Rao-Blackwellised PHD SLAM

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

School of Electrical & Electronic Engineering, Nanyang Technological University, Singapore
†School of Electrical & Electronic Engineering, University of Melbourne, Victoria 3010, Australia
{jsmullane,eadams}@ntu.edu.sg
†bnvo@unimelb.edu.au

***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 the 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.**

***Index Terms*— Bayesian SLAM, Random Finite Set (RFS), Feature-based Map, Probability Hypothesis Density (PHD), Point Process**

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 security patrolling [2], [3], [4], [5], [6], [7]. 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 [8]. Estimating a feature map, thus requires the joint estimation of the number of the features and their locations.

Current state-of-the-arts FB-SLAM solutions address two separate problems [5]:

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

This two-tiered approach to SLAM is sensitive to data association (DA) uncertainty [9]. This sensitivity is due to the fact that the current Bayesian SLAM framework does not in fully integrate DA uncertainty into the map estimation, specifically, the filtering step 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. Indeed, from an estimation viewpoint, the map is fundamentally a finite set (see section II-A). Using RFS theory, the FB-SLAM problem is 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 using 'brute force' implementations can be found in [10], [11], [12]. In this paper, a tractable 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 (RB) implementation of the PHD-SLAM filter is proposed based on the Gaussian mixture PHD filter for the map and the particle filter for the vehicle trajectory in a similar manner to FAST-SLAM. 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 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

The work of the second author is supported in part by discovery grant DP0880553 awarded by the Australian Research Council.

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 constructed by stacking features into a vector, and consider the following simplistic scenarios,

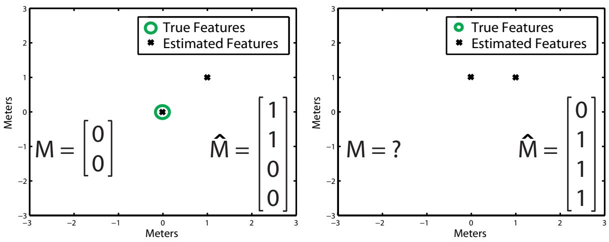

Fig. 1. Hypothetical scenario showing a fundamental inconsistency with vector representations of feature maps. How should the error be assigned when the estimates of the number of features in $\hat{\mathrm{M}}$ is incorrect?

A finite set representation of the map, $\mathcal{M}_k = \{m_k^1, \ldots, m_k^{N_k}\}$, where $m^1, \ldots, m^{N_k}$ are the N_k features present at time k, admits a mathematically consistent notion of estimation error since 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. The realisation that the map is set is not new. Indeed, in a number of influential works the map has been defined as a finite set [7]. However the mathematical tools for dealing with random finite sets were not yet available to the SLAM research community.

It should be noted that, while finite-sets naturally capture a feature map, a finite set map representation for autonomous grid-based frameworks [13], [14], is unnecessary, since the number of grid cells is known (*a priori* tessellation), and the order of the map states signifies their spatial location in the grid [12].

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 naturally represented by a finite set of readings, $\mathcal{Z}_k = \{z_k^1, z_k^2, \ldots, z_k^{\mathfrak{z}(k)}\}$.

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 [15], [16], based on a notion of integration and density that is consistent with point process theory [17].

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 $X_{0:k-1} = [X_0, X_1, \ldots, X_{k-1}]$ at time $k-1$, i.e.

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

Note that $FOV(X_{0:k-1}) = FOV(X_0) \cup FOV(X_1) \cup \cdots \cup FOV(X_{k-1})$. $\mathcal{M}_{k-1}$ 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,

$$\mathcal{M}_k = \mathcal{M}_{k-1} \cup \Big(FOV(X_k) \cap \bar{\mathcal{M}}_{k-1} \Big) \tag{2}$$

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 $\mathcal{M}_{k-1}$. Let the new features which have entered the FOV, i.e. the second term of eqn.(2), be modeled by the independent RFS, $\mathcal{B}_k(X_k)$. In this case, the RFS map transition density is given by,

$$f_{\mathcal{M}}(\mathcal{M}_k|\mathcal{M}_{k-1}, X_k) = \sum_{\mathcal{W} \subseteq \mathcal{M}_k} f_{\mathcal{M}}(\mathcal{W}|\mathcal{M}_{k-1}) f_{\mathcal{B}}(\mathcal{M}_k - \mathcal{W}|X_k) \tag{3}$$

where $f_{\mathcal{M}}(\cdot|\mathcal{M}_{k-1})$ is the transition density of the set of features that are in $FOV(X_{0:k-1})$ at time $k-1$ to time k, and $f_{\mathcal{B}}(\cdot|X_k)$ is the density of the RFS, $\mathcal{B}(X_k)$, of the new features that come through the field of view at time k.

