Mobile Robotics in a Random Finite Set Framework

John Mullane[1], Ba-Ngu Vo[2], Martin Adams[3], and Ba-Tuong Vo[2]

[1] Nanyang Technological University, Singapore, jsmullane@ntu.edu.sg
[2] University of Western Australia, Perth, Australia
ba-ngu.vo@uwa.edu.sg, ba-tuong.vo@uwa.edu.sg
[3] University of Chile, Santiago, Chile, martin@ing.uchile.cl

Abstract. This paper describes the Random Finite Set approach to Bayesian mobile robotics, which is based on a natural multi-object filtering framework, making it well suited to both single and swarm-based mobile robotic applications. By modeling the measurements and feature map as random finite sets (RFSs), joint estimates the number and location of the objects (features) in the map can be generated. In addition, it is shown how the path of each robot can be estimated if required. The framework differs dramatically from existing approaches since both data association and feature management routines are integrated into a single recursion. This makes the framework well suited to multi-robot scenarios due to the ease of fusing multiple map estimates from swarm members, as well as mapping robustness in the presence of other mobile robots which may induce false map measurements. An overview of developments thus far is presented, with implementations demonstrating the merits of the framework on simulated and experimental datasets.

Keywords: mobile robotics, Bayesian estimation, random finite sets, Probability Hypothesis Density.

1 Introduction

Mobile robotics is becoming an increasingly popular field of research, especially as embedded computing technology matures and gets ever more sophisticated. The goal of a mobile robot (or a swarm of mobile robots) is typically to travel through an unknown environment autonomously, while being continuously aware of both its surroundings and its position in relation to those surroundings. To measure the area of operation, robots are equipped with a suite of sensors such as lasers, radar, sonar or cameras which are inherently prone to sensing and data association error. In addition, the control inputs applied to a robots actuators introduce further uncertainty into the robot's position due to noise and kinematic modeling errors of the robot. Due to these multiple sources of uncertainty, Bayesian approaches have become widely popular [1], [2], [3], [4], which adopt probabilistic robot and sensor models and attempt to extract optimal estimates of the map and robots.

By far the most common approach to the problem is to use a random vector framework, in which the map and robot paths are modeled as random vectors containing positional information about the features and robot locations respectively [1]. While this model is the basis for the majority of existing mobile robotics algorithms, it requires independent data association and map management routines to respectively assign measurements to features and to estimate the number of features in the map [5], [6]. Recently, a new framework has been developed using Random Finite Set (RFS) models [7], [8], [9], which alleviates the need for independent routines and unifies the stochastic mobile robotics framework into a single Bayesian recursion. This new approach admits numerous benefits such as removal of data association, increased robustness to measurement error, integration of map management, straightforward fusion of multiple robot map estimates, expected map estimation and can be readily applied to single or multiple robot scenarios.

This paper advocates a fully integrated Bayesian framework for mobile robotics under DA uncertainty and unknown feature number. The key to this formulation is the representation of the map as a *finite set* of features. Indeed, from an estimation viewpoint, it is argued in section that the map is indeed a finite set and not a vector. Using Random Finite Set (RFS) theory, mobile robotics is then posed as a Bayesian filtering problem in which the posterior distribution of the set-valued map is propagated forward in time as measurements arrive. In the case of an unknown robot path, the joint density including the robot trajectory can be propagated. A tractable solution which propagates the first order moment of the map, its Probability Hypothesis Density (PHD), is presented. The PHD construct can also be interpreted in terms of occupancy maps [9], [10]. In this paper, both the map estimation from a known robot path and joint map/ trajectory estimation from an unknown robot path are examined separately. In particular, mapping robustness to multiple robots, which may interfere with the map building process, is demonstrated.

2 Background

Map estimation is closely related to the multi-target filtering problem, where the aim is to jointly estimate the time-varying number of targets (features) and their states from sensor measurements in the presence of data association uncertainty, detection uncertainty, clutter and noise. The first systematic treatment of this problem using random set theory was conceived by Mahler in 1994 [11], which later developed into Finite Set Statistics and the Probability Hypothesis Density (PHD) filter in 2003 [12]. A detailed treatment can be found in [13]. The mobile robotics problem was first formulated in an RFS framework in [7], with mapping and localisation algorithms presented in [8]. The approach modeled the joint vehicle trajectory and map as a single RFS, and recursively propagates its first order moment. Stemming from the popular FastSLAM algorithm [6], a factored approach to RFS-SLAM was proposed in [14] however vector-based approximations of RFS integrals were used which invalidate the approach. A theoretically correct factored approach was presented in [15], [9], which propagates the posterior PHDs of multiple trajectory-conditioned maps and the posterior distribution of the vehicle trajectory, with implementations on marine based radar data also appearing in [16].

