Multi-Fidelity Methods for Orbit Determination

Brandon A. Jones
The University of Texas at Austin

ABSTRACT

Applying full-fidelity physics to tracking space objects is unnecessary, and this paper considers an adaptive multi-fidelity uncertainty quantification approach to improve tractability in orbit determination via the particle-based Bernoulli filter. Multi-fidelity uncertainty propagation refers to leveraging a hierarchy of force models and/or propagation methods to reduce computation time for density function prediction. The stochastic collocation-based method leveraged in this work employs low-fidelity force models commensurate with the main problem in orbit propagation (two-body and J_2), autonomously identifies a subset of important points used to define nodes of a collocation surrogate, and performs high-fidelity propagation of those points to rectify the original surrogate and produce an approximate correction to the existing samples. This paper presents an application of the approach for propagating a particle ensemble in the Bernoulli filter, thereby reducing the computation time of the demanding prediction step. Runtime of the filter is reduced by a factor of 12 with no statistically significant change in filter performance. Further potential reductions in runtime are also discussed.

1. INTRODUCTION

Space Domain Awareness (SDA) requires propagation of a space-object catalog, with the state of each object represented by a Probability Density Function (PDF) describing statistical knowledge of position and velocity (for example). Propagated PDFs are required for key tasks of SDA systems, which includes conjunction assessment, sensor tasking, and object characterization. Prediction of the PDF requires two important elements: (i) a tractable approximation of the PDF, and (ii) selection of dynamic models employed to describe propagation. Generating accurate representations of a PDF is a non trivial task, resulting in recent research in, for example, methods of Gaussian Mixture Models (GMM) [1,2], higher-order state transition tensors [3,4], polynomial chaos [5], separated representations [6], and other approaches. However, little research exists to adaptively select the force- and measurement-model fidelity required for orbit-state PDF propagation and update. High-fidelity models require increased computation time, but not all orbit regimes and space objects require full-fidelity propagators. This is well understood in the specific case of selecting the degree and order of a spherical harmonic model for central-body gravity field modeling. Recently, the author of this paper presented a multi-fidelity approach to orbit uncertainty propagation that can leverage a hierarchy of increasing fidelity force models [7]. This enables a reduced runtime for both particle and Gaussian mixture representations of the orbit-state PDF. This paper leverages this method to reduce the computation cost of the prediction step in orbit determination.

The multi-fidelity uncertainty propagation approach leverages a hierarchy of orbit force model fidelities to correct a large ensemble of low-fidelity propagations. Using the subspace spanned by the low-fidelity samples, a set of points are selected for high-fidelity propagation. In this context, low-fidelity may be described by force model truncation, propagation via general perturbations, or numeric integration with an increased step size. A bi-fidelity approach combines the low-fidelity propagated states with a small number of samples (on the order of 10) using the highest-fidelity model, and generates a corrected stochastic collocation surrogate for propagating all samples. This method provides a reduced computational burden for the prediction step of Sequential Monte Carlo (SMC) and Gaussian Mixture Models (GMMs) when propagated via the unscented transform.

The goal of this effort is to reduce the computation burden of nonlinear filtering with applications to SDA, and this paper leverages the new approach to enable multi-fidelity orbit determination via SMC-based estimation. Recent efforts in nonlinear estimation for space objects focus on GMM-based filters for the sake of tractability (e.g., see [8,9]). A potential boon to space-object tracking is the use of SMC-based filters, but at the cost of increased runtime. While not their only advantage, such SMC approaches enable improved models for detection and survival probability, both of which are greatly simplified or require heuristic approaches to allow for their use in GMM filters (e.g., see [10, 11].

For cases where many low-fidelity propagations may be performed, the multi-fidelity approach makes SMC methods tractable with a reduced computation load. For example, thousands of low-fidelity propagations (e.g., two-body and J_2 perturbations) may be achieved via parallelization on a Graphical Processing Unit (GPU). This reduced runtime mitigates the existing limitations on SMC filters required for multi-target tracking. In theory, the approach also allows for multi-fidelity sensor models to be employed in the filter to approximate the measurement likelihood function, which will be developed in future work.

This paper presents the framework for multi-fidelity orbit determination via SMC methods, and demonstrates the efficacy of the approach with simulated observations. We focus on single-target tracking via a Bayes-optimal Bernoulli filter, which allows for incorporating clutter, missed detections, and target survival into a mathematically rigorous framework. While the Bernoulli filter offers such flexibility, this paper focuses on the changes in runtime and performance when leveraging the multi-fidelity propagator. A study of improved flexibility remains as part of a follow-up study. Performance of the multi-fidelity approach for particle prediction is quantified via estimated state accuracy, filter consistency, and runtime.

2. MULTI-FIDELITY PROBABILITY DENSITY FUNCTION PROPAGATION

Multi-fidelity propagation of the orbit-state PDF refers to leveraging orbit propagators of varying fidelity to reduce computation time for PDF prediction. For this paper, we leverage a bi-fidelity approach to reduce overall computation time when propagating a particle ensemble. Previous studies demonstrate the use of such a methods to reduce computation time with single-point propagation errors on the order of meters [7]. Such an approach may also be used to efficiently propagate a Gaussian mixture representation of the PDF via the unscented transform [7]. This section summarizes the approach used in this work for propagating the particle ensemble.