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,

$$f_{k|k-1}(\mathcal{M}_k, X_k|\mathcal{M}_{k-1}, X_{k-1}, U_k) = f_{\mathcal{M}}(\mathcal{M}_k|\mathcal{M}_{k-1}, X_k) f_X(X_k|X_{k-1}, U_k). \tag{4}$$

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

$$\mathcal{Z}_k = \bigcup_{m \in \mathcal{M}_k} \mathcal{D}_k(m, X_k) \cup \mathcal{C}_k(X_k) \tag{5}$$

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. This cannot be explicitly encapsulated by the classical vector measurement models. It is also assumed that $\mathcal{D}_k(m, X_k)$, and $\mathcal{C}_k(X_k)$ are independent RFSs conditional on X_k.

The RFS of the measurements generated by a feature at m is a Bernoulli RFS[1] 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.

Using Finite Set Statistics [16], the likelihood of the measurement $\mathcal{Z}_k$ is then given by,

$$g_k(\mathcal{Z}_k|X_k, \mathcal{M}_k) = \sum_{\mathcal{W} \subseteq \mathcal{Z}_k} f_{\mathcal{D}}(\mathcal{W}|\mathcal{M}_k, X_k) f_{\mathcal{C}}(\mathcal{Z}_k - \mathcal{W}|X_k) \quad (6)$$

with $f_{\mathcal{D}}(\cdot|\mathcal{M}_k, X_k)$ denoting the density of the RFS of observations, and $f_{\mathcal{C}}(\cdot|X_k)$ denoting the density of the clutter RFS, $\mathcal{C}_k$. It can be seen that this likelihood directly encapsulates the inherent measurement uncertainty, with $f_{\mathcal{D}}(\cdot|\mathcal{M}_k, X_k)$ considering detection uncertainty and measurement noises, and $f_{\mathcal{C}}(\cdot|X_k)$ modeling the spurious measurements. This density is typically *a priori* given as Poisson in number and uniform in space [2], [6]. 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 adopting an RFS representation.

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

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

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

$$p_{k|k-1}(\mathcal{M}_k, X_{1:k}) = f_X(X_k|X_{k-1}, U_k) \times \int f_{\mathcal{M}}(\mathcal{M}_k|\mathcal{M}_{k-1}, X_k) p_{k-1}(\mathcal{M}_{k-1}, X_{1:k-1}) \delta\mathcal{M}_{k-1}$$

$$p_k(\mathcal{M}_k, X_{1:k}) = \frac{g_k(\mathcal{Z}_k|X_k, \mathcal{M}_k) p_{k|k-1}(\mathcal{M}_k, X_{1:k})}{g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_0)} \quad (7)$$

where the δ implies a set integral. 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 (7) integrates uncertainty in DA and number of features into a single Bayesian filter and does not require separate DA step nor feature management, as are classically required [4], [5], [9]. 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 [16]. 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

For each closed set S and an RFS $\mathcal{M}$, let the number of points of $\mathcal{M}$ in S be denoted by $N_S(\mathcal{M})$. If the measure $\mathbb{E}[N_{(\cdot)}(\mathcal{M})]$ admits a density relative to the Lebesgue measure λ, then the PHD of an RFS $\mathcal{M}$ is the density [18],

$$v(m) = \frac{\mathbb{E}[N_{dm}(\mathcal{M})]}{\lambda(dm)}. \quad (8)$$

In other words, 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,

$$p(\mathcal{M}) = \frac{\prod_{m \in \mathcal{M}} v(m)}{\exp(\int v(m) dm)}. \quad (9)$$

