A Random Finite Set Conjugate Prior and Application to Multi-Target Tracking

Ba-Tuong Vo and Ba-Ngu Vo
School of Electrical, Electronic and Computer Engineering
The University of Western Australia, Crawley, WA 6009, Australia.
{ba-tuong.vo,ba-ngu.vo}@uwa.edu.au

***Abstract*—The objective of multi-object estimation is to simultaneously estimate the number of objects and their states from a set of observations in the presence of data association uncertainty, detection uncertainty, false observations and noise. This estimation problem can be formulated in a Bayesian framework by modeling the (hidden) set of states and set of observations as random finite sets (RFSs) where the model for the observation covers thinning, Markov shifts and superposition of false observations. A prior for the hidden RFS together with the likelihood of the realisation of the observed RFS gives the posterior distribution via the application of Bayes rule. We propose a new class of prior distribution and show that it is a conjugate prior with respect to the multi-target observation likelihood. This result is then applied to develop an analytic implementation of the Bayes multi-target filter for the class of linear Gaussian multi-target models.**

***Index Terms*—Random sets, Multi-Target Bayes filter, Multi-Bernoulli, Conjugate prior, Tracking**

I. INTRODUCTION

This paper investigates the problem of jointly estimating the number of objects and their states from a finite set of observations. In radar/sonar applications, the objective is to locate the targets from radar/sonar detections [1], [2]. In spatial statistics applications, for example, agriculture and forestry, it is of interest to study the underlying spatial distribution of points (plants) from point pattern observation extracted from aerial images [3], [4]. The major challenges in this estimation problem are

- detection uncertainty: each object may or may not generate an observation,
- clutter: observations are corrupted by spurious measurements (clutter) not originating from any object,
- data association uncertainty: there is no information on which object generated which observation.

This so-called multi-object estimation problem can be formulated in a Bayesian framework by modeling the (hidden) set of states and set of observations as random finite sets (RFSs). The model for the observed RFS covers thinning, Markov shifts and superposition of false observations. A prior for the hidden RFS together with the likelihood of the realization of the observed RFS gives the posterior distribution via the application of Bayes rule. In practice computing this posterior is intractable in general due the inherently combinatorial nature of the problem, and the usual complications associated with the curse of dimensionality [5]. A number of approximations have been proposed. The probability hypothesis density (PHD) [5] and Cardinalized PHD (CPHD) [6] updates are approximations based on moments and cardinality distributions, while the multi-Bernoulli update approximates the posterior distribution by a multi-Bernoulli RFS distribution [7].

In this paper we present a closed form solution to the posterior distribution for a class of multi-object prior distributions. In particular we introduce the notion of distinctly labeled RFS and propose a new class of prior distribution called generalized labeled multi-Bernoulli RFS (or generalized labeled iid cluster point process). We show that distinctly labeled multi-Bernoulli RFSs and Poisson point processes falls under the umbrella of generalized labeled multi-Bernoulli RFSs. More importantly we prove that the proposed generalized multi-Bernoulli point process is a conjugate prior with respect to the multi-target observation likelihood. Consequently, the posterior distribution can be expressed in closed form. This result is applied to develop an analytic implementation of the Bayes multi-target filter for linear Gaussian multi-target models.

II. BACKGROUND

In a multiple target system, the number of targets varies with time due to the appearance and disappearance of targets, and the number of measurements received at each time step does not necessarily match the number of targets due to missed detections and clutter. The objective of multiple target tracking is to jointly estimate the number of targets and their states from the accumulated observations. Suppose that at time k, there are $n(k)$ target states $x_{k,1},\ldots,x_{k,n(k)}$, each taking values in a state space $\mathbb{X} \subseteq \mathbb{R}^{n_x}$, and $m(k)$ observations $z_{k,1},\ldots,z_{k,m(k)}$ each taking values in an observation space $\mathbb{Z} \subseteq \mathbb{R}^{n_z}$. Define the *multiple target state* and *multiple target observation* respectively by the finite sets

$$X_k = \{x_{k,1},\ldots,x_{k,n(k)}\} \in \mathcal{F}(\mathbb{X}),$$
$$Z_k = \{z_{k,1},\ldots,z_{k,m(k)}\} \in \mathcal{F}(\mathbb{Z}),$$

where $\mathcal{F}(\mathbb{X})$ denotes the collection of finite sets of $\mathbb{X}$. The rationale behind this representation traces back to a fundamental consideration in estimation theory–estimation error [8].

A. Random Finite Sets

In the Bayesian estimation paradigm, the state and measurement are treated as realizations of random variables. Since the (multi-object) state is a finite set, the concept of a random finite set (RFS) is required to cast the multi-object estimation problem in the Bayesian framework. Intuitively, a *random finite set* (RFS), is a random (spatial) point pattern, e.g. measurements on a radar screen. What distinguishes an RFS from a random vector is that: the number of points is random; the points themselves are random and unordered. In essence, an RFS is simply a finite-set-valued random variable. Some examples of RFSs pertinent to development of our key results are given next.

