

Simultaneous Localization and Mapping for Spacecraft Rendezvous and Proximity Operations Using Random Finite Sets

DAVID E. GAYLOR, Member, IEEE
NASA Goddard Space Flight Center, Greenbelt, MD, USA

LAUREN KONITZER
NASA Goddard Space Flight Center, Greenbelt, MD, USA

MARK MORETTO
North Carolina State University, Raleigh, NC, USA

RICHARD LINARES
Massachusetts Institute of Technology, Cambridge, USA

Abstract— Future space missions require spacecraft to autonomously navigate non-cooperative environments for rendezvous and proximity operations (RPO). Current navigation filters used for RPO can have difficulty when optical sensors are used, due to complications with data association, high measurement uncertainty, and clutter. This paper provides an initial demonstration of the feasibility of using random finite set (RFS) filters for spacecraft relative navigation and pose estimation. Spacecraft relative navigation is formulated as a simultaneous localization and mapping (SLAM) problem, in which an observer spacecraft seeks to simultaneously estimate the location of features on a target object or spacecraft as well as its relative position, velocity and attitude. Simulated flash LIDAR measurements are processed using a Gaussian Mixture Probability Hypothesis Density (GMPHD) filter embedded in a particle filter to obtain a feature map of a target and a relative pose estimate between the target and observer over time. Results show that an RFS-based filter can successfully perform SLAM in a spacecraft relative navigation scenario.

This work was funded by the NASA Goddard Space Flight Center Internal Research and Development (IRAD) and NASA Pathways Internship programs. (Corresponding author: D. Gaylor)

David E. Gaylor is an Aerospace Engineer at NASA Goddard Space Flight Center, Greenbelt, MD (e-mail: david.gaylor@nasa.gov). Lauren Konitzer is an Aerospace Engineer at NASA Goddard Space Flight Center, Greenbelt, MD (e-mail: lauren.konitzer@nasa.gov). Mark Moretto is an Assistant Professor at North Carolina State University, Raleigh, NC (e-mail: mjmorett@ncsu.edu). Richard Linares is the Rockwell International Career Development Professor at the Massachusetts Institute of Technology, Cambridge, MA (e-mail: linaresr@mit.edu).

I. INTRODUCTION

Future space missions require autonomous navigation solutions for spacecraft rendezvous and proximity operations (RPO) to support a variety of missions including on-orbit servicing, active debris removal, and small body exploration. However, reliable solutions are difficult to achieve due to the challenges of operating in space. Noisy measurements cluttered with additional false detections can make a state estimation filter diverge if association between measurements and state space cannot be performed quickly and correctly.

A. Simultaneous Localization and Mapping

Spacecraft relative navigation can be formulated as a Simultaneous Localization and Mapping (SLAM) problem, where a robotic agent (e.g. the spacecraft) attempts to estimate a map of its environment and its pose (position, velocity, and attitude) within that environment at the same time. With feature-based SLAM methods, sensors typically measure features in the environment, for example the location of points or edges on the body of a rendezvous target using sensors such as cameras or LIDAR. Thus, the mapping part of SLAM becomes a multi-target tracking problem, as the observer is attempting to estimate the location of many features to construct a feature map of the other body.

Many methods have been proposed to solve the multi-target SLAM problem and are well summarized by surveys of current SLAM methods [1], [2]. These existing methods however tend to be limited by the fact that the underlying mathematics were originally intended for single-target tracking. To use the same methods for multi-target tracking, heuristic methods must be used to account for additional challenges introduced by multi-target tracking. Among the most significant of these challenges is the data association problem, where measurements must be matched to states in state space for navigation. False sensor returns can make it difficult to assign a measurement to a feature prior to filtering; some heuristic must be developed to assign measurements to expected features despite having extraneous measurements available. Conversely, a filtering algorithm must be able to account for a sensor failing to detect a feature that is expected, which commonly occurs with optical measurements due to difficult lighting conditions such as those found in space. With a low number of features and measurements, heuristic methods may suffice for solving these problems on a case by case basis. However, as the number of features and measurements increases, the data association problem becomes computationally challenging and the filters

that rely on the correct association can diverge if the association is incorrect. [3]

Previous studies have addressed SLAM methods for spacecraft relative navigation. Work by Augenstein as well as Sonnenburg et al., has demonstrated the problem can be solved using a SLAM formulation [4], [5]. However, both methods required significant computational resources for feature management and data association prior to the navigation filter; moreover, their formulations do not account for clutter and missed detections. Dennison, et al, successfully demonstrated their proposed architecture for autonomously performing SLAM, as well as estimating an asteroid's gravity field, rotational motion, and shape using an Unscented Kalman Filter (UKF) processing data from a swarm of spacecraft in a high fidelity simulation [6].

One of the most widely used feature-based SLAM algorithms is FastSLAM, where the feature map is estimated by an Extended Kalman Filter (EKF) and the corresponding trajectory and pose is estimated with a particle filter. [7] FastSLAM has been successful because of its lack of assumptions on the pose probability distribution, and real-time implementations, which are crucial for limited computation on-board a spacecraft [8]. However, the data association problem must still be solved before FastSLAM can be performed. A study by Cocaud and Kubota uses an algorithm based on an improved particle filtering method, intended for navigation and pinpoint landing on small bodies. [9] Their results address issues with drift in the particle filter that is commonly seen in FastSLAM, however their approach is reliant on heuristic methods for data association of features.