3 The Robotics RFS Filtering Framework

This section details both the RFS map estimation filter and the RFS SLAM filter.

3.1 The Robotic Mapping Problem

As was first proposed in [7] and [8], let $\mathcal{M}$ be the RFS representing the entire unknown environment comprising static map features and multiple mobile robots 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.

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

Given this representation, $\mathcal{M}_{k-1}$ evolves in time according to,

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

where $\bar{\mathcal{M}}_{k-1} = \mathcal{M} - \mathcal{M}_{k-1}$, i.e the set of features that are not in $\mathcal{M}_{k-1}$. If $f_{k|k-1}(\mathcal{M}_k|\mathcal{M}_{k-1}, X_{k-1})$ then represents the RFS feature map state transition density, the generalised Bayesian RFS robotic mapping recursion can be written [17],

$$p_{k|k-1}(\mathcal{M}_k|\mathcal{Z}_{0:k-1}, X_{0:k}) = \int f_{k|k-1}(\mathcal{M}_k|\mathcal{M}_{k-1}, X_k) \times p_{k-1}(\mathcal{M}_{k-1}|\mathcal{Z}_{0:k-1}, X_{k-1}) \delta\mathcal{M}_{k-1} \tag{3}$$

$$p_k(\mathcal{M}_k|\mathcal{Z}_{0:k}, X_{0:k}) = \frac{g_k(\mathcal{Z}_k|X_k, \mathcal{M}_k) p_{k|k-1}(\mathcal{M}_k|\mathcal{Z}_{0:k-1}, X_{0:k})}{\int g_k(\mathcal{Z}_k|X_k, \mathcal{M}_k) p_{k|k-1}(\mathcal{M}_k|\mathcal{Z}_{0:k-1}, X_{0:k}) \delta\mathcal{M}_k} \tag{4}$$

where $g_k(\mathcal{Z}_k|\cdot)$ denotes the likelihood of the RFS measurement and δ denotes a set integral. Integration over the map, requires integration over all possible feature maps (all possible locations *and* numbers of features).

3.2 The Joint Robotic Mapping and Localisation Problem

To jointly estimate the map and the robot trajectory, denoted by the random vector $X_{1:k}$, the posterior density of (4) can be modified to get,

$$p_k(\mathcal{M}_k, X_{1:k}|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0) = p_k(X_{1:k}|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0)p_k(\mathcal{M}_k|\mathcal{Z}_{0:k}, X_{0:k}). \tag{5}$$

where $U_{0:k-1}$ denotes the robot control inputs random vector. Note that the second term is exactly equivalent to the posterior of (4). The first term can be calculated via,

$$\begin{aligned} p_k(X_{1:k}|\mathcal{Z}_{0:k}, U_{0:k-1}, X_0) = g_k(\mathcal{Z}_k|\mathcal{M}_k, X_k) \times \frac{p_{k|k-1}(\mathcal{M}_k|\mathcal{Z}_{0:k-1}, X_{0:k})}{p_k(\mathcal{M}_k|\mathcal{Z}_{0:k}, X_{0:k})} \\ \times \frac{p_{k|k-1}(X_{1:k}|\mathcal{Z}_{0:k-1}, U_{0:k-1}, X_0)}{g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1})} \end{aligned} \tag{6}$$

Further details can be seen in [15], [16], [9].

3.3 First order Moment Approximation

As with the classical approaches, the previous Bayesian recursions are numerically intractable and sensible approximations are required. In this work, the predicted and posterior RFS maps of (3) and (4) are approximated by Poisson RFSs with PHDs $v_{k|k-1}(m|\mathcal{Z}_{0:k-1}, X_{0:k})$ and $v_k(m|\mathcal{Z}_{0:k}, X_{0:k})$. In essence, this approximation assumes that features are iid and the number of features is Poisson distributed, i.e.,

$$p_{k|k-1}(\mathcal{M}_k|\mathcal{Z}_{0:k-1}, X_{0:k}) \approx \frac{\prod\limits_{m \in \mathcal{M}_k} v_{k|k-1}(m|X_{0:k})}{\exp\left(\int v_{k|k-1}(m|X_{0:k})dm\right)} \tag{7}$$

$$p_k(\mathcal{M}_k|\mathcal{Z}_{0:k}, X_{0:k}) \approx \frac{\prod\limits_{m \in \mathcal{M}_k} v_k(m|X_{0:k})}{\exp\left(\int v_k(m|X_{0:k})dm\right)}. \tag{8}$$

