Kernel-Based Ensemble Gaussian Mixture Filtering for Orbit Determination with Sparse Data

Sehyun Yun[1,*], Renato Zanetti[2], Brandon A. Jones[2]

Aerospace Engineering and Engineering Mechanics Department, Cockrell School of Engineering, The University of Texas at Austin, Austin, Texas 78712

Abstract

In this paper, a modified kernel-based ensemble Gaussian mixture filtering (EnGMF) is introduced to produce fast and consistent orbit determination capabilities in a sparse measurement environment. The EnGMF is based on kernel density estimation (KDE) to combine particle filters and Gaussian sum filters. This work proposes using Silverman's rule of thumb to reduce the computational burden of KDE. Equinoctial orbital elements are used to improve the accuracy of the KDE bandwidth parameter in the modified EnGMF. A bi-fidelity approach to propagation and an adaptation algorithm for selecting the appropriate number of particles are also applied to the EnGMF to reduce the computational burden with an acceptable loss in accuracy for long time propagation. Through numerical simulation, the proposed implementation is compared to state-of-the-art approaches in terms of accuracy, consistency, and computational speed.

Keywords: Low Earth orbit constellations, Kernel density estimation, Particle filter, Gaussian mixture model, Bi-fidelity propagation, Adaptive algorithm

1. Introduction

In recent years, there has been an increasing interest in tracking an ever-growing number of space objects (SOs) for collision avoidance and space domain awareness

*Corresponding author

Email addresses: shyun@utexas.edu (Sehyun Yun), renato@utexas.edu (Renato Zanetti), brandon.jones@utexas.edu (Brandon A. Jones)

[1]Postdoctoral Fellow, Department of Aerospace Engineering and Engineering Mechanics

[2]Assistant Professor, Department of Aerospace Engineering and Engineering Mechanics

(DeMars et al., 2012; Pate-Cornell and Sachon, 2001; Castaings et al., 2015). As very large, low earth orbit (LEO) constellations are being developed and launched, the risk of collision in LEO keeps increasing because of a high density of SOs in this region. The high number and density of LEO SOs require accurate orbit determination and data association (National Research Council, 2012). Currently, only a limited number of radar-based surveillance sensors are available and used to estimate the state of an SO in LEO. The current approach to maintaining a LEO catalogue is not scalable to tens of thousands of spacecrafts. The solution of this problem is either adding more hardware (more tracking stations and/or clusters of supercomputers) or improving the computational efficiency of tracking and data association software used to maintain the catalogue.

A software-only solution is one in which the number of available measurements per SO is reduced because the current surveillance network is tasked to acquire data from many more SOs. It requires an efficient data association algorithm and an estimator able to extract as much information as possible from the sparse data. This paper addresses the latter, and proposes an accurate and computationally fast nonlinear estimation algorithm for orbit determination.

1.1. Overview of Estimators for Orbit Determination with Sparse Data Problem

For linear systems with linear measurements, the well-known Kalman filter (Kalman, 1960) provides a globally optimal solution, i.e., it extracts as much information from the data as possible (in a minimum mean square error sense). In the presence of nonlinearities (either in the dynamics, the measurements, or both), a nonlinear filter is able to produce a more accurate estimate than a linear one, i.e., extract more information from the data. Radar measurements of range, range-rate, and angles to an SO are inherently nonlinear. A nonlinear filter, therefore, will outperform a linear filter such as the extended Kalman filter (EKF) (Gelb, 1974) or unscented Kalman filter (UKF) (Julier and Uhlmann, 2004) even in the presence of near-linear dynamics.

To cope with the sparse data problem, this paper compares two nonlinear algorithms: the adaptive entropy-based Gaussian mixture information synthesis (AEGIS) (DeMars et al., 2013) and the kernel-based ensemble Gaussian mixture filtering (En-

GMF) (Anderson and Anderson, 1999). This work proposes new modifications to En-GMF to greatly improve its computational complexity. Two implementations of the UKF and the ensemble Kalman filter (EnKF) (Katzfuss et al., 2016) are also compared to the proposed approach, representing the SO with both Cartesian and equinoctial coordinates (Broucke and Cefola, 1972).

