

Large-Scale Space Object Tracking in a Proliferated LEO Scenario

Benjamin L. Reifler

*Department of Aerospace Engineering
and Engineering Mechanics
The University of Texas at Austin
Austin, TX, USA
benjamin.reifler@utexas.edu*

Andrey A. Popov

*Oden Institute for Computational
Engineering & Sciences
The University of Texas at Austin
Austin, TX, USA
andrey.a.popov@utexas.edu*

Brandon A. Jones

*Department of Aerospace Engineering
and Engineering Mechanics
The University of Texas at Austin
Austin, TX, USA
brandon.jones@utexas.edu*

Renato Zanetti

*Department of Aerospace Engineering
and Engineering Mechanics
The University of Texas at Austin
Austin, TX, USA
renato@utexas.edu*

Abstract— The proliferation of large satellite constellations in low Earth orbit (LEO) is dramatically increasing demand on existing systems for space domain awareness. The rapidly growing number of objects in LEO will reduce the average rate of observations per object, necessitating the development of multi-target algorithms that can handle higher levels of data sparsity without sacrificing computational efficiency. In this paper, we demonstrate that a multi-target filter that combines a number of useful capabilities is able to track and maintain custody of a simulated population of over 16,000 LEO objects without requiring the use of high-performance computing facilities. The filter is based on the generalized labeled multi-Bernoulli filter, and includes three previously-presented features: label space partitioning based on sensor fields of view, the ensemble Gaussian mixture filter (EnGMF), and bi-fidelity orbit uncertainty propagation. We also introduce a new, algorithmically simple method for adjusting the number of EnGMF particles to balance accuracy and computational efficiency, which we refer to as *progressive resampling*.

Index Terms—Space domain awareness, generalized labeled multi-Bernoulli, ensemble Gaussian mixture, multi-target, multi-sensor, multi-fidelity

I. INTRODUCTION

The proliferation of large satellite constellations in low Earth orbit (LEO) is dramatically increasing demand on existing systems for space domain awareness (SDA). Without the addition of new sensors, the rapidly growing number of space objects (SOs) in LEO will reduce the average rate of observations per SO, necessitating the development of multi-target tracking (MTT) algorithms that are specialized to the SDA problem and can handle higher levels of data sparsity without sacrificing computational efficiency. In this paper, a multi-target filter that combines a number of useful capabilities is used to track

a simulated population of LEO SOs, of which over 15,000 are based on real two-line elements (TLEs) and the remainder constitute a new large constellation comprising 1000 satellites. These SOs are tracked using a simulated heterogeneous network of ground-based sensors with limited fields of view (FOVs) based on the existing space surveillance network. The multi-target filter used in this work is based on the GLMB filter (GLMBF) [1]–[3], augmented with four features: label space partitioning based on sensor FOVs [4], the ensemble Gaussian mixture filter (EnGMF) [5], [6], bi-fidelity orbit uncertainty propagation [7], and a new, algorithmically simple method for adjusting the number of EnGMF particles to balance accuracy and computational efficiency.

The GLMBF is a closed-form solution to the Bayes multi-target filter recursion based on the family of δ -generalized labeled multi-Bernoulli (GLMB) random finite set (RFS) probability density functions (PDFs) [1]–[3]. The GLMBF assumes with Poisson distributed clutter and Bernoulli processes for birth and death. The GLMBF is a multiple hypothesis tracker, which enables it to preserve information that would be lost in a single hypothesis tracker, such as the probability hypothesis density (PHD) or labeled multi-Bernoulli (LMB) filter [8]–[10]. However, when used with a sensor network with a limited FOV, such that many objects are undetectable for most timesteps, multiple hypothesis trackers can experience runaway growth in the number of hypotheses. For the GLMBF, this can be mitigated by partitioning the label space based on each object’s estimated probability of being in each sensor’s FOV. The resulting filter is the FOV-partitioned GLMB filter (FP-GLMBF), which has been shown to enable tractable multiple hypothesis MTT with limited-FOV sensor networks [4], [11], [12].

The EnGMF is a single-target filter that enables accurate uncertainty propagation by parameterizing an object’s state PDF using a set of particles, but avoids particle depletion during measurement update by converting the particles into

This work was sponsored in part by DARPA under STTR contract number 140D0420C0062. The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

a conservative Gaussian mixture model (GMM) via kernel density estimation (KDE). It has been shown that the EnGMF performs well for tracking SOs over multiple orbital periods, especially when the state is represented by the equinoctial orbital elements [6], [13]. As an object's PDF is predicted over time, its particles tend to spread out due to perturbations, resulting in decreased accuracy, which negatively affects the FP-GLMBF's partition update and the EnGMF's measurement update. In this work, we mitigate this issue by adding new particles at a fixed rate per revolution, through a process we call *progressive resampling*.