1) Poisson RFSs: An RFS X on $\mathbb{X}$ is said to be *Poisson* with a given *intensity function* v (defined on $\mathbb{X}$) if its cardinality ,denoted as $|X|$, is Poisson distributed, with mean $\bar{N} = \int v(x)dx$, and for any finite cardinality, the elements x of X are independently and identically distributed (i.i.d.) according to the probability density $v(\cdot)/\bar{N}$ [9]. A Poisson RFS is completely characterized by its intensity function, also known in the tracking literature as the *Probability Hypothesis Density* (PHD). A Poisson RFS with intensity function v has probability density[1] (see [2] pp. 366).

$$\pi(X) = e^{-\bar{N}} v^X$$

where for any $h : \mathbb{X} \to \mathbb{R}$, $h^X \equiv \prod_{x \in X} h(x)$ with $h^{\emptyset} = 1$ by convention.

2) Bernoulli RFS: A *Bernoulli* RFS on $\mathbb{X}$ has probability $1 - r$ of being empty, and probability r of being a singleton whose (only) element is distributed according to a probability density p (defined on $\mathbb{X}$). The cardinality distribution of a Bernoulli RFS is a Bernoulli distribution with parameter r. The probability density of a Bernoulli RFS is given by (see [2] pp. 368)

$$\pi(X) = \begin{cases} 1 - r & X = \emptyset, \\ r \cdot p(x) & X = \{x\}. \end{cases}$$

3) Multi-Bernoulli RFS: A *multi-Bernoulli* RFS X on $\mathbb{X}$ is a union of a fixed number of independent Bernoulli RFSs $X^{(i)}$ with existence probability $r^{(i)} \in (0, 1)$ and probability density $p^{(i)}$ (defined on $\mathbb{X}$), $i = 1, ..., M$, i.e. $X = \bigcup_{i=1}^{M} X^{(i)}$. . A multi-Bernoulli RFS is thus completely described by the multi-Bernoulli parameter set $\{(r^{(i)}, p^{(i)})\}_{i=1}^{M}$. The mean cardinality of a multi-Bernoulli RFS is $\sum_{i=1}^{M} r^{(i)}$. Moreover, the probability density π is (see [2] pp. 368) $\pi(\emptyset) = \prod_{j=1}^{M} \left(1 - r^{(j)}\right)$ and

$$\pi(\{x_1, ..., x_n\}) = n! \ \pi(\emptyset) \sum_{\substack{\{i_1,...,i_n\} \\ \in \mathcal{F}_n(\{1,...,M\})}} \prod_{j=1}^{n} \frac{r^{(i_j)} p^{(i_j)}(x_j)}{1 - r^{(i_j)}}.$$

where $\mathcal{F}_n(\mathbb{X})$ denotes the collection of finite subsets of $\mathbb{X}$ with exactly n elements.

B. Multi-Target Bayes Filter

Mahler's Finite Set Statistics (FISST) provides powerful yet practical mathematical tools for dealing with RFSs [2], [5]. Using the FISST notion of integration and density the Bayes multiple target recursion propagates the posterior density of the multiple target state recursively in time, via the following prediction step and update step [2], [5], [6]:

$$\pi_{k|k-1}(X_k|Z_{1:k-1}) = \int f_{k|k-1}(X_k|X)\pi_{k-1}(X|Z_{1:k-1})\delta X,$$

$$\pi_k(X_k|Z_{1:k}) = \frac{g_k(Z_k|X_k)\pi_{k|k-1}(X_k|Z_{1:k-1})}{\int g_k(Z_k|X)\pi_{k|k-1}(X|Z_{1:k-1})\delta X},$$

where $f_{k|k-1}(\cdot|\cdot)$ and $g_k(\cdot|\cdot)$ are the multi-target transition and likelihood respectively, and integrals are by convention the *FISST set integrals* defined for a function $f : \mathcal{F}(\mathbb{X}) \to \mathbb{R}$ by

$$\int f(X)\delta X = \sum_{i=0}^{\infty} \frac{1}{i!} \int f(\{x_1, ..., x_i\}) dx_1 \cdots dx_i.$$

The multi-target transition and likelihood model target birth, death, detection uncertainty and clutter. The posterior multi-target density encapsulates all statistical information about the multi-target state given the observed data.

Despite its rigorous foundations and theoretical elegance ,the Bayes multiple target filter is generally intractable due the inherently combinatorial nature of multiple target probability densities, as well as the usual complications associated with the curse of dimensionality [2]. In addition, the multi-target Bayes filter in its current form does not provide tracks or target trajectories. Simply augmenting each target state with an identity or label does not solve the problem since the elements of a multi-target state do not necessarily have distinct labels, even though the elements themselves are distinct. In the next section, we introduce the notion of a labeled RFS to alleviate this problem. Moreover, we show that the Bayes labeled multi-target recursion admits a closed form solution.

III. Generalized Labeled Multi-Bernoulli RFS

A. Labeled RFS

This subsection introduces the notion of a labeled RFS. In applications such as multi-target tracking, apart from the states of the object, the identities of the objects are also required. This can be achieved by augmenting a label $\ell \in \mathbb{L}$, where $\mathbb{L}$ is the (discrete) space of labels, to the state $x \in \mathbb{X}$. However, care must be taken in the application of the multi-target Bayes recursion to ensure uniqueness of identities or distinctness of labels.

