Rao-Blackwellised PHD SLAM

Mullane, John; Vo, Ba-Ngu; Adams, Martin David

2010

Mullane, J., Vo, B. N., & Adams, M. D. (2010). Rao-Blackwellised PHD SLAM. IEEE International Conference on Robotics and Automation, 5410-5416.

<https://hdl.handle.net/10356/90756>

<https://doi.org/10.1109/ROBOT.2010.5509626>

© 2010 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: [DOI: <http://dx.doi.org/10.1109/ROBOT.2010.5509626>].

Downloaded on 26 Apr 2023 18:42:06 SGT

Rao-Blackwellised PHD SLAM

John Mullane, Ba-Ngu Vo, Martin D. Adams

Abstract—This paper proposes a tractable solution to feature-based (FB) SLAM in the presence of data association uncertainty and uncertainty in the number of features. By modeling the feature map as a random finite set (RFS), a rigorous Bayesian formulation of the FB-SLAM problem that accounts for uncertainty in the number of features and data association is presented. As such, the joint posterior distribution of the set-valued map and vehicle trajectory is propagated forward in time as measurements arrive. A first order solution, coined the PHD-SLAM filter, is derived, which jointly propagates the posterior PHD or intensity function of the map and the posterior distribution of the trajectory of the vehicle. A Rao-Blackwellised implementation of the PHD-SLAM filter is proposed based on the Gaussian mixture PHD filter for the map and a particle filter for the vehicle trajectory. Simulated results demonstrate the merits of the proposed approach, particularly in situations of high clutter and data association ambiguity.

I. INTRODUCTION

Following seminal developments in autonomous robotics [1], the problem of simultaneous localisation and mapping (SLAM) gained widespread interest, with numerous potential applications ranging from robotic planetary exploration to intelligent surveillance [2], [3], [4], [5]. This paper focusses on the Feature-based (FB) approach that decomposes physical environmental landmarks into parametric representations such as points, lines, circles, corners etc., known as features. FB maps are comprised of an unknown number of features at unknown spatial locations [6]. Estimating a feature map, thus requires the joint estimation of the number of the features and their locations.

Current state-of-the-art FB-SLAM solutions address two separate problems [3]:

- determining the measurement (to feature) association; and
- given the association, estimation of feature map and vehicle pose via stochastic filtering.

This two-tiered approach to SLAM is sensitive to data association (DA) uncertainty [7], because the current Bayesian SLAM framework does not fully integrate DA uncertainty (or uncertainty in feature / measurement number) into the map estimate. Specifically, the estimation component of the feature map assumes known DA. While a two-tiered approach is efficient and works well when DA uncertainty is low, it is not robust to high DA uncertainty e.g. in scenarios with high clutter and dense features and/or when the vehicle is moving/turning quickly. A SLAM solution that is robust

to DA under high clutter requires a framework that fully integrates DA uncertainty into the estimation of the map (and vehicle trajectory).

This paper advocates a fully integrated Bayesian framework for FB-SLAM under DA uncertainty and unknown number of features. The key to this formulation is the representation of the map as a finite set of features, which was first proposed in [8], and is further argued in this work from an estimation viewpoint (see section II-A). Using RFS theory, the FB-SLAM problem is then posed as a Bayesian filtering problem in which the joint posterior distribution of the set-valued map and vehicle trajectory are propagated forward in time as measurements arrive. The proposed Bayesian FB-SLAM framework allows for the joint, on-line estimation of the vehicle trajectory, the feature locations and the number of features in the map.

Preliminary studies of RFS-SLAM using ‘brute force’ implementations can be found in [8], [9]. In this paper, however, a tractable first order approximation, coined the (Probability Hypothesis Density) PHD-SLAM filter, is derived, which propagates the posterior PHDs of multiple trajectory-conditioned maps and the posterior distribution of the trajectory of the vehicle. A Rao-Blackwellised (RB) implementation of the PHD-SLAM filter is proposed based on the Gaussian mixture PHD filter for the map and a particle filter for the vehicle trajectory. Simulated results demonstrate the merits of the proposed approach, particularly in situations of high clutter and data association ambiguity.

II. BAYESIAN FEATURE-BASED SLAM

This section discusses the mathematical representation of the map and presents a Bayesian formulation of the FB-SLAM problem under uncertainty in DA and number of features. In particular, it is argued that fundamentally the map is a finite set and the concept of a random finite set is essential to a Bayesian FB-SLAM formulation.

A. Mathematical representation of the Feature Map

In the context of jointly estimating the number of features and their values, the collection of features, referred to as the feature map, is naturally represented as a finite set. The rationale behind this representation traces back to a fundamental consideration in estimation theory - estimation error. Without a meaningful notion of estimation error, estimation has very little meaning. Existing SLAM formulations do not admit a rigorous notion of mapping error despite the fact that it is equally as important as localisation error. To illustrate this point, recall that in existing SLAM formulations the map is

John Mullane and Martin D. Adams are with the school of Electrical & Electronic Engineering, Nanyang Technological University, Singapore. {jsmullane, eadams}@ntu.edu.sg

Ba-Ngu Vo is with the School of Electrical & Electronic Engineering, University of Melbourne, Australia. bnvo@unimelb.edu.au

constructed by stacking features into a vector, and consider the simplistic scenarios depicted in figure 1.