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

Consider now the joint posterior distribution of the map and the vehicle trajectory. From standard results on conditional expectation,

$$\mathbb{E}[N_S(\mathcal{M}_k)|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0] = \mathbb{E}[\mathbb{E}[N_S(\mathcal{M}_k)|X_{0:k}, \mathcal{Z}_{0:k}, U_{0:k-1}]] \quad (10)$$

Note that the 1st expectation on the RHS of the above equation is taken over all possible vehicle trajectories up to time k. The conditional expectation $\mathbb{E}[N_{(\cdot)}(\mathcal{M}_k)|X_{0:k}, \mathcal{Z}_{0:k}, U_{0:k-1}]$ is a measure (on the space of features) and its density (relative to the Lebesgue measure) is the posterior PHD of the map conditioned on the measurement history, the control history and the vehicle trajectory, i.e.

$$v_k(m|X_{0:k}, \mathcal{Z}_{0:k}, U_{0:k-1}) = \frac{\mathbb{E}[N_{dm}(\mathcal{M}_k)|X_{0:k}, \mathcal{Z}_{0:k}, U_{0:k-1}]}{\lambda(dm)}. \quad (11)$$

[1]The Bernoulli RFS is empty with a probability $1 - \alpha$ and is distributed according to a density p with probability α.

Similarly, the posterior expectation $\mathbb{E}[N_{(\cdot)}(\mathcal{M}_k)|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0]$ is a measure and its density is the posterior PHD of the map. Hence, it follows from (8), (10) and (11) that

$$v_k(m|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0) = \mathbb{E}\left[v_k(m|\mathcal{Z}_{0:k}, U_{0:k-1}, X_{0:k})\right] \tag{12}$$

In otherwords, the posterior PHD of the map is indeed the expectation of the trajectory-conditioned PHD of the map.

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 $v_k(\cdot|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0)$ 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 [19].

B. The RB PHD-SLAM recursion

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

$$\begin{aligned} p_k(\mathcal{M}_k, X_{1:k}|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0) = \\ p_k(X_{1:k}|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0)p_k(\mathcal{M}_k|\mathcal{Z}_{0:k}, X_{0:k}). \end{aligned} \tag{13}$$

Thus, the recursion for the joint map-trajectory posterior density according to (7) 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,

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

then,

$$\begin{aligned} p_{k|k-1}(\mathcal{M}_k|X_{0:k}) = \int f_{\mathcal{M}}(\mathcal{M}_k|\mathcal{M}_{k-1}, X_k)\times \\ p_{k-1}(\mathcal{M}_{k-1}|X_{0:k-1})\delta\mathcal{M}_{k-1} \end{aligned} \tag{14}$$

$$p_k(\mathcal{M}_k|X_{0:k}) = \frac{g_k(\mathcal{Z}_k|\mathcal{M}_k, X_k)p_{k|k-1}(\mathcal{M}_k|X_{0:k})}{g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_{0:k})} \tag{15}$$

$$\begin{aligned} p_k(X_{1:k}) = g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_{0:k})\times \\ \frac{f_X(X_k|X_{k-1}, U_{k-1})p_{k-1}(X_{1:k-1})}{g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1})}. \end{aligned} \tag{16}$$

The recursion defined by (14), (15), (16) is similar to FastSLAM [5], in the exploitation of the factorisation of the joint SLAM posterior of [20]. The difference is that the map and the measurements are random finite sets. Consequently, the propagation equations involve probability density 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 [15], then eqns.(14), (15) are approximated by propagating the corresponding PHD,

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

$$\begin{aligned} v_k(m|X_{0:k}) = v_{k|k-1}(m|X_{0:k})\Bigg[1 - P_D(m|X_k) + \\ \sum_{z\in\mathcal{Z}_k} \frac{\Lambda(m|X_k)}{c_k(z|X_k) + \int \Lambda(\zeta|X_k)v_{k|k-1}(\zeta|X_{0:k})d\zeta}\Bigg] \end{aligned} \tag{17}$$

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,

$P_D(m\|X_k)$	=	the probability of detecting a feature at m, from vehicle pose X_k.
$c_k(z\|X_k)$	=	PHD of the clutter RFS $\mathcal{C}_k$ in eqn.(5) at time k.