Definition: For any $\tilde{X} \subset \mathbb{X} \times \mathbb{L}$, let $\mathcal{L}(\tilde{X})$ denote the set of labels of $\tilde{X}$, i.e. $\mathcal{L}(\tilde{X}) \triangleq \{\ell : (x, \ell) \in \tilde{X}\}$. A labeled RFS with state space $\mathbb{X}$ and (discrete) label space $\mathbb{L}$ is an RFS on $\mathbb{X} \times \mathbb{L}$ such that for each realization $\tilde{X}$ we have $|\mathcal{L}(\tilde{X})| = |\tilde{X}|$.

Note that the set integral for a function $\tilde{f} : \mathcal{F}(\mathbb{X} \times \mathbb{L}) \to \mathbb{R}$ is defined by

$$\int \tilde{f}(\tilde{X})\delta\tilde{X} = \sum_{i=0}^{\infty} \frac{1}{i!} \int \sum_{(\ell_1,...\ell_n) \in \mathbb{L}^n} \tilde{f}(\{(x_1, \ell_1), ..., (x_i, \ell_i)\}) dx_1 \cdots dx_i.$$

[1]For simplicity, in this paper, we shall not distinguish a FISST set derivative of a belief functional and a probability density. While the former is not a probability density [2], it is, equivalent to a probability density relative to the distribution of a Poisson RFS with unit intensity (see [10]).

and that a labeled RFS is related to its unlabeled version via

$$\pi(\{x_1,...,x_n\}) = \sum_{(\ell_1,...\ell_n)\in\mathbb{L}^n} \tilde{\pi}(\{(x_1,\ell_1),...,(x_n,\ell_n)\})$$

Example 1: A labeled Poisson RFS $\tilde{X}$ with state space $\mathbb{X}$ and label space $\mathbb{L}$= $\{\alpha_i : i \in \mathbb{N}\}$ is a Poisson RFS X on $\mathbb{X}$ with intensity v tagged/augmented with labels from $\mathbb{L}$. A sample from a labeled Poisson RFS with intensity v can be generated by the following procedure:

$\tilde{X} = \emptyset$;
$n \sim Poiss(\langle v, 1\rangle)$;
for $i = 1 : n$
  $x \sim v(\cdot)/\langle v, 1\rangle$;
  $\tilde{x} = (x, \alpha_i)$;
  $\tilde{X} = \tilde{X} \cup \{\tilde{x}\}$;
end;

The set of unlabeled elements generated by the above procedure is a Poisson RFS. However, the set of labeled elements is not necessarily a Poisson RFS on $\mathbb{X}\times$ $\mathbb{L}$. Indeed the density of a labeled Poisson RFS is given by

$$\tilde{\pi}(\{(x_1,\ell_1),...,(x_n,\ell_n)\}) = \frac{\delta_{\mathbb{L}(n)}(\{\ell_1,...\ell_n\})}{e^{\langle v,1\rangle}n!}\prod_{i=1}^{n} v(x_i)$$

where $\mathbb{L}(n) = \{\alpha_1, ..., \alpha_n\}$, $\delta_Y(X) = 1$ if $X = Y$ and zero otherwise. It can be easily verified that $\int \tilde{\pi}(\tilde{X})\delta\tilde{X} = 1$.

Example 2: A labeled multi-Bernoulli $\tilde{X}$ with state space $\mathbb{X}$ and label space $\mathbb{L}$ is a multi-Bernoulli RFS X on $\mathbb{X}$ augmented with labels corresponding to the successful Bernoulli's. A sample from a labeled multi-Bernoulli RFS is generated from the multi-Bernoulli parameters $\{(r^{(i)}, p^{(i)})\}_{i=1}^{M}$, by the following procedure:

$\tilde{X} = \emptyset$;
for $i = 1, ..., M$
  $u \sim U[0, 1]$;
  if $u \le r^{(i)}$,
    $x \sim p^{(i)}$;
    $\tilde{x} = (x, \alpha_i)$;
  end;
  $\tilde{X} = \tilde{X} \cup \{\tilde{x}\}$;
end;

It is clear that the above procedure always generates a finite set of augmented states with distinct labels and that the set of unlabeled elements is a multi-Bernoulli. However, the set of labeled elements is not necessarily a multi-Bernoulli RFS on $\mathbb{X}\times$ $\mathbb{L}$. The probability density of a labeled multi-Bernoulli is given by

$$\tilde{\pi}(\{(x_1,\ell_1),...,(x_n,\ell_n)\}) = \delta_n(|\{\ell_1,...\ell_n\}|) \times$$
$$1_{\mathbb{L}(M)}(\{\ell_1,...\ell_n\})\prod_{i=1}^{M}\left(1-r^{(i)}\right)\prod_{j=1}^{n}\frac{r^{(\ell_j)}p^{(\ell_j)}(x_j)}{1-r^{(\ell_j)}}$$

where $1_Y(X) = 1$ if $X \subseteq Y$ and 0 otherwise. Again it can be verified that $\int \tilde{\pi}(\tilde{X})\delta\tilde{X} = 1$.