For linear measurements and dynamics, the UKF reduces to the Kalman filter and is the optimal out of all linear estimators regardless of the probability distributions (in a minimum mean square error, MMSE, sense). The UKF is typically more robust to nonlinearities than the EKF (Julier and Uhlmann, 2004) but can still fail to produce an adequate estimate in the case of high nonlinearities. The nonlinearities of orbital dynamics are easily mitigated by choosing to represent the SO's state with an appropriate set of orbital elements, for example equinoctial elements. Changes in these elements, specifically the angle quantity, are linear, and variations due to nonlinear effects are relatively small. This choice of coordinates, therefore, allows for accurate and computationally inexpensive long time propagations of the mean and covariance matrix (Junkins et al., 1996; Sabol et al., 2010), for example when using the unscented transformation (UT). The price to pay for linear dynamics is typically an even more nonlinear measurement model, which may cause UKF divergence in a scarce-measurement environment, as shown in the numerical results section of this paper. In measurement-rich environments, when long propagations are followed by dense measurements arcs, a batch least-squares approach is often the preferred orbit determination solution (Tapley et al., 2004), as it allows to extract more information from nonlinear measurements than linear sequential filters. After processing the measurement batch, the mean and covariance of the estimate can be propagated with the UT to start a new iteration. Batch least squares does not provide full information about the probability distribution function and it only returns the mean and covariance matrix and the underlying distribution is typically assumed Gaussian, hence they work best when many measurements are available such that after incorporating them all the resulting uncertainty is close to Gaussian. Nonlinear recursive filters, on the other hand, approximate the optimal MMSE estimator, which has the lowest square estimation error (on average), and provides a full description of the underlying PDFs. AEGIS and our proposed modification

to EnGMF are two examples of nonlinear filters.

While this work focuses on the estimation problem, adding data association to a single-target nonlinear filter is a problem of considerable current interest. Data association and collision detection benefit from full knowledge of the PDF, which can be approximated with AEGIS and the EnGMF but inevitably results in a Gaussian assumption for linear estimators. As long as the PDF remains approximately Gaussian after measurement batches are incorporated, linear filters produce excellent performance. Hughes and Alfriend (2020), for example, assumes that the initial orbit determination solution is an estimate with a Gaussian distribution, and employs modified equinoctial elements to propagate the state and associate a sequence of observations to an SO using the Mahalanobis distance.

The AEGIS method is based on the standard Gaussian sum filter (GSF) (Sorenson and Alspach, 1971; Alspach and Sorenson, 1972). The GSF is a nonlinear estimator for nonlinear systems and it has been applied to SO tracking applications (Horwood and Poore, 2011; Horwood et al., 2012). To deal with multimodal and non-Gaussian distributions, the GSF approximates the probability density functions (PDF) as a Gaussian mixture model (GMM). The GSF provides a nearly optimal solution when enough components are taken and each Gaussian component has a small enough covariance matrix such that the nonlinear dynamic and measurement functions can be accurately approximated to linear functions in the support of each Gaussian component. In the presence of a Gaussian prior and a nonlinear measurement, the GSF outperforms linear filters when the prior is approximated by many Gaussians of smaller covariance such that the measurement is approximately linear in their support.

One of the limitations of the standard GSF is that the weights of the Gaussian components remain the same during nonlinear propagations. Several studies recently have been proposed to address this issue and improve the standard GSF algorithm to better account for nonlinear dynamics (DeMars et al., 2013; Terejanu et al., 2011; Tuggle and Zanetti, 2018). One of these approaches is AEGIS, which splits the Gaussian components to reduce the effects of nonlinearities of a dynamical system during the prediction of state uncertainty (DeMars et al., 2013).

Another approach to nonlinear filtering is sequential importance sampling with re-

sampling (SISR), commonly known as particle filters (PFs) (Arulampalam et al., 2004). PFs are known to suffer from degeneracy with near-deterministic dynamics, i.e., with little process noise. As orbital dynamics is well characterized, a particle filter implementation of orbit determination inevitably requires low process noise. Modifications have been investigated to improve the standard SISR methods such as the bootstrap particle filter (BPF), auxiliary particle filter (APF), and regularized particle filter (RPF) (Doucet et al., 2001) by combining particle filters and GSF (Anderson and Anderson, 1999; Yun and Zanetti, 2019; Raihan and Chakravorty, 2018; Yun and Zanetti, 2020; Liu et al., 2016). For example, the sequential Monte Carlo filtering with Gaussian mixture model (SMCGMM) proposed in Yun and Zanetti (2019) assumes that each particle of the pre-propagation distribution to be a Gaussian component having a zero or small covariance matrix. Raihan and Chakravorty (2018) and Yun and Zanetti (2020) integrate a PF with a clustering algorithm (e.g., K-means algorithm or expectation-maximization (EM) algorithm) to approximate the prior distribution with a GMM. Although clustering to form the GMM provides an accurate solution for a highly nonlinear system, it is computationally expensive and not of practical use for tracking large LEO constellations.