Multi-fidelity orbit uncertainty propagation provides high-accuracy state PDF prediction without requiring each particle in the ensemble to be propagated with the full high-fidelity dynamics model, improving computational efficiency [7]. This is important when tracking many SOs. A set of particle state histories are computed using a low-fidelity dynamics model, then stochastic collocation [14] is used to adjust the final state of each particle to more closely match the high-fidelity model, based on a relatively small number of particles that are repropagated with the high-fidelity dynamics model.

This paper builds on work we have presented in previous papers [4], [15], [16]. In [15], we applied the FP-GLMBF and EnGMF to a population of 8000 SOs, with a fixed number of particles per track and without multi-fidelity propagation. In [16], an unpartitioned GLMBF and EnGMF were applied to a small cluster of objects in geostationary orbit (GEO), with bi-fidelity propagation and a particle adaptation method based on convergence assessment. In [4], which is currently under review, we provide a more detailed derivation of the FP-GLMBF, including a new measurement-driven birth (MDB) model, and assess the filter's performance in scenarios that include spontaneous birth, death, and spawning. In this paper, we apply the FP-GLMBF and EnGMF to a population of 16,297 LEO SOs, with bi-fidelity propagation and a new method for adding particles over time. Unlike in [16], the objects in this scenario receive asynchronous measurement updates due to FOV limitations.

II. ALGORITHM DESCRIPTION

This section describes the multi-target filter system from bottom to top, beginning with the particle propagation method, followed by the single target filter, resampling method, and high-level multi-target filter.

A. Multi-Fidelity Orbit Propagation

In order to balance computational efficiency and propagation accuracy, we use a bi-fidelity method based on stochastic collocation. For PDF prediction, we will propagate N particles. In the multi-fidelity approach, these particles are propagated from one simulation step to the next via a cheap low-fidelity dynamics model. When we are going to extract the estimated PDF or perform measurement gating or update, a high-fidelity correction is performed by identifying a set of *important samples*, which are repropagated using a more expensive high-fidelity model.

This correction is performed via stochastic collocation [7], [14]: First, the initial state of each particle and its low-fidelity state at the last M steps are concatenated, resulting in the $(M + 1)d \times N$ matrix $X^L(\Xi) = [x^L(\xi_1) \dots x^L(\xi_N)]$, where $\Xi = \{\xi_i\}_{i=1}^N$ is the set of random inputs, $d = 6$ is the dimension of the state space,

$$x^L(\xi_i) = \begin{bmatrix} x^L(0, \xi_i) \\ x^L(\Delta t_{k-M}, \xi_i) \\ \vdots \\ x^L_k(\Delta t_k, \xi_i) \end{bmatrix}, \quad (1)$$

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. The number of important samples is limited to $r \leq (M + 1)d$. Each column of $X^L(\Xi)$ may be approximated using the surrogate

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

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

$$\bar{\Xi} = \arg \min_{\Xi} \inf_{y \in \text{span}(X^L(\Xi))} \|x^L(\xi) - y\|_{\infty}, \quad (3)$$

which is achieved using a greedy algorithm based on the pivoted Cholesky decomposition [7], [17]. 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

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

Our low-fidelity model solves the orbital equations of motion for the point mass gravity of the Earth and its J_2 perturbation [18] via the fourth-order Runge–Kutta (RK4) method with a fixed step size of 10 s. The J_2 perturbation is necessary in the low-fidelity model for this scenario because its effect is significant in LEO. We also tested an analytical low-fidelity model based on Vinti theory [19], which again included point mass and J_2 gravity, but we found that it was slower in this scenario than the RK4-based model. We hypothesize that this may be due to the majority of low-fidelity propagator calls in this scenario being short arcs of 60 s.

B. Ensemble Gaussian Mixture Filtering

The EnGMF provides the accuracy of a particle filter (PF) while avoiding the problem of particle depletion, enabling it to function with a much smaller number of particles than a PF, reducing overall computational complexity. The EnGMF is parameterized by a set of particles, which are initialized by drawing N random samples from an initial PDF. In this work, the particles are defined in equinoctial orbital elements instead of the typical Cartesian state space to keep the dynamics nearly linear even in the presence of perturbations [6], [13]. The first

equinoctial element is typically the semi-major axis (SMA) of the object's orbit, but we replace it with the mean motion $n = \sqrt{\mu_\oplus/a^3}$, where μ_\oplus is the standard gravitational parameter of the Earth and a is the SMA, in order to make the problem more linear.

The PDF is predicted in time by simply predicting each particle. When performing a measurement update or extracting the PDF's estimated mean and covariance, the particles are combined into a GMM via KDE: each particle becomes the mean of a GMM component with weight N^{-1} and conservative covariance

$$B_S = \beta_S P, \quad (5)$$

where β_S is a bandwidth parameter found using Silverman's rule and P is the sample covariance of the N particles [6], [20].