B. Random Finite Sets for SLAM

The random finite set (RFS) approach to multi-target tracking is an emerging alternative to traditional methods. In the RFS formulation, the collection of individual targets is treated as a set-valued state, and the collection of individual observations is treated as a set-valued observation. Modeling set-valued states and set-valued observations as RFSs allows the problem of dynamically estimating multiple targets in the presence of clutter and association uncertainty to be directly addressed in a Bayesian filtering framework.

The first RFS filter, the Probability Hypothesis Density (PHD) filter, was proposed for multi-target tracking by Ronald Mahler in the early 2000's [10]. Mullane, et al., demonstrated a feature-based map used for SLAM is fundamentally a finite set, and is thus more accurately represented and estimated with RFS methods [3]. Their results show that in the presence of uncertain dynamics and measurements

and high degrees of clutter, the PHD filter can outperform EKF-based mapping filters.

Several recent studies have also leveraged the ability of RFS-based tracking methods for space-based multi-target tracking applications. McCabe and DeMars used RFS-based filters in feature-based robotic mapping for planetary landing, successfully mapping the features of a lunar environment during a descent trajectory given a known spacecraft pose [11]. Their work compares the multi-target tracking performance of several RFS-based filters for this purpose. Studies have also performed and evaluated the use of RFS-based filters for space object tracking, resulting in an integrated framework for successful identification and tracking of space objects [12]. These works have shown that RFS-based methods such as the PHD filter are a promising approach for handling the challenges associated with multi-target tracking in space.

The objective of this paper is to provide an initial demonstration of the feasibility of using random finite set (RFS) filters for spacecraft relative navigation and pose estimation. We formulate spacecraft relative navigation as a SLAM problem and use an RFS-based filter for mapping combined with a particle filter for pose estimation. In this paper, we simulate processing flash LIDAR measurements using a Gaussian Mixture Probability Hypothesis Density (GMPHD) filter embedded in a particle filter to obtain a feature map of a target and a relative pose estimate between the target and observer.

II. THEORY

A. Random Finite Sets for Multi-Target Tracking

In multi-target tracking, we wish to jointly estimate the number of targets and their states. We are given a set of noisy, potentially cluttered measurements. Assume that at time step k , there exists a set of $N(k)$ targets and $M(k)$ measurements. In general there is no specific order in which it is known that these measurements and states are associated, thus we represent them as random finite sets given by:

$$X_k = \{x_{k,1}, \dots, x_{k,N(k)}\} \quad (1)$$

$$Z_k = \{z_{k,1}, \dots, z_{k,M(k)}\} \quad (2)$$

In the RFS framework, the sets X_k and Z_k are a multi-target state and multi-target observation, respectively. X_k and Z_k are random finite sets, variables that can be characterized by a probability distribution and a family of joint probability densities of all the elements of X_k and Z_k [13]. Additionally, the cardinality $|\cdot|$, or number of elements in a random finite set becomes a variable that can be estimated, and does not have to remain the same over time.

This formulation allows more possibilities for the time evolution of the states contained in X_k . For a given multi-target state X_{k-1} , each element $x_{k-1} \in X_{k-1}$ either continues to exist at the next time step k with probability $p_{S,k}$ or ceases to exist with probability $1 - p_{S,k}$. Alternatively, a target which does not exist at time $k - 1$ may begin to exist at time step k independently of the targets existing at time step $k - 1$. We say that this target has been birthed. Similarly, a new target may begin to exist at time step k by being “spawned” from a target which existed at time step $k - 1$. In real scenarios, both of these events may occur because a target is entering the region of state space of interest after previously being outside it, or a target that was initially estimated to be one target becomes better resolved and is more likely to be several targets that are closely spaced. Thus at a time step k , a general target state is represented as

$$X_k = \left[\bigcup_{\zeta \in X_{k-1}} S_{k|k-1}(\zeta) \right] \cup \left[\bigcup_{\zeta \in X_{k-1}} \Gamma_{k|k-1}(\zeta) \right] \cup B_k \quad (3)$$

where X_k is composed of a union of a set of surviving targets $S_{k|k-1}$ each with transition density $f_{k|k-1}(x_k|x_{k-1})$, a set of newly birthed targets B_k , and a set of spawned targets $\Gamma_{k|k-1}$. Note that Γ and B are entirely general and their forms are determined by the specific scenario in which they are used.

Similarly, the RFS measurement model accounts for probabilities of detection and the possibility of clutter. A target $x_k \in X_k$ can either be detected with probability $p_{D,k}$, or not detected (missed) with probability $1 - p_{D,k}$. Additionally, the framework of RFS allows for the concept of clutter measurements, which we represent as an additional RFS K_k of false detections which do not originate from a target. Thus, our set of measurements at a time step k is denoted by

$$Z_k = K_k \cup \left[\bigcup_{x \in X_k} \Theta_k(x) \right] \quad (4)$$

where Z_k is composed of a union of sets of actual measurements $\Theta_k(x)$ which occur with probability density $g_k(z_k|x_k)$ and clutter measurements K_k . K_k is also entirely general and determined by the specific scenario involved.

