

SLAM with Single Cluster PHD Filters

Chee Sing Lee, Daniel E. Clark, Joaquim Salvi

Abstract—Recent work by Mullane, Vo, and Adams has re-examined the probabilistic foundations of feature-based Simultaneous Localization and Mapping (SLAM), casting the problem in terms of filtering with random finite sets. Algorithms were developed based on Probability Hypothesis Density (PHD) filtering techniques that provided superior performance to leading feature-based SLAM algorithms in challenging measurement scenarios with high false alarm rates, high missed detection rates, and high levels of measurement noise. We investigate this approach further by considering a hierarchical point process, or single-cluster multi-object, model, where we consider the state to consist of a map of landmarks conditioned on a vehicle state. Using Finite Set Statistics, we are able to find tractable formulae to approximate the joint vehicle-landmark state based on a single Poisson multi-object assumption on the predicted density. We describe the single-cluster PHD filter and the practical implementation developed based on a particle-system representation of the vehicle state and a Gaussian mixture approximation of the map for each particle. Synthetic simulation results are presented to compare the novel algorithm against the previous PHD filter SLAM algorithm. Results presented indicate a superior performance in vehicle and map landmark localization, and comparable performance in landmark cardinality estimation.

Keywords: Simultaneous Localization and Mapping, probability hypothesis density filtering, doubly-stochastic processes, estimation.

I. INTRODUCTION

Work in SLAM has traditionally focused on repurposing single target tracking techniques such as the Extended Kalman Filter [1] to estimate a single combined state vector describing the pose of the vehicle and the map features. As a result of this approach, additional data association mechanisms are required to correctly identify which portion of the state vector generates an individual measurement, or whether the measurement originated from an actual target at all. A number of different algorithms have been developed, including gated nearest neighbor [2], joint compatibility branch and bound [3], and 1-Point RANSAC [4].

In contrast to the data association management strategies, the multitarget Bayes filter [5] extends the same principles which form the foundation of the single-object filtering techniques to estimate multi-object states. By reformulating the estimation problem into one where the number of targets is also uncertain, it inherently separates true measurements from clutter, and determines the most likely association between measurements and targets. However, for a non-trivial

number of targets, this filter is computationally unattainable. Instead of propagating the full multi-object posterior, only the first-order moment of the multi-target probability density is propagated in order to achieve tractability. The resulting algorithm is known as the Probability Hypothesis Density (PHD) filter [6], due to Stein and Winter's theory of evidence accrual [7]. Practical implementations of this filter have been presented, which represent the PHD with particles [8] or a Gaussian mixture [9]. The Cardinalized PHD filter [10] is an extension of the PHD filter that propagates the cardinality distribution alongside the first-order moment. This allows the multi-target state to be generalized from a Poisson process to an independently and identically distributed (i.i.d.) process. A Gaussian mixture implementation for the CPHD has also been proposed [11]. Recently, Mullane, Vo and Adams proposed a SLAM algorithm which combines a Rao-Blackwellized particle filter with the Gaussian mixture PHD filter for estimation of map landmarks [12]–[14]. In comparison to more established SLAM methods, it is most similar to the FastSLAM family of algorithms [15], but performs much better when faced with increased levels of measurement clutter. In that work, the authors arrived at their implementation by making a number of approximations based on assumed map cardinality. The PHD based SLAM algorithm presented is based on a different approximation, one which deals with the underlying statistical process governing the existence of map features.

The remainder of the paper is organized as follows: Section I discusses relevant past work in the field. Sections II and III describe the proposed methodology and implementation respectively. Section IV details simulation experiments and discusses the results thereof. Finally, Section V offers conclusions and sets a direction for further study.

II. SINGLE CLUSTER PROCESS SLAM

In Simultaneous Localization And Mapping (SLAM), a compact and convenient representation of the environment is a map of point features. These are obtained by abstracting measurements from onboard sensors such as sonar or cameras. Such feature abstraction does not always provide correct results, especially under adverse environmental conditions, and these imperfections are manifested as false positives (clutter) or negatives (missed detections). Therefore, it is desirable to develop SLAM algorithms that are able to cope with poor quality measurements, as robust feature extraction may not always be available. In this work, we arrive at such an algorithm by first modeling the SLAM problem as a doubly-stochastic point process, and then applying a PHD filter for the estimation of doubly-stochastic processes.

Chee Sing Lee and Joaquim Salvi are with the Computer Vision and Robotics Group (VICOROB), University of Girona, Girona, Spain cslee@eia.udg.edu, qsalvi@eia.udg.edu

Daniel Clark is with the Joint Research Institute in Signal and Image Processing, Heriot-Watt University, Edinburgh, UK d.e.clark@hw.ac.uk

Point processes model the random occurrence of events in a given interval of some arbitrary space [16]. Their applications span many disciplines, including epidemiology, ecology, finance, and telephone queuing.