This PHD approximation has been proven to be effective in multi-target tracking [13]. As shown in [12], [17], the recursion of (3) and (4) can equally be reduced to a *predictor corrector* form. The PHD predictor equation is,

$$v_{k|k-1}(m|X_{0:k}) = v_{k-1}(m|X_{0:k-1}) + b(m|X_k) \tag{9}$$

where $b(m|X_k)$ is the PHD of the new feature RFS, $\mathcal{B}(X_k)$. The PHD corrector equation is then,

$$\begin{aligned} v_k(m|X_{0:k}) = v_{k|k-1}(m|X_{0:k}) \Bigg[1 - P_D(m|X_k) + \\ \sum_{z \in \mathcal{Z}_k} \frac{\Lambda(m|X_k)}{c_k(z|X_k) + \int \Lambda(\zeta|X_k) v_{k|k-1}(\zeta|X_{0:k}) d\zeta} \Bigg] \end{aligned} \tag{10}$$

where, $P_D(m|X_k)$ is the probability of detecting a feature at m, $c_k(z|X_k)$ is the PHD of the clutter RFS and $\Lambda(m|X_k) = P_D(m|X_k)g_k(z|m, X_k)$. The PHD recursion is far more numerically tractable than propagating the RFS map densities of (4). In addition, the recursion can be readily extended to incorporate multiple sensors / swarms of robots by sequentially updating the map PHD with the measurement from each robot. In addition, using the same Poisson approximations, (6) can be readily evaluated. A graphical depiction of a PHD approximated by a Gaussian Mixture before and after an update is shown in figure 1.

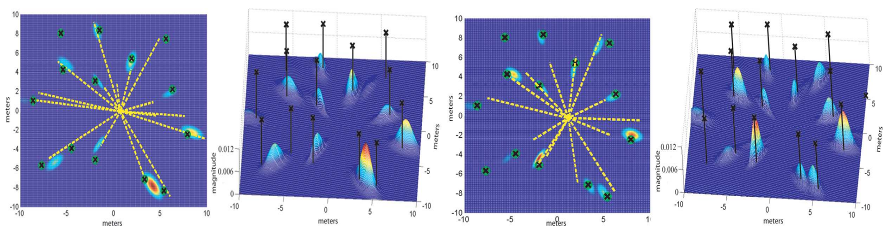

Fig. 1. An example of a map PHD superimposed on the true map represented by black dots. The peaks of the PHD represent locations with highest concentration of expected number of features. The PHD on the left is at time $k-1$, and that on the right is at time k.

The PHD construct is also beneficial to the multi-robot submapping problem [18], where the global map may be recovered by simply adding (and averaging the over-lap) the PHD sub-maps.

4 Filter Implementations

This section outlines a Gaussian Mixture implementation of the proposed filters. For the Robotic Mapping filter of section 3.1,let the predicted map PHD,

$$v_{k-1}(m|X_{k-1}) = \sum_{j=1}^{J_{k-1}} \eta_{k-1}^{(j)} \mathcal{N}\big(m; \mu_{k-1}^{(j)}, P_{k-1}^{(j)}\big) \tag{11}$$

which is a mixture of J_{k-1} Gaussians, with $\eta_{k-1}^{(j)}$, $\mu_{k-1}^{(j)}$ and $P_{k-1}^{(j)}$ being the corresponding predicted weights, means and covariances respectively for the j^{th} Gaussian component of the map PHD. Let the new feature intensity, $b(m|\mathcal{Z}_{k-1}, X_k)$, also be a Gaussian mixture of the form,

$$b(m|\mathcal{Z}_{k-1}, X_k) = \sum_{j=1}^{J_{b,k}} \eta_{b,k}^{(j)} \mathcal{N}\big(m; \mu_{b,k}^{(j)}, P_{b,k}^{(j)}\big) \tag{12}$$

where, $J_{b,k}$ defines the number of Gaussians in the new feature intensity at time k and $\eta_{b,k}^{(j)}$, $\mu_{b,k}^{(j)}$ and $P_{b,k}^{(j)}$ are the corresponding components. The predicted intensity is therefore also a Gaussian mixture,