The goal of multi-target filtering is to obtain a posterior density of the multi-target state X_k given a multi-target observation Z_k . The posterior density can be calculated using a Bayesian recursion given by

$$p_{k|k-1}(X_k|Z_{1:k-1}) = \int f_{k|k-1}(X_k|X) p_{k-1}(X|Z_{1:k-1}) \mu_s(dX) \quad (5)$$

$$p_k(X_k|Z_{1:k}) = \frac{g_k(Z_k|X_k) p_{k|k-1}(X_k|Z_{1:k-1})}{\int g_k(Z_k|X) p_{k|k-1}(X|Z_{1:k-1}) \mu_s(dX)} \quad (6)$$

where $f_{k|k-1}(\cdot|\cdot)$ is a transition density from one state to another, $g_k(\cdot|\cdot)$ is the observation likelihood of a measurement given a state; μ_s is an appropriate reference measure on the collection of all finite subsets of state space [14] and $Z_{1:k-1}$ is the prior history of measurements. Unfortunately this recursion is computationally intractable due to the set integrals required.

B. Gaussian Mixture PHD Filter

To approximate the multi-target posterior density in equations 5 and 6, we choose to propagate the posterior intensity, which is the first order statistical moment of the multi-target state. This simplification forms the basis of the Probability Hypothesis Density (PHD) filter [10]. As a result, the integral of the intensity function over a particular region of state space is equal to the number of targets expected to exist in that region of state space. This means that the number of expected targets at each time step can be propagated jointly. In other words, we find that $\hat{N}(k) = \int \nu(x) dx$, where the intensity function $\nu(x)$ is referred to as the probability hypothesis density and is integrated over the whole RFS space, allowing for estimation of the cardinality of the multi-target state set.

A number of additional assumptions are made to obtain a tractable PHD recursion from equations 5-6. Full derivations and proof are detailed by Vo and Ma in [14] and will not be reproduced here. To use the PHD recursion in a tractable, closed form manner, the following assumptions must be made: each target follows a linear Gaussian dynamics model, each sensor follows a linear Gaussian measurement model, and the intensities of the birth and spawn sets are modeled as Gaussian mixtures. This means that we assume each individual target follows a similar model as is typically seen with a classic Kalman filter:

$$f_{k|k-1}(x|\zeta) = \mathcal{N}(x; F_{k-1}\zeta, Q_{k-1}) \quad (7)$$

$$g_k(z|x) = \mathcal{N}(z; H_k x, R_k) \quad (8)$$

where in general $\mathcal{N}(\cdot; m, P)$ represents a Gaussian distribution with mean m and covariance P ; Q_{k-1} is process noise covariance on the dynamics, F_{k-1} is the state transition matrix, H_k is the measurement model observation matrix, and R_k is the covariance of

the measurement noise. Additionally, the Gaussian mixture birth and spawn intensities become:

$$\beta_k(x) = \sum_{i=1}^{J_{\beta,k}} w_{\beta,k}^{(i)} \mathcal{N}(x; m_{\beta,k}^{(i)}, P_{\beta,k}^{(i)}) \quad (9)$$

$$\gamma_{k|k-1}(x|\zeta) = \sum_{j=1}^{J_{\gamma,k}} w_{\gamma,k}^{(j)} \mathcal{N}(x; F_{\gamma,k}^{(j)} \zeta + d_{\gamma,k-1}^{(j)}, Q_{\gamma,k}^{(j)}) \quad (10)$$

where J, w, P are given model parameters that determine the shape of the corresponding intensities and can be chosen depending on the scenario, allowing for the inclusion of *a priori* knowledge of the targets if desired or available.

Finally, we assume that the posterior intensity at $k-1$ is also a Gaussian mixture, which continues to be a Gaussian mixture at the next time step:

$$v_{k-1}(x) = \sum_{i=1}^{J_{k-1}} w_{k-1}^{(i)} \mathcal{N}(x; m_{k-1}^{(i)}, P_{k-1}^{(i)}) \quad (11)$$

The end result of using these assumptions is the following computationally tractable recursion, where the predicted intensity at time k is also a Gaussian mixture consisting of the sum of surviving, spawned, and birthed targets:

$$v_{k|k-1}(x) = p_{S,k} \sum_{j=1}^{J_{k-1}} w_{k-1}^{(j)} \mathcal{N}(x; m_{S,k|k-1}^{(j)}, P_{S,k|k-1}^{(j)}) + v_{\gamma,k|k-1}(x) + \beta_k(x) \quad (12)$$

$$m_{S,k|k-1}^{(j)} = F_{k-1} m_{k-1}^{(j)}$$

$$P_{S,k|k-1}^{(j)} = Q_{k-1} + F_{k-1} P_{k-1}^{(j)} F_{k-1}^T$$

where the intensity of spawned targets depends on the set of previously existing targets:

$$v_{\gamma,k|k-1}(x) = \sum_{j=1}^{J_{k-1}} \sum_{\ell} w_{k-1}^{(j)} w_{\gamma,k}^{(\ell)} \mathcal{N}(x; m_{\gamma,k|k-1}^{(j,\ell)}, P_{\gamma,k|k-1}^{(j,\ell)}) \quad (13)$$

$$m_{\gamma,k|k-1}^{(j,\ell)} = F_{\gamma,k-1}^{(\ell)} m_{k-1}^{(j)} + d_{\gamma,k-1}^{(\ell)}$$

$$P_{\gamma,k|k-1}^{(j,\ell)} = Q_{\gamma,k-1}^{(\ell)} + F_{\gamma,k-1}^{(\ell)} P_{k-1}^{(j)} (F_{\gamma,k-1}^{(\ell)})^T$$