In order to incorporate process noise, when the GMM is constructed, a linearized process noise model is used to inflate the sample covariance before applying Silverman's rule. In Cartesian coordinates, the process noise transition matrix is

$$\Gamma(\Delta t) = \begin{bmatrix} \Delta t I_3 \\ \frac{1}{2} \Delta t^2 I_3 \end{bmatrix}, \quad (6)$$

where Δt is the time since the last measurement update, or since track initialization if there has been no update, and I_3 denotes the 3×3 identity matrix. Each particle is transformed from equinoctial elements to Cartesian coordinates in the Earth-centered inertial (ECI) reference frame, then a random process noise sample with covariance $\Gamma(\Delta t) Q \Gamma^T(\Delta t)$, where Q is the process noise covariance, is added, and the particle is transformed back to equinoctial elements. Note that the particles with noise added are only used to compute the sample covariance, and are not stored.

After adding noise to all particles, we compute the sample covariance and apply Silverman's rule to conservatively construct a GMM from the set of particles. Each particle is used as the mean of a GMM component with weight N^{-1} and a covariance of

$$B_S = h_S \left(\frac{4}{d+2} \right)^{\frac{2}{d+4}} N^{-\frac{2}{d+4}} \tilde{P}, \quad (7)$$

where h_S is a tunable parameter and \tilde{P} is the sample covariance including process noise. For this work, we set $h_S = 1$, which results in the optimal bandwidth matrix B_S (in terms of the mean integrated squared error (MISE)) when the sampling distribution is Gaussian [20].

In this work, the resulting GMM is updated via the square-root unscented Kalman filter (SR-UKF) update [21], which is more accurate for nonlinear measurement models than the linearization-based extended Kalman filter (EKF) update, but comes at a higher computational cost. Because FOV limitations mean that measurement updates are relatively rare in this scenario, the cost of measurement updates was not considered a significant factor in achieving the desired performance. To apply the SR-UKF update to a GMM, the single-Gaussian

SR-UKF update is applied to each component individually. The component weights are updated by

$$w_i^+ = \frac{q_i w_i^-}{\sum_{j=1}^N q_j w_j^-}, \quad (8)$$

where w_i^- is the prior weight of component i , w_i^+ is its posterior weight, and q_i is its measurement likelihood. The sum in the denominator of (8) is the overall measurement likelihood, which is needed for multi-target data association.

After each measurement update, a new set of particles are sampled from the posterior PDF. Because of the relatively small number of particles used in this work, outliers in resampling could have negative effects on accuracy. Therefore, rejection sampling is used to ensure that no particle is more than 3σ from the GMM component from which it was sampled.

C. Progressive Resampling

In order to balance the need for sufficient particle density with computational efficiency, we add new particles to each EnGMF PDF over time. In this approach, each PDF is initialized with $N = N_{\min}$ particles, and the corresponding GMM in equinoctial elements is computed and stored. As the PDF is predicted, new particles are drawn from the stored prior at some rate $\frac{dN}{d\tau}$, where τ is a time-like variable, until $N = N_{\max}$. Each new particle is predicted to each time in the bi-fidelity time history up through the current time, so the state history is always fully populated for high-fidelity correction. After each measurement update, the number of particles is reset to N_{\min} during resampling and the prior GMM is replaced. To save memory, the prior GMM is replaced with the version parameterized by N_{\min} particles, instead of the GMM produced by the SR-UKF update, which may be much larger. The values of N_{\min} , N_{\max} , and $\frac{dN}{d\tau}$ are all tunable parameters.

Fig. 1 shows an illustration of progressive resampling with a one-dimensional PDF. Particles are added from top to bottom in this example to aid explanation, while the actual method uses random sampling. The particle addition rate is such that one particle is added for each prediction step. At time t_0 , the initial PDF is parameterized by two particles (the two filled circles). Without progressive resampling, these two particles become an increasingly poor parameterization of the PDF over time, with the final PDF at t_3 lacking its lower mode entirely. When the PDF is predicted from t_0 to t_1 , a third particle is sampled from the PDF at t_0 (the top hollow circle). Then, all three particles are predicted to t_1 . This is repeated for the prediction steps from t_1 to t_2 and t_2 to t_3 , resulting in a better distribution of particles at each step.

In this work, the time-like variable τ is based on the number of revolutions (orbits) since the last full resampling (i.e., initialization or the last measurement update). Because the orbital period of an SO changes over time due to perturbations, τ is accumulated from timestep k to $k+1$ as follows:

$$\tau_{k+1} = \tau_k + (t_{k+1} - t_k) \frac{\hat{n}_k}{2\pi}, \quad (9)$$

where \hat{n}_k is the SO's estimated mean motion at time t_k . Because particles can only be added in integer amounts, when $\tau \frac{dN}{d\tau}$ is




Fig. 1. Illustration of progressive resampling in a nonlinear system.

greater than some positive integer ΔN , the number of particles is increased to $N + \Delta N$ and τ is decremented by $\Delta N \frac{d\tau}{dN}$.

