Efficient Cislunar Multi-Target Tracking with Adaptive Multi-Fidelity Propagation

Benjamin L. Reifler
The University of Texas at Austin
Brandon A. Jones
The University of Texas at Austin

ABSTRACT

Cislunar space presents new challenges for space situational awareness (SSA). Choatic dynamics make orbit uncertainty propagation difficult, requiring significant computation time for sufficient accuracy and often resulting in rapid growth in uncertainty. The size of the cislunar regime, lunar and solar exclusion zones, and regions of poor illumination all contribute to sparsity of observations. As the population of objects in cislunar space grows, so will the need for efficient multi-target tracking (MTT) algorithms that can mitigate these issues. In order to reduce the computational cost of track prediction while maintaining accuracy, this paper presents a method for adaptive bi-fidelity orbit uncertainty propagation. The proposed method is applied in simulated cislunar SSA scenarios and is shown to provide an improvement in computational efficiency compared to the original bi-fidelity method without compromising tracking accuracy.

1. INTRODUCTION

Increasing human activity in cislunar space presents new challenges for space situational awareness (SSA) [1]. In particular, chaotic dynamics make orbit uncertainty propagation difficult, requiring significant computation time for sufficient accuracy and often resulting in rapid growth in uncertainty [2, 3]. The size of the cislunar regime, lunar and solar exclusion zones, and regions of poor illumination all contribute to sparsity of observations [4, 5]. As the population of objects in cislunar space grows, so will the need for efficient multi-target tracking (MTT) algorithms that can mitigate these issues. The computational cost of a multi-target filter's prediction step is generally proportional to the cost of predicting a single object's probability density function (PDF) times the number of hypothesized tracks. Therefore, total filter runtime is increasingly dependent on the efficiency of orbit uncertainty propagation as the number of objects increases.

In order to reduce the computational cost of track prediction while maintaining accuracy, this paper presents a method for adaptive multi-fidelity orbit uncertainty propagation. The proposed method is applied to the ensemble Gaussian mixture filter (EnGMF), a particle-based single-target filter that has been shown to enable accurate space object tracking with sparse observations [6, 7, 8]. We use the EnGMF with adaptive multi-fidelity propagation in a GLMB filter (GLMBF), which is a labeled random finite set (RFS)-based multi-target filter [9, 10, 11], and assess the effect of our adaptive propagation method on tracking accuracy and computational cost in simulated cislunar SSA scenarios. This work builds on previous research into the application of (non-adaptive) multi-fidelity propagation to the cislunar regime [12].

Multi-fidelity orbit uncertainty propagation balances the accuracy of high-fidelity propagation and the speed of low-fidelity propagation to increase the computational efficiency of particle-based filters [13]. For example, bi-fidelity propagation works by propagating all particles with a relatively cheap low-fidelity model, identifying a small set of important particles, propagating these particles again with the more costly high-fidelity model, and correcting the low-fidelity particles to better match the distribution of the high-fidelity particles via stochastic collocation [14]. The accuracy of multi-fidelity propagation can also be increased by using a time history of states for each particle to construct the multi-fidelity surrogate. By selecting the correct low- and high-fidelity models for a given scenario, computation time may be minimized without significantly decreasing prediction accuracy.

Approved for public release; distribution is unlimited. Public Affairs release approval #AFRL-2024-4281.

The varied orbital dynamics and regimes over cislunar and near-lunar space introduce new challenges for multi-fidelity uncertainty propagation not present in the near-Earth case. Depending on an object's location in cislunar space, it is possible to select the optimal pair of dynamics models from a set of models to achieve a desired level of accuracy while minimizing computational cost. For example, if the object is close to the surface of the Moon, a detailed non-spherical gravity field model may be necessary, but if the object is farther from the surface, a point-mass model may be sufficient. If the object is far from both the Earth and Moon, the optimal low-fidelity model may be the circular restricted three-body problem (CRTBP) in a rotating reference frame or a model that includes point-mass gravity of the Earth and Moon.

In this paper, we apply our adaptive multi-fidelity propagation method to simulated SSA scenarios in which objects pass through a variety of orbital regimes with different dynamics. We use these scenarios to quantify the effect of the proposed method on multi-target filter runtime and tracking accuracy, with the goal of improving our ability to efficiently track the growing cislunar population.

2. BACKGROUND

2.1 Multi-Fidelity Orbit Uncertainty Propagation

Our approach to multi-fidelity orbit uncertainty propagation uses stochastic collocation to significantly reduce the computational cost of propagating a set of particles with only a slight reduction in accuracy [13]. This paper focuses on bi-fidelity propagation. To apply the bi-fidelity method, we define two dynamics models: a high-fidelity model suitable for state estimation and a low-fidelity model that is computationally cheaper to evaluate but less accurate. For example, for propagation in Earth orbit, the low-fidelity model may include only two-body dynamics and the J_2 perturbation, while the high-fidelity model includes drag, solar radiation pressure (SRP), and a more detailed model of Earth's gravity field. The set of particles representing a space object (SO)'s state-space PDF is propagated using the low fidelity model. When a high-fidelity estimate is required (e.g., for estimate extraction, to enable a measurement update, or for data association), a high-fidelity correction is performed. High-fidelity correction works by identifying a set of *important samples*, which are propagated again from the time of the last measurement update using the high-fidelity model. Then, stochastic collocation is used to adjust the low-fidelity samples to match the distribution of the high-fidelity samples [14].