Then, the measurement updated posterior is also a Gaussian mixture:

$$v_k(x) = (1 - p_{D,k}) v_{k|k-1}(x) + \sum_{z \in Z_k} v_{D,k}(x; z) \quad (14)$$

where the intensity of detected targets is:

$$v_{D,k}(x; z) = \sum_{j=1}^{J_{k|k-1}} w_k^{(j)}(z) \mathcal{N}(x; m_{k|k}^{(j)}(z), P_{k|k}^{(j)}) \quad (15)$$

$$m_{k|k}^{(j)}(z) = m_{k|k-1}^{(j)} + K_k^{(j)}(z - H_k m_{k|k-1}^{(j)})$$

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

$$K_k^{(j)} = P_{k|k-1}^{(j)} H_k^T (H_k P_{k|k-1}^{(j)} H_k^T + R_k)^{-1}$$

and where the weight of each element of the Gaussian Mixture is given by:

$$w_k^{(j)}(z) = \frac{p_{D,k} w_{k|k-1}^{(j)} q_k^{(j)}(z)}{\kappa_k(z) + p_{D,k} \sum_{\ell=1}^{J_{k|k-1}} w_{k|k-1}^{(\ell)} q_k^{(\ell)}(z)}$$

$$q_k^{(j)}(z) = \mathcal{N}(z; H_k m_{k|k-1}^{(j)}, R_k + H_k P_{k|k-1}^{(j)} H_k^T)$$

Thus, equations 12-15 combined give a tractable, closed-form solution for the GMPHD filter. This recursion takes in noisy, cluttered measurements and extracts an intensity function which represents the location of likely target positions in state space. The form presented here is similar to the Kalman filter, but it is more general due to the nuances introduced by the RFS formulation.

1. Extended Kalman Filter Approximations

Similar to how the Kalman filter can be modified to incorporate nonlinear dynamics and measurement models, Vo and Ma [14] have shown that the PHD filter can also use nonlinear dynamics and measurements by locally linearizing the state transition matrix and observation matrix. These approximations for the state transition and observation model derivative matrices are used. In this paper, as the problem defined in Section III utilizes nonlinear dynamics and measurement models.

C. Rao-Blackwellized Particle Filter

The GMPHD filter in equations 12-15 is used to perform the mapping duties of SLAM, similar to the approach used by Vo and Ma [14]. In the context of a SLAM application, the targets tracked by the filter are the features which have been identified and extracted from images of the target spacecraft taken by the observer. To perform the localization aspect of SLAM, a method must be used to obtain a relative pose estimate of the observer based on the feature map multi-target state provided by the PHD filter. Mullane et al., successfully achieved this by using an outer Rao-Blackwellized particle filter (RB-PF) [3]. To fully perform SLAM with an RFS-based filter, we also use a RB-PF to recursively obtain a maximum *a posteriori* pose estimate.

In this formulation, each particle of the particle filter represents an individual estimate of the pose of the observer as well as a corresponding estimate of the map. The map comes from the PHD filter which each particle utilizes independently, thus the feature map belonging to each particle is conditioned on an independent estimate of the pose. Figure 1 depicts

the relationships between the GMPHD filter and the particle filter. The combined filter structure is referred to as the Rao-Blackwellized GMPHD (RBGM-PHD) filter in the remainder of this paper. Further details of the filter implementation are provided in [16].




Fig. 1. Simplified diagram depicting the flow of data between different parts of the filter [17].

1. Importance Weighting

For this method to succeed, there needs to be a way to determine which particle has the “best” joint estimate of feature map and observer pose. To achieve this, an importance weight $p_l(Z_k | X_k^{[l]}, Z_{1:k-1})$ is calculated for each l th particle. In other words, a weight can be strategically assigned to each particle based on how well the observations in Z_k match with the map conditioned pose associated with each particle. The basis of importance weighting begins by assuming that a posterior density of states x and measurements z can generally be approximated with N_p particles as

$$p(x_k | z_{1:k}) \approx \sum_{l=1}^{N_p} \eta_k^{[l]} \delta(x_k | x_k^{[l]}) \quad (16)$$

where the weight $\eta_k^{[l]}$ from Bayes’ Rule of the l th particle is calculated recursively as

$$\eta_k^{[l]} \propto \eta_{k-1}^{[l]} \frac{p(z_k | x_k^{[l]}) p(x_k^{[l]} | x_{k-1}^{[l]})}{\pi(x_k | x_{k-1}, z_k)} \quad (17)$$

This approximation of the posterior approaches the true posterior density as $N_p \rightarrow \infty$. Thus, the weights of the particles in equation 17 must be calculated to sequentially update the weight of a particle based on newly obtained measurements at each time step. Further assumptions must be made to make this a tractable computation. In this paper,

Single Cluster (SC) importance weighting [18] is used.

Lee argues that SLAM is fundamentally a single cluster process, where the daughter process of the map is conditioned on the parent process of the vehicle pose. Thus by assuming the measurements obtained are a result of Poisson point processes, a single cluster update for the PHD filter can be obtained. This argument and subsequent derivation of how it affects the calculation of the multi-target likelihood is presented in [19] and is not reproduced here.

Using these arguments, the right hand side of equation 17 becomes computable and the updated weight $\eta_k^{[l]}$ of the l th particle can be calculated from the previous map PHD $v_k^{-[l]}$ as:

$$\eta_k^{[l]} = \exp \left(\sum_{r=1}^{N(k)-[l]} w_k^{r,[l]} \right) \times \prod_{z \in Z_k} \left(\kappa(z) + p_{D,k} \sum_{r=1}^{N(k)-[l]} p(z | \mathcal{N}_k^{r,[l]}, x_{0:k}^{[l]}) w_k^{r,[l]} \right) \eta_{k-1}^{[l]} \quad (18)$$

where $\mathcal{N}_k^{r,[l]}$ is the r th Gaussian in the intensity $v_k^{-[l]}$ of targets in the PHD filter associated with particle l and $x_{0:k}$ is the history of states. This is a computationally tractable form of equation 17, and it represents a multi-target likelihood, or how likely the measurements “agree” with each particle’s pose-conditioned feature map. Moreover, it does not require data association for computation of the multi-target likelihood.