The products $\prod_{i=1}^{M}\left(1-r^{(i)}\right)\prod_{j=1}^{n}\frac{r^{(\ell_j)}p^{(\ell_j)}(x_j)}{1-r^{(\ell_j)}}$ can be written as

$$\prod_{i=1}^{M}\left(\sum_{j=1}^{n}\delta_{\alpha_i}(\ell_j)r^{(i)}p^{(i)}(x_j) + \left(1-1_{\{\ell_1,...\ell_n\}}(\alpha_i)\right)(1-r^{(i)})\right)$$

Instead of the parameter set $\{(r^{(i)}, p^{(i)})\}_{i=1}^{M}$ and corresponding labels α_i, we can generalize this to an arbitrary finite parameter set $\{(r^{(\theta)}, p^{(\theta)}) : \theta \in \Theta\}$ and labels $\alpha(\theta)$ where $\alpha : \Theta \to \mathbb{L}$ is 1-1 mapping. The density can be written in terms of a product over the parameter set Θ as follows

$$\tilde{\pi}(\tilde{X}) = \delta_{|\tilde{X}|}(|\mathcal{L}(\tilde{X})|)1_{\alpha(\Theta)}(\mathcal{L}(\tilde{X}))\left[\Phi(\tilde{X};\cdot)\right]^{\Theta}$$

where

$$\begin{aligned}\Phi(\tilde{X};\theta) &= \sum_{(x,\ell)\in\tilde{X}}\delta_{\alpha(\theta)}(\ell)r^{(\theta)}p^{(\theta)}(x) \\ &\quad + \left(1-1_{\mathcal{L}(\tilde{X})}(\alpha(\theta))\right)(1-r^{(\theta)})\end{aligned}$$

B. Generalized Labeled Multi-Bernoulli RFS

Definition: A generalized labeled multi-Bernoulli RFS is an RFS on $\mathbb{X}\times$ $\mathbb{L}$ with probability density of the form

$$\tilde{\pi}(\tilde{X}) = \delta_{|\tilde{X}|}(|\mathcal{L}(\tilde{X})|)\sum_{c\in\mathbb{C}}w^{(c)}\left(\mathcal{L}(\tilde{X})\right)\left[\tilde{p}^{(c)}\right]^{\tilde{X}}$$

where $\mathbb{C}$ is a discrete index set, $w^{(c)}$ and $\tilde{p}^{(c)}$ satisfy

$$\begin{aligned}\int \tilde{p}^{(c)}(x,\ell)dx &= 1 \\ \sum_{c\in\mathbb{C}}\sum_{L\subseteq\mathbb{L}}w^{(c)}(L) &= 1,\end{aligned}$$

Remark: $\delta_{|\tilde{X}|}(|\mathcal{L}(\tilde{X})|)$ is the distinct-label indicator of $\tilde{X}$, i.e. takes on the value 1 if the labels of $\tilde{X}$ are distinct and zero otherwise. Hence, the RFS defined above is a labeled RFS.

Remark: The cardinality distribution of a generalized labeled multi-Bernoulli RFS is given by

$$\rho(n) = \sum_{c\in\mathbb{C}}\sum_{L\in\mathcal{F}_n(\mathbb{L})}w^{(c)}(L)$$

It can be shown that the labeled Poisson and labeled Multi-Bernoulli RFSs are special cases of generalized labeled multi-Bernoulli RFSs.

IV. Closed Form Bayes Multi-Target Recursion

In this section we show that the generalized labeled multi-Bernoulli density is closed under the Bayes multi-target prediction and update. In other words, if we start with a generalized labeled multi-Bernoulli initial prior, then all subsequent multi-target prediction and update densities are also generalized labeled multi-Bernoulli.

As mentioned earlier, for the multi-target Bayes filter to provide tracks, we augment each target state with an identity or label. To address the uniqueness of labels we evoke the labeled RFS developed in the previous section. The labeling scheme used in this work is as follows. Each target is identified by a label $\ell = (k, i)$, where k is the time of birth and i is the index of the state of the target when it is born. Targets born from labeled Poisson or labeled multi-Bernoulli have unique labels as shown previously. Using the following notations for the space of labels,

$$\begin{aligned} \mathbb{L}_k &= \{k\} \times \mathbb{N}, \\ \mathbb{L}_{0:k} &= \mathbb{L}_{0:k-1} \cup \mathbb{L}_k, \\ \mathbb{L}_k(M) &= \{k\} \times \{1, ...M\}, \end{aligned}$$

a target born at time k, has state $\tilde{x} \in \mathbb{X}\times\mathbb{L}_k$ and a multi-target state at time k, is a finite subset of $\mathbb{X}\times\mathbb{L}_{0:k}$.

A. Multi-Target Prediction

The multi-target dynamical model involves thinning, Markov shifts and superposition of new targets.

Given a multi-target state $\tilde{X}'$ at time $k-1$, each target $(x', \ell') \in \tilde{X}'$ either dies with probability $q_{S,k|k-1}(x', \ell')$ and takes on the value $\emptyset$, or continues to exist at time k with probability $p_{S,k|k-1}(x', \ell') = 1 - q_{S,k|k-1}(x'_i, \ell'_i)$ and moves to a new state (x, ℓ) with probability density $f_{k|k-1}(x|x', \ell')\delta_\ell(\ell')$. (note that the label or identity of the target is preserved in the transition). Assuming that, conditional on $\tilde{X}'$, each of the transition of the target kinematic state are mutually independent, then the set of surviving targets at time k is modeled by the labeled multi-Bernoulli RFS with parameter set $\{(r^{(x',\ell')}, p^{(x',\ell')}) : (x', \ell') \in \tilde{X}'\}$ where $r^{(x',\ell')} = p_{S,k|k-1}(x', \ell')$, $p^{(x',\ell')} = f_{k|k-1}(\cdot|x', \ell')$, and labeling function $\alpha : \tilde{X}' \to \mathbb{L}_{0:k-1}$ defined by $\alpha(x', \ell') = \ell'$. The multi-target density for surviving targets at time k given the multi-target state $\tilde{X}'$ at time $k-1$ is thus