In our SLAM algorithm, we propose that the existence of map features be modeled as a point process conditioned on the vehicle pose, which is also a stochastic process. This conditionality relationship is known as a *doubly-stochastic process*. We refer to this process as a *single cluster point process*, where the vehicle state is the *parent process*, and the map state is the *daughter process*. Our algorithm propagates the parent process and the first-order moment of the daughter process, known as its Probability Hypothesis Density (PHD), or intensity.

Let \mathbf{X}_k be the random vector that represents the vehicle state, and \mathbf{M}_k be the Random Finite Set (RFS) that represents the location of map features, which exist in the space $\mathcal{X} \subseteq \mathbb{R}^{n_m}$.

$$\mathbf{X}_k = [x_{k,1} \dots x_{k,n_x}] \quad (1)$$

$$\mathbf{M}_k = \{\mathbf{m}_{k,1} \dots \mathbf{m}_{k,\nu_k}\} \in \mathcal{F}(\mathcal{X}) \quad (2)$$

Where $\mathcal{F}(\mathcal{X})$ is the set of all finite subsets of \mathcal{X} . In addition, the vehicle receives measurements which are represented as an RFS \mathbf{Z}_k taken from the measurement space $\mathcal{Z} \subseteq \mathbb{R}^{n_z}$.

$$\mathbf{Z}_k = \{\mathbf{z}_{k,1} \dots \mathbf{z}_{k,\mu_k}\} \in \mathcal{F}(\mathcal{Z}) \quad (3)$$

This RFS is the union of measurements that originate from true targets, and measurements generated by a Poisson false alarm process whose PHD is $\kappa(\mathbf{z}) = \lambda U(\mathcal{Z})$, where λ is the Poisson rate parameter and $U(\mathcal{Z})$ is the uniform distribution over \mathcal{Z} . The measurement model $\mathbf{z} = h(\mathbf{m}, \mathbf{X})$ relates the measurements to landmark locations and the vehicle pose.

Let $p_k(\mathbb{X}_k) = p_k(\mathbf{X}_k, \mathbf{M}_k)$ be the joint posterior probability distribution of the vehicle state \mathbf{X}_k and multi-object map state \mathbf{M}_k at time step k . From a methodological perspective, the estimation problem to solve remains the same as other SLAM formulations: the sequential Bayesian estimation of $p_k(\mathbb{X}_k)$. For brevity, let $\mathbb{X} = \mathbb{X}_k$ and $\mathbb{X}' = \mathbb{X}_{k-1}$. The multiobject Bayes prediction and update formulas are:

$$p_{k|k-1}(\mathbb{X}|\mathbf{Z}_{1:k-1}) = \int \pi_k(\mathbb{X}|\mathbb{X}') p_{k-1}(\mathbb{X}') \delta \mathbb{X}' \quad (4)$$

$$p_k(\mathbb{X}|\mathbf{Z}_{1:k}) = \frac{g_k(\mathbf{Z}_k|\mathbb{X}) p_{k|k-1}(\mathbb{X}|\mathbf{Z}_{1:k-1})}{\int g_k(\mathbf{Z}_k|\mathbb{X}) p_{k|k-1}(\mathbb{X}|\mathbf{Z}_{1:k-1}) \delta \mathbb{X}} \quad (5)$$

The difference between the above equations and those of the single object Bayes filter is that because \mathbf{M}_k is a random finite set rather than a random vector, the denominator of (5) becomes a set integral [6], which involves an infinite sum over all possible joint distributions of different cardinalities. These integrals make implementation of the multiobject Bayes filter intractable for a variable number of targets. Instead, we propagate the first moment of the multiobject probability distribution. In order to make the resulting algorithm tractable, we make an assumption on the predicted daughter process, namely that it can be approximated with a Poisson spatial point process. This assumption was made in the derivation of the single-cluster PHD filter [17].