2. Resampling

In addition to importance weighting, a particle filter requires periodic resampling to maintain particle hypothesis in regions of the posterior that have higher probability. The primary reasons for this are to prevent degenerate particle estimates with unreasonably high variance and to improve the exploration of state space, enabling more efficient filtering. In this paper, Low Variance Resampling (LVR) [20] was chosen as the resampling technique because it is relatively simple yet rigorously systematic and thus lends itself well to applications in which the “best” sampling distribution isn’t necessarily known or may be complicated and difficult to resolve. Moreover, LVR schemes are of complexity $\mathcal{O}(N_p)$, as compared to naive approaches which are typically of complexity $\mathcal{O}(N_p \log N_p)$. In LVR, particle hypotheses have a certain probability of being deleted and reset to be more similar to higher weighted particles, based on their original weight [21].

The important properties of LVR are: the probability that a particle gets resampled is directly pro-

portional to its weight, if resampling does not take place, the particle maintains the same weight, and particles may be “roughened” after being resampled, meaning that their properties are varied slightly from the particles from which they were resampled in an attempt to maintain particle diversity and efficient exploration of state space. The magnitude of roughening which occurs is determined by the user. In this paper, particles are resampled at a given time step only if the “effective” number of particles, defined as $\sum \eta_i^2$, falls below a chosen threshold which is tunable by the user.

3. Pose Estimation

The relative pose estimate at each time step is updated by extracting the pose of the particle with the highest weight prior to resampling, i.e. the maximum *a posteriori* estimate. With the inclusion of a particle filter for pose estimation the overall method employed in this paper resembles that of FastSLAM, where the EKF portion for mapping has been replaced with a GMPHD filter.

III. SPACECRAFT RELATIVE NAVIGATION SIMULATION

This section describes the spacecraft navigation scenario that was simulated to test the proposed RFS-SLAM filter.

A. Test Case Scenario

Figure 2 shows a general spacecraft RPO scenario, including relevant coordinate frames and defining vectors. The observer spacecraft (also commonly referred to as the chaser or deputy spacecraft) is collecting measurements of features located on a target spacecraft (also commonly referred to as the chief spacecraft.) The features being observed are physical details that are “optically significant,” such as edges or corners of a solar panel, or craters and rocks on an asteroid.




Fig. 2. Relationship between observer and target coordinate frames and vector definitions.

1. Target Properties

The target object is assumed to be spherical and covered with randomly distributed features. It has a radius of 50 meters and has 20 randomly distributed features on its surface. This size was chosen to be between a typical satellite and a small asteroid. For this initial demonstration, the features are chosen from a uniform distribution constrained to the surface of the sphere – however, the filter used in this paper is aware of this constraint.

B. Assumptions

Several assumptions are being made to simplify the models used for this initial feasibility study. The measurement frame (i.e. the camera frame) is the same as the observer body fixed frame to avoid including another rotation from an external measurement frame to the observer body frame. A spherical target was chosen to eliminate the need for occlusion computations. The motion of the spherical target includes the general dynamics of features rotating in and out of view of the observer. Clutter measurements are drawn from a uniform distribution in the field of view of the sensor independently at each time step. In other words, clutter measurements are non-persistent. Persistent clutter could be used instead without loss of generality in the algorithms used, and it has been demonstrated by past studies. [22] And the noise on the measurements is Gaussian.

C. System Dynamics

The target is in a circular orbit about the Earth, and the observer is in close proximity to the target, i.e. within less than 1km range. Thus, the dynamics of the observer relative to the target can be approximated by the linear Clohessy-Wiltshire (CW) equations.

To simplify the dynamic equations, the motion is constrained such that the attitude of the target is fixed in the CW frame, similar to a nadir pointing spacecraft. We also assume that the observer is perfectly tracking the target in its field of view, which is a reasonable assumption given an appropriate control system for pointing. In other words, the center of the target is assumed to be located in the center of the images taken by the observer. This negates the need to estimate angular velocity, since the angular velocity of the observer body frame with respect to the Hill body frame can then be directly calculated as:

$$\omega = \frac{\mathbf{r} \times \dot{\mathbf{r}}}{|\mathbf{r}|^2} \quad (19)$$

The relative attitude of the observer and target are expressed using Modified Rodrigues Parameters (MRPs) as they have one easily avoidable singular-

ity, compared to multiple singularities common in an Euler angle description of the attitude dynamics.

D. Measurement Model

In the scenarios studied in this paper, the observer spacecraft collects flash LIDAR measurements of the target spacecraft or body. These measurements are corrupted with Gaussian-distributed noise after being simulated from the truth trajectory. Additionally, extraneous measurements which do not originate from features on the target are appended to the noisy measurements which did originate from features on the target. Although the algorithms used in this paper are, in theory, fully capable of handling missed detections, this was not simulated.

We define that a measurement of an extracted feature consists of angular position in pixel coordinates on an image as well as an associated range for that pixel position. It is important to note that the measurements have no other identifying information, i.e. the filter doesn't receive any information about which feature a particular measurement originated from.