This work leverages the method of multi-fidelity uncertainty quantification proposed in [12, 13] and augmented for the case of orbit-state uncertainty propagation [7]. Specifically, we begin with many low-fidelity samples $\boldsymbol{x}^L(\boldsymbol{\xi}) \in \mathbb{X}$ where $\mathbb{X}$ is the single-target state space, and generate a surrogate approximation based on a subset of high-fidelity samples $\boldsymbol{x}^H(\boldsymbol{\xi}) \in \mathbb{X}$ to produce approximate points $\widehat{\boldsymbol{x}}^H(\boldsymbol{\xi})$. Let Ξ denote the set of random inputs $\{\boldsymbol{\xi}_i\}_{i=1}^m$ that uniquely define particles in a given ensemble. For example, $\boldsymbol{\xi}$ may be a vector of realized random numbers from a given pseudo-random number generator, and their mapping to $\boldsymbol{x}^H$ is problem dependent. Then

$$\boldsymbol{X}^H(\Xi) \equiv \begin{bmatrix} \boldsymbol{x}^H(\boldsymbol{\xi}_1) & \cdots & \boldsymbol{x}^H(\boldsymbol{\xi}_m) \end{bmatrix} \in \mathbb{R}^{n\times m} \tag{1}$$

is the matrix of m high-fidelity propagated samples given the set Ξ. These samples then define a space

$$\mathbb{X}^H(\Xi) \equiv \text{span}\left(\boldsymbol{X}^H(\Xi)\right) = \text{span}\begin{bmatrix} \boldsymbol{x}^H(\boldsymbol{\xi}_1) & \cdots & \boldsymbol{x}^H(\boldsymbol{\xi}_m) \end{bmatrix} \subseteq \mathbb{X}, \tag{2}$$

that is also denoted as a function of Ξ. Similar definitions hold for the low-fidelity samples, e.g., $\boldsymbol{X}^L(\Xi)$ and $\mathbb{X}^L(\Xi)$. The method of stochastic collocation (e.g., see Chapter 20 of [14]) produces a surrogate

$$\boldsymbol{x}^H(\boldsymbol{\xi}) \approx \widehat{\boldsymbol{x}}^H(\boldsymbol{\xi}) = \sum_{\ell=1}^{r} c_\ell(\boldsymbol{\xi})\, \boldsymbol{x}^H(\overline{\boldsymbol{\xi}}_\ell), \tag{3}$$

where $r \ll m$ is the number of terms in the expansion, c_ℓ are the coefficients, and $\overline{\boldsymbol{\xi}}_\ell$ are the random vector corresponding to the collocation points. The approach to multi-fidelity propagation leveraged here uses low-fidelity samples to generate c_ℓ and r evaluations of the high-fidelity propagator. This approach requires three key assumptions:

- $\boldsymbol{X}^L(\Xi)$ allows for identifying the r samples required for Eq. (3),
- The r high-fidelity samples produce a sufficiently accurate basis for $\boldsymbol{x}^H(\boldsymbol{\xi}) \approx \widehat{\boldsymbol{x}}^H(\boldsymbol{\xi})$, and
- Coefficients c_ℓ of the expansion

$$\boldsymbol{x}^L(\boldsymbol{\xi}) \approx \widehat{\boldsymbol{x}}^L(\boldsymbol{\xi}) = \sum_{\ell=1}^{r} c_\ell(\boldsymbol{\xi})\, \boldsymbol{x}^L(\overline{\boldsymbol{\xi}}_\ell) \tag{4}$$

  are sufficiently accurate to be leveraged in Eq. (3).

Coefficients c_ℓ and the important points $\overline{\boldsymbol{\xi}}_\ell$ are jointly computed and identified in a single algorithm. Let the set of random inputs corresponding to important samples be denoted

$$\overline{\Xi} \equiv \left\{\overline{\boldsymbol{\xi}}_\ell\right\}_{\ell=1}^{r}, \tag{5}$$

and are identified to minimize the distance

$$\text{dist}\left(\boldsymbol{x}^L(\boldsymbol{\xi}), \mathbb{X}^L(\overline{\Xi})\right) \equiv \inf_{\boldsymbol{y}\in\mathbb{X}^L(\overline{\Xi})} \|\boldsymbol{x}^L(\boldsymbol{\xi}) - \boldsymbol{y}\| \tag{6}$$

between points in $\boldsymbol{X}^L(\Xi)$ and the space $\mathbb{X}^L(\overline{\Xi})$.