Fig. 1. Hypothetical scenario showing a fundamental inconsistency with vector representations of feature maps. If M is the true map, how should the error be assigned when the number of features in the map estimate, $\hat{M}$, is incorrect?

A finite set representation of the map, $\mathcal{M}_k = \{m_k^1, \dots, m_k^{N_k}\}$, where $m_k^1, \dots, m_k^{N_k}$ are the N_k features present at time k , admits a mathematically consistent notion of estimation error since the distance between sets is a well understood concept. In contrast, stacking individual features into a single vector does not admit a satisfactory notion of error as illustrated in Figure 1.

For the most common sensor models considered in SLAM, the order in which sensor readings are recorded at each sampling instance bears no significance. Moreover, the number of measurements, $\mathfrak{Z}(k)$, at any given time is not fixed due to detection uncertainty, spurious measurements and unknown feature number. Thus, this type of measurement may also be represented as a finite set of readings, $\mathcal{Z}_k = \{z_k^1, z_k^2, \dots, z_k^{\mathfrak{Z}(k)}\}$.

B. The Bayes FB-SLAM Filter

In the Bayesian estimation paradigm, the state/parameter and measurement are treated as realizations of random variables. Since the map (and the measurement) is a finite set, the concept of a random finite set is essential for Bayesian map estimation. In essence, a *random finite set* (RFS) is simply a finite-set-valued random variable. Similar to random vectors, the probability density (if it exists) is a very useful descriptor of an RFS, especially in filtering and estimation. However, the space of finite sets does not inherit the usual Euclidean notion of integration and density. Hence, standard tools for random vectors are not appropriate for random finite sets. Mahler's Finite Set Statistics (FISST) provides practical mathematical tools for dealing with RFSs [10], [11], based on a notion of integration and density that is consistent with point process theory [12].

Let $\mathcal{M}$ be the RFS representing the entire unknown map and let $\mathcal{M}_{k-1}$ be the RFS representing the subset of the map that has passed through the field-of-view (FOV) of the on-board sensor with trajectory $X_{0:k-1} = [X_0, X_1, \dots, X_{k-1}]$ at time $k-1$, i.e.

$$\mathcal{M}_{k-1} = \mathcal{M} \cap \text{FOV}(X_{0:k-1}). \quad (1)$$

Note that $\text{FOV}(X_{0:k-1}) = \text{FOV}(X_0) \cup \text{FOV}(X_1) \cup \dots \cup \text{FOV}(X_{k-1})$. $\mathcal{M}_{k-1}$ therefore represents the set on the

space of features which intersects with the union of individual FOVs, over the vehicle trajectory up to and including time $k-1$. Given this representation, $\mathcal{M}_{k-1}$ evolves in time according to,

$$\mathcal{M}_k = \mathcal{M}_{k-1} \cup \left(\text{FOV}(X_k) \cap \bar{\mathcal{M}}_{k-1} \right) \quad (2)$$

where $\bar{\mathcal{M}}_{k-1} = \mathcal{M} - \mathcal{M}_{k-1}$ (note the difference operator used here is the set difference), i.e the set of features that are not in $\mathcal{M}_{k-1}$. Modeling the vehicle dynamics by the standard Markov process with transition density $f_X(X_k|X_{k-1}, U_k)$, where U_k denotes the control input at time k , the joint transition density of the map and the vehicle pose can be written as,

$$f_{k|k-1}(\mathcal{M}_k, X_k | \mathcal{M}_{k-1}, X_{k-1}, U_k) = f_{\mathcal{M}}(\mathcal{M}_k | \mathcal{M}_{k-1}, X_k) f_X(X_k | X_{k-1}, U_k). \quad (3)$$

where $f_{\mathcal{M}}(\mathcal{M}_k | \mathcal{M}_{k-1}, X_k)$ denotes the RFS map transition density, modelled by eqn.(2).

The measurement $\mathcal{Z}_k$ received by the vehicle with pose X_k , at time k , can be modeled by

$$\mathcal{Z}_k = \bigcup_{m \in \mathcal{M}_k} \mathcal{D}_k(m, X_k) \cup \mathcal{C}_k(X_k) \quad (4)$$

where $\mathcal{D}_k(m, X_k)$ is the RFS of measurements generated by a feature at m and $\mathcal{C}_k(X_k)$ is the RFS of the spurious measurements at time k . Therefore $\mathcal{Z}_k$ consists of a random number, $\mathfrak{Z}(k)$, of measurements, whose order of appearance has no physical significance. The RFS of the measurements generated by a feature at m is a Bernoulli RFS¹ given by, $\mathcal{D}_k(m, X_k) = \emptyset$ with probability $1 - p_D(m|X_k)$ and $\mathcal{D}_k(m, X_k) = \{z\}$ with probability density $p_D(m|X_k)g_k(z|m, X_k)$. For a given robot pose X_k , $p_D(m|X_k)$ is the probability of the sensor detecting a feature at m , and when conditioned on detection, $g_k(z|m, X_k)$ is the likelihood that a feature at m generates the measurement z . The RFS $\mathcal{C}_k(X_k)$ represents the spurious measurements registered, which may be dependent on the vehicle pose, X_k . It is interesting to note that by characterising clutter through a discrete distribution (Poisson) and density function (uniform), FB-SLAM related literature that probabilistically consider clutter are in fact implicitly adopting an RFS model [4], [5].