Other examples of hybrid PF/GSF algorithms include Anderson and Anderson (1999) and Liu et al. (2016), which approximate each propagated particle as a Gaussian component with a non-zero covariance matrix calculated by bandwidth selection for kernel density estimation (KDE). KDE is a non-parametric technique to estimate the PDF of a random variable (Scott, 1992). The KDE algorithm with a Gaussian kernel is similar to the EM clustering algorithm in that they construct a GMM using the particles. However, in the KDE algorithm, every particle is considered as a Gaussian component to establish a GMM whereas EM clustering algorithm groups several particles into each Gaussian mixture components. An adaptable bandwidth selection suffers from a high computational cost similar to the clustering algorithms presented in Raihan and Chakravorty (2018) and Yun and Zanetti (2020).

1.2. Contributions of This Work

In this paper, a modification to the EnGMF algorithm is proposed to efficiently track SOs in LEO with short and sparse observation data. A key element of the EnGMF algorithm is the determination of the covariance matrix of each Gaussian component in a GMM. The covariance matrix is determined by the bandwidth parameter of a kernel function. Although the optimal bandwidth parameter can be obtained using a data-driven method (Park and Marron, 1990; Jones et al., 1996), this approach is computationally expensive. Alternatively, we can compare the simulation results of a system using a range of the bandwidth parameter (Liu et al., 2016) and tune this parameter according to the system. In this paper, we propose an approach to achieve a near-optimal bandwidth parameter with a low computational cost for orbit determination with sparse observation data. We achieve this by computing the bandwidth of a Gaussian kernel in the KDE algorithm with Silverman's rule of thumb (Silverman, 1986) to reduce the KDE computational burden. Moreover, a bi-fidelity approach to propagation (Jones and Weisman, 2019) and an adaptive algorithm are applied to the modified EnGMF to reduce its computational cost. The bi-fidelity approach uses both high- and low-fidelity models to maximize the accuracy of orbit uncertainty propagation, while minimizing computational cost for the propagation of the EnGMF. The choice of the number of particles is also one key parameter of the EnGMF to reduce the computational cost. Thus, an adaptive algorithm for selecting an appropriate number of particles is used for the EnGMF based on convergence assessment (Elvira et al., 2017). The contributions of this paper can be summarized as follows:

- A modification to the EnGMF algorithm is proposed to efficiently track SOs in LEO with short and sparse observation data.
- A bi-fidelity propagation approach is included in the EnGMF algorithm to reduce its computational cost.
- An algorithm to autonomously adapt the number of EnGMF particles is proposed to reduce its computational cost.
- The proposed EnGMF is compared to the UKF, EnKF, and AEGIS method

through numerical simulation in terms of accuracy, consistency, and computational speed.

1.3. Outline

The remainder of this paper is organized as follows. First, the dynamics and measurement models are described and the coordinate systems are presented. Then, the two nonlinear estimation techniques, the AEGIS and a modified EnGMF, are introduced in section III and an adaptive EnGMF with bi-fidelity propagation algorithm is presented in section IV. In section V, simulation results are shown using the proposed algorithms followed by some concluding remarks on the methodology and results.

2. System Models

2.1. Dynamics Model

The inertial position and velocity of an SO are denoted by $\boldsymbol{r}^I = [r_x\ r_y\ r_z]^{\mathrm{T}}$ and $\boldsymbol{v}^I = [v_x\ v_y\ v_z]^{\mathrm{T}}$. The orbital dynamics of an SO in Earth-Centered Inertial (ECI) coordinates are given by

$$\begin{bmatrix} \dot{\boldsymbol{r}}^I \\ \dot{\boldsymbol{v}}^I \end{bmatrix} = \begin{bmatrix} \boldsymbol{v}^I \\ -\frac{\mu}{r^3}\boldsymbol{r}^I + \boldsymbol{a}^I_{NS} + \boldsymbol{a}^I_{3B} + \boldsymbol{a}^I_{drag} + \boldsymbol{a}^I_{srp} \end{bmatrix} \tag{1}$$

where μ is the Earth's gravitational parameter and r is the Euclidean norm of $\boldsymbol{r}^I$. $\boldsymbol{a}^I_{NS}$ is the gravitational perturbation due to non-spherical effect of the Earth gravity, $\boldsymbol{a}^I_{3B}$ indicates the third-body perturbations of the Moon and the Sun, and $\boldsymbol{a}^I_{drag}$ and $\boldsymbol{a}^I_{srp}$ represent the acceleration perturbation due to atmospheric drag and solar radiation pressure (SRP), respectively. For this study, the EGM2008 (Pavlis et al., 2008) gravity model is used for the Earth and 70 $\times$ 70 degrees and order are applied for gravity modeling, and the planetary and lunar ephemeris DE430 (Folkner et al., 2009) is selected to compute the location of the Moon and the Sun. The 70 $\times$ 70 gravity field is selected to match one traditional definition of high fidelity propagation for LEO objects in the context of space situational awareness (Weedon and Cefola, 2010), but can be varied based on required prediction accuracy and computation constraints.