$$\tilde{f}_{S,k|k-1}(\tilde{X}|\tilde{X}') = \delta_{|\tilde{X}|}(|\mathcal{L}(\tilde{X})|)1_{\mathcal{L}(\tilde{X}')}(\mathcal{L}(\tilde{X}))\Big[\Phi_{k|k-1}(\tilde{X};\cdot)\Big]^{\tilde{X}'}$$

where

$$\begin{aligned} \Phi_{k|k-1}(\tilde{X}; x', \ell') = &\sum_{(x,\ell)\in\tilde{X}} \delta_{\ell'}(\ell) p_{S,k|k-1}(x', \ell') f_{k|k-1}(x|x', \ell') \\ &+ \Big(1 - 1_{\mathcal{L}(\tilde{X})}(\ell')\Big) q_{S,k|k-1}(x', \ell') \end{aligned}$$

Given a prior $\tilde{\pi}_{k-1}$ on the multi-target state at time $k-1$, the density of the surviving multi-target state at time k is given by the marginal

$$\tilde{\pi}_{S,k|k-1}(\tilde{X}) = \int \tilde{f}_{S,k|k-1}(\tilde{X}|\tilde{X}')\tilde{\pi}_{k-1}(\tilde{X}')\delta\tilde{X}'.$$

Proposition: If the prior for the multi-target state at time $k-1$ is a generalized labeled multi-Bernoulli of the form,

$$\tilde{\pi}_{k-1}(\tilde{X}') = \delta_{|\tilde{X}'|}(|\mathcal{L}(\tilde{X}')|) \sum_{c\in\mathbb{C}} w_{k-1}^{(c)} \left(\mathcal{L}(\tilde{X}')\right) \left[\tilde{p}_{k-1}^{(c)}\right]^{\tilde{X}'}$$

Then the surviving multi-target state at time k is also a generalized labeled multi-Bernoulli with density given by

$$\begin{aligned} \tilde{\pi}_{S,k|k-1}(\tilde{X}) = \delta_{|\tilde{X}|}(|\mathcal{L}(\tilde{X})|) \sum_{c\in\mathbb{C}} \sum_{L'\subseteq\mathbb{L}_{0:k-1}} 1_{L'}(\mathcal{L}(\tilde{X})) \times \\ w_{k-1}^{(c)}(L')\Big[q_{S,k|k-1}^{(c)}\Big]^{L'} \left[\frac{K_{S,k|k-1}^{(c)}}{q_{S,k|k-1}^{(c)}}\right]^{\mathcal{L}(\tilde{X})} \Big[\tilde{p}_{S,k|k-1}^{(c)}\Big]^{\tilde{X}} \end{aligned}$$

where

$$\begin{aligned} \tilde{p}_{S,k|k-1}^{(c)}(x, \ell) &= \frac{\left\langle f_{k|k-1}(x|\cdot, \ell), p_{S,k|k-1}(\cdot, \ell)\tilde{p}_{k-1}^{(c)}(\cdot, \ell)\right\rangle}{K_{S,k|k-1}^{(c)}(\ell)} \\ K_{S,k|k-1}^{(c)}(\ell) &= \int\left\langle f_{k|k-1}(x|\cdot, \ell), p_{S,k|k-1}(\cdot, \ell)\tilde{p}_{k-1}^{(c)}(\cdot, \ell)\right\rangle dx \\ q_{S,k|k-1}^{(c)}(\ell') &= \left\langle q_{S,k|k-1}(\cdot, \ell'), \tilde{p}_{k-1}^{(c)}(\cdot, \ell')\right\rangle \end{aligned}$$

The new born target set at time k is modeled as labeled Poisson RFS with intensity γ_k. and label space $\mathbb{L}_k$. the multi-target density of the new born targets are

$$f_{B,k}(\tilde{X}) = \delta_{\mathbb{L}_k(|\tilde{X}|)}(\mathcal{L}(\tilde{X})\}) \frac{e^{-\langle\gamma_k,1\rangle}}{|\tilde{X}|!} \prod_{(x,\ell)\in\tilde{X}} \gamma_k(x)$$

Since, the multi-target state at time k is the union of the surviving targets and new born targets, the multi-target transition density is the convolution

$$\begin{aligned} \tilde{f}_{k|k-1}\left(\tilde{X}|\tilde{X}'\right) &= \sum_{W\subseteq\tilde{X}} \tilde{f}_{B,k}(\tilde{X} - W)\tilde{f}_{S,k|k-1}(W|\tilde{X}') \\ &= \tilde{f}_{B,k}(\tilde{X} - \tilde{X}_S)\tilde{f}_{S,k|k-1}(\tilde{X}_S|\tilde{X}') \end{aligned}$$

where $\tilde{X}_S = \{(x, \ell) \in \tilde{X} : \mathcal{L}(\tilde{X}) \subset \mathbb{L}_{0:k-1}\}$. This can be verified by noting that $f_{S,k|k-1}(W|\tilde{X}')$ in the convolution sum is zero, except possibly for $W \subseteq \tilde{X}_S$, moreover, if $W \subset \tilde{X}_S, W \neq \emptyset$ then $f_{B,k}(\tilde{X} - W)$ is zero because $\tilde{X} - W$ has labels belonging to $\mathbb{L}_{0:k-1}$. Note that since the label space $\mathbb{L}_k$ of the birth targets and the label space $\mathbb{L}_{0:k-1}$ of the surviving targets are mutually exclusive, the superposition of the birth targets and surviving targets is indeed a labeled RFS (this is not true in general).