Furthermore, from (9), $p_{k|k-1}(\mathcal{M}_k|X_{0:k})$ and $p_k(\mathcal{M}_k|X_{0:k})$ are approximated by,

$$\begin{aligned} p_{k|k-1}(\mathcal{M}_k|X_{0:k}) &\approx \frac{\prod_{m\in\mathcal{M}_k} v_{k|k-1}(m|X_{0:k})}{\exp\left(\int v_{k|k-1}(m|X_{0:k})dm\right)} \\ p_k(\mathcal{M}_k|X_{0:k}) &\approx \frac{\prod_{m\in\mathcal{M}_k} v_k(m|X_{0:k})}{\exp\left(\int v_k(m|X_{0:k})dm\right)}. \end{aligned}$$

Subsequently, from the mapping recursion of eqn.(15) and setting $\mathcal{M}_k = \emptyset$, it can be shown that the measurement likelihood in the vehicle trajectory recursion of eqn.(16) can be evaluated as,

$$\begin{aligned} g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, X_{0:k}) = \prod_{z\in\mathcal{Z}_k} c_k(z)\times \\ \exp\left(\hat{N}_k - \hat{N}_{k|k-1} - \int c_k(z)dz\right). \end{aligned} \tag{18}$$

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 [5]. 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.(15), while a particle filter is then used to propagate the posterior density of the vehicle trajectory of eqn.(16). 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 $X_{0:k-1}^{(i)} = [X_0, X_1^{(i)}, X_2^{(i)}, \ldots, X_{k-1}^{(i)}]$ 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 prior map PHD for the i^{th} particle, $v_{k-1}^{(i)}(\cdot|X_{k-1}^{(i)})$, be a Gaussian mixture of the form,

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

which is a mixture of $J_{k-1}^{(i)}$ 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. Let the new feature intensity for the particle, $b(\cdot|\mathcal{Z}_{k-1}, X_k^{(i)})$, from the sampled pose, $X_k^{(i)}$ at time k also be a Gaussian mixture of the form,

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

where, $J_{b,k}^{(i)}$ defines 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. This is analogous to the proposal distribution in the particle filter and provides an initial estimate of the new features entering the map.

The predicted intensity is therefore also a Gaussian mixture,

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

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.(17), that the posterior map PHD, $v_k(\cdot|X_k^{(i)})$ is then also a Gaussian mixture given by,

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

The components of the above equation are given by,

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

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

where, $q^{(i,j)}(z, X_k) = \mathcal{N}\big(z; H_k \mu_{k|k-1}^{(i,j)}, S_k^{(i,j)}\big)$. The components are obtained from the standard EKF update equations,

$$\mu_{k|k}^{(i,j)} = \mu_{k|k-1}^{(i,j)} + K_k^{(i,j)}(z - H_k \mu_{k|k-1}^{(i,j)})$$

$$P_{k|k}^{(i,j)} = [I - K_k^{(i,j)} \nabla H_k] P_{k|k-1}^{(i,j)}$$

$$K_k^{(i,j)} = P_{k|k-1}^{(i,j)} \nabla H_k^T [S_k^{(i,j)}]^{-1}$$

$$S_k^{(i,j)} = R_k + \nabla H_k P_{k|k-1}^{(i,j)} \nabla H_k^T$$

with ∇H_k being the Jacobian of the measurement equation with respect to the landmarks estimated location. The clutter RFS, C_k, is assumed Poisson distributed [2], [6] in number and uniformly spaced over the mapping region. Therefore the clutter intensity is given by, $c(z) = \lambda_c \mathcal{U}(z)$, where λ_c is the average number of clutter measurements and $\mathcal{U}(\cdot)$ denotes a uniform distribution on the measurement space. As with other feature-based GM implementations [21], pruning and merging operations are required to curb the explosive growth in the number of Gaussian components of the posterior map PHD. These operations are carried out as in [18].