D. FOV-Based Label-Partitioned GLMB Filtering

This section provides a brief overview of RFSs, then describes the FP-GLMBF algorithm as it is used in this work. The full algorithm and its derivation can be found in [4]. In this section, lower case 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. The subscript k to indicate the current time is omitted and the subscript $k + 1$ is abbreviated to a subscript plus sign. Let $\mathcal{F}(\mathbb{X})$ denote the set of all finite subsets of \mathbb{X} . We also define the multi-object exponential $f^X = \prod_{x \in X} f(x)$ and Kronecker delta

$$\delta_Y[X] = \begin{cases} 1, & X = Y, \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

Given some vector space \mathbb{X} and discrete label space \mathbb{L} , an RFS is a random variable on $\mathcal{F}(\mathbb{X})$ and a labeled RFS is a random variable on $\mathcal{F}(\mathbb{X} \times \mathbb{L})$ [22]. Given $\mathbf{x} = (x, l) \in \mathbb{X} \times \mathbb{L}$, let $\text{lab } \mathbf{x} = l$ denote the projection of the label-augmented state space $\mathbb{X} \times \mathbb{L}$ onto its label space \mathbb{L} . We can now define the distinct label indicator $\Delta(\mathbf{X}) = \delta_{|\mathbf{X}|} [|\text{lab } \mathbf{X}|]$, which will be used ensure that a labeled RFS realization will not contain duplicate labels.

The GLMBF and FP-GLMBF are constructed using RFS PDFs from the δ -GLMB and LMB families. However, LMB PDFs are used to represent spontaneous birth and spawning, neither of which appear in our scenario, so they are not relevant to this paper. A δ -GLMB PDF has the form

$$\boldsymbol{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{(I, \xi) \in \mathcal{F}(\mathbb{L}) \times \Xi} w^{(I, \xi)} \delta_I[\text{lab } \mathbf{X}] \left(p^{(\xi)} \right)^{\mathbf{X}}, \quad (11)$$

where Ξ is a discrete space, (I, ξ) is a component of the PDF, $w^{(I, \xi)}$ is its associated weight, and $p^{(\xi)}(\cdot, l)$ is the state-space PDF of an object with label l given index ξ . In MTT, each component is a data association hypothesis, with I being the set of objects that exist, ξ being the combined data association history, $p^{(\xi)}$ being the tracks for the existing objects, and $w^{(I, \xi)}$ being the estimated probability that the hypothesis is true.

To construct the partition of the label space \mathbb{L} at timestep $k + 1$, we first define the instantaneous FOV of each of our S

sensors as a subset of the state space \mathbb{X} . If the sensors' FOVs do not overlap, this creates a partition of the state space:

$$\mathcal{V}_+ = \left\{ V_+^{(i)} \right\}_{i=0}^S, \quad (12)$$

where $V_+^{(0)} = \mathbb{X} - V_+$ is the region outside all sensors' FOVs and $V_+ = \bigcup_{i=1}^S V_+^{(i)}$ is the network's combined FOV. Then, we can take a number of random samples from the weighted sum of all predicted tracks $p_+^{(\xi)}(\cdot, l)$ with label l and check if they are within $V_+^{(i)}$ to numerically compute $p_V^{(i)}(l)$, the predicted probability that object l lies within the i th group of the partition. By assigning each object l to group $L_+^{(i)}$ such that $p_V^{(i)}(l)$ is maximized, we construct a partition of the label space:

$$\mathcal{L}_+ = \left\{ L_+^{(i)} \right\}_{i=0}^S. \quad (13)$$

Note that the partition definition is more complicated when object birth and spawning are considered [4], but that is beyond the scope of this work.

To apply the partition to a GLMB density or a set of partitioned GLMB PDFs, we first split all the GLMB PDFs that contain labels in multiple groups by removing all tracks for objects in another group and merging any duplicate hypotheses that are created (i.e., summing their weights). Then, all PDFs in each group are merged by combining their track lists and taking the Cartesian product of their hypothesis lists, with each new hypothesis's weight being equal to the product of the weights of the hypotheses that were combined to create it.

As described in [4], the group of objects outside all sensors' FOVs only need to have their tracks predicted, without modifying their data association hypotheses. This means that we can improve computational efficiency and avoid unnecessary truncation by not merging GLMB PDFs in this group into a single PDF. Furthermore, by assuming state-independent probabilities of detection and survival conditioned on object l being in the i th sensor's FOV, denoted $P_D^{(i)}(l)$ and $P_S(l)$, the FP-GLMBF equations can be simplified by assuming probabilities of detection

$$p_D^{(\xi)}(l) = p_V^{(i)}(l | \xi) P_D^{(i)}(l) \quad (14)$$

and survival

$$p_S^{(\xi)}(l) = 1 - p_V^{(i)}(l | \xi) (1 - P_S(l)), \quad (15)$$

where $p_V^{(i)}(l | \xi)$ is the probability of an object lying within the sensor's FOV given association history ξ . The probability $p_V^{(i)}(l | \xi)$ is computed by drawing random samples from the predicted track for object l given history ξ and checking if they are inside $V_+^{(i)}$.