The flash LIDAR measurements are simulated from the simulated feature locations using a pinhole projection camera model for the pixel coordinates. The range measurement of a feature is determined by the Euclidean distance from the optical center $\mathbf{C}$ of the sensor to the feature located on the target at point $\mathbf{P}$. The sensor frame is defined such that the u and v axes are aligned with the rows and columns of the detector and the third axis points along the boresight of the optics. The pixel coordinates of the feature image $\mathbf{p}$ are defined in the sensor frame as $\{u, v, w\}$. The measurements of a feature are thus defined as $\mathbf{y} = [u \ v \ \rho]$, where ρ is defined as:

$$\rho = \|\mathbf{X} - \mathbf{X}_c\| \quad (20)$$

where $\mathbf{X}$ is the feature location and $\mathbf{X}_c$ is the sensor location in the same frame, in this case the Hill frame previously defined as $\{\mathbf{B}\}$. To calculate u and v a pinhole camera projection is defined with the following standard equations:

$$\mathbf{x} = w \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} wu \\ wv \\ w \end{bmatrix} \quad (21)$$

where $\mathbf{x}$ is calculated as:

$$\mathbf{x} = \begin{bmatrix} fm_u & s & P_u \\ 0 & fm_v & P_v \\ 0 & 0 & 1 \end{bmatrix} [R][I_{3 \times 3} - \mathbf{X}_c] \begin{bmatrix} \mathbf{X} \\ 1 \end{bmatrix} \quad (22)$$

where $[R]$ is the rotation matrix from the Hill frame to the sensor frame, f is the focal length of the detector or the distance between the sensor plane and image plane, s can be used to account for any skew in the optics if desired, P_u and P_v are the

location in pixels of the optical center of the sensor projected onto the image plane, and m_u and m_v are the inverse sizes of the pixels in each dimension. [23] These relationships fully dictate the base measurement model used to generate non-corrupted measurements.

Noise and $n_{clutter}$ clutter measurements are added after truth measurements are calculated. The noise is generated with a zero-mean Gaussian distribution with a block diagonal covariance $[\sigma_u^2, \sigma_v^2, \sigma_\rho^2]$. Clutter events were Poisson distributed and clutter measurements were uniformly distributed across the field of view. If a measurement with noise is outside the the sensor field of view, then it is disregarded prior to the filter. The constants used to simulate these data are included in Table I. These values were chosen to produce an angle of view of approximately 14° , to closely mimic the specifications of the flash LIDAR on-board Raven, an experimental payload on the International Space Station which performs autonomous tracking of rendezvous targets. [24].

TABLE I
Measurement model parameters.

Parameter	Value	Units
f	25	mm
P_u, P_v	128	$[]$
s	0	$[]$
m_u, m_v	20480	m^{-1}
n_{col}, n_{row}	256	$[]$
σ_u, σ_v	1	$[]$
σ_ρ	10	m
$n_{clutter}$	10	$[]$

IV. RESULTS AND DISCUSSION

The purpose of this paper is to demonstrate that an RFS-based filter can perform SLAM for spacecraft relative navigation. The mapping results of the PHD filter are analyzed by observing how closely the estimated map matches the actual map of visible features at each time step, both in terms of how many features are tracked and their proximity to the actual feature locations. Pose estimation is analyzed by comparing the estimated pose with the true pose.

A. Test Case Descriptions

The measurement models and dynamics used to generate the test cases are described in Sections II and III. Two test cases are used to demonstrate the results of the filter. The orbit for the first test case is shown in Figure 3; the blue curve is the trajectory, the sphere is the target body, and the red x's on

the target body are the locations of the features to be tracked. This is a periodic orbit around the target body that could be used during an inspection phase of a servicing mission or the mapping phase of a small body exploration mission. This orbit was chosen to investigate the long-term behavior of the filter, on the order of several hours. In this case, the observer periodically encounters the same features as ones it has observed in the past. Currently, no mechanism is in place to have the filter store features that have been tracked in the past; features are simply deleted from the tracked set once they go out of view. However in the future, this test case would be ideal for implementing a loop-closure method, i.e. storing old features and identifying if they are encountered later in the trajectory.




Fig. 3. Trajectory for Case 1, the periodic orbit [17].

A fundamental problem with SLAM implementations is dealing with sparse maps. Thus, a second test case was chosen to investigate how the filter performs in the event that few or no measurements originate from a true feature for a period of time, i.e. no features on the target body are visible to the sensor. The relative orbit of the observer with respect to the target is shown in Figure 4 and is referred to as a walking safety ellipse (WSE). This type of relative orbit may also be used for the inspection and rendezvous phase of a servicing mission [25].

The WSE brings the observer to within 30m of the surface of the target and as far away as 650m. The trajectory shown is over a time period of 15,000 seconds to see the behavior of the orbit, though this full time span is not shown in the results. The sensor periodically sees no features on the target body for at least 1 time step, leading to a sparse feature map.

In both test cases, the birth model $\gamma_k(x)$ for the PHD filter is given to be a Gaussian Mixture with means located at the true feature locations and covariance of 10 meters. Usually birth distributions would be defined as a distribution mixture around the edge of the field of view, however, this can be computationally intractable since for every additional birth distribution, another potential target would be propagated through every prediction and update. Thus, we assumed some knowledge of the




Fig. 4. Trajectory for Case 2, the walking safety ellipse [17].

a-priori state of your target set, to avoid this potential computational issue. Since feature spawning is not expected to occur in this scenario, no spawn intensity function is specified. All time steps are in units of 50 seconds. Each test case is given 10 uniformly distributed clutter measurements at each time step, in addition to the noisy measurements which originate from true features. For the RB-PF implementation, 100 particles are used. The pose hypotheses of the particles are initially distributed slightly offset from the true pose state with a covariance corresponding to a standard deviation of 0.1 meters in the position coordinates, 0.1 cm/s in the velocity coordinates, and $1e-3$ in the dimensionless attitude coordinates. This covariance was chosen to allow for a relatively small initial particle dispersion similar to what may be expected from an external estimate of the pose from other sensors.