B. The Vehicle Trajectory

The proposed filter adopts a particle approximation of the posterior vehicle trajectory, $p_k(X_{1:k})$, which is sampled/re-sampled as follows:

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

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

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

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

Step 2: *Resampling Step*

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

Setting the proposal density to be the vehicle transition density,

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

which can be evaluate in closed form according to eqn.(18), 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 to classical vector-based FastSLAM [5].

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 [5] 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 input variance of 0.25, steering input variance of 9, range measurement variance of 2 and bearing measurement variance of 12.5, $P_D = 0.95$, $\lambda_c = 5$, using a sensor with 15m maximum range and a 360^o field of view. Measurement noise was inflated to hinder data association in the vector-based filter. 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.

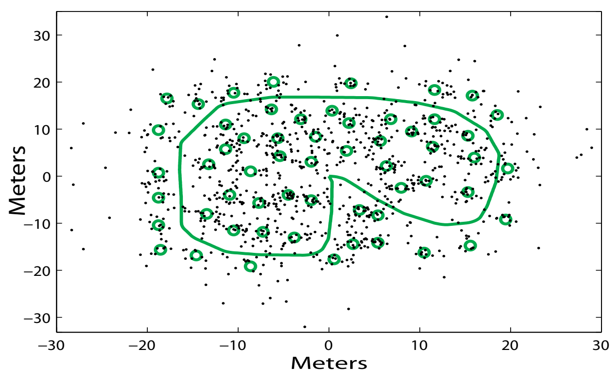

Fig. 2. *The simulated environment showing point features (circles) and the vehicle trajectory (line). A sample measurement plotted from the ground truth trajectory is also shown (black points).*

Figure 2 shows the simple ground truth trajectory and map for the trial. A sample result from a single Monte Carlo (MC) 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 error, 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 popular feature management methods [4], [5].

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

Given that the feature map is a set, to explicitly quantify the map estimation error, a mathematically consistent set error metric [11], [12], [22] can be adopted which jointly evaluates the error in both feature location *and* feature number estimates. The metric optimally assigns each feature

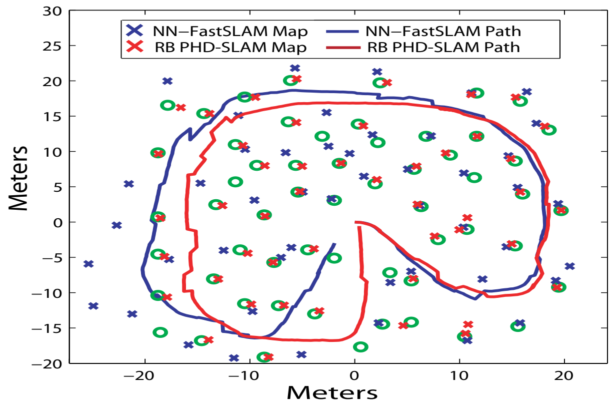

Fig. 3. *Graphical representation of the posterior FB-SLAM estimate from each filter. The robustness of the proposed approach in high measurement noise is evident.*

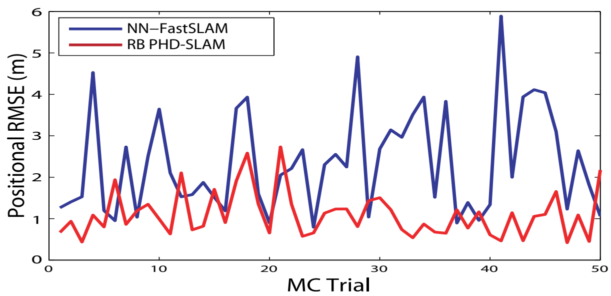

Fig. 4. *Comparison of the positional RMSE over 50 MC trials in the simulated environment. In the presence of large data association uncertainty and clutter, a marked improvement in trajectory estimation using the proposed filter is noticeable.*

estimate to its ground truth feature through the Hungarian 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.