Generating $\overline{\Xi}$ to minimize Eq. (6) is (generally) not tractable, and we employ a greedy algorithm to generate $\overline{\Xi}$. As detailed in [13], a solution to the pivoted Cholesky decomposition

$$\left[\boldsymbol{X}^L\right]^T \boldsymbol{G}^L \left[\boldsymbol{X}^L\right] = \boldsymbol{A}^T \boldsymbol{L}\boldsymbol{L}^T \boldsymbol{A} \tag{7}$$

ranks the distance of points $\boldsymbol{x}^L(\boldsymbol{\xi}_i)$ in Eq. (6) and provides a means for computing c_ℓ. In Eq. (7), $\boldsymbol{A}$ is a pivot matrix that orders samples based on distance to $\mathbb{X}^L(\Xi)$, and $\boldsymbol{G}^L$ is the Gramian matrix generated by the low-fidelity samples. For details on generating $\boldsymbol{L}$ and $\boldsymbol{A}$, see Algorithm 1 of [13]. Given $\boldsymbol{L}$, then

$$\boldsymbol{L}\boldsymbol{L}^T\boldsymbol{c} = \boldsymbol{g}, \tag{8}$$

where

$$g_\ell = \left\langle \boldsymbol{x}^L(\boldsymbol{\xi}), \boldsymbol{x}^L(\overline{\boldsymbol{\xi}}_\ell)\right\rangle, \qquad \ell = 1, \ldots, r, \tag{9}$$

the vector $\boldsymbol{c} \in \mathbb{R}^r$ includes the expansion coefficients c_ℓ, and $< \cdot, \cdot >$ denotes the discrete inner product. Solving Eqs. (7)-(9) requires a fixed value of r. For the sake of simplicity in this work, we assume $r = 12$. Previous work developed a method for autonomously selecting r to achieve a given accuracy, and typically resulted in $r = 11$ or 12 for the case of orbit uncertainty propagation.

3. SINGLE-TARGET BERNOULLI FILTER

The Bernoulli filter is a Bayes-optimal single-target filter that incorporates possible target birth/death and factors that influence the measurement update such as detection probability and clutter measurements [15]. Note that this work does not leverage a multi-target tracking framework such as the Bayes-optimal approach in [16]. This presentation does not derive the Bernoulli filter, but the interested reader may review [15] and [17].

The Bernoulli filter incorporates knowledge of target existence in the definition of the single-target state PDF. This state is described as a Random Finite Set (RFS) $\mathcal{X}$ with maximum cardinality zero, i.e., the target does not exist if $\mathcal{X} = \varnothing$ and otherwise has estimated state $\boldsymbol{x} \in \mathbb{X}$ and $\mathcal{X} = \{\boldsymbol{x}\}$. Let q be the target's existence probability. Then the state PDF is

$$p(\mathcal{X}) = \begin{cases} 1-q, & \mathcal{X} = \varnothing, \\ q \cdot p(\boldsymbol{x}), & \mathcal{X} = \{\boldsymbol{x}\}, \\ 0, & \text{otherwise}, \end{cases} \tag{10}$$

where $p(\boldsymbol{x})$ is the PDF for the single-target state $\boldsymbol{x}$. The Bayes-optimal filter then predicts the state PDF from time t_{k-1} to t_k via the Chapman-Kolmogorov equation

$$p_{k|k-1}\left(\mathcal{X}_k\right) = \int f\left(\mathcal{X}_k|\mathcal{X}\right) p_{k-1}\left(\mathcal{X}\right) \, \delta\mathcal{X}, \tag{11}$$

where $f()$ is the transition kernel and the integral denotes the set integral as seen in RFS-based filters (e.g., see [15]). The transition kernel incorporates both the probability of object survival $p_S(\boldsymbol{x})$ and evolution of the estimated state $\boldsymbol{x}$. This state is then updated via Bayes theorem

$$p_{k|k}\left(\mathcal{X} \mid \mathcal{Z}\right) \propto \varphi_k\left(\mathcal{Z} \mid \mathcal{X}\right) p_{k|k-1}\left(\mathcal{X}\right) \tag{12}$$

where $\mathcal{Z}$ is a measurement RFS and $\varphi()$ is the RFS measurement likelihood function. Note that $|\mathcal{Z}| \geq 0$ with variations in the number of elements due to missed detections and/or clutter measurements. The measurement likelihood incorporates detection probability $p_D(\boldsymbol{x})$ and clutter in the measurement update. The remainder of this section outlines the Bernoulli filter implementation via Sequential Monte Carlo (SMC).

SMC filters approximate the state PDF via

$$p(\boldsymbol{x}) = \sum_{i=1}^{N} \omega^{(i)} \delta_{\boldsymbol{x}^{(i)}}(\boldsymbol{x}) \tag{13}$$

where $\omega^{(i)} \geq 0$ are the weights of particles $\boldsymbol{x}^{(i)}$, $\sum \omega^{(i)} = 1$, and $\delta()$ is the common Kronecker delta function. The SMC-based implementation of the Bernoulli filter recursively updates the existence probability q and the weights $\omega^{(i)}$ of the particles. A rigorous treatment of SMC-based estimators may be found in [18], and this work leverages particle regularization for resampling.