The initial relative states used to generate these trajectories are given in Table II. The features are randomly distributed onto the target body and measurements are simulated.

TABLE II
Test case initial conditions.

Test Case	1	2
x (m)	-200	0
y (m)	0	0
z (m)	0	-500
$\dot{x}$ (m/s)	0	0.1
$\dot{y}$ (m/s)	0.4241	-0.01
$\dot{z}$ (m/s)	0	0.1

B. Test Case 1: Periodic Orbit

Figure 5 shows the results of the RBGMPHD filter mapping step for the first case, the periodic

orbit. This case was chosen to be slightly less strict about allowing features to be accepted past the pruning step. From the first time step, the map is fully tracked; as time progresses, the features continue to be observed over a long period of time.




Fig. 5. Sequential snapshots of the estimated feature map for the full RBGMPHD filter used on Case 1 [17].

Figure 6 depicts the number of true features visible to the observer compared to the number of features obtained by the PHD filter. Additionally, the less strict pruning of the features is seen as the filter consistently overestimates how many features are visible. Further tuning would make it possible to improve these estimates, as well as implementing the merging procedure specified by Vo and Ma [14]. [14]




Fig. 6. Comparison between the true number of visible features vs. how many features the RBGMPHD filter has estimated for Case 1 [17].

Figure 7 shows the position and velocity errors of the highest weight particle compared to the filter 3σ uncertainties, which are calculated from the sample covariance of all the particle poses. The black line shows the difference between the true and estimated pose, and the red dashed line shows the 3σ uncertainties of the 100 pose hypotheses. Note that there is no sample covariance calculated for the Euler angles, as these have been converted from MRPs. In the MRP formulation, the concept of an additive sample covariance is not valid. The periodic nature of the orbit is seen in the filter covariance uncertainties, as the sample covariance periodically increases and decreases, particularly in the Z coordinate which

may be slightly less observable due to the planar nature of the orbit. In some dimensions, the filter appears to converge on the correct pose as time progresses. These results show that the RBGMPHD filter has potential for accurate and stable pose estimation over a long period of time.




Fig. 7. Comparison between sample covariance 3σ uncertainties vs. the estimate error for Case 1 [17].

C. Test Case 2: Walking Safety Ellipse

Figure 8 shows the results of the RBGMPHD filter for the highest weight particle at 4 selected time steps. The PHD filter is not only determining the feature map from the first time step, but the map continues to be updated as features move in and out of the field of view as the observer sees different sides of the target. Moreover, when measurements become sparse such as at 1400 seconds, the filter is able to recover and maintain a stable map estimate.




Fig. 8. Sequential snapshots of the estimated feature map from the RBGMPHD filter over time for Case 2 [17].

Figure 9 depicts the number of true features visible to the observer compared to the number of features obtained by the PHD filter. This figure shows that the estimated number of features roughly tracks the truth.




Fig. 9. Comparison between the true number of visible features vs. how many features the RBGMPPH filter has estimated for Case 2 [17].

Figure 10 shows the position and velocity errors compared to the sample covariance 3σ uncertainties of all the particle poses, as well as the Euler angle differences between the true and estimated relative attitude. These results show that the covariance uncertainties are reasonable for the amount of estimation error, and the pose is being tracked to within a few meters in relative position and cm/s in relative velocity, and generally less than one degree of error in relative attitude. These results show that even in the presence of a sparse feature map with features frequently moving in and out of the field of view of the sensor, the RBGMPPH formulation can successfully perform SLAM, although we note that the birth model is a Gaussian Mixture with means located at the true feature locations and covariance of 10 meters.




Fig. 10. Comparison between sample covariance 3σ uncertainties vs. the position and velocity errors for Case 2 [17].

V. CONCLUSIONS

Spacecraft relative navigation can be formulated as a SLAM problem where the mapping consists of tracking a number of features on the target spacecraft or body. We proposed using an RFS-based filter for mapping combined with a particle filter for pose estimation to perform SLAM. The proposed method avoids some of the issues with current methods, such as the need for a separate heuristic data association algorithm to determine a mapping between measurements and features prior to filtering. Additionally, the RFS framework inherently accounts for measurement issues which occur with optical navigation in space, such as clutter measurements,

and the possibility that features move in and out of the field of view of a sensor over time.

The test cases used to test the RBGMPPH filter provided a challenging dynamics and measurement environment for the filter; namely an abnormally low signal to noise ratio (i.e. many clutter measurements, at times a ratio of 1:10). Realistically, the first use of an RFS-based filter on-board a spacecraft may be in less difficult circumstances. For example the initial pose of the particle filter may be initialized by another estimate of the spacecraft pose from other sensors. The majority of the challenges seen with the RBGMPPH filter came from particle filter tuning issues and sensitivity to pose estimate errors. Despite this, careful selection of a resampling method can allow the particle filter to converge estimates even when the estimation errors grow large. Thus, the results presented in this paper demonstrate that use of an RFS-based filter is feasible for spacecraft relative navigation. Future studies could investigate the ability of this method to handle higher numbers of features, robustness to reduction in number of tracked features and larger apriori errors, and increased fidelity of dynamics and measurement models.