The Bayesian FB-SLAM recursion is next outlined. Let $p_k(\mathcal{M}_k, X_{0:k} | Z_{1:k}, U_{1:k}, X_0)$ denote the joint posterior density of the map $\mathcal{M}_k$, and the vehicle trajectory $X_{0:k}$. For clarity of exposition, the following abbreviations shall be adhered to,

$$p_{k|k-1}(\mathcal{M}_k, X_{0:k}) = p_{k|k-1}(\mathcal{M}_k, X_{0:k} | Z_{0:k}, U_{0:k-1}, X_0) \\ p_k(\mathcal{M}_k, X_{0:k}) = p_k(\mathcal{M}_k, X_{0:k} | Z_{0:k}, U_{0:k-1}, X_0)$$

¹The Bernoulli RFS is empty with a probability $1 - \epsilon$ and is distributed according to a density π with probability ϵ .

The recursion for a static feature map is then given as follows,

$$\begin{aligned}
p_{k|k-1}(\mathcal{M}_k, X_{0:k}) &= f_X(X_k|X_{k-1}, U_k) \times \\
&\int f_{\mathcal{M}}(\mathcal{M}_k|\mathcal{M}_{k-1}, X_k) p_{k-1}(\mathcal{M}_{k-1}, X_{1:k-1}) \delta \mathcal{M}_{k-1} \\
p_k(\mathcal{M}_k, X_{0:k}) &= \frac{g_k(\mathcal{Z}_k|X_k, \mathcal{M}_k) p_{k|k-1}(\mathcal{M}_k, X_{0:k})}{g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_0)}
\end{aligned} \tag{5}$$

where δ implies a set integral, and $g_k(\mathcal{Z}_k|X_k, \mathcal{M}_k)$ denotes the likelihood of the measurement $\mathcal{Z}_k$ given the pose X_k and map $\mathcal{M}_k$. The joint posterior density encapsulates all statistical information about the map and vehicle pose, that can be inferred from the measurements and control history up to time k . The Bayesian FB-SLAM recursion (5) integrates uncertainty in DA and number of features into a single Bayesian filter and does not require any separate DA steps nor feature management, as are classically adopted [2], [3], [7]. As with the standard Bayes filter, the above recursion is computationally intractable in general. The following section therefore investigates tractable approximations to the Bayes FB-SLAM filter.

III. THE PHD-SLAM FILTER

Since the full Bayes FB-SLAM filter is numerically intractable, it is necessary to look for tractable but principled approximations. The probability hypothesis density (PHD) approach which propagates the 1st order moment of the posterior multi-target RFS has proven to be both powerful and effective in multi-target filtering [11]. However, this technique cannot be directly applied to FB-SLAM which propagates the joint posterior density of the map and the vehicle trajectory. This section derives a recursion that jointly propagates the posterior PHD of the map and the posterior density of the vehicle trajectory.

A. The Posterior PHD of the Map

The integral of the PHD v over a set S gives the expected number of points of $\mathcal{M}$ that are in S . If the RFS $\mathcal{M}$ is Poisson, i.e. the number of points is Poisson distributed and the points themselves are independently and identically distributed, then the probability density of $\mathcal{M}$ can be constructed exactly from the PHD [13],

$$p(\mathcal{M}) = \frac{\prod_{m \in \mathcal{M}} v(m)}{\exp(\int v(m) dm)}. \tag{6}$$

In this sense, the PHD can be thought of as a 1st moment approximation of the probability density of an RFS.

In addition, the PHD construct allows an alternative notion of expectation for maps. A salient property of the PHD construct in map estimation is that the posterior PHD of the map is indeed the expectation of the trajectory-conditioned PHDs. More concisely,

$$v_k(m) = \mathbb{E}[v_k(m|X_{0:k})], \tag{7}$$

where the expectation is taken over the vehicle trajectory $X_{0:k}$. This result follows from standard properties of the PHD (intensity function) of an RFS, see for example classical texts such as [13]. Such an averaging property for map estimates is not available in existing SLAM approaches.

Apart from being a first order approximation of the posterior density of the map, the posterior PHD plays a vital role in the map estimation process itself. Given the joint posterior of the map and the trajectory, an estimate of vehicle trajectory can be computed by marginalising over the map to obtain the posterior of the vehicle trajectory and take the mean. It is well-known that the posterior mean of the vehicle trajectory is Bayes optimal. While the posterior density of the map can be obtained by marginalising over the vehicle trajectory, the expectation of the map is not defined. Fortunately, a Bayes optimal estimator for the map can be obtained using the posterior PHD $v_k(m|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0)$ of the map by integrating the posterior PHD to obtain the estimated number of features $\hat{N}_k$ and then finding the $\hat{N}_k$ highest the local maxima of the posterior PHD [14].