Given the existence probability q_{k-1} and single-target state PDF $p_{k-1}(\boldsymbol{x})$ at t_k, then

$$q_{k|k-1} = p_B\left(1 - q_{k-1}\right) + q_{k-1} \sum_{i=1}^{N} \omega^{(i)}\, p_S\left(\boldsymbol{x}^{(i)}\right) \tag{14}$$

where p_B is the probability that the target is "born" in the surveillance region and $p_S()$ is the state-dependent survival probability. While presented here for the sake of completeness, object birth is not considered in this paper and $p_B = 0$. For the case where p_S is not dependent on $\boldsymbol{x}$, i.e., $p_S(\boldsymbol{x}) = p_S$, then the second term simplifies to $q_{k-1}p_S$. Particles are predicted by drawing a sample $\boldsymbol{x}_k^{(i)} \sim \rho(\boldsymbol{x}_k|\boldsymbol{x}_{k-1}^{(i)}, \mathcal{Z})$ where $\rho(\boldsymbol{x}_k|\boldsymbol{x}_{k-1}^{(i)}, \mathcal{Z})$ is the *proposal density*. While different forms of the proposal density exist, this work leverages the common assumption

$$\rho(\boldsymbol{x}_k|\boldsymbol{x}_{k-1}^{(i)}, \mathcal{Z}) = \pi_{k|k-1}(\boldsymbol{x}_k|\boldsymbol{x}_{k-1}) \tag{15}$$

where $\pi_{k|k-1}()$ is the (single-) target state transition kernel. Hence, predicted particles are simply the propagated particles $\boldsymbol{x}_{k-1}^{(i)}$ from t_{k-1} to t_k. With this simplification and ignoring birth, then

$$\omega^{(i)} = \frac{p_S(\boldsymbol{x}^{(i)})\, r_{k-1}}{r_{k|k-1}}\, \omega_{k-1}^{(i)}, \qquad i = 1, \ldots, N \tag{16}$$

See [17] for a presentation of the Bernoulli filter without this simplification via the proposal density.

The measurement update in the Bernoulli filter seeks to update the existence probability and the particles weights given the measurement set $\mathcal{Z}$. For the SMC-based filter, then

$$I_1 = \sum_{i=1}^{N} p_D(\boldsymbol{x}_{k|k-1}^{(i)})\omega_{k|k-1}^{(i)} \tag{17}$$

$$I_2(\boldsymbol{z}) = \sum_{i=1}^{N} p_D(\boldsymbol{x}_{k|k-1}^{(i)}) g_k(\boldsymbol{z}|\boldsymbol{x}_{k|k-1}^{(i)})\omega_{k|k-1}^{(i)} \tag{18}$$

$$\Delta_k = I_1 - \sum_{\boldsymbol{z}\in\mathcal{Z}} \frac{I_2(\boldsymbol{z})}{\lambda\, c(\boldsymbol{z})} \tag{19}$$

where $g_k()$ is the single-target measurement likelihood function, λ is the mean number of clutter measurements per scan (assuming a Poisson distribution for cardinality), and $c(\boldsymbol{z})$ is the density of clutter measurements. The updated existence probability is then

$$q_{k|k} = \frac{1 - \Delta_k}{1 - \Delta_k \cdot q_{k|k-1}}. \tag{20}$$

and the updated weights are

$$\omega_{k|k}^{(i)} = \frac{\tilde{\omega}_{k|k}^{(i)}}{\sum_{i=1}^{N} \tilde{\omega}_{k|k}^{(i)}} \tag{21}$$

$$\tilde{\omega}_{k|k}^{(i)} = \left[1 - p_D(\boldsymbol{x}_{k|k-1}^{(i)}) + p_D(\boldsymbol{x}_{k|k-1}^{(i)}) \sum_{\boldsymbol{z}\in\mathcal{Z}} \frac{g_k(\boldsymbol{z}|\boldsymbol{x}_{k|k-1}^{(i)})}{\lambda\, c(\boldsymbol{z})}\right] \omega_{k|k-1}^{(i)}. \tag{22}$$

To improve robustness of the filter, we resample using regularization (e.g., see Table 3.6 of [18]) after each measurement update. This promotes particle diversity, and results in equally weighted particles for the *a posteriori* PDF. In this approach, the empirical mean and covariance matrix are generated using the available particles. A slight deviation is applied to each particle generated via resampling based on the filter state variance and a fixed-width scale factor derived from the dimensionality of the state space. Future work will instead consider resampling as part of the prediction step with the measurement likelihood $\varphi(\mathcal{Z} \mid \boldsymbol{x}_{k|k-1}^{(i)})$ as a proposal density. Given the existing surrogate generated as part of the propagation, the multi-fidelity approach reduces the cost of resampling based on the prediction residuals, thereby improving robustness and (theoretically) reducing N.

4. SIMULATION RESULTS

4.1. SCENARIO DESCRIPTION

To demonstrate the improvement in computational efficiency for tracking space objects, we consider the case of a low-Earth orbiting spacecraft observed by a small radar sensor network. Each station in the network generates azimuth, elevation, range, and range-rate measurements. Table 1 provides the true spacecraft state at epoch, and Fig. 1 illustrates the groundtrack compared to the simulated sensor network. Note the ground stations are isolated to locations in the Northern/Western hemisphere. Fig. 2 presents the possible observation times for each of the stations over the 16 hour simulation time, and there is almost a six-hour gap between the initial time and first possible detection. Note that only one measurement is possibly generated per satellite pass and is randomly selected over the indicated time interval.