The predicted multi-target density from time $k-1$ to k is

$$\begin{aligned} \tilde{\pi}_{k|k-1}(\tilde{X}) &= \int \tilde{f}_{k|k-1}(\tilde{X}|\tilde{X}')\tilde{\pi}_{k-1}(\tilde{X}')\delta\tilde{X}' \\ &= \tilde{f}_{B,k}(\tilde{X} - \tilde{X}_S)\tilde{\pi}_{S,k|k-1}(\tilde{X}_S) \end{aligned}$$

B. Multi-Target Bayes Update

A given state $(x, \ell) \in \tilde{X}$ is either detected with probability $p_D(x)$ and generates an observation z with likelihood $g(z|x)$, or missed with probability $q_D(x) = 1 - p_D(x)$, i.e. each state $x \in X$ generates a Bernoulli RFS $\Theta(x)$ with $r = p_D(x)$ and $p(\cdot) = g(\cdot|x)$. In addition, the sensor also receives a set of false alarms or clutter which can be modeled as a Poisson RFS K with intensity function $\kappa(\cdot)$. In essence, the observation RFS is the superposition of Poisson clutter with thinned and Markov-shifted measurements.

Assuming that, conditional on the multi-target state $\tilde{X}$, the constituent RFSs are mutually independent, the multi-target measurement likelihood is given by [2]

$$\tilde{g}(Z|\tilde{X}) = e^{-\langle\kappa,1\rangle}\kappa^Z q_D^{\tilde{X}} \sum_\theta \left[\psi_Z(\cdot;\theta)\right]^{\tilde{X}}$$

where the mapping $\theta : \mathcal{L}(\tilde{X}) \to \{0, 1, ..., |Z|\}$, such that $\theta(\ell) = \alpha(\ell') > 0$ implies $\ell = \ell'$, and

$$\psi_Z(x, \ell; \theta) = \begin{cases} \frac{p_D(x)}{q_D(x)} \frac{g(z_{\theta(\ell)}|x)}{\kappa(z_{\theta(\ell)})} & \theta(\ell) > 0 \\ 1 & \theta(\ell) = 0 \end{cases}$$

Proposition: Suppose that the prior distribution is a generalized labeled multi-Bernoulli of the form,

$$\tilde{\pi}(\tilde{X}) = \delta_{|\tilde{X}|}(|\mathcal{L}(\tilde{X})|) \sum_{c \in \mathbb{C}} w^{(c)} \left(\mathcal{L}(\tilde{X})\right) \left[\tilde{p}^{(c)}\right]^{\tilde{X}}$$

Then, under the multi-target measurement model, the posterior distribution is also a generalized multi-Bernoulli given by

$$\pi(\tilde{X}|Z) = \delta_{|\tilde{X}|}(|\mathcal{L}(\tilde{X})|) \sum_{c \in \mathbb{C}} \sum_{\theta} \tilde{w}^{(c,\theta)}(\mathcal{L}(\tilde{X})|Z) \left[\tilde{p}^{(c,\theta)}(\cdot|Z)\right]^{\tilde{X}}$$

where

$$\tilde{w}^{(c,\theta)}(\mathcal{L}(\tilde{X})|Z) \propto w^{(c)}(\mathcal{L}(\tilde{X})) \left[K_D^{(c,\theta)}(\cdot|Z)\right]^{\mathcal{L}(\tilde{X})}$$

$$\tilde{p}^{(c,\theta)}(x,\ell|Z) = \frac{\tilde{p}^{(c)}(x,\ell) q_D(x,\ell) \psi_Z(x,\ell;\theta)}{K_D^{(c,\theta)}(\ell|Z)}$$

$$K_D^{(c,\theta)}(\ell|Z) = \left\langle \tilde{p}^{(c)}(\cdot,\ell) q_D(\cdot,\ell), \psi_Z(\cdot,\ell;\theta) \right\rangle$$

V. Demonstration

Consider a 10 target scenario on the region $[-1000, 1000]m \times [-1000, 1000]m$. Targets move with constant velocity as shown in Figure 1.

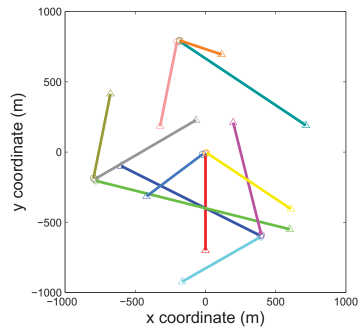

Figure 1. Target trajectories. Start/Stop positions given by ○/△.