The procedure for high-fidelity correction begins by concatenating the last M low-fidelity propagated states for each particle, resulting in the $(M+1)\,d \times N$ matrix $X^L(\Xi) = \left[x^L(\xi_1) \quad \dots \quad x^L(\xi_N)\right]$, where $d = 6$ is the dimension of the translational state space, N is the number of particles, $\Xi = \{\xi_i\}_{i=1}^N$ is the set of random inputs,

$$x^L(\xi_i) = \begin{bmatrix} x^L(\Delta t_{k-M}, \xi_i) \\ \vdots \\ x_k^L(\Delta t_k, \xi_i) \end{bmatrix}, \tag{1}$$

x^L is the low-fidelity state, and Δt_k is the difference between time t_k and the time of the last measurement update, or the time of track initialization if there has been no update yet. Our implementation propagates particles backward in time M timesteps using the low-fidelity model when needed, instead of storing their histories, to reduce memory usage. This makes high-fidelity correction more computationally expensive, but the size of the effect depends on the number of particles and the computational complexity of the low-fidelity model. The number of important samples is limited to $r \leq (M+1)\,d$, so including the M previous steps helps ensure that enough important samples may be used to achieve the desired accuracy. Each column of $X^L(\Xi)$ may be approximated using the surrogate

$$x^L(\xi_i) \approx \sum_{j=1}^{r} c_j(\xi_i)\, x^L\left(\bar{\xi}_j\right), \tag{2}$$

where $c_j(\xi_i)$ are a set of coefficients and $\bar{\xi}_j$ are the random inputs of the important samples. Next, $c_j(\xi_i)$ and $\bar{\xi}_j$ are found by solving the optimization problem

$$\bar{\Xi} = \arg\min_{\Xi} \inf_{y \in \operatorname{span}\left(X^L(\Xi)\right)} \left\| x^L(\xi) - y \right\|, \tag{3}$$

using a greedy algorithm based on the pivoted Cholesky decomposition [13, 15]. Finally, the high-fidelity values $x^H(\bar{\xi}_j)$ for the important samples are computed by repropagating with the high-fidelity model, and the corrected samples are given by

$$x^H(\xi_i) \approx \sum_{j=1}^{r} c_j(\xi_i) x^H(\bar{\xi}_j) \, . \tag{4}$$

2.2 The Ensemble Gaussian Mixture Filter

The EnGMF enables accurate but efficient nonlinear estimation. The filter parameterizes the estimated state PDF as a set of particles and uses kernel density estimation (KDE) to avoid particle depletion, allowing it to function with fewer particles than a particle filter (PF) [6, 7]. The EnGMF is initialized by drawing N random samples from the initial PDF. The PDF is predicted by propagating each particle forward in time. To update the PDF, first, we compute the sample covariance P of the particles. Then, Silverman's rule is used to compute the bandwidth parameter β_S, and the particles are converted to a Gaussian mixture model (GMM), with each particle becoming the mean of a GMM component with weight N^{-1} and covariance $B_S = \beta_S P$ [16, 7]. The GMM is then updated using an appropriate GMM-based filter and N new particles are sampled from the updated GMM. The KDE GMM is also used to extract the estimated mean and covariance, as opposed to using the sample mean and covariance directly.

In this paper, the GMM update is performed using the square root unscented Kalman filter (SR-UKF) update [17], to enable accurate nonlinear update with fewer particles. This is intended to reduce the cost of prediction and high-fidelity correction, but does increase the computational cost of the measurement update per particle. The EnGMF generally performs better for orbit determination when its particles' states are parameterized by equinoctial orbital elements, due to improved linearity when applying Silverman's rule [18, 7]. However, because equinoctial elements are designed to represent two-body orbits, here we instead use Cartesian coordinates.

2.3 The Generalized Labeled Multi-Bernoulli Filter

The GLMBF is a labeled RFS-based multi-target filter. An RFS may be thought of as a set of random vectors whose cardinality is also a random variable. Random finite sets enable Bayesian estimation of the states of unknown numbers of objects. The GLMBF is a closed-form solution to the Bayes mulit-target filter recursion [9, 10]. This paper uses the joint predict–update formulation of the GLMBF [11].