$$v_{k|k-1}(m|X_k) = \sum_{j=1}^{J_{k|k-1}} \eta_{k|k-1}^{(j)} \mathcal{N}\big(m; \mu_{k|k-1}^{(j)}, P_{k|k-1}^{(j)}\big) \tag{13}$$

which consists of $J_{k|k-1} = J_{k-1} + J_{b,k}$ Gaussians representing the union of the prior map intensity, $v_{k-1}(m|X_{k-1})$, and the proposed new feature intensity, according to (9). Since the measurement likelihood is also of Gaussian form, it follows from (10) that the posterior map PHD, $v_k(m|X_k)$ is then also a Gaussian mixture given by,

$$v_k(m|X_k) = v_{k|k-1}(m|X_k)\Bigg[1 - P_D(m|X_k) + \sum_{z \in \mathcal{Z}_k} \sum_{j=1}^{J_{k|k-1}} v_{G,k}^{(j)}(z, m|X_k)\Bigg]. \tag{14}$$

The components of the above equation are given by,

$$v_{G,k}^{(j)}(z, m|X_k) = \eta_k^{(j)}(z|X_k)\mathcal{N}(m; \mu_{k|k}^{(j)}, P_{k|k}^{(j)}) \tag{15}$$

$$\eta_k^{(j)}(z|X_k) = \frac{P_D(m|X_k)\eta_{k|k-1}^{(j)} q^{(j)}(z, X_k)}{c(z) + \sum_{\ell=1}^{J_{k|k-1}} P_D(m|X_k)\eta_{k|k-1}^{(\ell)} q^{(\ell)}(z, X_k)} \tag{16}$$

where, $q^{(j)}(z, X_k) = \mathcal{N}\big(z; H_k \mu_{k|k-1}^{(j)}, S_k^{(j)}\big)$. 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. The clutter RFS, C_k, is assumed Poisson distributed [2] in number and uniformly spaced over the mapping region and Gaussian management methods are carried out as in [19].

For the Robotic Mapping filter of section 3.2, the location density of (6) can be propagated via Particle Filtering techniques [6], [15]. If the vehicle transition density is chosen as the proposal distribution, the weighting for the i^{th} particle becomes,

$$\widetilde{w}_k^{(i)} = g_k(\mathcal{Z}_k|\mathcal{Z}_{0:k-1}, \widetilde{X}_{0:k}^{(i)}) w_{k-1}^{(i)}. \tag{17}$$

Then for each hypothesised robot trajectory particle, an independent robotic mapping filter of equations (17)-(22) is executed.

5 Results & Analysis

The benchmark algorithms used in the analysis are the EKF-based mapping filters [4] and the FastSLAM localisation [6] algorithm with maximum likelihood data association, using a mutual exclusion constraint and a 95% χ^2 confidence gate. An arbitrary simulated feature map and vehicle trajectory are used as shown in figure 2. Existing approaches to mobile robotics typically deal with interfering measurements through ‘feature management’ routines, primarily via the landmark’s (feature’s) quality (LQ) [4] or a binary Bayes filter [6]. These operations are typically independent of the main filtering update, whereas the proposed approach unifies feature management, data association and state filtering into a single Bayesian update giving it a more robust performance in the presence of multiple mobile robots.

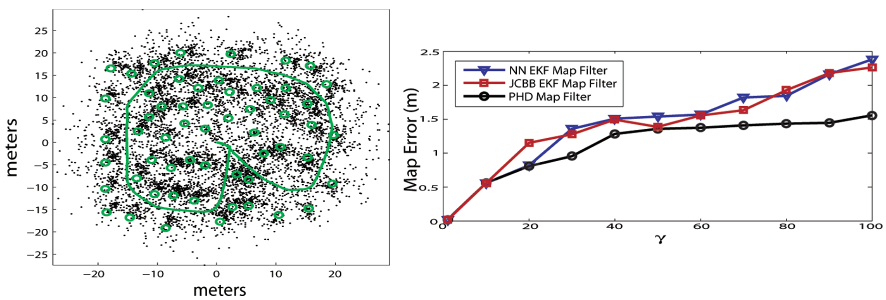

Fig. 2. *Left: The simulated environment showing point features (green circles). A sample measurement history plotted from the robot trajectory (green line) is shown. Right: Comparison of mapping error vs. measurement noise for the proposed filters and classical vector EKF solutions.*

Figure 2 also shows the filter performance in increasing measurement noise. Figure 3 shows a comparison of the map estimate for each filter at differing noise inflation values, γ. In a mobile robot swarm, mapping robots should exclude the moving robots from the static feature map. As such, figure 3 also depicts the suitability of the proposed approach to swarm based mapping, as the approach is robust to an increasing number of other robots. The proposed mapping framework performs well in the presence of increased measurement uncertainty and other mobile robots.