The prior filtering density for the group corresponding to the i th sensor's FOV after repartitioning at timestep $k + 1$ is

$$\boldsymbol{\pi}^{(i)}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{(I, \xi) \in \mathcal{F}(L_+^{(i)}) \times \Xi} w^{(I, \xi)} \delta_I[\text{lab } \mathbf{X}] \left(p^{(\xi)} \right)^{\mathbf{X}}, \quad (16)$$

where Ξ is the set of histories of maps $\theta : \mathbb{L} \rightarrow \{0, 1, \dots\}$, assigning measurements from each sensor i to labels in its

corresponding group $L^{(i)}$ at time k . If $\theta(l) > 0$, it gives the index of the measurement assigned to object l at time k , and if $\theta(l) = 0$, no measurement is assigned to object l at that time. With the simplifying assumptions described in the previous paragraph, the predicted and updated PDF is

$$\begin{aligned} \pi_+(X_+ | Z_+^{(i)}) &\propto \Delta(X_+) \sum_{I, \xi, I_+, \theta_+} w^{(I, \xi)} w_+^{(I, \xi, I_+, \theta_+)} \left(Z_+^{(i)} \right) \delta_{I_+} [\text{lab } X_+] \\ &\quad \times \left(p_+^{(\xi, \theta_+)} \left(\cdot | Z_+^{(i)} \right) \right)^X, \end{aligned} \quad (17)$$

where $I_+ \in L_+^{(i)}$, $\theta_+ \in \Theta_+$, the set of maps at $k+1$, and

$$\begin{aligned} w_+^{(I, \xi, I_+, \theta_+)} \left(Z_+^{(i)} \right) &= \left(p_S^{(\xi)} \right)^{I \cap I_+} \left(1 - p_S^{(\xi)} \right)^{I - I_+} \\ &\quad \times \left(\psi_+^{(\xi, \theta_+)} \left(\cdot | Z_+^{(i)} \right) \right)^{I_+}, \end{aligned} \quad (18)$$

$$p_+^{(\xi, \theta_+)} \left(x_+, l | Z_+^{(i)} \right) = \frac{p_+^{(\xi)} \left(x_+, l \right) \psi_+^{(\theta_+(l))} \left(x_+, l | Z_+^{(i)} \right)}{\psi_+^{(\xi, \theta_+(l))} \left(l | Z_+^{(i)} \right)}, \quad (19)$$

$$p_+^{(\xi)} \left(x_+, l \right) = \left\langle f_{S,+} \left(x_+ | \cdot, l \right), p^{(\xi)} \left(\cdot, l \right) \right\rangle, \quad (20)$$

$$\psi_+^{(\xi, j)} \left(l | Z_+^{(i)} \right) = \left\langle p_+^{(\xi)} \left(\cdot, l \right), \psi_+^{(j)} \left(\cdot, l | Z_+^{(i)} \right) \right\rangle, \quad (21)$$

$$\begin{aligned} \psi_+^{(j)} \left(x_+, l | Z_+^{(i)} \right) &= \delta_0[j] \left(1 - p_D^{(\xi)} \left(l \right) \right) + \left(1 - \delta_0[j] \right) \\ &\quad \times \frac{p_D^{(\xi)} \left(l \right) g^{(i)} \left(z_{+,j}^{(i)} | x_+, l \right)}{\kappa^{(i)} \left(z_{+,j}^{(i)} \right)}, \end{aligned} \quad (22)$$

where $Z_+^{(i)}$ is the set of measurements $z_{+,j}^{(i)}$ collected by sensor i , $g^{(i)}$ is the sensor's measurement likelihood function, $\kappa^{(i)}$ is its clutter intensity, and $f_{S,+}$ is a track's predicted PDF given its earlier state.

Finally, the predicted and updated GLMB density is truncated by computing the cost $c(l, j) = -\log \left(\eta^{(I, \xi)} \left(l, j \right) \right)$, for each individual association j of object l , where

$$\eta^{(I, \xi)} \left(l, j \right) = \begin{cases} 1 - p_S^{(\xi)} \left(l \right), & j < 0, \\ p_S^{(\xi)} \left(l \right) \psi_+^{(\xi, j)} \left(l | Z_+^{(i)} \right), & j \geq 0, \end{cases} \quad (23)$$

$j < 0$ indicates that object l does not exist, $j = 0$ indicates that it exists but was not detected, and $j > 0$ indicates that it exists and generated the j th measurement. These costs are assembled into a cost matrix, and the resulting ranked assignment problem is solved via Murty's algorithm or a Gibbs sampler-based stochastic algorithm to truncate the hypothesis list by finding the highest-weighted new GLMB components [3].

The estimated multi-target state is extracted from each GLMB PDF by taking its maximum *a posteriori* cardinality estimate and then the set of tracks from its highest-weighted hypothesis with that cardinality. The overall multi-target state estimate is simply the union of the estimates from each GLMB density.