Table 1. True spacecraft state in orbital elements

Element	Value
Semimajor Axis	7778.1363 km
Eccentricity	0.01
Inclination	80.0°
Right Ascension	100.0°
Arg. of Periapsis	0.0°
True Anomaly	0.0°

Propagation of particles in the Bernoulli filter leverages the multi-fidelity propagation framework presented in Section 2. The low- and high-fidelity force models are described in Table 2 with additional force model parameters provided in Table 3. Note that the low-fidelity propagator is the so-called *main problem* for orbit propagation, while the high-fidelity models include additional perturbations and increased fidelity. The Bernoulli filter assumes $p_S = 0.99$ and $p_D = 0.99$. Future work will stress the filter by considering missed detections and advanced models for $p_D(\boldsymbol{x})$, but this work focuses on the gains in computation time with the multi-fidelity framework.

The Bernoulli filter uses a Cartesian definition of the state. The *a priori* state PDF is Gaussian with covariance matrix

$$\overline{\boldsymbol{P}} = \begin{bmatrix} (0.1 \text{ km})^2 \cdot \mathbb{I}_{3\times 3} & \boldsymbol{0} \\ \boldsymbol{0} & (10^{-4} \text{ km/s})^2 \cdot \mathbb{I}_{3\times 3} \end{bmatrix}. \tag{23}$$

The initial mean is based on a random sample of a Gaussian with a mean vector defined by the true state derived from Table 1. This yields

$$\overline{\boldsymbol{x}}_0 = \begin{bmatrix} \boldsymbol{r} \\ \boldsymbol{v} \end{bmatrix}, \qquad \boldsymbol{r} = \begin{bmatrix} -1337.006 \\ 7583.346 \\ 0.028 \end{bmatrix} \text{ km}, \qquad \boldsymbol{v} = \begin{bmatrix} -1.236505 \\ -0.218029 \\ 7.120749 \end{bmatrix} \text{ km/s}. \tag{24}$$

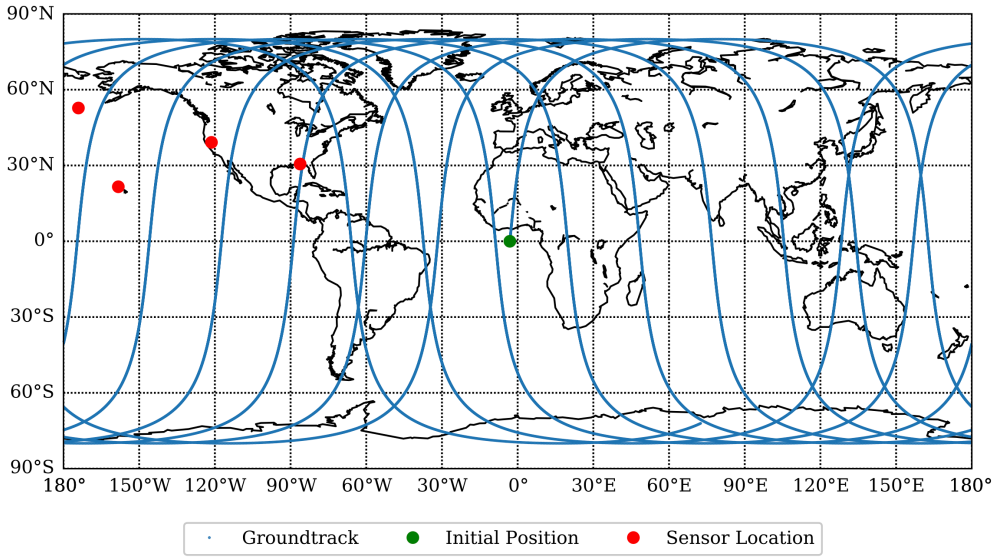

Fig. 1. Groundtrack for simulated satellite and radar ground stations

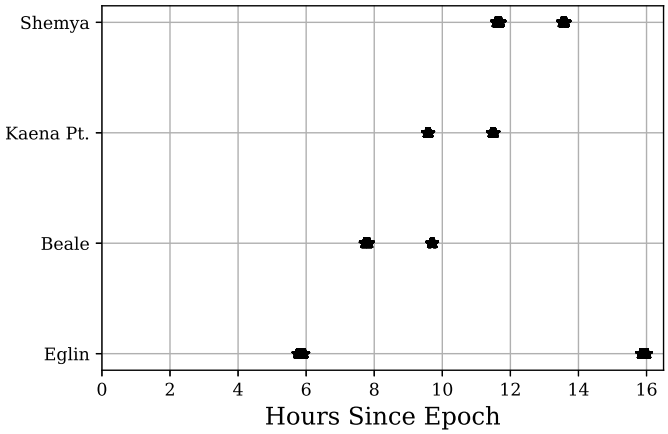

Fig. 2. Times of satellite visibility for each ground station

Table 2. Force models for low- and high-fidelity models

Model	**Low-Fidelity**	**High-Fidelity**
Central Body Gravity	Two-Body and J_2	EGM2008, 70×70
Third-Body Perturbations	None	DE430 Sun and Moon [19]
Solar Radiation Pressure	None	Cannonball
Atmospheric Drag	None	Cannonball
Coordinate System Reduction	None	IAU2006 [20]

