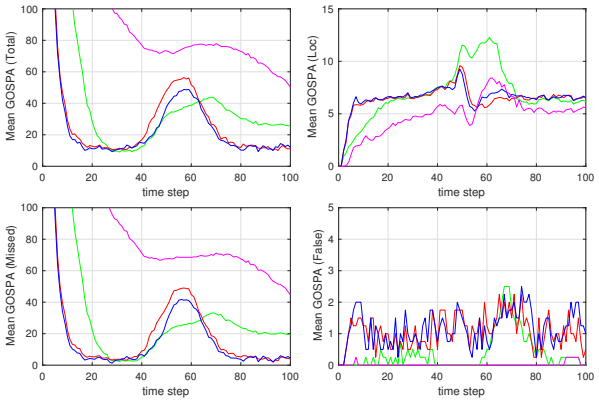

(d) $\lambda^{fa} = 30$, $P_k^d = 0.75$

Fig. 5: Filtering performance evaluation of δ-GLMB with joint prediction and update (green), LMB (magenta), PMBM with recycling (red), and PMB (Murty) with recycling (blue) using GOSPA in a challenging scenario with coalescence.

[2] S. Blackman and R. Popoli, "Design and analysis of modern tracking systems(book)," *Norwood, MA: Artech House*, 1999.

[3] R. P. Mahler, *Statistical multisource-multitarget information fusion*. Artech House, Inc., 2007.

[4] ——, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, 2003.

[5] B.-N. Vo and W.-K. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4091–4104, 2006.

[6] R. Mahler, "PHD filters of higher order in target number," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 43, no. 4, 2007.

[7] B.-T. Vo, B.-N. Vo, and A. Cantoni, "Analytic implementations of the cardinalized probability hypothesis density filter," *IEEE Trans. Signal Process.*, vol. 55, no. 7, pp. 3553–3567, 2007.

[8] ——, "The cardinality balanced multi-target multi-Bernoulli filter and its implementations," *IEEE Trans. Signal Process.*, vol. 57, no. 2, pp. 409–423, 2009.

[9] 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.

[10] 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.

[11] B. N. Vo, B.-T. Vo, and H. Hoang, "An efficient implementation of the generalized labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, 2016.

[12] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3246–3260, 2014.

[13] 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.

[14] A. F. García-Fernández, J. Williams, K. Granström, and L. Svensson, "Poisson multi-Bernoulli mixture filter: direct derivation and implementation," *arXiv preprint arXiv:1703.04264*, 2017.

[15] J. L. Williams, "An efficient, variational approximation of the best fitting multi-Bernoulli filter," *IEEE Trans. Signal Process.*, vol. 63, no. 1, pp. 258–273, 2015.

[16] A. S. Rahmathullah, Á. F. García-Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *2017 20th International Conference on Information Fusion*, Xi'an, P.R. China, Jul. 2017.

[17] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3447–3457, 2008.

[18] K. G. Murty, "An algorithm for ranking all the assignments in order of increasing cost," *Operations research*, vol. 16, no. 3, pp. 682–687, 1968.

[19] J. Williams and R. Lau, "Approximate evaluation of marginal association probabilities with belief propagation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 50, no. 4, pp. 2942–2959, 2014.

[20] D. Eppstein, "Finding the k shortest paths," *SIAM Journal on computing*, vol. 28, no. 2, pp. 652–673, 1998.

[21] J. L. Williams, "Hybrid Poisson and multi-Bernoulli filters," *2012 15th International Conference on Information Fusion*, pp. 1103–1110, 2012.

[22] L. Svensson, D. Svensson, M. Guerriero, and P. Willett, "Set JPDA filter for multitarget tracking," *IEEE Transactions on Signal Processing*, vol. 59, no. 10, pp. 4677–4691, 2011.

[23] C. Hue, J.-P. Le Cadre, and P. Pérez, "Tracking multiple objects with particle filtering," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 38, no. 3, pp. 791–812, 2002.