The use of bi-fidelity orbit uncertainty propagation has two implications for the FP-GLMBF implementation. First, we

compute $p_V^{(i)}(l)$ for all l using only their low-fidelity state PDFs, in order to avoid the expense of high-fidelity correction. We then perform high-fidelity correction for all tracks for objects in groups in the FOV to enable accurate measurement gating and measurement update. Second, the tracks obtained during estimate extraction must undergo high-fidelity correction to ensure accuracy, but these tracks are then discarded to avoid increasing error by replacing their initial PDFs with bi-fidelity approximations.

III. SCENARIO DESCRIPTION

The simulation covers a period of three days, beginning on April 18, 2021 at 0:00 UTC, with a timestep size of 1 min. The simulated SO population consists of two sub-populations. The first is a hypothetical large satellite constellation comprising 1000 satellites whose orbits all have an SMA of 7528.1363 km and an inclination of 53° , forming a shell around the Earth between $\pm 53^\circ$ latitude. The second is the set of 15,302 SOs that were included in the TLE catalog on the start date and had eccentricity less than 0.25 and mean motion greater than 11.25 rev/day, excluding five objects that would reenter the Earth's atmosphere within the three day period. This results in a total population of 16,297 SOs.

Key orbital elements for these objects are shown in Fig. 2. As the figure shows, LEO objects tend to have highly inclined orbits, resulting in changing measurement opportunities for ground-based sensors as the Earth rotates.




Fig. 2. Key orbital elements for the population used in this scenario.

The true dynamics for the simulation include a 16×16 spherical harmonics-based model of the Earth's gravity field (EGM2008) [23], the point mass gravity of the Sun and Moon, and atmospheric drag. The ballistic coefficient for all objects is a constant 0.021 kg/m^2 . The true dynamics are solved using a

variable-step RK45 integrator. The filter uses the true model as its high-fidelity model (i.e., there is no simulated mismodeling), but it still assumes a process noise covariance of $Q = \sigma_a I_3$, where $\sigma_a = 10^{-12} \text{ km/s}^2$ to account for error due to bi-fidelity propagation. The true initial states for all objects are sampled from their initial tracks in the FP-GLMBF to ensure statistical consistency.

For progressive resampling, the EnGMF has $N_{\min} = 100$, $\frac{dN}{d\tau} = 20 \text{ rev}^{-1}$, and $N_{\max} = 500$. For bi-fidelity propagation, the filter stores the initial state and five most recent timesteps for each particle.

The sensor network comprises ten ground-based radars, each with a conical FOV of half-angle 1° . The locations, field of regard (FOR) definitions, and noise statistics for these sensors are shown in Tables I–III. A sensor’s FOR is the union of all its possible FOVs, which are typically constrained by topography and sensor pointing hardware. All sensors provide measurements of topocentric azimuth, elevation, and range, and all sensors except for #9 also provide measurements of range-rate. All sensors have a conditional probability of detection $P_D = 0.98$ for objects in their FOV, while the filter believes that $P_D = 0.96$. The sensors are tasked at each simulation step to point at the object in their FOV with the longest time since it was last in the FOV of any sensor, whether or not it was successfully detected at that time. In order to generate the task schedule before beginning the simulation, the true states of the SOs are used to determine sensor pointing directions.

TABLE I
SENSOR LOCATIONS.

#	Latitude (deg)	Longitude (deg)	Altitude (m)
1	30.572460	-86.2146900	73.6404
2	54.361800	-0.670100	211.1952
3	41.752500	-70.538056	109.9920
4	39.136070	-121.350720	142.5200
5	76.570278	-68.299167	370.7752
6	48.724722	-97.899722	379.7776
7	-7.906630	-14.402580	58.3280
8	64.300222	-149.190956	164.0092
9	9.395390	167.479130	-23.6192
10	42.617440	-71.490890	156.4660

TABLE II
SENSOR FOR LIMITS.

#	Azimuth (deg)		Elevation (deg)		Range (km)
	Min.	Max.	Min.	Max.	Max.
1	120	240	1.0	105	40,744
2	0	360	3.0	85	4,828
3	347	227	3.0	85	5,556
4	126	6	3.0	85	5,556
5	297	177	3.0	85	4,828
6	313	63	1.9	45	3,300
7	0	360	-5.0	90	59,264
8	184	64	3.0	85	4,828
9	0	360	0.0	90	90,000
10	0	360	0.5	90	90,000

TABLE III
SENSOR NOISE STANDARD DEVIATIONS.