In the single-cluster Poisson process, the prior PHD of the joint vehicle landmark state can be factorized as follows,

$$D_{k-1}(\mathbf{X}, \mathbf{m}) = s_{k-1}(\mathbf{X}) \times \tilde{D}_{k-1}(\mathbf{m}|\mathbf{X}) \quad (6)$$

where $s_{k-1}(\mathbf{X})$ is the intensity of the parent process, and $\tilde{D}_{k-1}(\mathbf{m}|\mathbf{X})$ is the intensity of the daughter process. The single cluster PHD prediction equation is

$$D_{k|k-1}(\mathbf{X}, \mathbf{m}) = \int s_{k-1}(\mathbf{X}') \pi_{k|k-1}(\mathbf{X}|\mathbf{X}') \tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}') d\mathbf{X}' \quad (7)$$

$$\tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}') = \gamma_{k|k-1}(\mathbf{m}|\mathbf{X}') + \int D_{k-1}(\mathbf{m}'|\mathbf{X}') \tilde{\pi}_{k|k-1}(\mathbf{m}|\mathbf{m}'; \mathbf{X}') d\mathbf{m}' \quad (8)$$

$\pi_{k|k-1}(\mathbf{X}|\mathbf{X}')$ is the Markov transition density for the parent process, and $\tilde{\pi}_{k|k-1}(\mathbf{m}|\mathbf{m}', \mathbf{X}')$ is the conditional Markov transition density for the daughter process; $\gamma_{k|k-1}(\mathbf{m}|\mathbf{X})$ is the PHD for the daughter birth process. The single-cluster PHD update can be separated into a parent update and a daughter update.

$$s_k(\mathbf{X}) = \frac{L_{\mathbf{Z}_k}(\mathbf{X}) s_{k|k-1}(\mathbf{X})}{\int L_{\mathbf{Z}_k}(\mathbf{X}) s_{k|k-1}(\mathbf{X}) d\mathbf{X}} \quad (9)$$

$$\tilde{D}_k(\mathbf{m}|\mathbf{X}) = (1 - p_D(\mathbf{m}|\mathbf{X})) \tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}) + \sum_{\mathbf{z} \in \mathbf{Z}_k} \frac{p_D(\mathbf{m}|\mathbf{X}) \tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}) g_k(\mathbf{z}|\mathbf{m}, \mathbf{X})}{\kappa_k(\mathbf{z}) + \int p_D(\mathbf{m}|\mathbf{X}) \tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}) g_k(\mathbf{z}|\mathbf{m}, \mathbf{X}) d\mathbf{m}} \quad (10)$$

Where $g_k(\mathbf{z}|\mathbf{m}, \mathbf{X})$ is the single-object measurement likelihood, and $L_{\mathbf{Z}_k}(\mathbf{X})$ is the multi-object measurement likelihood, both conditional on the vehicle state. The multi-object likelihood is defined as

$$L_{\mathbf{Z}_k}(\mathbf{X}) = \exp \left\{ - \int p_D(\mathbf{m}|\mathbf{X}) \tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}) d\mathbf{m} \right\} \times \prod_{\mathbf{z} \in \mathbf{Z}_k} \left(\kappa_k(\mathbf{z}) + \int p_D(\mathbf{m}|\mathbf{X}) \tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}) g_k(\mathbf{z}|\mathbf{m}; \mathbf{X}) d\mathbf{m} \right) \quad (11)$$

Note that the equations pertaining to the daughter (8)(10) are identical to the conventional PHD filter prediction and update equations [6]. These update equations define the primary difference between the Single-Cluster PHD filter and the previously proposed Rao-Blackwellized PHD filter,

III. IMPLEMENTATION

We implement the Single-Cluster PHD filter using a Dirac mixture model for the PHD of the parent. Each component of the parent mixture model is associated with a Gaussian mixture model which represents the PHD of the daughter process conditioned on that particular parent component. At each iteration of the filter, we begin with the following prior

PHDs:

$$s_{k-1}(\mathbf{X}) = \sum_{i=1}^{N_{k-1}} \eta_{k-1}^{(i)} \delta(\mathbf{X} - \mathbf{X}_{k-1}^{(i)}) \quad (12)$$

$$D_{k-1}^{(i)}(\mathbf{m}|\mathbf{X}) = \sum_{j=1}^{J_{k-1}^{(i)}} w_{k-1}^{(j|i)} \mathcal{N}(\mathbf{m}; \mu_{k-1}^{(j|i)}, \mathbf{P}_{k-1}^{(j|i)}) \quad (13)$$

The notation $\mathcal{N}(m; \mu, \mathbf{P})$ is used to denote a Gaussian distribution with mean vector μ and covariance matrix \mathbf{P} , and $\delta(\mathbf{X} - \alpha)$ denotes the Dirac delta distribution centered at α . Mixture models may be conveniently represented by the set of their parameters, so we may alternatively express the prior PHDs like so:

$$s_{k-1}(\mathbf{X}) = \{\eta_{k-1}^{(i)}, \mathbf{X}_{k-1}^{(i)}\}_{i=1}^{N_{k-1}} \quad (14)$$

$$D_{k-1}^{(i)}(\mathbf{m}|\mathbf{X}) = \{w_{k-1}^{(j|i)}, \mu_{k-1}^{(j|i)}, \mathbf{P}_{k-1}^{(j|i)}\}_{j=1}^{J_{k-1}^{(i)}} \quad (15)$$