By assuming an unknown vehicle trajectory, and applying random control inputs, this section analyses the proposed joint mapping and localisation filter in comparison to the popular FastSLAM algorithm [6]. Figure 4 demonstrates the ability of the trajectory estimation filter to estimate the location of the robot at each time step, showing less error than conventional methods. In terms of estimating the map from an unknown path, figure 5 show how the RFS framework and proposed filters maintain accurate estimates of the number of features in the map and their locations, even in the presence of false measurements from other robots and mapping noise. Experimental results based on a 77GHz millimeter wave radar mounted on an autonomous robot are shown in 6.

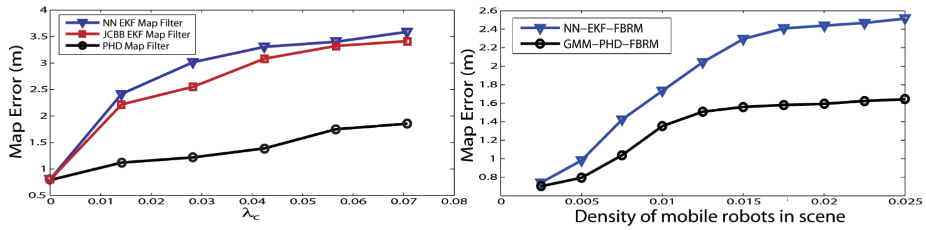

Fig. 3. *Left: Feature mapping error vs. clutter density for vector based NN-EKF and JCBB-EKF approaches and the proposed PHD framework, with the PHD approach seen to perform well in high clutter. Right: Comparison of the map estimation error in the presence of increasing densities per square meter of mobile robots.*

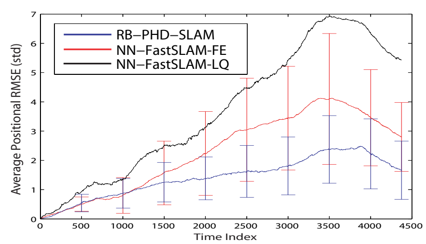

Fig. 4. *The mean and standard deviation of the expected trajectory estimates of the proposed RFS approach versus that of FastSLAM over 50 MC runs. LQ refers to an implementation with the 'landmark quality' method of [4].*

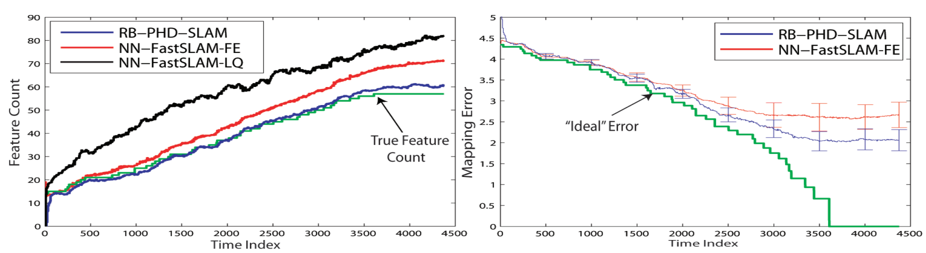

Fig. 5. *Left: The average estimated number of features in the map vs. ground truth for each approach. The feature number estimate from the proposed approach can be seen to closely track that of the ground truth. Right: A comparative plot of the mean and standard deviation of the map estimation error vs. time. Note that the 'ideal' error converges to zero, an important property for robotic mapping filters.*

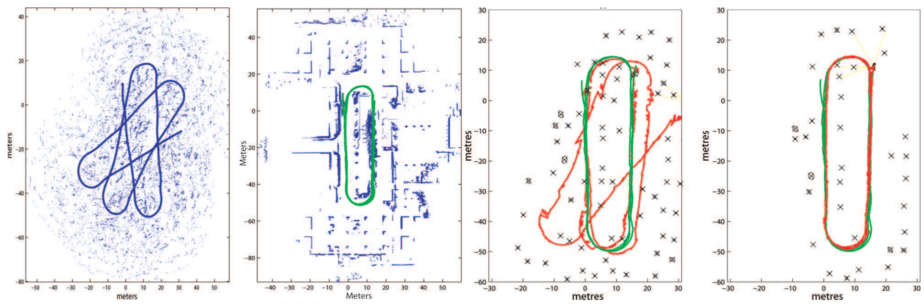

Fig. 6. *A: Raw radar measurements and noise vehicle path. B: The scan map plotted from the GPS path. C: Posterior Estimate from FastSLAM, D: Posterior Estimate from PHD-SLAM.*