#	Az. (deg)	El. (deg)	Range (km)	Range-rate (km/s)
1	0.0470	0.0462	0.1623	0.0002
2	0.0215	0.0189	0.0175	0.0002
3	0.0370	0.0313	0.0165	0.0023
4	0.0264	0.0238	0.0177	0.0024
5	0.0481	0.0329	0.0143	0.0013
6	0.0175	0.0129	0.0244	0.0014
7	0.0176	0.0135	0.0804	0.0294
8	0.0462	0.0469	0.0194	0.0027
9	0.0129	0.0114	0.0110	N/A
10	0.0073	0.0081	0.0054	0.0001

The FP-GLMBF begins with one estimated track for each SO, each in its own group within the initial label space partition. Each track is initialized by propagating its catalog-derived initial state backward in time to $t = -24 \text{ h}$, then using the result as the mean of a Gaussian PDF. The covariance of this PDF is aligned with the mean’s velocity direction, with position and velocity standard deviations of 100 m and 1 cm/s along the velocity direction and 1 m and 0.1 mm/s along any perpendicular axis. This Gaussian PDF is then sampled to initialize the EnGMF. The EnGMF PDF is predicted forward to $t = -12 \text{ h}$, and updated once with a geocentric measurement of right ascension, declination, range, and range-rate. These measurements have standard deviations of 0.1 deg for the angles, 0.1 km for range, and 0.01 km/s for range-rate. Finally, the PDF is predicted to $t = 0$ and updated with another measurement from the same sensor. This process is necessary because the task of constructing a realistic initial PDF for an SO is non-trivial due to nonlinearity.

Clutter returns are not generated in this scenario, and the FP-GLMBF assumes a uniform clutter intensity of $\kappa = 10^{-9}$. Each group of the partition is limited to a maximum of 100 hypotheses, and hypotheses with weights less than 10^{-6} are pruned. The ranked assignment problem in the FP-GLMBF joint predict–update step is solved via Murty’s algorithm. The measurement gating distance is 6. The probability of an object being inside a sensor’s FOV is computed using 50 random samples. Because no object can actually die during the simulation, the conditional probability of survival is set conservatively to $P_S = 0.9999$.

IV. RESULTS

This section discusses the results of one run of the simulation, performed on a 2.4–3.9 GHz CPU with 192 concurrent threads. The total runtime is 10 h, as shown in Fig. 3. The final average number of particles per track is 237.3.

The overall tracking performance, quantified using the optimal subpattern assignment (OSPA) metric [24], is shown in Fig. 4. This metric uses cutoff 15 km and order 2. These values are recorded once every 2 h of simulated time.

In Fig. 4, we see that the location error rises slightly near the beginning of the simulation, before dropping to a lower value. This rise is due to the track initialization procedure, which included an update for all tracks at $t = 0$, which is not consistent

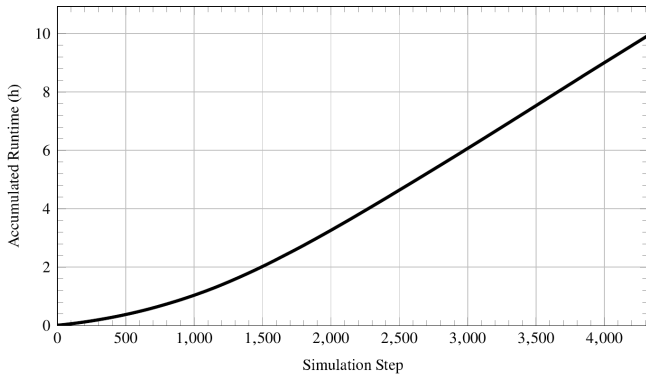


Fig. 3. Total accumulated filtering time over the course of the simulation.




Fig. 4. Tracking error over time. Note that the cutoff is 15 km.

with the actual timing of observations in the simulation. Fig. 5 shows that the average observation rate is a little more than one per day, and some objects are observed much less frequently.




Fig. 5. Observation rate statistics over time.

In this trial, the filter only ever believes that one object has died, and it manages to reacquire the object with a better estimate of its state, as shown in Fig. 6. This is responsible for the temporary increase in cardinality error visible in Fig. 4. Note that while this object's initial position error exceeds its estimated

overall 3σ bound, this is actually the case for 12% of all objects and is not necessarily responsible for the filter's temporary belief that it has disappeared. The fact that the object is later reacquired shows the value of the multiple hypothesis tracking provided by the FP-GLMBF, which was able to continue tracking this object even while its highest-weighted hypothesis said that it had died.




Fig. 6. Error and uncertainty for the object that is temporarily believed to have died. Note that its position error initially exceeds the estimated overall 3σ bound, but is within the bound after reacquisition.

At the end of the simulation, 75 objects (0.46% of the population) have position errors exceeding the 15 km cutoff. Of these, six appear to be due to label switching and are not reflected in the OSPA distance. Notably, most of the 75 objects are observed three times throughout the simulation, which is close to the average observation rate. However, their average first observation occurs at $t = 0.91$ day, which may be too late to begin accurate data association without additional observations. The simplistic sensor tasking algorithm used here may have missed earlier opportunities to observe these objects and may have also chosen suboptimal follow-up observation times. This could be improved by using a less greedy, information-theoretic approach to sensor tasking [25].