A. Map Prediction

The prediction for the daughter process is the same as that for the standard Gaussian mixture PHD filter.

$$\tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) = \tilde{D}_{b,k|k-1}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) + \tilde{D}_{S,k|k-1}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) \quad (16)$$

The first term is a Gaussian mixture corresponding to births, or the appearance of new features in the map.

$$\tilde{D}_{b,k|k-1}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) = \sum_{j=1}^{J_{b,k|k-1}} w_{b,k|k-1}^{(j)} \mathcal{N}(\mathbf{m}; \mu_{b,k|k-1}^{(j)}, \mathbf{P}_{b,k|k-1}^{(j)}|\mathbf{X}) \quad (17)$$

The parameters for the birth Gaussians are derived from the measurement scan of the previous time step. The second term in the prediction represents the propagation of features that survive from the prior map:

$$\tilde{D}_{S,k|k-1}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) = p_S \sum_{j=1}^{J_{S,k|k-1}} w_{S,k|k-1}^{(j)} \mathcal{N}(\mathbf{m}; \mu_{S,k|k-1}^{(j)}, \mathbf{P}_{S,k|k-1}^{(j)}|\mathbf{X}) \quad (18)$$

$$w_{S,k|k-1}^{(j)} = w_{k-1}^{(j)} \quad (19)$$

$$\mu_{S,k|k-1}^{(j)} = f(\mu_{k-1}^{(j)}) \quad (20)$$

$$\mathbf{P}_{S,k|k-1}^{(j)} = \mathbf{F}_k \mathbf{P}_{k-1}^{(j)} \mathbf{F}_k^T + \mathbf{Q}_k \quad (21)$$

B. Vehicle Prediction

The sampling property of the Dirac delta function means that substitution of (12) and (13) into (7) results in the following sum:

$$D_{k|k-1}(\mathbf{X}, \mathbf{m}) = \sum_{i=0}^{N_{k-1}} \pi_{k|k-1}(\mathbf{X}|\mathbf{X}_{k-1}^{(i)}) \tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) \quad (22)$$

The Markov transition density is then approximated by sampling M particles from it. The result is a new Dirac mixture for the parent, containing $N_{k-1} \times M$ components. Because the predicted maps $\tilde{D}_{k|k-1}(\mathbf{m}|\mathbf{X}_i')$ do not depend on the current predicted vehicle pose, each of the M parent

components that originate from the same $\mathbf{X}_{k-1}^{(i)}$ can be assigned identical copies of the predicted map.

C. Measurement Update

Like the prediction, the measurement update for the daughter process mirrors that of the standard GM-PHD filter.

$$\tilde{D}_k^{(i)}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) = (1 - p_D(\mathbf{m}|\mathbf{X}_{k|k-1}^{(i)})) \tilde{D}_{k|k-1}^{(i)}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) + \sum_{\mathbf{z} \in \mathbf{Z}_k} \tilde{D}_{D,k}^{(i)}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) \quad (23)$$

where

$$\tilde{D}_{D,k}^{(i)}(\mathbf{m}|\mathbf{X}_{k-1}^{(i)}) = \sum_{j=1}^{J_{k|k-1}} \frac{w_k^{(j|i)} \mathcal{N}(\mathbf{m}; \mu_k^{(j|i)}, \mathbf{P}_k^{(j|i)})}{\kappa_k(\mathbf{z}) + \sum_{l=1}^{J_{k|k-1}} w_k^{(l|i)}} \quad (24)$$

$$w_k^{(j|i)} = p_D(\mu_{k|k-1}^{(j|i)}|\mathbf{X}_{k|k-1}^{(i)}) \times g_k(\mathbf{z}|\mu_{k|k-1}^{(j|i)}; \mathbf{X}_{k|k-1}^{(i)}) w_{k|k-1}^{(j|i)} \quad (25)$$

$$\mu_k^{(j|i)} = \mu_{k|k-1}^{(j|i)} - K_k^{(j|i)}(\mathbf{z} - \hat{\mathbf{z}}_k^{(j|i)}) \quad (26)$$

$$\mathbf{P}_k^{(j|i)} = (\mathbf{I} - \mathbf{K}_k^{(j|i)} \mathbf{J}_k^{(j|i)}) \mathbf{P}_{k|k-1}^{(j|i)} \quad (27)$$

$$\mathbf{J}_k^{(j|i)} = \left. \frac{\partial}{\partial \mathbf{m}} h(\mathbf{m}, \mathbf{X}) \right|_{\mathbf{m}=\mu_{k|k-1}^{(j|i)}, \mathbf{X}=\mathbf{X}_{k|k-1}^{(i)}} \quad (28)$$