2.3.1 Notation

In this section, lowercase letters denote vectors, capital letters denote sets, blackboard bold letters represent spaces, and calligraphic letters denote sets of sets. Bold symbols indicate the use of label-augmented vectors. For variables that change over time, the subscript k to indicate time is not written and the subscript $k+1$ to indicate the following timestep is abbreviated to a subscript plus sign. The multi-object exponential is defined as $f^X = \prod_{x \in X} f(x)$, the Kronecker delta is defined as

$$\delta_Y[X] = \begin{cases} 1, & X = Y, \\ 0, & \text{otherwise}, \end{cases} \tag{5}$$

the indicator function is defined as

$$1_Y(X) = \begin{cases} 1, & X \subseteq Y, \\ 0, & \text{otherwise}, \end{cases} \tag{6}$$

and $1_Y(x) = 1_Y(\{x\})$. The function $\mathcal{F}$ denotes the set of all finite subsets of a space. A label-augmented RFS is defined on $\mathbb{X} \times \mathbb{L}$, where $\mathbb{X}$ is the state space and $\mathbb{L}$ is a discrete label space. This means that each element in a realization of a labeled RFS is of the form $\boldsymbol{x} = (x, l)$, where x is the state and l is the label. In this work, object labels are of the form $l = (i, k, j)$, where i is the index of the sensor that first detected the object, k is the timestep at which the object's initial track was created, and j is the object's unique index out of all objects first detected by sensor i at time k. A labled RFS realization may not contain duplicate labels, and this is enforced using the distinct label indicator

$$\Delta(\boldsymbol{X}) = \delta_{|\boldsymbol{X}|}[|\text{lab}(\boldsymbol{X})|] \, , \tag{7}$$

where $\text{lab}(x, l) = l$ denotes the projection of the label-augmented state space $\mathbb{X} \times \mathbb{L}$ onto its discrete label space $\mathbb{L}$ and $\text{lab}(\boldsymbol{X}) = \{\text{lab}(\boldsymbol{x}_1), \ldots, \text{lab}(\boldsymbol{x}_n)\}$.

2.3.2 Random Finite Set PDFs

The GLMBF is constructed using two families of labeled-RFS PDFs: the δ-generalized labeled multi-Bernoulli (GLMB) and labeled multi-Bernoulli (LMB) distributions [9].

A δ-GLMB RFS PDF may be parameterized by components $(I,\xi) \in \mathcal{F}(\mathbb{L}) \times \Xi$, where Ξ is a discrete space, with associated weights $w^{(I,\xi)}$. For MTT, each component typically represents a data association hypothesis, with I being a set of objects that may exist, ξ being their combined measurement association history, and $w^{(I,\xi)}$ being the estimated probability that the hypothesis is true. A δ-GLMB RFS density is of the form

$$\boldsymbol{\pi}(\boldsymbol{X}) = \Delta(\boldsymbol{X}) \sum_{(I,\xi)\in\mathcal{F}(\mathbb{L})\times\Xi} w^{(I,\xi)} \delta_I[\mathrm{lab}(\boldsymbol{X})] \left(p^{(\xi)}\right)^{\boldsymbol{X}} , \tag{8}$$

where $p^{(\xi)}(\cdot,l)$ is an object's state-space PDF given association history ξ and label l.

Instead of representing a mixture of multiple data association hypotheses, an LMB PDF represents a single hypothesis that contains a probability of existence and a state-space PDF for each object. An LMB PDF may be parameterized by $\left\{\left(r^{(l)}, p^{(l)}\right)\right\}_{l\in\mathbb{L}}$, where $r^{(l)}$ is the probability that object l exists and $p^{(l)}$ is its state-space PDF. The PDF of an LMB RFS has the form

$$\boldsymbol{\pi}(\boldsymbol{X}) = \Delta(\boldsymbol{X})\, w(\mathrm{lab}(\boldsymbol{X}))\, p^{\boldsymbol{X}} , \tag{9}$$

where $p(x,l) = p^{(l)}(x)$ and

$$w(I) = \prod_{l\in\mathbb{L}} \left(1 - r^{(l)}\right) \prod_{l\in I} \frac{1_{\mathbb{L}}(l)\, r^{(l)}}{1 - r^{(l)}} . \tag{10}$$

2.3.3 Joint Predict–Update Recursion

Given the initial filtering density in (8) at timestep k, the predicted and updated density at timestep $k+1$ is given by

$$\boldsymbol{\pi}_+(\boldsymbol{X}_+) \propto \Delta(\boldsymbol{X}_+) \sum_{I,\xi,I_+,\theta_+} w^{(I,\xi)} w_+^{(I,\xi,I_+,\theta_+)}(Z_+)\, \delta_{I_+}[\mathrm{lab}(\boldsymbol{X}_+)] \left(p_+^{(\xi,\theta_+)}(\cdot \mid Z_+)\right)^{\boldsymbol{X}_+} , \tag{11}$$

where $I \in \mathcal{F}(\mathbb{L})$, $\xi \in \Xi$, $I_+ \in \mathcal{F}(\mathbb{L}_+)$, and $\theta_+ \in \Theta_+$, where $\mathbb{L}_+ = \mathbb{L} \cup \mathbb{B}_+$, $\mathbb{B}_+$ is the space of object labels that could be born at this time, and Θ_+ is the set of maps $\theta_+ : \mathbb{L}_+ \to \{0 : |Z_+|\}$ assigning measurements in the current measurement set Z_+ to object labels, where $\theta_+(l) = 0$ implies that label l is not assigned a measurement. Additionally,