B. The RB PHD-SLAM recursion

Using standard conditional probability, the joint posterior density of the map and the trajectory can be decomposed as

$$\begin{aligned}
p_k(\mathcal{M}_k, X_{0:k}|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0) &= \\
p_k(X_{0:k}|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0) p_k(\mathcal{M}_k|\mathcal{Z}_{0:k}, X_{0:k}).
\end{aligned} \tag{8}$$

Thus, the recursion for the joint map-trajectory posterior density according to (5) is equivalent to jointly propagating the posterior density of the map conditioned on the trajectory and the posterior density of the trajectory. If, as before for compactness,

$$\begin{aligned}
p_{k|k-1}(\mathcal{M}_k|X_{0:k}) &= p_{k|k-1}(\mathcal{M}_k|\mathcal{Z}_{0:k-1}, X_{0:k}) \\
p_k(\mathcal{M}_k|X_{0:k}) &= p_k(\mathcal{M}_k|\mathcal{Z}_{0:k}, X_{0:k}) \\
p_k(X_{0:k}) &= p_k(X_{0:k}|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0)
\end{aligned}$$

then,

$$\begin{aligned}
p_{k|k-1}(\mathcal{M}_k|X_{0:k}) &= \int f_{\mathcal{M}}(\mathcal{M}_k|\mathcal{M}_{k-1}, X_k) \times \\
&p_{k-1}(\mathcal{M}_{k-1}|X_{0:k-1}) \delta \mathcal{M}_{k-1}
\end{aligned} \tag{9}$$

$$p_k(\mathcal{M}_k|X_{0:k}) = \frac{g_k(\mathcal{Z}_k|\mathcal{M}_k, X_k) p_{k|k-1}(\mathcal{M}_k|X_{0:k})}{g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_{0:k})} \tag{10}$$

$$\begin{aligned}
p_k(X_{0:k}) &= g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_{0:k}) \times \\
&\frac{f_X(X_k|X_{k-1}, U_{k-1}) p_{k-1}(X_{1:k-1})}{g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1})}.
\end{aligned} \tag{11}$$

The recursion defined by (9), (10), (11) is similar to FastSLAM [3], in the exploitation of the factorisation of the joint SLAM posterior. The difference is that the map and the measurements are random finite sets. Consequently, the propagation equations involve probability densities of random finite sets and marginalisation over the map involves set integrals.

Abbreviating $v_{k|k-1}(m|X_{0:k}) = v_{k|k-1}(m|Z_{0:k-1}, X_{0:k})$ and $v_k(m|X_{0:k}) = v_k(m|Z_{0:k}, X_{0:k})$, and following the trajectory-conditioned PHD mapping filter in [9], eqns.(9), (10) are approximated by propagating the corresponding PHD [10],

$$v_{k|k-1}(m|X_{0:k}) = v_{k-1}(m|X_{0:k-1}) + b(m|X_k)$$

$$v_k(m|X_{0:k}) = v_{k|k-1}(m|X_{0:k}) \left[1 - P_D(m|X_k) + \sum_{z \in \mathcal{Z}_k} \frac{\Lambda(m|X_k)}{c_k(z|X_k) + \int \Lambda(\zeta|X_k) v_{k|k-1}(\zeta|X_{0:k}) d\zeta} \right] \quad (12)$$

where $\Lambda(\cdot|X_k) = P_D(\cdot|X_k)g_k(z|\cdot, X_k)$, $b(m|X_k)$ is the PHD of the new feature RFS, $\mathcal{B}(X_k)$, discussed previously in section II-B and,

$$\begin{aligned} P_D(m|X_k) &= \text{the probability of detecting a feature at } m, \text{ from vehicle pose } X_k. \\ c_k(z|X_k) &= \text{PHD of the clutter RFS } \mathcal{C}_k \text{ in eqn.(4) at time } k. \end{aligned}$$

To evaluate eqn.(11), the term $g_k(\mathcal{Z}_k|Z_{0:k-1}, X_{0:k})$ involves the set integration,

$$g_k(\mathcal{Z}_k|Z_{0:k-1}, X_{0:k}) = \int p(\mathcal{Z}_k, \mathcal{M}_k|Z_{0:k-1}, X_{0:k}) \delta M_k$$

which is numerically intractable in general. Moreover, the EKF approximation given the vector map in FastSLAM [3] cannot be used since they are two fundamentally different quantities and it is not known how they are even related.

However, recall from (6) that $p_{k|k-1}(\mathcal{M}_k|X_{0:k})$ and $p_k(\mathcal{M}_k|X_{0:k})$ are approximated by,

$$p_{k|k-1}(\mathcal{M}_k|X_{0:k}) \approx \frac{\prod_{m \in \mathcal{M}_k} v_{k|k-1}(m|X_{0:k})}{\exp \left(\int v_{k|k-1}(m|X_{0:k}) dm \right)} \quad (13)$$

$$p_k(\mathcal{M}_k|X_{0:k}) \approx \frac{\prod_{m \in \mathcal{M}_k} v_k(m|X_{0:k})}{\exp \left(\int v_k(m|X_{0:k}) dm \right)}. \quad (14)$$

Subsequently, from the mapping recursion of eqn.(10) and setting $\mathcal{M}_k = \emptyset$, it can be shown that the measurement likelihood in the vehicle trajectory recursion of eqn.(11) can be evaluated as,

$$g_k(\mathcal{Z}_k|Z_{0:k-1}, X_{0:k}) = \prod_{z \in \mathcal{Z}_k} c_k(z) \times \exp \left(\hat{N}_k - \hat{N}_{k|k-1} - \int c_k(z) dz \right). \quad (15)$$

This is an important result, which allows for the likelihood of the measurement conditioned on the trajectory (but not the map), to be calculated in closed-form, as opposed to using approximations [3]. This is exploited in the following section describing the filter implementation.