The primary non-gravitational force acting on SOs in low earth orbit (LEO) is the drag force. The drag acceleration due to atmospheric density relies upon the drag coefficient, the cross-sectional area of an SO face perpendicular to velocity vector, and the height of an SO above the Earth's surface. The acceleration due to drag is then given by

$$\boldsymbol{a}^I_{drag} = -\frac{1}{2} C_d \frac{A}{m} \rho_d \|\boldsymbol{v}_{rel}\| \boldsymbol{v}_{rel} \tag{2}$$

where C_d is the drag coefficient, m and A are the mass and cross-sectional area of the SO, respectively, ρ_d is the atmospheric density at altitude of the SO, $\|\cdot\|$ means the Euclidean norm, and $\boldsymbol{v}_{rel}$ is the atmosphere-relative velocity vector. For computing the atmospheric density, the exponential density model is employed in this study (Vallado, 1997).

The acceleration due to SRP depends on the shape of an SO and the cannonball model, i.e., spherical object, is assumed in this paper. The acceleration perturbation due to SRP is then given by

$$\boldsymbol{a}^I_{srp} = -\frac{S_F A C_R \boldsymbol{u}^I_{sun}}{mc} \tag{3}$$

where S_F is the solar flux, m is the mass of the SO, c is the speed of light, C_R is the coefficient of reflectivity, u^I_{sun} is the unit vector pointing from the SO to the Sun in the ECI frame.

2.2. Measurement Model

The four-dimensional measurement vector $\boldsymbol{y}_k$ at time t_k contains range, range-rate, right ascension and declination of the observed SO as seen by a ground-based radar sensor. The range ρ_{range} and range-rate $\rho_{rangerate}$ measurements are calculated along the line of sight (LOS) between the ground-based radar and the SO. The relative position vector $\boldsymbol{\rho}^I = [\rho_x \ \rho_y \ \rho_z]^{\mathrm{T}}$ between the SO and a ground station $\boldsymbol{r}^I_S$ coordinatized in ECI is given by:

$$\boldsymbol{\rho}^I = \boldsymbol{r}^I - \boldsymbol{r}^I_S \tag{4}$$

The error-free range measurement is given:

$$\rho_{range} = \|\boldsymbol{\rho}^I\| = \|\boldsymbol{r}^I - \boldsymbol{r}_S^I\| \tag{5}$$

By differentiating Eq. (5) with respect to time, the error-free range-rate measurement is obtained as follows:

$$\rho_{rangerate} = \frac{\left(\boldsymbol{r}^I - \boldsymbol{r}_S^I\right)^{\mathrm{T}} \left(\boldsymbol{v}^I - \boldsymbol{v}_S^I\right)}{\|\boldsymbol{r}^I - \boldsymbol{r}_S^I\|} \tag{6}$$

where $\boldsymbol{v}_S^I$ is the time rate of change of the ground station position vector with respect to the inertial frame.

Along with the range and range-rate, angle data in the form of right ascension α and declination δ are measured to estimate the states of the SO. The error-free angle observation equations are described as follows:

$$\alpha = \tan^{-1}\left(\frac{\rho_y}{\rho_x}\right), \quad \delta = \sin^{-1}\left(\frac{\rho_z}{\|\boldsymbol{\rho}^I\|}\right) \tag{7}$$

All measurements are corrupted by zero-mean, Gaussian noise. In this study, light travel time delay and measurement biases are not considered.

2.3. Coordinate systems

The dynamic equations of the SO presented above are expressed in Cartesian coordinates, which results in nonlinear differential equations. Alternatively, equinoctial orbital elements (Broucke and Cefola, 1972) offer a near-linear dynamics. The Keplerian motion is exactly linear, and nonlinearities arise only due to perturbations such as non-central gravity and drag. The equinoctial orbital elements are expressed as func-

tions of the Keplerian orbital elements as follows:

$$\begin{aligned} a &= a \\ h &= e \sin(\omega + \Omega) \\ k &= e \cos(\omega + \Omega) \\ \lambda_0 &= M_0 + \omega + \Omega \\ p &= \tan(i/2) \sin(\Omega) \\ q &= \tan(i/2) \cos(\Omega) \end{aligned} \tag{8}$$

where a is the semi-major axis, e is the eccentricity, i is the inclination, Ω is the longitude of the ascending node, ω is the argument of periapsis, and M_0 is the mean anomaly.