ACKNOWLEDGMENTS

A preliminary version of this research was presented at the AAS/AIAA Space Flight Mechanics Meeting in 2019 [17]. The authors would like to thank Dr. Martin Adams and Felipe Inostroza at the Universidad de Chile for their advice and consultation.

REFERENCES

- [1] N. Dhiman, D. Deodhare, and D. Khemani. Where am i? creating spatial awareness in unmanned ground robots using slam: A survey. *Sadhana Academy Proceedings*, vol. 40, no. 5, pp. 1385–1433, 2015.
- [2] G. Bresson, Z. Alsayed, L. Yu, and S. Glaser. Simultaneous localization and mapping: A survey of current trends in autonomous driving. *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 3, pp. 194–220, 2017.
- [3] 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*. Berlin: Springer-Verlag, 2011.
- [4] S. Augenstein. Monocular pose and shape estimation of moving targets for autonomous rendezvous and docking. *Doctoral Thesis, Stanford University*, 2011. [Online]. Available: <https://purl.stanford.edu/xf316zh6614>
- [5] A. Sonnenburg, M. Tkocz, and K. Janschek. EKF-slam based approach for spacecraft rendezvous navigation with unknown target spacecraft. *International Federation of Automatic Control Proceedings*, vol. 43, no. 15, pp. 339–344, 2010.
- [6] K. Dennison, N. Stacey, and S. D’Amico.

- Autonomous asteroid characterization through nanosatellite swarming
IEEE Transactions on Aerospace and Electronic Systems, vol. 59, no. 4, pp. 4604–4624, 2023.
- [7] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit
Fastslam: A factored solution to the simultaneous localization and mapping problem
Proceedings of the National Conference on Artificial Intelligence, vol. 18, pp. 593–598, 2002.
- [8] E. Eade and T. Drummond
Scalable monocular slam
IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 1, pp. 469–476, 2006.
- [9] C. Cocard and T. Kubota
Slam-based navigation scheme for pinpoint landing on small celestial body
Advanced Robotics, vol. 26, no. 15, pp. 1747–1770, 2012.
- [10] R. Mahler
Multitarget bayes filtering via first-order multitarget moments
IEEE Transactions on Aerospace and Electronic Systems, vol. 39, no. 4, pp. 1152–1178, 2003.
- [11] J. S. McCabe and K. J. DeMars
Feature-based robotic mapping with generalized labeled multi-bernoulli filters for planetary landers
AIAA/AAS Astrodynamics Specialist Conference, 2016.
- [12] J. S. McCabe, K. J. DeMars, and C. Frueh
Integrated detection and tracking for multiple space objects
Advances in the Astronautical Sciences, vol. 155, pp. 1035–1052, 2015.
- [13] R. Mahler
Statistical Multisource-Multitarget Information Fusion. Boston, MA: Artech House, 2007.
- [14] B.-N. Vo and W.-K. Ma
The gaussian mixture probability hypothesis density filter
IEEE Transactions on Signal Processing, vol. 54, no. 11, pp. 4091–4104, 2006.
- [15] R. Mahler
A survey of phd filter and cphd filter implementations
Proceedings of SPIE, vol. 6567, no. 1, pp. 0–12, 2007.
- [16] L. Schlenker
Spacecraft relative navigation using random finite sets
Master's thesis, University of Minnesota, 2019. [Online]. Available: <http://hdl.handle.net/11299/206138>
- [17] L. Schlenker, M. Moretto, D. Gaylor, and R. Linares
Simultaneous localization and mapping for spacecraft rendezvous and proximity operations using random finite sets
Advances in the Astronautical Sciences, vol. 168, no. IV, pp. 3401–3420, 2019.
- [18] K. Leung, F. Inostroza, and M. Adams
Multifeature-based importance weighting for the phd slam filter
IEEE Transactions on Aerospace and Electronic Systems, vol. 52, no. 6, pp. 2697–2714, 2016.
- [19] C. S. Lee, D. Clark, and J. Salvi
Slam with dynamic targets via single-cluster phd filtering
IEEE Selected Topics in Signal Processing, vol. 7, no. 3, pp. 543–552, 2013.
- [20] T. Li, M. Bolic, and P. M. Djuric
Resampling methods for particle filtering
IEEE Signal Processing Magazine, vol. 5, pp. 70–86, 2015.
- [21] S. Thrun, W. Burgard, and D. Fox
Probabilistic Robotics. Boston, MA: MIT Press, 2005.
- [22] J. Mullane, B.-N. Vo, B.-T. Vo, and M. Adams
Mobile robotics in a random finite set framework
International Conference in Swarm Intelligence, vol. 6279, no. 2, pp. 519–528, 2011.
- [23] J. Christian
Fundamentals of vision based navigation
NASA Engineering and Safety Center, Tech. Rep., 2014.
- [24] M. Strube, R. Henry, E. Skelton, J. V. Eepoel, N. Gill, and R. McKenna
Raven: An on-orbit relative navigation demonstration using international space station visiting vehicles
Advances in the Astronautical Sciences, vol. 154, pp. 659–672, 2015.
- [25] M. Vavrina, E. Skelton, K. DeWeese, B. Naasz, D. Gaylor, and C. D'Souza
Safe rendezvous trajectory design for the restore-1 mission
Advances in the Astronautical Sciences, vol. 168, no. IV, pp. 3649–3668, 2019.