$$w_+^{(I,\xi,I_+,\theta_+)}(Z_+) = \left(r_{B,+}\right)^{\mathbb{B}_+\cap I_+} \left(1 - r_{B,+}\right)^{\mathbb{B}_+ - I_+} \left(\bar{p}_S^{(\xi)}\right)^{I\cap I_+} \left(1 - \bar{p}_S^{(\xi)}\right)^{I - I_+} \left(\bar{\psi}_+^{(\xi,\theta_+(\cdot))}(\cdot \mid Z_+)\right)^{I_+} , \tag{12}$$

$$p_+^{(\xi,\theta_+)}(x_+, l \mid Z_+) = \frac{\bar{p}_+^{(\xi)}(x_+, l)\, \psi_+^{(\theta_+(l))}(x_+, l \mid Z_+)}{\bar{\psi}_+^{(\xi,\theta_+(l))}(l \mid Z_+)} , \tag{13}$$

$$\bar{p}_+^{(\xi)}(x_+, l) = 1_{\mathbb{B}_+}(l)\, p_{B,+}(x_+, l) + 1_{\mathbb{L}}(l)\, \frac{\left\langle p_S(\cdot,l)\, f_{S,+}(x_+ \mid \cdot, l), p^{(\xi)}(\cdot,l)\right\rangle}{\bar{p}_S^{(\xi)}(l)} , \tag{14}$$

$$\bar{p}_S^{(\xi)}(l) = \left\langle p^{(\xi)}(\cdot,l), p_S(\cdot,l)\right\rangle , \tag{15}$$

$$\bar{\psi}_+^{(\xi,j)}(l \mid Z_+) = \left\langle \bar{p}_+^{(\xi)}(\cdot,l), \psi_+^{(j)}(\cdot, l \mid Z_+)\right\rangle , \tag{16}$$

$$\psi_+^{(j)}(x_+, l \mid Z_+) = \delta_0[j]\,(1 - p_D(x_+, l)) + (1 - \delta_0[j])\, \frac{p_D(x_+, l)\, g\left(z_{+,j} \,\middle|\, x_+, l\right)}{\kappa\left(z_{+,j}\right)} , \tag{17}$$

where $r_{B,+}(l)$ is the probability that object l is born, $p_{B,+}(\cdot,l)$ is the single-target PDF for the newborn object, p_S is the probability of survival from one step to the next, p_D is the probability of detection, $f_{S,+}(x_+ \mid \cdot, l)$ is the surviving object transition density, g is the measurement likelihood, and κ is the expected clutter intensity.

The number of possible hypotheses grows exponentially over time. To maintain computational tracatability, the set of new hypotheses (I,ξ,I_+,θ_+) in (11) resulting from prior hypothesis (I,ξ) is truncated using a ranked assignment algorithm and the following cost matrix:

$$\left[\begin{array}{ccc|ccc|ccc} c(l_1,-1) & \cdots & \infty & c(l_1,0) & \cdots & \infty & c(l_1,1) & \cdots & c(l_1,|Z_+|) \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \infty & \cdots & c(l_M,-1) & \infty & \cdots & c(l_M,0) & c(l_M,1) & \cdots & c(l_M,|Z_+|) \end{array}\right], \tag{18}$$

where $c(l,j) = -\log\Big(\eta^{(I,\xi)}(l,j)\Big)$ and

$$\eta^{(I,\xi)}(l,j) = \begin{cases} 1 - r_{B,+}^{(l)}, & j < 0 \wedge l \in \mathbb{B}_+, \\ r_{B,+}^{(l)}\tilde{\psi}_+^{(\xi,j)}(l \mid Z_+)\,, & j \geq 0 \wedge l \in \mathbb{B}_+, \\ 1 - \bar{p}_S^{(\xi)}(l)\,, & j < 0 \wedge l \in I, \\ \bar{p}_S^{(\xi)}(l)\tilde{\psi}_+^{(\xi,j)}(l \mid Z_+)\,, & j \geq 0 \wedge l \in I. \end{cases} \tag{19}$$

If an object is assigned to the left block of the cost matrix, it does not exist, if it is assigned to the center block, it exists but was not detected at time $k+1$, and if it is assigned to the right block, it exists and produced measurement $z_{+,j}$. The classic approach to solve this ranked assignment problem is to use Murty's algorithm [19, 9, 10], but for large cost matrices, a Gibbs sampler-based approach is more efficient [11].

3. ADAPTIVE MULTI-FIDELITY ORBIT UNCERTAINTY PROPAGATION

This paper proposes an improvement to multi-fidelity orbit uncertainty propagation that is designed to further reduce its computational cost without reducing its accuracy. Recall from Section 2.1 that the bi-fidelity method uses a pair of dynamics models, one low- and one high-fidelity, to provide a particular balance of computational cost and prediction accuracy. The proposed approach takes advantage of the fact that different orbital regimes may have different optimal pairs of dynamics models, depending on the desired level of accuracy. For example, accurate prediction in the near-Earth regime typically requires accounting for acceleration due to non-spherical gravity, at least in the high-fidelity model, while these terms may be neglected when an object is farther away from the Earth. The same holds for the near-Moon regime.