$$\mathbf{K}_k^{(j|i)} = \mathbf{P}_{k|k-1}^{(j|i)} \mathbf{J}_k^{(j|i)} \mathbf{S}_k^{(j|i), -1} \quad (29)$$

$$\mathbf{S}_k^{(j|i)} = \mathbf{J}_k^{(j|i)} \mathbf{P}_{k|k-1}^{(j|i)} \mathbf{J}_k^{(j|i), T} + \mathbf{R}_k \quad (30)$$

$$\hat{\mathbf{z}}_k^{(j|i)} = h(\mu_{k|k-1}^{(j|i)}, \mathbf{X}_{k|k-1}^{(i)}) \quad (31)$$

Let $FOV_k(\mathbf{X}_k) \in \mathcal{X}$ be the vehicle's sensor field of view at time k , dependent on the current vehicle location. Assuming a constant probability of detection p_D , we have:

$$p_D(\mathbf{m}|\mathbf{X}_{k|k-1}^{(i)}) = \begin{cases} p_D & \text{if } \mathbf{m} \in FOV_k(\mathbf{X}_{k|k-1}^{(i)}) \\ 0 & \text{otherwise} \end{cases} \quad (32)$$

Consequently, for landmarks outside of the field of view, the updated feature will be identical to the predicted one because only the first term of (23) will be non-zero. This means that only the features within $FOV_k(\mathbf{X}_{k|k-1}^{(i)})$ need to be updated, and remaining feature estimates can be propagated forward untouched. In order to perform the measurement update for the parent process, we must first compute the multi-object measurement likelihood.

$$L_{\mathbf{Z}_k}(\mathbf{X}^i) = \exp \left\{ - \sum_{j=1}^{J_{k|k-1}} p_D(\mu_{k|k-1}^{(j|i)}|\mathbf{X}_{k|k-1}^{(i)}) w_{k|k-1}^{(j|i)} \right\} \times \prod_{\mathbf{z} \in \mathbf{Z}_k} \left(\kappa_k(\mathbf{z}) + \sum_{j=1}^{J_{k|k-1}} w_k^{(j|i)} \right) \quad (33)$$

With this likelihood in hand, the weights of the Dirac mixture can be updated:

$$\eta_k^{(i)} = \frac{L_{\mathbf{Z}_k}(\mathbf{X}^i)}{\sum_{l=1}^{N_{k|k-1}} L_{\mathbf{Z}_k}(\mathbf{X}^l)} \eta_{k-1}^{(i)} \quad (34)$$

At this point we have the updated posterior parent and daughter PHDs. However, some steps need to be taken to

manage the computational complexity of the filter. During the prediction for the parent, each component in the Dirac mixture is “shotgunned” into M new components, resulting in a new mixture containing $M \times N_{k-1}$ components. Left unchecked, the size of the parent process mixture would grow exponentially with every time step. To curb this growth, we prune the mixture to the N_{k-1} components with the highest weights. The Gaussian mixtures for the daughter process also have the potential for this exponential growth, as the measurement update generates $|\mathbf{Z}| + 1$ new Gaussians for each component in the predicted mixture. Many of these come from low-likelihood measurement associations and contribute little to the updated PHD. We employ Salmond’s clustering method for Gaussian mixture reduction [18] to eliminate or merge these low-weight components to generate a more manageable approximation of the PHD.

It is important to remark that this implementation is not the only way to realize the single-cluster PHD SLAM algorithm. For example, in [17], a single-cluster PHD filter with a Gaussian parent was proposed. However, we have decided on this implementation because the Dirac mixture should be more able to capture non-linear motion models.

IV. SYNTHETIC RESULTS

The Single-Cluster PHD SLAM algorithm was validated with simulations on synthetic data. The simulated vehicle conforms to the Ackerman steering motion model used in [19], and the sensor input consists of range and bearing measurements to point features in the environment. The standard deviation of the zero-mean Gaussian odometry noise was 2 m/s for the velocity, and 5° for the steering angle. The measurement noise had a std. deviation of 1 m and 2° for the range and bearing components respectively. A probability of detection $p_D = 0.95$ was used to cull the true feature measurements, and false alarm measurements were added to the sensor inputs, at an average of $\lambda = 5$ per scan. Figure 1 depicts an example scenario generated with these parameters.