IV. FILTER IMPLEMENTATION

Following the description of the proposed RB-PHD-SLAM filter in the previous section, a Gaussian mixture (GM) PHD filter is used to propagate the trajectory-conditioned posterior PHD of the map of eqn.(10), while a particle filter is then used to propagate the posterior density of the vehicle trajectory of eqn.(11). As such, let the PHD-SLAM density at time $k-1$ be represented by a set of L particles,

$$\left\{ w_{k-1}^{(i)}, X_{0:k-1}^{(i)}, v_{k-1}^{(i)}(\cdot|X_{0:k-1}^{(i)}) \right\}_{i=1}^L,$$

where $X_{0:k-1}^{(i)} = [X_0, X_1^{(i)}, X_2^{(i)}, \dots, X_{k-1}^{(i)}]$ is the i^{th} hypothesised vehicle trajectory and $v_{k-1}^{(i)}(\cdot|X_{0:k-1}^{(i)})$ is its map PHD. The filter then proceeds to approximate the posterior density by a new set of weighted particles,

$$\left\{ w_k^{(i)}, X_{0:k}^{(i)}, v_k^{(i)}(\cdot|X_{0:k}^{(i)}) \right\}_{i=1}^L,$$

as follows:

A. The Per-particle GM PHD Feature Map

Let the new feature intensity for the particle, $b(\cdot|X_{k-1}^{(i)})$, from the sampled pose, $X_k^{(i)}$ at time k be a Gaussian mixture of the form,

$$b(m|X_{k-1}^{(i)}, X_k^{(i)}) = \sum_{j=1}^{J_{b,k}^{(i)}} \eta_{b,k}^{(i,j)} \mathcal{N}(m; \mu_{b,k}^{(i,j)}, P_{b,k}^{(i,j)})$$

where, $J_{b,k}^{(i)}$ is the number of Gaussians in the new feature intensity at time k and $\eta_{b,k}^{(i,j)}$, $\mu_{b,k}^{(i,j)}$ and $P_{b,k}^{(i,j)}$ are the corresponding components. The prior map PHD for the i^{th} particle, $v_{k-1}^{(i)}(\cdot|X_{k-1}^{(i)})$, is a Gaussian mixture of the form,

$$v_{k-1}(m|X_{k-1}^{(i)}) = \sum_{j=1}^{J_{k-1}^{(i)}} \eta_{k-1}^{(i,j)} \mathcal{N}(m; \mu_{k-1}^{(i,j)}, P_{k-1}^{(i,j)})$$

where $J_{k-1}^{(i)}$ denotes the number of Gaussians, with $\eta_{k-1}^{(i,j)}$, $\mu_{k-1}^{(i,j)}$ and $P_{k-1}^{(i,j)}$ being their corresponding predicted weights, means and covariances respectively. The predicted intensity is therefore also a Gaussian mixture,

$$v_{k|k-1}(m|X_k^{(i)}) = \sum_{j=1}^{J_{k|k-1}^{(i)}} \eta_{k|k-1}^{(i,j)} \mathcal{N}(m; \mu_{k|k-1}^{(i,j)}, P_{k|k-1}^{(i,j)})$$

which consists of $J_{k|k-1}^{(i)} = J_{k-1}^{(i)} + J_{b,k}^{(i)}$ Gaussians representing the union of the prior map intensity, $v_{k-1}(\cdot|X_{k-1}^{(i)})$, and the proposed new feature intensity. Since the measurement likelihood is also of Gaussian form, it can be seen from eqn.(12), that the posterior map PHD, $v_k(\cdot|X_k^{(i)})$ is then also

a Gaussian mixture given by,

$$v_k(m|X_k^{(i)}) = v_{k|k-1}(m|X_k^{(i)}) \left[1 - P_D(m|X_k^{(i)}) + \sum_{z \in \mathcal{Z}_k} \sum_{j=1}^{J_{k|k-1}^{(i)}} v_{G,k}^{(i,j)}(z, m|X_k^{(i)}) \right].$$

where,

$$v_{G,k}^{(i,j)}(z, m|X_k^{(i)}) = \eta_k^{(i,j)}(z|X_k^{(i)}) \mathcal{N}(m; \mu_{k|k}^{(i,j)}, P_{k|k}^{(i,j)})$$

$$\eta_k^{(j)}(z|X_k^{(i)}) = \frac{P_D(m|X_k^{(i)}) \eta_{k|k-1}^{(i,j)} q^{(i,j)}(z, X_k^{(i)})}{c(z) + \sum_{\ell=1}^{J_{k|k-1}^{(i)}} P_D(m|X_k^{(i)}) \eta_{k|k-1}^{(i,\ell)} q^{(i,\ell)}(z, X_k^{(i)})}$$

$$q^{(i,j)}(z, X_k) = \mathcal{N}(z; H_k \mu_{k|k-1}^{(i,j)}, S_k^{(i,j)})$$

with ∇H_k being the Jacobian of the measurement equation with respect to the feature's estimated location. The terms $\mu_{k|k}$, $P_{k|k}$ and S_k can be obtained using any standard filtering technique such as EKF or UKF. In this paper, the EKF updates are adopted. Gaussian pruning and merging operations are carried out as in [15].