Given a set of dynamics models $F = \{f_1, f_2, \ldots\}$, we seek a map $m : \mathbb{X} \to F \times F$ such that $\big(f^L, f^H\big) = m(\hat{\boldsymbol{x}})$, where $\hat{\boldsymbol{x}}$ is an object's mean state estimate at time k and the functions f^L and f^H are the low- and high-fidelity models, respectively, that solve the following optimization problem:

$$\min_{(f^L,f^H)\in F\times F} t_R\Big(f^L, f^H\Big) \quad \text{subject to} \quad D\Big(p_+ \,\Big\|\, \tilde{p}_+\Big(f^L, f^H\Big)\Big) < \varepsilon_D\,, \tag{20}$$

where $t_R\big(f^L, f^H\big)$ is the real-time computational cost of choosing model pair $\big(f^L, f^H\big)$ (including both prediction and correction), D is a divergence metric, p_+ is the object's state-space PDF predicted from k to $k+1$ using a full-fidelity dynamics model, $\tilde{p}_+\big(f^L, f^H\big)$ is the object's PDF predicted using f^L and corrected using f^H, and ε_D is a user-defined value. In order to achieve the desired improvement in computational efficiency, it is obviously impractical to solve the optimization problem online. Instead, simulations should be performed in advance to compute an approximation of the map m given F. Note that if $f^L = f^H$ over the time interval between two measurement updates, there is no need to perform high-fidelity correction, which may result in significant time savings.

As a proof-of-concept, the implementation used to generate the results in the following section uses a simple definition for m, comprising three position-based regions: one near the Earth, one near the Moon, and one not near either body. A position in space is defined as near a body, in this case, if it is inside a sphere of some set radius centered on the body. The precise radii used in this work were found by a coarse search for the smallest multiple of R which resulted in no significant decrease in tracking accuracy compared to the original bi-fidelity method, where R is the radius of either body. Maintaining the same average tracking accuracy was used as a proxy for satisfying the divergence constraint ε_D in (20).

4. SIMULATION

This section describes a multi-target tracking simulation used to assess the tracking error and time cost of the proposed adaptive orbit propagation algorithm.

4.1 Scenario Description

The simulation considers four small clusters of SOs, each containing five objects. Cluster 1 is in a distant retrograde orbit (DRO), Cluster 2 is in a near-rectilinear halo orbit (NRHO), Cluster 3 is in a lunar transfer orbit (LTO), and Cluster 4 is heading for a low flyby of the Moon. The objects in each cluster all have the same initial state PDF, with the respective mean given in Table 1 and diagonal covariance $10^{-8} I_6$ in the normalized units (LU^2, LU^2/TU^2) in the rotating barycentric frame. Figure 1 shows the trajectories of the mean objects for each cluster over the ten-day simulation. The nominal trajectory for Cluster 4, the flyby, reaches a minimum altitude of 133 km, assuming a spherical Moon, but due to random initialization, the actual flyby altitude may be lower or higher for each object.

Table 1: Initial means for the objects in each cluster, expressed in normalized units (LU, LU/TU) in the rotating barycentric frame.

Cluster	Mean					
	x	y	z	$\dot{x}$	$\dot{y}$	$\dot{z}$
1	0.806	0.000	0.000	0.000	0.519	0.000
2	1.022	0.000	−0.182	0.000	−0.103	0.000
3	−0.112	0.000	0.000	2.194	−3.440	0.000
4	0.949	−0.019	0.304	−0.006	0.064	0.003

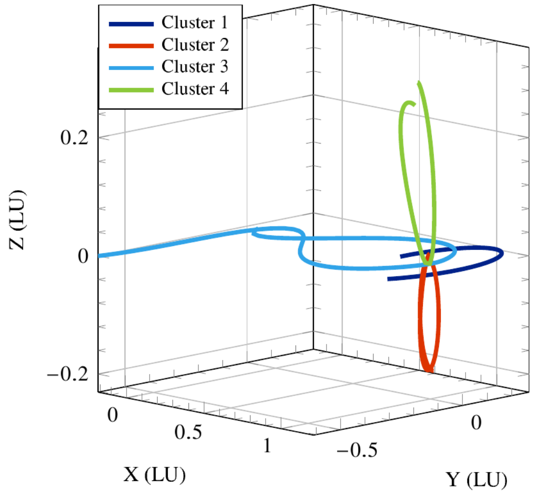

Fig. 1: Plot of mean trajectories for each cluster in the rotating barycentric frame.