Using a fixed ground truth map and trajectory, 50 Monte Carlo runs were executed, with both odometry and sensor inputs regenerated for each run. The filter used 64 components for the parent mixture, with a shotgun factor of $M = 4$ in the prediction. Birth components were computed from the previous timestep’s vehicle location and measurements, and a constant birth weight of $w_b = 0.01$ was used. The expected *a posteriori* state estimate was taken for the vehicle pose, while the maximum *a posteriori* state estimate was taken for the map. The estimated number of map features, \hat{N}_k , was determined by summing the weights in the Gaussian mixture of the most highly weighted particle, and the \hat{N}_k most highly weighted Gaussians were taken as the map estimate.

For comparison, the Rao-Blackwellized PHD SLAM algorithm in [14] was also implemented, using the single feature map assumption for the particle weight update. As the RB-PHD SLAM has already been shown to out-perform more well established methods such as the Extended Kalman Filter and FastSLAM, these were not included in the simulations.

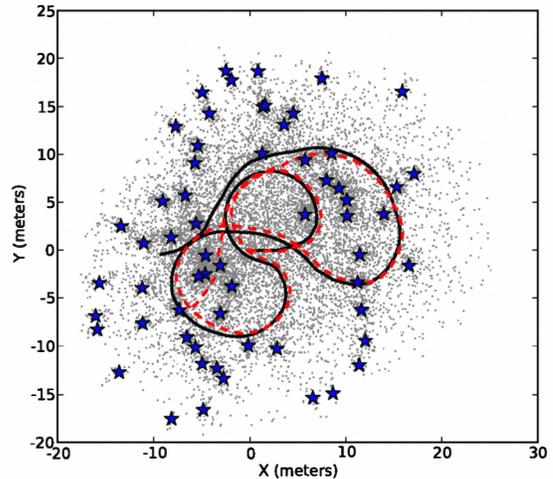


Fig. 1: Illustration of the scenario used for simulations, showing the true vehicle trajectory (black solid), dead-reckoning trajectory (red dashed), true map landmarks (blue stars), and cumulative sensor measurements (gray dots)

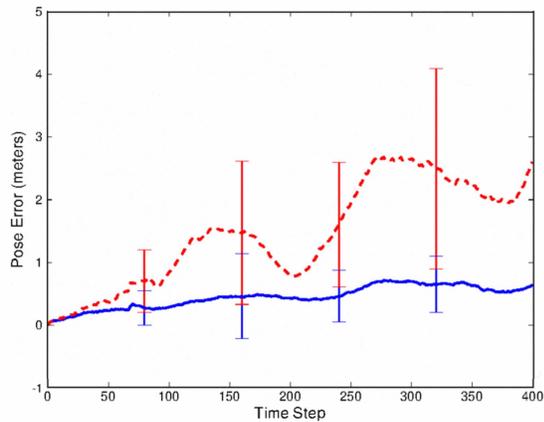
Both algorithms were implemented in C with the CUDA 4.0 Parallel Computing Toolkit, and the simulation was run with an NVIDIA Tesla C2070 GPU.

Performance in vehicle localization was evaluated using Euclidean distance, while mapping performance was evaluated using the OSPA metric [20], with order $p = 1$ and cutoff $c = 5$ meters. Results are shown in Figures 2 and 3.

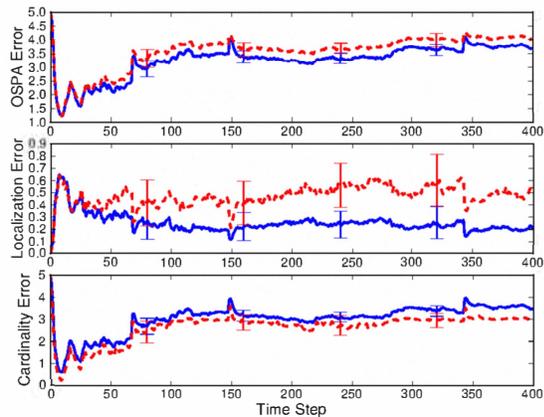
V. DISCUSSION AND CONCLUSIONS

The results presented in this work indicate that the single-cluster PHD filter provides a viable solution to the SLAM problem in scenarios with high measurement clutter. Compared to previous PHD SLAM methods, it achieves superior performance in localizing both the vehicle and map landmarks, while estimation of landmark cardinality remains comparable. With the Dirac mixture implementation for the parent, the implementation of the single-cluster PHD SLAM algorithm is closely related to the RB-PHD SLAM, the difference being the weight updates for the vehicle pose particles. It is likely that the improved performance of the algorithm stems from the approximations made in order to achieve tractable update formulae. In RB-PHD-SLAM, Poisson assumptions were made on both the prior and posterior, and a further assumption was made on the number of features. In the single-cluster PHD SLAM, only a Poisson approximation was made on the prior.