Table 3. Force model and satellite parameters

Parameter	Value
Satellite Mass	500 kg
Satellite Drag and SRP Area	1 m^2
C_D	2.0
C_R	1.5
Epoch Time	January 3, 2010 23:58:53.82 UTC

The filter uses 10^6 particles, which was selected to prevent degeneracy. Particles are generated by random sampling of a Gaussian distribution with mean $\overline{\boldsymbol{x}}_0$ and covariance $\overline{\boldsymbol{P}}$.

All runtime tests are executed on a Dell PowerEdge R730 with dual 20-core Xeon E-2698 v4 processors, 128 GB of memory, and running RedHat 7.2. Propagation software is written in C and compiled using the GNU gcc compiler version 4.8.5. Filter-specific software and the multi-fidelity framework are written in Python 2.7. When enabled, parallelization of particle propagation leverages the ipyparallel module, version 6.2.1 with 25 engines. When conducting runtime tests, no other processes are running on the computer.

4.2. PROPAGATION ACCURACY

Before demonstrating efficacy of the multi-fidelity approach for orbit determination, this section presents the propagation accuracy for the first $\sim$6 hour measurement gap. To aid in interpreting runtime comparisons, parallelization is disabled for these tests. Hence, to improve tractability, only 10^5 samples are propagated.

Table 4. Runtime and RMS differences for multi-fidelity and full-fidelity propagation of 10^5 samples

Case	**Runtime [s]**	**RMS Position Difference [m]**
Low-Fidelity	3.6	834.17
Multi-Fidelity	7.6	0.022
Full-Fidelity	627.0	–

Table 4 includes the accuracy of the proposed approach and compares the runtimes for the low-, multi-, and full-fidelity methods. This is based on the propagation from the epoch time to the time of the first measurement at approximately six hours (see Fig. 2). No parallelization is enabled for this runtime test. The multi-fidelity approach reduces runtime by a factor of $\sim$82 when compared to full-fidelity propagation, and with negligible change in propagated position. Uncorrected particles using the low-fidelity propagator exhibit 834 m errors. While the multi-fidelity approach produces centimeter-level differences in the propagated state, note that errors on the order of tens of meters have been seen in other, longer duration tests [7].

Figure 3 presents normalized histograms of the propagation errors and Gaussian fits of the data for each of the components of position. In this case, the distribution of the errors is approximately Gaussian distributed with near-zero mean. Such a property of the errors is not guaranteed (in general), and future work will seek to predict the distribution of error in the multi-fidelity samples.

4.3. ORBIT DETERMINATION ACCURACY

Finally, this section presents the accuracy and tractability of the Bernoulli filter with a multi-fidelity propagator when compared to a full-fidelity model. As discussed previously, we use 10^6 particles with propagation parallelized over 25 cores.

Figs. 4 and 5 present the filter accuracy and PDF variance for the two propagation methods. The mean and variance are empirically determined based on the particles after resampling (equally weighted). For both the position and velocity, the variance appears to be more pessimistic when using the multi-fidelity propagation. While not included here for the sake of conciseness, running the filters with variations in the random number generator seed value produce variations in the 3σ boundaries. A Monte Carlo analysis (i.e., running the filter multiple times) is required to fully compare the uncertainty of the two filters, but differences between the two results appear to be within the statistical noise. For all cases, estimated state error (defined as mean state minus the truth) are within the $\pm 3\sigma$ boundaries.

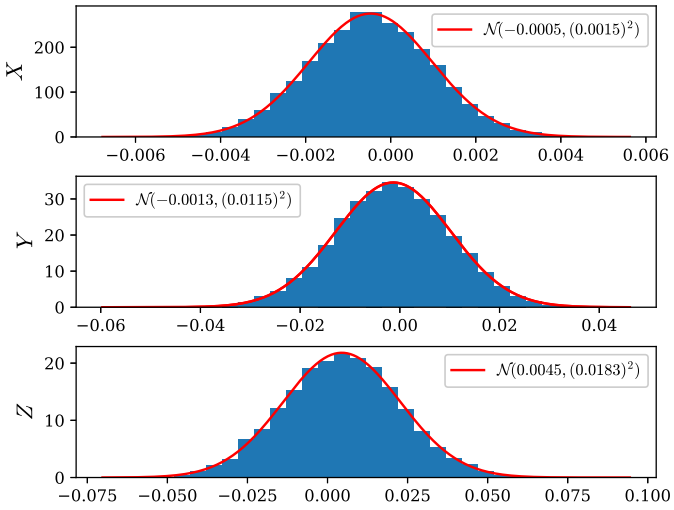

Fig. 3. Normalized histogram of propagation errors (in meters) using the multi-fidelity framework and the empirically determined Gaussian distribution of deviations