6 Conclusion

This paper shows that from a fundamental estimation viewpoint that a feature-based map is a finite set and subsequently presented a Bayesian filtering formulation as well as a tractable solution for the feature-based mobile robotics problem. The framework outlined here presents a new direction of research for the multiple mobile robot community, which naturally encapsulates the inherent system uncertainty. Both a mapping-only and joint robot trajectory / map filter was introduced and analysed. In contrast to existing frameworks, the RFS approach to mobile robotics jointly estimates the number of features in the map as well as their individual locations in the presence of data association uncertainty and clutter. It was also shown that this Bayesian formulation admits a number of optimal Bayes estimators for mobile robotics problems. Analysis was carried out both in a simulated environment through Monte Carlo trials, demonstrating the robustness of the proposed filter, particulary in the presence of large data association uncertainty and clutter, illustrating the merits of adopting an RFS approach a swarm-based robotics applications.

7 Acknowledgements

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. Smith, R., Self, M., Cheeseman, P.: Estimating uncertain spatial relationships in robotics. Autonomous Robot Vehicles (1990) 167–193

2. Makarsov, D., Durrant-Whyte, H.: Mobile vehicle navigation in unknown environments: a multiple hypothesis approach. IEE Proceedings of Contr. Theory Applict. **vol. 142** (July 1995)
3. Thrun, S.: Particle filter in robotics. In: Uncertainty in AI (UAI). (2002)
4. Dissanayake, G., Newman, P., Durrant-Whyte, H., Clark, S., Csorba, M.: A solution to the simultaneous localization and map building (SLAM) problem. IEEE Transactions on Robotic and Automation **17**(3) (June 2001) 229–241
5. Guivant, J., Nebot, E., Baiker, S.: Autonomous navigation and map building using laser range sensors in outdoor applications. Journal of Robotic Systems **17**(10) (October 2000) 565–583
6. Montemerlo, M., Thrun, S., Siciliano, B.: FastSLAM: A Scalable Method for the Simultaneous Localization and Mapping Problem in Robotics. Springer (2007)
7. Mullane, J., Vo, B., Adams, M., Wijesoma, W.: A random set formulation for bayesian SLAM. In: proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, France (September 2008)
8. Mullane, J., Vo, B., Adams, M., Wijesoma, W.: 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)
9. Mullane, J., Vo, B., Adams, M., Vo, B.: A random finite set approach to bayesian SLAM. IEEE Transactions on Robotics (To Appear.)
10. Mullane, J., Vo, B., Adams, M., Vo, B.: Random Finite Sets for Robot Mapping & SLAM. Springer (To Appear.)
11. Mahler, R.: Global integrated data fusion. Proc. 7th Nat. Symp. on Sensor Fusion **1** (1994) 187–199
12. Mahler, R.: Multi-target bayes filtering via first-order multi-target moments. IEEE Transactions on AES **4**(39) (October 2003) 1152–1178
13. Mahler, R.: Statistical Multisource Multitarget Information Fusion. Artech House (2007)
14. Kaylan, B., Lee, K., Wijesoma, W.: FISST-SLAM: Finite set statistical approach to simultaneous localization and mapping. International Journal of Robotics Research **29**(10) (September 2010. Published online first in Oct, 2009.) 1251–1262
15. Mullane, J., Vo, B., Adams, M.: Rao-blackwellised PHD SLAM. In: proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Alaska, USA (May 2010)
16. Mullane, J., Keller, S., Rao, A., Adams, M., Yeo, A., Hover, F., Patrikalakis, N.: X-band radar based SLAM in singapores off-shore environment. In: proceedings of the 11th IEEE ICARCV, Singapore (December 2010)
17. Vo, B., Singh, S., Doucet, A.: Sequential monte carlo methods for multi-target filtering with random finite sets. IEEE Transactions on Aerospace and Electronic Systems **41**(4) (October 2005) 1224–1245
18. Shoudong, H., Zhan, W., Dissanayake, G.: Sparse local submap joining filter for building large-scale maps. IEEE Transactions on Robotics **24**(5) (May 2008) 1121–1130
19. Vo, B., Ma, W.: The gaussian mixture probability hypothesis density filter. IEEE Transactions on Signal Processing **54**(11) (November 2006) 4091–4104