To illustrate these differences, a small experiment was performed. Using the same set of sensor inputs, the measurement updates for the two methods were executed on identical predicted states. Figure 4 shows the resulting parent distributions. It is apparent that the SC-PHD SLAM update generates a significantly more focused parent distribution than the RB-PHD-SLAM update. This suggests that the multi-object likelihood used in the single-cluster derivation is more discriminating, and is better able to concentrate the parent particles about the true vehicle position. We



(a) Vehicle Pose Error



(b) Map OSPA Error, with localization and cardinality components

Fig. 2: MC simulations results, with $1-\sigma$ bounds indicated. SC-PHD SLAM: blue, solid. RB-PHD SLAM: red, dashed

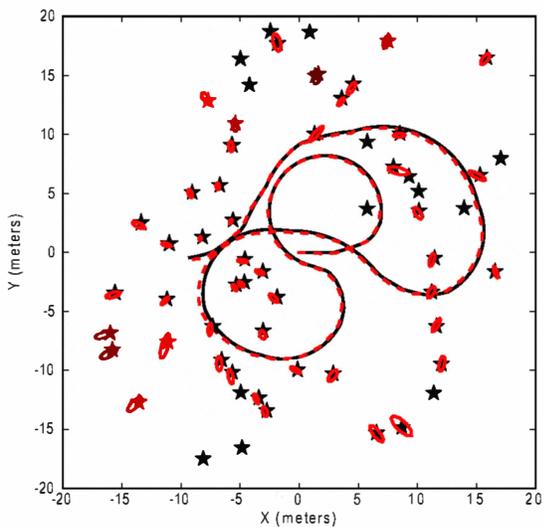


Fig. 3: Example of resulting trajectories and map from SC-PHD SLAM (shown in red)

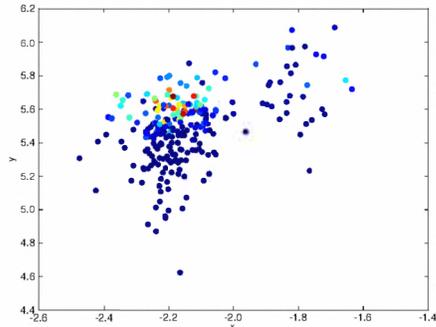
would also like to remark about computational performance of the PHD filter SLAM. As implemented in this work, the running time for both algorithms was on the order of several minutes per Monte Carlo run. At this point, online application appears out of the question. However, as more computational resources become available in the future, these methods are likely to become increasingly relevant. Moreover, because the PHD update is applied only to map features within the vehicle’s field of view, the SC-PHD SLAM will scale suitably to longer missions. A number of different research avenues are open to continue this work. We feel that the synthetic simulations presented here sufficiently validate the SC-PHD SLAM algorithm, but further work will reinforce these results using real world datasets. The Poisson assumption on the map process introduces a high variance in the cardinality estimate. We believe that relaxing this assumption to one of an i.i.d. process, by substituting the single-cluster PHD filter with a single-cluster CPHD filter, will improve the cardinality estimation.

VI. ACKNOWLEDGMENTS

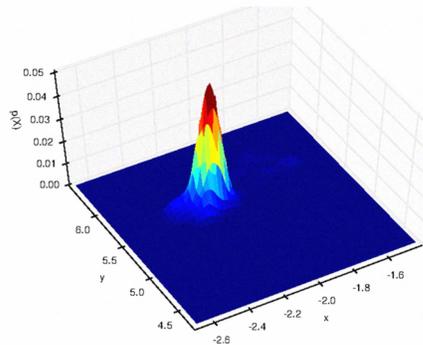
Chee Sing Lee is granted with a PhD FI Scholarship of the Catalan Government. Daniel Clark is a Royal Academy of Engineering/ EPSRC Research Fellow. This work has been supported by the FP7-ICT-2011-7 project PANDORA—Persistent Autonomy through Learning, Adaptation, Observation and Re-planning (Ref 288273) funded by the European Commission. Joaquim Salvi is supported by the project RAIMON—Autonomous Underwater Robot for Marine Fish Farms Inspection and Monitoring (Ref CTM2011-29691-C02-02) funded by the Spanish Ministry of Science and Innovation.