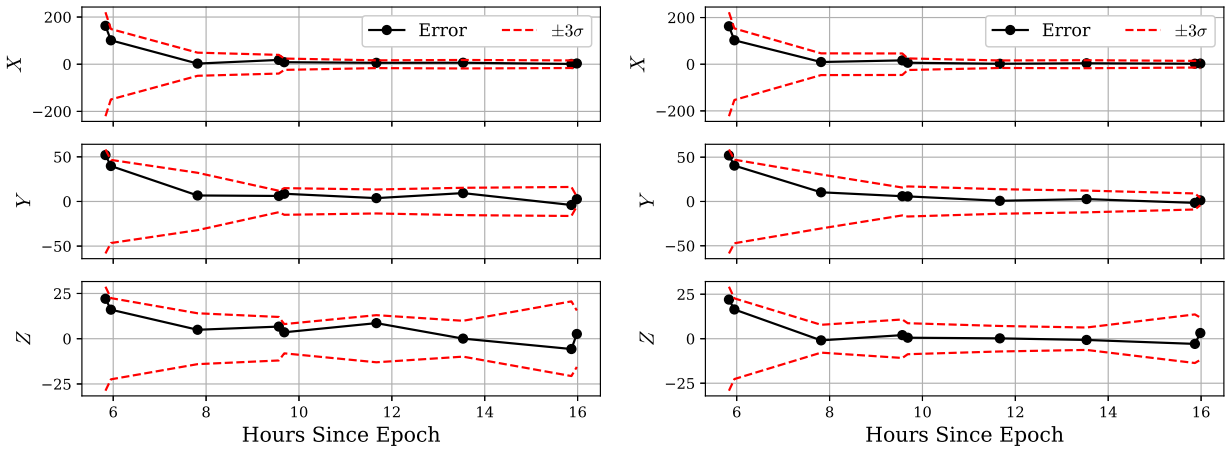

Fig. 4. Filter position errors (in meters) and $\pm 3\sigma$ covariance envelopes starting with the first observation. Left image is for the multi-fidelity propagator and the right image is for the full-fidelity model.

Table 5. Runtime [s] breakdown for filter execution with 10^6 particles

Component	**Multi-Fidelity**	**Full-Fidelity**
Prediction	29.5	970.5
Update	54.5	52.9
Total	84.0	1023.4

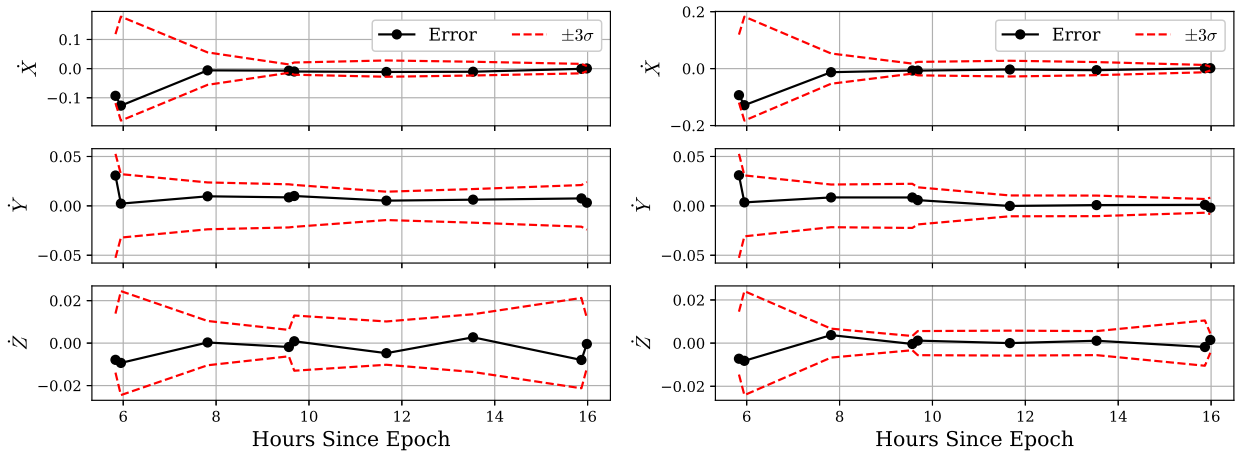

Fig. 5. Filter velocity errors (in meters/second) and $\pm 3\sigma$ covariance envelopes starting with the first observation. Left image is for the multi-fidelity propagator and the right image is for the full-fidelity model.

Table 5 summarizes the filter runtime when leveraging the multi- and high-fidelity propagators. As stated previously, the filter implementation leverages Python instead of a compiled language. Absolute runtime will be improved after porting to a compiled language (e.g., C/C++), but the results in Table 5 illustrate relative improvement with the change in propagation method. In the case of multi-fidelity propagation of particles (and 25 cores), the measurement update requires more computation time. Note that the update may also be parallelized, but this was not done here. The multi-fidelity approach reduces the total runtime by a factor of $\sim$12.