B. The Vehicle Trajectory

The proposed filter adopts a particle approximation of the posterior vehicle trajectory, $p_k(X_{0:k})$, which is sampled/resampled as follows:

At time $k \geq 1$, **Step 1: Sampling Step**

- For $i = 1, \dots, L$, sample $\tilde{X}_k^{(i)} \sim q(\cdot)$ and set

$$\tilde{w}_k^{(i)} = \frac{g_k(\mathcal{Z}_k | \mathcal{Z}_{0:k-1}, \tilde{X}_{0:k}^{(i)}) f_X(\tilde{X}_k^{(i)} | X_{k-1}^{(i)}, U_{k-1})}{q(\tilde{X}_k^{(i)} | \cdot)} w_{k-1}^{(i)}.$$

- Normalise weights: $\sum_{i=1}^L \tilde{w}_k^{(i)} = 1$.

Step 2: Resampling Step

- Resample $\{\tilde{w}_k^{(i)}, \tilde{X}_{0:k}^{(i)}\}_{i=1}^L$ to get $\{w_k^{(i)}, X_{0:k}^{(i)}\}_{i=1}^L$.

The vehicle transition is chosen as the proposal density then,

$$\tilde{w}_k^{(i)} = g_k(\mathcal{Z}_k | \mathcal{Z}_{0:k-1}, \tilde{X}_{0:k}^{(i)}) w_{k-1}^{(i)}$$

which can be evaluated in closed form according to eqn.(15), where,

$$\hat{N}_{k|k-1}^{(i)} = \sum_{j=1}^{J_{k|k-1}^{(i)}} \eta_{k|k-1}^{(i,j)} \text{ and } \hat{N}_k^{(i)} = \sum_{j=1}^{J_k^{(i)}} \eta_k^{(i,j)}.$$

The following section presents results and analysis of the proposed RB-PHD-SLAM filter, and compares it to classical vector-based FastSLAM [3].

V. RESULTS & ANALYSIS

This section details results and analysis from trials carried out using a simple simulated dataset, depicted in figure 2. For comparative purposes, the benchmark algorithm used in the analysis is the FastSLAM [3] algorithm with maximum likelihood data association, using mutual exclusion constraint and a 95% χ^2 confidence gate. Both filters use 50 particles to approximate the trajectory density. Nominal parameters for the trials were: velocity standard deviation (std.) of 0.5m/s, steering std. of 3°, range std. of 1.4m and bearing std. of 3.5°, $P_D = 0.95$, $\lambda_c = 5$, using a sensor with 15m maximum range and a 360° field of view. The clutter RFS, C_k , is assumed Poisson distributed [5], [4] in number and uniformly spaced over the mapping region, specifically, $c(z) = \lambda_c \mathcal{U}(z)$, where λ_c is the average number of clutter measurements and $\mathcal{U}(\cdot)$ denotes a uniform distribution on the measurement region. For both filters, the trajectory of highest weight is chosen as the vehicle path estimate, with its corresponding map being used to estimate the features, using an existence threshold of 0.5.

Fig. 2. The simulated environment showing point features (circles) and the vehicle trajectory (line). The measurements over the ground truth trajectory are also shown (black points).

Figure 2 shows the simple ground truth trajectory and map for the trial. A sample result from a single trial is then depicted in figure 3. An improved vehicle trajectory and feature map estimate is evident. Given that the RB-PHD-SLAM filter incorporates data association and feature number uncertainty into its Bayesian recursion, it is more robust to large sensing errors, as it does not rely on hard measurement-feature assignment decisions. Furthermore, it jointly estimates the number of features and their locations, alleviating the need for separate feature management [2], [3].

The trajectory RMSE over the estimated trajectory for 50 independent trials is presented in figure 4, further demonstrating the reduced trajectory estimation error of the proposed filter.

Given that a set representation is used for the map, to explicitly quantify the map estimation error, a mathematically consistent set error metric [9], [16] can be adopted which jointly evaluates the error in both feature location and feature number estimates. The metric optimally assigns each feature estimate to its ground truth feature through the Hungarian

Fig. 3. Graphical representation of the posterior FB-SLAM estimate from each filter, using the trajectory and map of the highest weighted particle.

Fig. 4. Comparison of the trajectory RMSE for 50 runs, with approximate means plotted (dashed). In the presence of large data association uncertainty and clutter, a marked improvement in trajectory estimation using the proposed filter is noticeable.

assignment algorithm and evaluates an error distance, while penalising for under/over estimating the correct number of features. Figure 5 plots the map estimation error metric for each MC trial. The metric mathematically quantifies the mapping error depicted in figure 3, allowing for the comparison of feature map estimates. The results demonstrate the improved map estimate (both in terms of feature number and location) from the proposed RB-PHD-SLAM filter under difficult sensing conditions.

Fig. 5. Comparison of the feature map estimation error, with approximate means plotted (dashed). For comparing the mapping error, the posterior map estimate of FastSLAM is ‘interpreted’ as a set.

Figure 6 presents the mean and standard deviation of the estimated vehicle trajectory’s RMSE over 50 Monte Carlo

trials carried out at increasing levels of measurement noise. The merits of the proposed Bayesian SLAM framework and RB-PHD-SLAM filter are verified, as its encapsulation of data association and feature/measurement number uncertainty into a single update (as opposed to the common two-tiered approach), increases its robustness to situations which may corrupt data-association reliant approaches.

Fig. 6. Mean and Standard Deviation of the estimated vehicle trajectory’s RMSE at increasing levels of measurement noise.