Single target prediction and update is performed via the EnGMF, with a constant number of particles $N = 200$. Tracking information is provided to the filter by a single sensor located at the center of the Moon, which records Moon-centered inertial (MCI) measurements of right ascension and declination, and their respective rates. The sensor produces measurements every 12 h after the beginning of the simulation, with probability of detection $P_D = 0.95$. The angle and angle rate measurements have constant standard deviations 0.1 arcsec and 0.001 arcsec/s, respectively. Clutter is not modeled, and the multi-target filter assumes a constant clutter intensity $\kappa = 10^{-9}$. The filter also assumes probability of survival near one for all objects, does not model spontaneous birth, prunes hypotheses with weights less than 10^{-3}, and truncates its hypothesis list via Murty's algorithm to a maximum of 50 hypotheses. This limit on the number of updated hypotheses is allocated to each prior GLMB component in proportion to the square root of its weight.

The true dynamics model includes the point-mass gravity of the Earth, Moon, and Sun, and a 70 × 70 spherical harmonics models for the Earth and Moon's non-spherical gravity (EGM2008 and LP150Q) [20, 21]. Non-spherical gravity terms are only computed for whichever body is closest at the beginning of the propagation timespan. For bi-fidelity propagation, the high-fidelity model is the same as the truth model but with only a 10×10 gravity field model. The low-fidelity model includes only the point-mass gravity of the Earth, Moon, and Sun. Note that a 10×10 gravity field may not be sufficient for accurate propagation in the low lunar orbit (LLO) regime [22, p. 150].

For adaptive bi-fidelity propagation, the near-Earth regime is defined as a sphere around the Earth with a radius of 10 Earth radii. The near-Moon regime is defined as a sphere around the Moon with a radius of 10 Moon radii. The low-fidelity model described in the previous paragraph is used in all three regions. When near either body, the high-fidelity model includes the point-mass gravity of the Earth, Moon, and Sun, and a 10×10 non-spherical gravity field for the near body. When not near either body, the high-fidelity model is identical to the low-fidelity model.

The simulation timestep is 1 h and all models are integrated using an RK45 variable-step integrator. With either bi-fidelity propagation method, high-fidelity correction is performed only before each multi-target measurement update.

4.2 Results

Five scenarios are constructed using the four SO clusters described in the previous section: a scenario with all 20 objects, a scenario only considering the five objects in Cluster 1, a scenario only considering Cluster 2, and so on. Each set of results reflects the average of 100 Monte Carlo runs with each of the four propagation methods (low-fidelity, high-fidelity, bi-fidelity, and adaptive bi-fidelity). The observations, initial SO states, and initial EnGMF particles are randomized for each set of four runs.

The simulation is implemented in C++ with the maximum safe level of optimization enabled. Each run is performed on a single thread. The simulation runs are performed on a computer with six 3.2 GHz CPUs. The computer can support a total of 12 concurrent threads, but only a maximum of four runs are performed simultaneously, to avoid running out of memory or biasing timing results by overtaxing the system.

Figure 2 shows the average optimal subpattern assignment (OSPA) tracking error with each propagation method [23], with cutoff 100 km and order 2, for the scenario with all objects. The low-fidelity results are significantly worse than the other three methods. This is partially due to the fact that some objects' state-space PDF estimates tend to diverge with only low-fidelity propagation. The four methods have average OSPA errors 19.08 km, 15.59 km, 15.93 km, and 15.69 km respectively. This shows that, in this case, neither original bi-fidelity propagation nor our new adaptive bi-fidelity propagation yields a significant decrease in tracking accuracy.

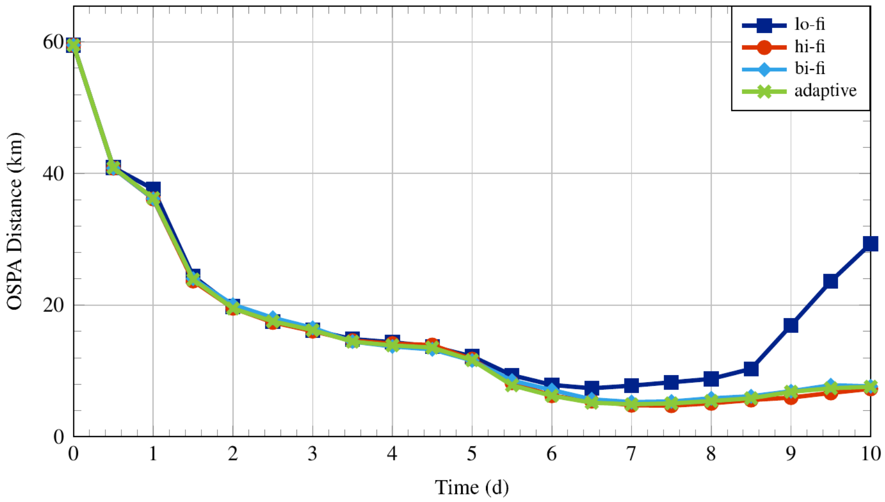

Fig. 2: Average OSPA tracking error with each propagation method, computed after each measurement update.