REFERENCES

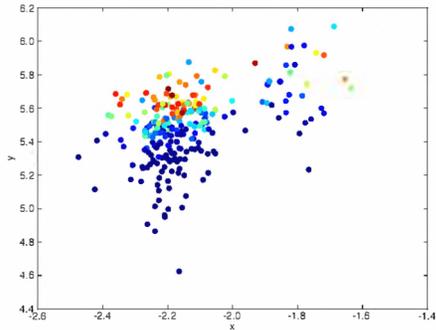
- [1] M. Dissanayake, P. Newman, S. Clark, H. Durrant-Whyte, and M. Csorba, “A solution to the simultaneous localization and map building (SLAM) problem,” *Robotics and Automation, IEEE Transactions on*, vol. 17, no. 3, pp. 229–241, Jun. 2001.
- [2] Y. Bar-Shalom, *Tracking and data association*. San Diego, CA, USA: Academic Press Professional, Inc., 1987.
- [3] J. Neira and J. Tardós, “Data association in stochastic mapping using the joint compatibility test,” *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 890–897, December 2001.
- [4] J. Civera, O. G. Grasa, A. J. Davison, and J. M. M. Montiel, “1-point RANSAC for extended kalman filtering: Application to real-time structure from motion and visual odometry,” *J. Field Robot.*, vol. 27, pp. 609–631, September 2010. [Online]. Available: <http://dx.doi.org/10.1002/rob.v27:5>
- [5] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, Inc., 2007.
- [6] R. Mahler, “Multitarget bayes filtering via first-order multitarget moments,” *IEEE Trans. Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, October 2003.
- [7] M. C. Stein and C. Winter, “An additive theory of bayesian evidence accrual,” Analysis and Assessment Division, Los Alamos National Laboratory, Tech. Rep. ADA364591, 1987. [Online]. Available: <http://handle.dtic.mil/100.2/ADA364591>
- [8] B.-N. Vo, S. Singh, and A. Doucet, “Sequential monte carlo methods for multitarget filtering with random finite sets,” *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 41, no. 4, pp. 1224–1245, oct. 2005.
- [9] B.-N. Vo and W.-K. Ma, “The gaussian mixture probability hypothesis density filter,” *Signal Processing, IEEE Transactions on*, vol. 54, no. 11, pp. 4091–4104, 2006.



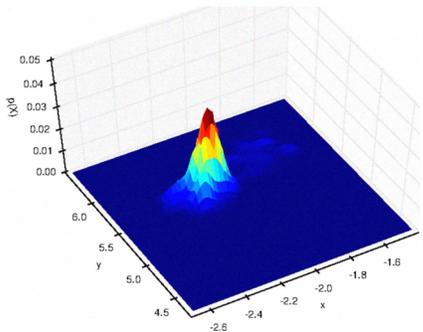
(a) SC-PHD SLAM particles



(b) SC-PHD SLAM smoothed distribution



(c) RB-PHD SLAM particles



(d) RB-PHD SLAM smoothed distribution

Fig. 4: Comparison of updated vehicle distributions. Smoothing was performed using a Gaussian kernel with a bandwidth $h = 0.05$

- [10] R. Mahler, "PHD filters of higher order in target number," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 43, no. 4, pp. 1523–1543, 2007.
- [11] B.-T. Vo, B.-N. Vo, and A. Cantoni, "Analytic implementations of the cardinalized probability hypothesis density filter," *IEEE Transactions on Signal Processing*, pp. 3553–3567, 2007.
- [12] J. Mullane, B.-N. Vo, and M. D. Adams, "Rao-blackwellised PHD SLAM," in *ICRA'10*, 2010, pp. 5410–5416.
- [13] J. Mullane, B.-N. Vo, M. D. Adams, and B.-T. Vo, "A random-finite-set approach to bayesian SLAM," *IEEE Transactions on Robotics*, vol. PP, no. 99, pp. 1–15, 2011.
- [14] J. Mullane, B.-N. Vo, M. Adams, and B.-T. Vo, *Random Finite Sets for Robot Mapping and SLAM - New Concepts in Autonomous Robotic Map Representations*, ser. Springer Tracts in Advanced Robotics. Springer, 2011, vol. 72.
- [15] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit, "FastSLAM: A factored solution to the simultaneous localization and mapping problem," in *Proceedings of the AAAI National Conference on Artificial Intelligence*. Edmonton, Canada: AAAI, 2002.
- [16] D. J. Daley and D. Vere-Jones, *An introduction to the theory of point processes. Vol. I*, 2nd ed., ser. Probability and its Applications (New York). New York: Springer-Verlag, 2003, elementary theory and methods.
- [17] A. Swain and D. Clark, "The single-group phd filter: an analytic solution," in *International Conference on Data Fusion*, 2011.
- [18] D. Salmond, "Mixture reduction algorithms for point and extended object tracking in clutter," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 45, no. 2, pp. 667–686, april 2009.
- [19] J. E. Guivant and E. M. Nebot, "Optimization of the simultaneous localization and map-building algorithm for real-time implementation," *IEEE Transactions on Robotics*, vol. 17, no. 3, pp. 242–257, 2001.
- [20] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Transactions on Signal Processing*, vol. 56, no. 8-1, pp. 3447–3457, 2008.