V. CONCLUSION

We have demonstrated that our filter can track and maintain custody of more than 16,000 space objects in low Earth orbit without requiring the use of high-performance computing facilities. This was achieved by combining the FOV-partitioned GLMB filter (FP-GLMBF), ensemble Gaussian mixture filter (EnGMF), multi-fidelity orbit propagation, and our new progressive resampling method. Progressive resampling allowed the single-target EnGMF to provide accurate estimated probability density functions (PDFs) to the multi-target FP-GLMBF, even for objects that had not been observed recently, while reducing overall memory usage. We also showed that, in this scenario, an object's low-fidelity state PDF could be used for the FP-GLMBF's partition update without needing to perform high-fidelity correction.

ACKNOWLEDGMENT

The authors thank S. Yun and N. Ravago for their work on our previous papers and their code that helped inform the implementation of this filter.

REFERENCES

- [1] B.-T. Vo and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Transactions on Signal Processing*, vol. 61, no. 13, Jul. 2013.
- [2] B.-N. Vo, B.-T. Vo, and D. Phung, "Labeled random finite sets and the Bayes multi-target tracking filter," *IEEE Transactions on Signal Processing*, vol. 62, no. 24, Dec. 2014.
- [3] B.-N. Vo, B.-T. Vo, and H. G. Hoang, "An efficient implementation of the generalized labeled multi-Bernoulli filter," *IEEE Transactions on Signal Processing*, vol. 65, no. 8, Apr. 2017.
- [4] B. L. Reifler and B. A. Jones, "Tractable generalized labeled multi-Bernoulli filtering with limited fields of view," *IEEE Transactions on Aerospace and Electronic Systems*, 2023, under review.
- [5] J. L. Anderson and S. L. Anderson, "A monte carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts," *Monthly Weather Review*, vol. 127, no. 12, pp. 2741–2758, Dec. 1999.
- [6] S. Yun, R. Zanetti, and B. A. Jones, "Kernel-based ensemble gaussian mixture filtering for orbit determination with sparse data," *Advances in Space Research*, vol. 69, no. 12, pp. 4179–4197, Apr. 2022.
- [7] B. A. Jones and R. Weisman, "Multi-fidelity orbit uncertainty propagation," *Acta Astronautica*, vol. 155, pp. 406–417, Oct. 2019.
- [8] R. P. S. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003.
- [9] 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, Nov. 2006.
- [10] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The labeled multi-bernoulli filter," *IEEE Transactions on Signal Processing*, vol. 62, no. 12, Jun. 2014.
- [11] B. L. Reifler and B. A. Jones, "Tracking of large satellite constellations using a partitioned GLMB filter," *Proceedings of the 2019 AAS/AIAA Astrodynamics Specialist Conference*, Aug. 2019, paper no. AAS 19-674.
- [12] —, "Tracking spawning events in cislunar space using a label-partitioned GLMB filter," *Proceedings of the 2022 AAS/AIAA Astrodynamics Specialist Conference*, Aug. 2022, paper no. AAS 22-540.
- [13] R. A. Broucke and P. J. Cefola, "On the equinoctial orbital elements," *Celestial Mechanics*, vol. 5, pp. 303–310, 1972.
- [14] D. Xiu, "Stochastic collocation methods: A survey," in *Handbook of Uncertainty Quantification*, R. G. Ghanem, D. Higdon, and H. Owhadi, Eds. Springer International Publishing, 2017, ch. 20, pp. 699–716.
- [15] B. L. Reifler, S. Yun, B. A. Jones, and R. Zanetti, "Multi-target ensemble Gaussian mixture tracking with sparse observations," *Proceedings of the 2021 Advanced Maui Optical and Space Surveillance Technologies Conference*, Sep. 2021.
- [16] S. Yun, N. Ravago, B. L. Reifler, R. Zanetti, and B. A. Jones, "Generalized labeled multi-Bernoulli filter with kernel-based ensemble Gaussian mixture filtering for orbit determination with sparse data," *Proceedings of the 2022 Advanced Maui Optical and Space Surveillance Technologies Conference*, Sep. 2022.
- [17] X. Zhu, A. Narayan, and D. Xiu, "Computational aspects of stochastic collocation with multifidelity models," *SIAM/ASA Journal on Uncertainty Quantification*, vol. 2, no. 1, pp. 444–463, 2014.
- [18] H. Schaub and J. L. Junkins, *Analytical Mechanics of Space Systems*. AIAA, 2003.
- [19] J. P. Vinti, "New method of solution for unretarded satellite orbits," *Journal of Research of the National Bureau of Standards*, vol. 63B, no. 2, pp. 105–116, Jun. 1959.
- [20] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*. Routledge, 1998.
- [21] R. van der Merwe and E. 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.
- [22] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House, Inc., 2007.
- [23] K. Pavlis, Nikolaos, S. A. Holmes, S. C. Kenyon, and J. K. Factor, "The development and evaluation of the Earth Gravitational Model 2008 (EGM2008)," *Journal of Geophysical Research*, vol. 117, no. B04406, Apr. 2012.
- [24] 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, pp. 3447–3457, Aug. 2008.
- [25] N. Ravago and B. A. Jones, "Risk-aware sensor scheduling and tracking of large constellations," *Advances in Space Research*, vol. 68, pp. 2530–2550, 2021.