The kinematic target state is a vector of planar position and velocity $x_k = [\, p_{x,k}, p_{y,k}, \dot{p}_{x,k}, \dot{p}_{y,k} \,]^T$. Measurements are noisy vectors of planar position only $z_k = [\, z_{x,k}, z_{y,k} \,]^T$. The single-target state space model is a linear Gaussian transition density $f_{k|k-1}(x|\zeta) = \mathcal{N}(x; F_k\zeta, Q_k)$ and linear Gaussian likelihood $g_k(z|x) = \mathcal{N}(z; H_k x, R_k)$ with parameters

$$F_k = \begin{bmatrix} I_2 & \Delta I_2 \\ 0_2 & I_2 \end{bmatrix} \quad Q_k = \sigma_\nu^2 \begin{bmatrix} \frac{\Delta^4}{4} I_2 & \frac{\Delta^3}{2} I_2 \\ \frac{\Delta^3}{2} I_2 & \Delta^2 I_2 \end{bmatrix}$$

$$H_k = \begin{bmatrix} I_2 & 0_2 \end{bmatrix} \qquad R_k = \sigma_\varepsilon^2 I_2$$

where I_n and 0_n denote the $n \times n$ identity and zero matrices respectively, $\Delta = 1s$, $\sigma_\nu = 5m/s^2$ and $\sigma_\varepsilon = 10m$. The survival and detection probabilities are $p_{S,k} = 0.99$ and $p_{D,k} = 0.98$. Births follow a Poisson RFS with intensity $\gamma_k(x) = \sum_{i=1}^4 w_\gamma \mathcal{N}(x; m_\gamma^{(i)}, P_\gamma)$,where $w_\gamma = 0.03$, $m_\gamma^{(1)} = [\, 0,0,0,0 \,]^T$, $m_\gamma^{(2)} = [\, 400,-600,0,0 \,]^T$, $m_\gamma^{(3)} = [\, -800,-200,0,0 \,]^T$, $m_\gamma^{(4)} = [\, -200,800,0,0 \,]^T$, and $P_\gamma = \mathrm{diag}([\, 10,\ 10,\ 10,\ 10 \,]^T)^2$. Clutter follows a Poisson RFS with uniform spatial density and a mean rate of 60 returns per scan and average intensity of $\lambda_k = 1.5 \times 10^{-5} m^{-2}$. Figure 2 shows the measurements for x and y coordinates versus time.

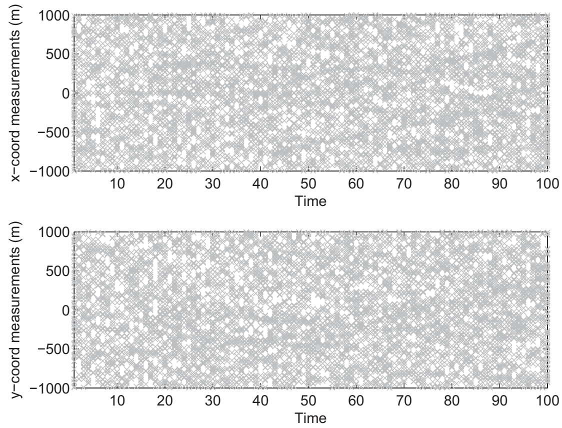

Figure 2. Measurements for x and y coordinates versus time.

The multi-target Bayes (MTB) filter is implemented via the previously derived closed form solution. The (continuous) kinematic component of the target state, and hence the individual target tracks, are represented as Gaussian mixtures which are predicted and updated using the standard Gaussian sum filter equations. The (discrete) label component of the target states are explicitly expanded and represented as a grid or table, where each point corresponds to a particular component in the generalized labeled Multi-Bernoulli density. The multi-target prediction and update steps are then calculated on component by component basis, using Murty's algorithm [11] to generate predicted and updated components respectively, in order of decreasing weights post prediction and update. To ensure tractability of the filter, the predicted and updated multi-target density must be truncated, retaining only the components with the highest weights. This is achieved by employing Murty's algorithm to calculate newly predicted and updated components in decreasing order of predicted and updated weights and stopping at a determined threshold. For each existing component, the number of new components calculated and stored in each forward propagation is set to be proportional to the weight of the original component, subject to a maximum of 100 terms. The resultant posterior at each time step is then further truncated to a maximum of 100 terms.

Figure 3 shows the filter estimates for single sample run in x and y coordinates versus time. It can be seen that the filter initiates and terminates tracks with a very short delay, and generally produces accurate estimates of the individual target states and the total target numbers, with very few dropped or false tracks. A direct performance comparison with the GM-CPHD filter [12], [13] is also undertaken. Figure 4 compares the true and estimated cardinality for the two filters versus time. In this scenario, the MTB filter vastly outperforms the CPHD filter in estimating the number of targets. Figure 5 compares the OSPA penalty ($p = 1$, $c = 100m$) [14] for the two filters, further confirming that the MTB filter generally outperforms the CPHD filter. The breakup of the OSPA into localization and cardinality components shown in Figure 6 further confirm a noticeable improvement in both the localization and cardinality performance.