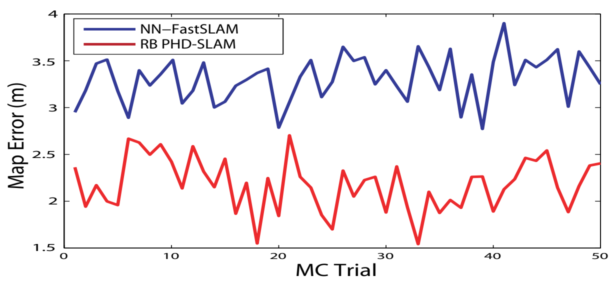

Fig. 5. *Comparison of the feature map estimation error. For the purpose of evaluating the mapping error, the posterior map estimate of FastSLAM is 'interpreted' as a set.*

Figure 6 presents the mean and standard deviation of the estimated vehicle trajectory's RMSE over 50 MC 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.

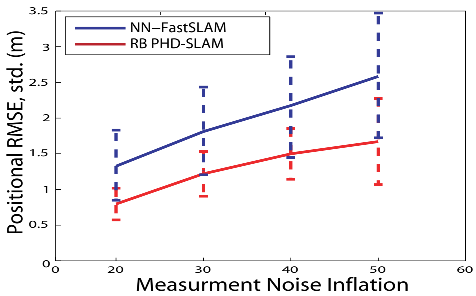

Fig. 6. *Mean and Standard Deviation of the estimated vehicle trajectory's RMSE at increasing levels of measurement noise.*

VI. Conclusion

This paper presented both a Bayesian recursion and tractable solution for the feature-based SLAM problem. The filter jointly propagates and estimates the vehicle trajectory, number of features in the map as well as their individual locations in the presence of data association uncertainty and clutter. The key to the approach is to adopt the natural 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. A closed form solution for the trajectory weighting was also presented, alleviating the need for approximation, which is commonly used. Simulated results demonstrated the robustness of the proposed filter, particulary in the presence of large data association uncertainty and clutter.

References

[1] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. *Autonomous Robot Vehicles*, pages 167–193, 1990.

[2] D. Makarsov and H.F. Durrant-Whyte. Mobile vehicle navigation in unknown environments: a multiple hypothesis approach. *IEE Proceedings of Contr. Theory Applict.*, vol. 142, July 1995.

[3] J. Guivant, E. Nebot, and S. Baiker. Autonomous navigation and map building using laser range sensors in outdoor applications. *Journal of Robotic Systems*, 17(10):565–583, October 2000.

[4] 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.

[5] M. Montemerlo, S. Thrun, and B. Siciliano. *FastSLAM: A Scalable Method for the Simultaneous Localization and Mapping Problem in Robotics*. Springer, 2007.

[6] 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.

[7] H.F. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping: Part I. *IEEE Robotics and Automation Magazine*, 13(2):99–110, June 2006.

[8] S. Thrun. Particle filter in robotics. In *Uncertainty in AI (UAI)*, 2002.

[9] 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.

[10] 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.

[11] 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.

[12] J. Mullane. *Autonomous Navigation: On Issues Concerning Measurement Uncertainty*. PhD thesis, Dept. of EEE, Nanyang Technological University, Singapore, 2009.

[13] M.C. Martin and H.P. Moravec. Robot evidence grids. Technical Report CMU-RI-TR-96-06, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, March 1996.

[14] J. Mullane, M. Adams, and W.S. Wijesoma. Robotic mapping using measurement likelihood filtering. *International Journal of Robotics Research*, 2(28):172–190, 2009.

[15] R. Mahler. Multi-target bayes filtering via first-order multi-target moments. *IEEE Transactions on AES*, 4(39):1152–1178, October 2003.

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

[17] 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.

[18] 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.

[19] R. Mahler. Implementation and application of PHD/CPHD filters. In *Int'l Conf. on Information Fusion*, Seattle, WA, 2009.

[20] R. Murphy. Bayesian map learning in dynamic environments. In *Proc. Conf. Neural Inf. Process. Syst.*, pages 1015–1021, Colorado, 1999.

[21] H.F. Durrant-Whyte, S. Majumder, M. de Battista, and S. Scheding. A bayesian algorithm for simultaneous localisation and map building. In *The Tenth International Symposium of Robotics Research (ISRR)*, Victoria, Australia, 2001.

[22] D. Schuhmacher, B.T. Vo, and B.N. Vo. A consistent metric for performance evaluation of multi-object filters (to appear). *IEEE Transactions on Signal Processing*, 86(8):3447–3457, 2008.