5. CONCLUSIONS AND FUTURE WORK

This paper presented the use of a multi-fidelity framework for orbit-state propagation in the SMC-Bernoulli filter for tracking a single space object via ground-based radar. The multi-fidelity framework combines the use of low- and high-fidelity propagators to reduce the overall computation time of a particle ensemble with some, in this case negligible, loss of accuracy. This produces over one order of magnitude reduction in runtime for the particle filter. Differences in the filter's solution are within the statistical error of the estimated posterior PDF. Previous discussions highlighted specific items that require follow-up study. Additionally, we plan to demonstrate the use of multi-fidelity methods for more computationally expensive measurement models, e.g., more realistic models with light-time correction and atmospheric effects. The proposed approach makes particle filtering for tracking space objects more tractable for cases that require such flexibility in the tracking algorithm.

REFERENCES

1. Joshua T. Horwood, Nathan D. Aragon, and Aubrey B. Poore. Gaussian sum filters for space surveillance: Theory and simulations. *Journal of Guidance, Control, and Dynamics*, 34(6):1839–1851, November - December 2011.

2. Kyle J. DeMars, Robert H. Bishop, and Moriba K. Jah. Entropy-based approach for uncertainty propagation of nonlinear dynamical systems. *Journal of Guidance, Control, and Dynamics*, 36(4):1047–1057, July-August 2013.

3. Ryan S. Park and Daniel J. Scheeres. Nonlinear mapping of Gaussian statistics: Theory and applications to spacecraft trajectory design. *Journal of Guidance, Control, and Dynamics*, 29(6):1367–1375, November - December 2006.

4. Kohei Fujimoto, Daniel J. Scheeres, and K. Terry Alfriend. Analytical nonlinear propagation of uncertainty in the two-body problem. *Journal of Guidance, Control, and Dynamics*, 35(2):497–509, March-April 2012.

5. Brandon A. Jones, Alireza Doostan, and George H. Born. Nonlinear propagation of orbit uncertainty using non-intrusive polynomial chaos. *Journal of Guidance, Control, and Dynamics*, 36(2):430–444, March - April 2013.

6. Marc Balducci, Brandon A. Jones, and Alireza Doostan. Orbit uncertainty propagation and sensitivity analysis with separated representations. *Celestial Mechanics and Dynamical Astronomy*, 129(1-2):105–136, 2017.

7. Brandon A. Jones and Ryan Weisman. Multi-fidelity orbit uncertainty propagation. In *1st IAA Conference on Space Situational Awareness*, Orlando, FL, November 13-15 2017.

8. Kyle J. DeMars, Islam I. Hussein, Carolin Frueh, Moriba Jah, and R. Scott Erwin. Multiple-object space surveillance tracking using finite-set statistics. *Journal of Guidance, Control, and Dynamics*, 38(9):1741–1756, 2015.

9. Brandon A. Jones and Ba-Ngu Vo. A labeled multi-Bernoulli filter for space object tracking. In *2015 AAS/AIAA Space Flight Mechanics Meeting*, AAS 15-413, Williamsburg, VA, January 11-15 2015.

10. Ba-Ngu Vo and Wing-Kin Ma. The Gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, November 2006.

11. Brandon A. Jones, Noble Hatten, Nicholas Ravago, and Ryan P. Russell. Ground-based tracking of geosynchronous space objects with a GM-CPHD filter. In *Proceedings of the 2016 Advanced Maui Optical and Space Surveillance Technologies Conference*, Wailea, Maui, Hawaii, September 20-23 2016.

12. Akil Narayan, Claude Gittelson, and Dongbin Xiu. A stochastic collocation algorithm with multifidelity models. *SIAM Journal on Scientific Computing*, 36(2):A495–A521, 2014.

13. 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.

14. Roger G. Ghanem, David Higdon, and Houman Owhadi, editors. *Handbook of Uncertainty Quantification*. Springer International Publishing, Switzerland, 2017.

15. Ronald P. S. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, Boston, Massachusetts, 2007.

16. Ba-Tuong Vo and Ba-Ngu Vo. Labeled random finite sets and multi-object conjugate priors. *IEEE Transactions on Signal Processing*, 61(13):3460–3475, July 2013.

17. Branko Ristic, Ba-Tuong Vo, Ba-Ngu Vo, and Alfonso Farina. A tutorial on Bernoulli filters: Theory, implementation and applications. *IEEE Transactions on Signal Processing*, 61(13):3406–3430, July 2013.

18. Branko Ristic, Sanjeev Arulampalam, and Neil Gordon. *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House, Boston, Massachusetts, 2004.

19. William M. Folkner, James G. Williams, Dale H. Boggs, Ryan S. Park, and Petr Kuchynka. The planetary and lunar ephemerides DE430 and DE431. IPN Progress Report 42-196, Jet Propulsion Laboratory, California Institute of Technology, http://ipnpr.jpl.nasa.gov/progress_report/42-196/196C.pdf, February 2009.

20. Gérard Petit and Brian Luzum. IERS conventions (2010). IERS Technical Note 36, International Earth Rotation and Reference Systems Service (IERS), Frankfurt am Main, Germany, 2010.