Figure 3 shows the computational cost of the multi-target filter at each simulation timestep. As expected, these results show that both the original and adaptive bi-fidelity methods roughly match the computational cost of low-fidelity propagation when only predicting the multi-target density. When a measurement update occurs, the cost of filtering increases for all four methods, but we can see that the adaptive method results in smaller spikes than the original bi-fidelity method. The biggest advantage of bi-fidelity propagation (original or adaptive) is observed around $t = 8.5\,\mathrm{d}$, which corresponds to the LTO objects (Cluster 3) returning to the near-Earth regime. This causes a wide spike in high-fidelity cost, but a narrower and shorter spike in the bi-fidelity costs. The average total filter runtimes for the four methods are 25.72 s, 115.05 s, 34.97 s, 30.94 s. This indicates that adaptive bi-fidelity propagation reduces filtering time by an average of 12% compared to the original bi-fidelity method in this case.

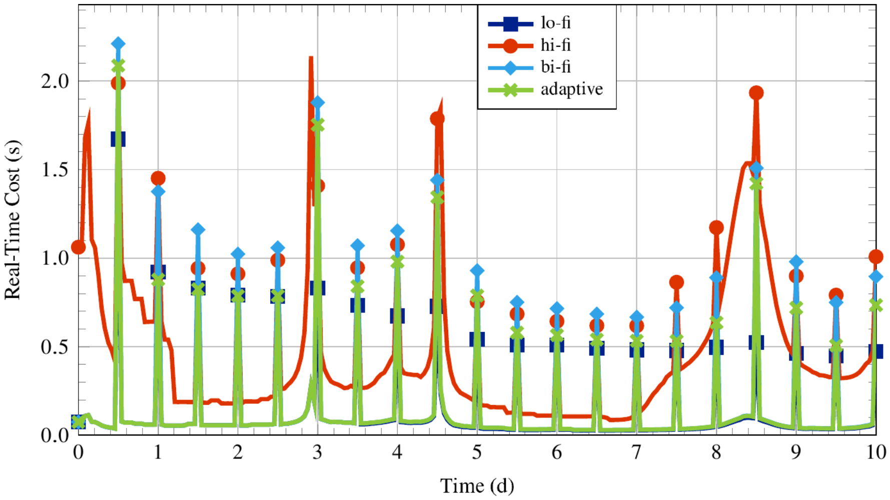

Fig. 3: Average runtime cost of multi-target filter with each propagation method. Markers correspond to measurement updates.

Tables 2 and 3 summarize the results for the other four scenarios, each only considering the objects in a single cluster. These results show that the adaptive bi-fidelity propagation causes a negligible change in tracking error compared to the original bi-fidelity method. Furthermore, the new propagation method reduces total runtime cost by 11–17% across the five scenarios, with the greatest improvement in Cluster 1. This is because the DRO never passes close enough to either body to require high-fidelity correction, meaning that adaptive bi-fidelity propagation has the same computational cost as low-fidelity propagation in this case.

Table 2: Average OSPA tracking error (km) for each propagation method and scenario. "Change" is relative increase from original bi-fidelity to adaptive bi-fidelity.

Method	All Clusters	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Lo-Fi	19.08	15.45	11.36	15.23	17.04
Hi-Fi	15.59	15.75	8.74	9.52	14.69
Bi-Fi	15.93	15.66	8.96	9.51	14.72
Adaptive	15.69	15.26	9.12	9.67	15.13
Change	−2%	−3%	2%	2%	3%

Table 3: Average total runtime cost (s) of multi-target filter for each propagation method and scenario. "Change" is relative decrease from original bi-fidelity to adaptive bi-fidelity.

Method	All Clusters	Cluster 1	Cluster 2	Cluster 3	Cluster 4
Lo-Fi	25.72	6.25	9.26	14.77	8.28
Hi-Fi	115.05	12.61	43.25	121.97	28.97
Bi-Fi	34.97	7.47	15.89	21.61	11.14
Adaptive	30.94	6.22	14.06	18.12	9.87
Change	12%	17%	12%	16%	11%

5. CONCLUSION AND FUTURE WORK

We have presented an adaptive bi-fidelity method for orbit uncertainty propagation. The method uses a predetermined map to select the optimal pair of low- and high-fidelity dynamics models for bi-fidelity propagation based on an object's estimated state. We have shown via numerical simulation that the method provides an improvement in computational efficiency compared to the original bi-fidelity method without compromising tracking accuracy. The greatest improvement has been observed in orbit families that spend significant periods of time not close to either the Earth or Moon. Fortunately, this includes popular families for cislunar mission planning, including distant retrograde orbits (DROs), near-rectilinear halo orbits (NRHOs), and lunar transfer orbits (LTOs).

As a proof-of-concept for the proposed method, these results provide motivation for the development of a more rigorous map for cislunar orbit uncertainty propagation. This map could draw on a more varied set of models, such as the circular restricted three-body problem (CRTBP) in the rotating frame and models including additional forces such as solar radiation pressure (SRP). Another possible area for further development of the adaptive bi-fidelity method may be to replace estimate-based model selection with per-particle model selection. This could result in a partition of the particle set based on the model selection map. Particle partitioning may be necessary to use the proposed method in scenarios with much larger gaps between detections, where a single object's state-space probability density function (PDF) may have significant mass in multiple regimes.