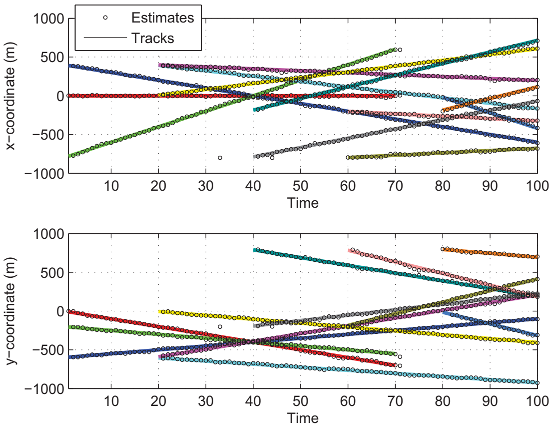

Figure 3. Estimates and tracks for x and y coordinates versus time.

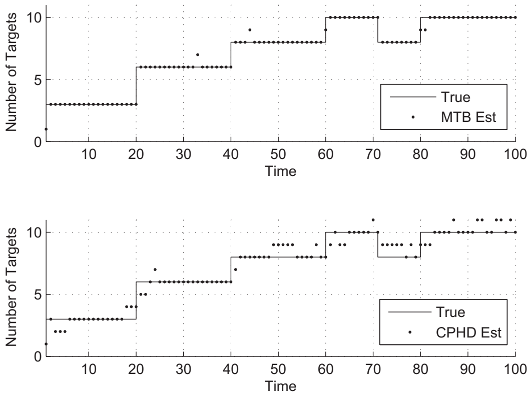

Figure 4. Number of targets versus time for MTB and CPHD filters.

VI. Conclusion

This paper proposes the first random finite set conjugate prior corresponding to the multi-target likelihood in the multi-target Bayes filter. The resultant recursion propagates an exact solution to the multi-target posterior recursively in time. The solution is built upon the concept of a labeled random finite set, and formally incorporates the propagation and estimation of unique track labels within the filtering framework. A tractable implementation is demonstrated for linear Gaussian multi-target models, by use of the Gaussian sum filter to propagate individual target tracks, and more importantly by application of Murty's algorithm to dynamically truncate posterior components, thereby ensuring a polynomial time complexity. Preliminary results indicate that the proposed MTB filter performs remarkably well, and significantly outperforms the CPHD filter in both localization and cardinality estimation.

Acknowledgements

Dr. Ba Tuong Vo is a recipient of the Australian Research Council Postdoctoral Fellowship DP0989007. Prof. Ba Ngu Vo is a recipient of the Australian Research Council Future Fellowship FT0991854.

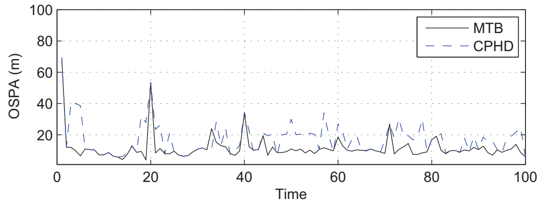

Figure 5. OSPA distance ($p = 1$, $c = 100m$) for MTB and CPHD filters.

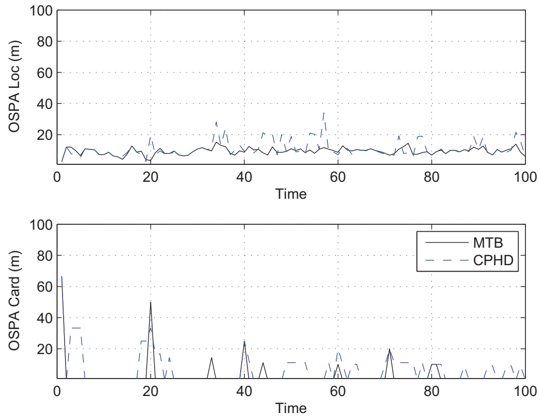

Figure 6. OSPA components for MTB and CPHD filters.

References

[1] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Academic Press, San Diego, 1988.

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

[3] J. Lund and E. Thonnes, "Perfect simulation and inference for point processes given noisy observations", *Computational Statistics*, Vol. 19, pp. 317–336, 2004.

[4] J. Moller and R. Waagepeterson, *Modern Statistics for Spatial Point Processes*, Scandinavian Journal of Statistics, Vol. 34, pp. 643–684, 2006.

[5] R. Mahler, "Multi-target Bayes filtering via first-order multi-target moments," *IEEE Trans. Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, 2003.

[6] R. Mahler, "PHD filters of higher order in target number," *IEEE Trans. Aerospace & Electronic Systems*, Vol. 43, No. 3, July 2007.

[7] B.-T. Vo, B.-N. Vo, and A. Cantoni, "The cardinality balanced multi-target multi-Bernoulli filter and its implementations," *IEEE Trans. Signal Processing*, Vol. 57, No. 2, pp. 409–423, Feb. 2009.

[8] B.-N. Vo, B.-T. Vo, N.-T. Pham and D. Suter, "Joint Detection and Estimation of Multiple Objects from Image Observations," *IEEE Trans. Signal Processing*, Vol. 58, No. 10, pp. 5129–5241, 2010.

[9] D. Daley and D. Vere-Jones, *An introduction to the theory of point processes*. Springer-Verlag, 1988.

[10] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multi-target filtering with random finite sets," *in IEEE Trans. Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1224–1245, 2005.

[11] K. G. Murty, *An Algorithm for Ranking all the Assignments in Order of Increasing Cost*, Operations Research, Vol. 16, No. 3. (1968), pp. 682-687.

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

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

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