A. Extensions of RB-PHD-SLAM

The PHD approach detailed in this paper can be readily extended to included non-Poisson RFSs via the Cardinalised PHD construct [17] and the multiBernoulli recursion [18]. The trajectory conditioned measurement likelihood can be calculated exactly for each map RFS approximation.

1) *RB-CPHD-SLAM*: Let the density of the map RFS (eqn.14) be approximated by an i.i.d. cluster RFS,

$$p_k(\mathcal{M}_k|X_{0:k}) \approx \frac{|\mathcal{M}_k|! \rho_k(|\mathcal{M}_k|) \prod_{m \in \mathcal{M}_k} v_k(m|X_{0:k})}{\left(\int v_k(m|X_{0:k}) dm \right)^{|\mathcal{M}_k|}}$$

which is completely characterised by its cardinality distribution, ρ , and PHD, v , and where $\mathbb{E}[\rho_k] = \int v_k(m) dm$. The trajectory conditioned map recursion of eqn.(10) may then be approximated by a CPHD filter [17], and by evaluating the trajectory conditioned measurement likelihood in eqn.(10) using the i.i.d. cluster RFS,

$$g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_{0:k}) \approx \frac{|\mathcal{Z}_k|! \rho_k(|\mathcal{Z}_k|) \kappa_k^{\mathcal{Z}_k}}{\left(\int c_k(z|X_k) dz \right)^{|\mathcal{Z}_k|}} \times \frac{\rho_{k|k-1}(\emptyset|X_{0:k})}{\rho_k(\emptyset|X_{0:k})}$$

or for $\mathcal{M}_k = \bar{m}$, where $\bar{m}$ is a feature chosen according to a given strategy, and assuming Poisson clutter,

$$\begin{aligned} g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_{0:k}) &\approx \\ (1 - P_D(\bar{m}|X_k)) \kappa_k^{\mathcal{Z}_k} + P_D(\bar{m}|X_k) &\sum_{z \in \mathcal{Z}_k} \kappa_k^{\mathcal{Z}_k - z} g_k(z|\bar{m}, X_k) \\ &\frac{\exp \left(\int c_k(z|X_k) dz \right)}{\rho_k(1|X_{0:k}) v_k(\bar{m}|X_{0:k}) \int v_k(m|X_{0:k}) dm} \\ &\times \frac{\rho_{k|k-1}(1|X_{0:k}) v_k(\bar{m}|X_{0:k}) \int v_k(m|X_{0:k}) dm}{\rho_k(1|X_{0:k}) v_k(\bar{m}|X_{0:k}) \int v_k(m|X_{0:k}) dm} \end{aligned}$$

Given that the CPHD filter propagates the *distribution* of the number of features as opposed to just its mean (for the PHD approach of this paper), it is anticipated that the map, and subsequently the trajectory, estimates from RB-CPHD-SLAM would be remarkably improved in comparison to RB-PHD-SLAM.

2) *RB-MeMber-SLAM*: A multi-Bernoulli RFS, being simply a set of Bernoulli RFSs mentioned previously in section II-B, is completely described by the multi-Bernoulli parameter set $\{(\epsilon^{(i)}, \pi^{(i)})\}_{i=1}^m$, its probability density p is $p(\emptyset) = \prod_{j=1}^m (1 - \epsilon^{(j)})$ and,

$$p(\{m_1, \dots, m_n\}) = p(\emptyset) \sum_{1 \leq i_1 \neq \dots \neq i_n \leq m} \prod_{j=1}^n \frac{\epsilon^{(i_j)} p^{(i_j)}(m_j)}{1 - \epsilon^{(i_j)}}.$$

Thus, if the RFS map density is approximated by a multi-Bernoulli RFS, $p_k(\mathcal{M}_k | \mathcal{Z}_{0:k}, X_{0:k}) = \{(\epsilon_k^{(j)}(X_{0:k}), \pi_k^{(j)}(m | X_{0:k}))\}_{j=1}^{m_k(X_{0:k})}$, then the trajectory conditioned map recursion may be approximated by a MeMber filter [18], and as before with $\mathcal{M}_k = \emptyset$,

$$g_k(\mathcal{Z}_k | \mathcal{Z}_{0:k-1}, X_{0:k}) \approx \frac{\kappa_k^{\mathcal{Z}_k}}{\exp(\int c_k(z | X_k) dz)} \times \frac{\prod_{j=1}^{m_{k|k-1}(X_{0:k})} (1 - \epsilon_{k|k-1}^{(j)}(X_{0:k}))}{\prod_{j=1}^{m_k(X_{0:k})} (1 - \epsilon_k^{(j)}(X_{0:k}))}.$$