6. ACKNOWLEDGMENT

This material is based on research sponsored by the Air Force Research Laboratory (AFRL) under agreement number FA9453-21-2-0064. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of AFRL and/or the U.S. Government. AFRL Public Release Case Number: AFRL-2024-4281.

7. REFERENCES

[1] M. J. Holzinger, C. C. Chow, and P. Garretson. *A Primer on Cislunar Space*. Air Force Research Laboratory, 2021.

[2] Michael R. Thompson, Nathan P. Ré, Cameron Meek, and Bradley Cheetham. Cislunar orbit determination and tracking via simulated space-based measurements. *Proceedings of the 2021 Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2021.

[3] Sam Wishnek, Marcus J. Holzinger, and Patrick Handley. Robust cislunar initial orbit determination. *Proceedings of the 2021 Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2021.

[4] Mark Bolden, Timothy Craychee, and Erin Griggs. An evaluation of observing constellation orbit stability, low signal-to-noise, and the too-short-arc challenges in the cislunar domain. *Proceedings of the 2020 Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2020.

[5] Surabhi Bhadauria and Carolin Frueh. Optical observation regions in cislunar space using the bi-circular restricted four body problem geometry. *Proceedings of the 2022 Advanced Maui Optical and Space Surveillance Technologies Conference*, September 2022.

[6] Jeffrey L. Anderson and Stephen L. Anderson. A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. *Monthly Weather Review*, 127(12):2741–2758, December 1999.

[7] Sehyun Yun, Renato Zanetti, and Brandon A. Jones. Kernel-based ensemble Gaussian mixture filtering for orbit determination with sparse data. *Advances in Space Research*, 69(12):4179–4197, April 2022.

[8] Benjamin L. Reifler, Andrey A. Popov, Brandon A. Jones, and Renato Zanetti. Large-scale space object tracking in a proliferated LEO scenario. *Proceedings of the 26th International Conference on Information Fusion*, June 2023.

[9] Ba-Tuong Vo and Ba-Ngu Vo. Labeled random finite sets and multi-object conjugate priors. *IEEE Transactions on Signal Processing*, 61(13), July 2013.

[10] Ba-Ngu Vo, Ba-Tuong Vo, and Dinh Phung. Labeled random finite sets and the Bayes multi-target tracking filter. *IEEE Transactions on Signal Processing*, 62(24), December 2014.

[11] Ba-Ngu Vo, Ba-Tuong Vo, and Hung Gia Hoang. An efficient implementation of the generalized labeled multi-Bernoulli filter. *IEEE Transactions on Signal Processing*, 65(8), April 2017.

[12] Trevor N. Wolf, Enrico M. Zucchelli, and Brandon A. Jones. Multi-fidelity uncertainty propagation for objects in cislunar space. *Proceedings of the 2022 AIAA SciTech Forum*, January 2022.

[13] Brandon A. Jones and Ryan Weisman. Multi-fidelity orbit uncertainty propagation. *Acta Astronautica*, 155:406–417, October 2019.

[14] Dongbin Xiu. Stochastic collocation methods: A survey. In Roger G. Ghanem, David Higdon, and Houman Owhadi, editors, *Handbook of Uncertainty Quantification*, chapter 20, pages 699–716. Springer International Publishing, 2017.

[15] Xueyu Zhu, Akil Narayan, and Dongbin Xiu. Computational aspects of stochastic collocation with multifidelity models. *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):444–463, 2014.

[16] Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. Routledge, 1998.

[17] Rudolph van der Merwe and Eric A. Wan. The square-root unscented Kalman filter for state and parameter-estimation. *Proceedings of the 2001 IEEE International Conference on Acoustics, Speech, and Signal Processing*, May 2001.

[18] R. A. Broucke and P. J. Cefola. On the equinoctial orbital elements. *Celestial Mechanics*, 5:303–310, 1972.

[19] Katta G. Murty. An algorithm for ranking all the assignments in order of increasing cost. *Operations Research*, 16(3):682–687, 1968.

[20] K. Pavlis, Nikolaos, Simon A. Holmes, Steve C. Kenyon, and John K. Factor. The development and evaluation of the Earth Gravitational Model 2008 (EGM2008). *Journal of Geophysical Research*, 117(B04406), April 2012.

[21] A. S. Konopliv, S. W. Asmar, E. Carranza, W. L. Sjogren, and D. N. Yuan. Recent gravity models as a result of the Lunar Prospector mission. *Icarus*, 150:1–18, 2001.

[22] Sean M. McArdle. *Gravity Modeling for Lunar Orbits*. PhD thesis, The University of Texas at Austin, 2022.

[23] Dominic Schuhmacher, Ba-Tuong Vo, and Ba-Ngu Vo. A consistent metric for performance evaluation of multi-object filters. *IEEE Transactions on Signal Processing*, 56(8):3447–3457, August 2008.