Equivalently, using $\mathcal{M}_k = \bar{m}$,

$$g_k(\mathcal{Z}_k | \mathcal{Z}_{0:k-1}, X_{0:k}) \approx \frac{(1 - P_D(\bar{m} | X_k)) \kappa_k^{\mathcal{Z}_k} + P_D(\bar{m} | X_k) \sum_{z \in \mathcal{Z}_k} \kappa_k^{\mathcal{Z}_k - z} g_k(z | \bar{m}, X_k)}{\exp(\int c_k(z | X_k) dz)} \times \frac{1}{\Gamma} \prod_{j=1}^{m_{k|k-1}(X_{0:k})} (1 - \epsilon_{k|k-1}^{(j)}(X_{0:k})) \times \frac{1}{\Gamma} \sum_{j=1}^{m_{k-1}(X_{0:k-1})} \frac{\epsilon_{k|k-1}^{(j)}(X_{0:k}) \pi_{k|k-1}^{(j)}(\bar{m} | X_{0:k})}{1 - \epsilon_{k|k-1}^{(j)}(X_{0:k})}$$

where,

$$\Gamma = \prod_{j=1}^{m_k(X_{0:k})} (1 - \epsilon_k^{(j)}(X_{0:k})) \sum_{j=1}^{m_k(X_{0:k})} \frac{\epsilon_k^{(j)}(X_{0:k}) \pi_k^{(j)}(\bar{m} | X_{0:k})}{1 - \epsilon_k^{(j)}(X_{0:k})}$$

Using a Bernoulli RFS to represent each feature allows for the joint encapsulation of its existence probability and location in a single pdf, in contrast to existing SLAM methods [2], [3], [4]. It is expected that RB-MeMber-SLAM would perform well in the presence of highly non-linear process and/or measurement models.

VI. CONCLUSION

This paper argues that from a fundamental estimation perspective, the feature map is more appropriately represented as a finite set. A Bayesian recursion and tractable solution for the feature-based SLAM problem were then presented. The filter jointly propagates and estimates the vehicle trajectory, number of features in the map as well as their individual locations while subject to data association uncertainty and clutter. The key to the approach is to adopt a finite-set representation of the map and to use the tools of finite-set-statistics to cast the problem into the Bayesian paradigm. A Rao-Blackwellised implementation of the filter

was outlined, in which the PHD of the map was propagated using a Gaussian mixture PHD filter, and a particle filter propagated the vehicle trajectory density, admitting an exact solution for the trajectory conditioned measurement likelihood. Simulated results demonstrated the robustness of the proposed filter, particularly in the presence of large data association uncertainty and clutter. The presented framework admits numerous extensions for future work.

VII. ACKNOWLEDGMENTS

This research is funded in part by the Singapore National Research Foundation through the Singapore-MIT Alliance for Research and Technology CENSAM. The second author is supported in part by discovery grant DP0880553 awarded by the Australian Research Council.

REFERENCES

- [1] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. *Autonomous Robot Vehicles*, pages 167–193, 1990.
- [2] G. Dissanayake, P. Newman, H.F. Durrant-Whyte, S. Clark, and M. Csorba. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Transactions on Robotic and Automation*, 17(3):229–241, June 2001.
- [3] M. Montemerlo, S. Thrun, and B. Siciliano. *FastSLAM: A Scalable Method for the Simultaneous Localization and Mapping Problem in Robotics*. Springer, 2007.
- [4] W.S. Wijesoma, L.D.L Perera, and M.D. Adams. Toward multidimensional assignment data association in robot localization and mapping. *IEEE Transactions on Robotics*, 22(2):350–365, April 2006.
- [5] D. Makarsov and H.F. Durrant-Whyte. Mobile vehicle navigation in unknown environments: a multiple hypothesis approach. *IEE Proceedings of Contr. Theory Appl.*, vol. 142, July 1995.
- [6] S. Thrun. Particle filter in robotics. In *Uncertainty in AI (UAI)*, 2002.
- [7] J. Niera and J.D. Tardos. Data association in stochastic mapping using the joint compatibility test. *IEEE Transactions on Robotics and Automation*, 17(6):890–897, December 2001.
- [8] J. Mullane, B.N. Vo, M. Adams, and W.S. Wijesoma. A random set formulation for bayesian slam. In *proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, France, September 2008.
- [9] J. Mullane, B.N. Vo, M. Adams, and W.S. Wijesoma. A random set approach to slam. In *proceedings of the IEEE International Conference on Robotics and Automation (ICRA) workshop on Visual Mapping and Navigation in Outdoor Environments*, Japan, May 2009.
- [10] R. Mahler. Multi-target bayes filtering via first-order multi-target moments. *IEEE Transactions on AES*, 4(39):1152–1178, October 2003.
- [11] R. Mahler. *Statistical Multisource Multitarget Information Fusion*. Artech House, 2007.
- [12] B.N. Vo, S. Singh, and A. Doucet. Sequential monte carlo methods for multi-target filtering with random finite sets. *IEEE Transactions on Aerospace and Electronic Systems*, 41(4):1224–1245, October 2005.
- [13] D. Stoyan, W.S. Kendall, and J. Mecke. *Stochastic Geometry and Its Applications*. John Wiley and Sons, Inc, New York, second edition, 1995.
- [14] R. Mahler. Implementation and application of PHD/CPHD filters. In *Int'l Conf. on Information Fusion*, Seattle, WA, 2009.
- [15] B.N. Vo and W.K. Ma. The gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, November 2006.
- [16] D. Schuhmacher, B.T. Vo, and B.N. Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE Transactions on Signal Processing*, 86(8):3447–3457, 2008.
- [17] B.T. Vo, B.N. Vo, and A. Cantoni. Analytic implementations of the cardinalized probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 55(7):3553–3567, July 2007.
- [18] B.T. Vo, B.N. Vo, and A. Cantoni. The cardinality balanced multi-target multi-bernoulli and its implementations. *IEEE Transactions on Signal Processing*, 57(2):409–423, February 2009.