3. Estimation Techniques

This section reviews the AEGIS and introduces the proposed nonlinear estimation algorithms to cope with the sparse data problem: a modified EnGMF.

3.1. Adaptive Entropy-based Gaussian Mixture Information Synthesis

AEGIS uses an entropy-based method to detect nonlinearity of a dynamical system during the prediction of state uncertainty and then applies a splitting technique to decrease the approximation error caused by truncating the nonlinear functions of the system to low-order. The AEGIS method is based on the standard GSF which is a nonlinear estimator. In the GSF, non-Gaussian PDFs are approximated as a GMM as follows:

$$p(\boldsymbol{x}) = \sum_{i=1}^{N} \omega^{(i)} n(\boldsymbol{x}; \boldsymbol{\mu}^{(i)}, P^{(i)}) \tag{9}$$

where $\boldsymbol{x}$ is a random variable, $p(\boldsymbol{x})$ is the PDF of $\boldsymbol{x}$, N is the number of all Gaussian components, $n(\boldsymbol{x}|\boldsymbol{\mu}, P)$ represents the Gaussian PDF with mean $\boldsymbol{\mu}$ and covariance P; and $\boldsymbol{\mu}^{(i)}$, $P^{(i)}$, and $\omega^{(i)}$ are the means, covariance matrices, and weights of the i^{th} Gaussian component. The PDF normalization and positivity properties lead to the

following constraints on the weights:

$$\omega^{(i)} \geq 0, \ \forall i \qquad \sum_{i=1}^{N} \omega^{(i)} = 1 \tag{10}$$

The performance of the GSF mainly depends on both the number and the weights of the components of a GMM; however, both of them are held constant during the propagation step. To improve the standard GSF algorithm to better adapt to nonlinearities of the system, the AEGIS approach allows for the modification of the Gaussian components over the propagation step based on two main mechanisms.

The first step of AEGIS is to monitor the nonlinearity of the dynamics using a property derived from the differential entropy for linearized dynamical systems. The differential entropy of a continuous random variable $\boldsymbol{x}$ is defined as follows (DeMars et al., 2013):

$$H(\boldsymbol{x}) = -\int_S p(\boldsymbol{x}) \log\left(p(\boldsymbol{x})\right) d\boldsymbol{x} = E\{-\log\left(p(\boldsymbol{x})\right)\} \tag{11}$$

where S is the support set. In this paper, all logarithms are assumed to be natural. The analytic solution of the differential entropy for a multivariate Gaussian distribution is then expressed as follows:

$$H(\boldsymbol{x}) = \frac{1}{2} \log |2\pi e P| \tag{12}$$

where P is the covariance matrix and $|\cdot|$ represents the matrix determinant. By taking a derivative with respect to time for Eq. (12), the time rate of the differential entropy for a linearized dynamical system is obtained as follows (DeMars et al., 2013):

$$\dot{H}(\boldsymbol{x}) = \text{trace}\{F\left(\boldsymbol{\mu}(t), t\right)\} \tag{13}$$

Therefore, the entropy value for a linearized system can be calculated by numerically integrating Eq. (13) with an appropriate initial condition, which requires only the evaluation of the trace of the dynamics Jacobian. On the other hand, a nonlinear determination of the differential entropy can be evaluated via Eq. (12) by a nonlinear implementation of the integration of the covariance matrix; for example, unscented transformation

is one of the most popular and effective methods for moment evaluation. Any deviation between the linear and nonlinear values of the entropy then indicates that nonlinearity is impacting the Gaussian component. As a result, the difference between the linearized and nonlinear predictions of the entropy can be monitored without the full solution to both the linearized and nonlinear predictors. In other words, when the difference between these values of entropy exceeds a preassigned threshold, a splitting algorithm is applied to the Gaussian component during a propagation. A smaller threshold leads to more frequent splitting during the propagation.

Once the nonlinear effects have been detected from the first step, a splitting algorithm is applied to mitigate the effects by replacing a Gaussian component with several Gaussian components. For the univariate case, each Gaussian component can be decomposed into 3 components using splitting libraries which are shown in Table 1. The splitting technique for a univariate case with splitting library can be then extended to the multivariate case by considering the principal directions of the covariance matrix. The details of the algorithm are explained in DeMars et al. (2013). After the propagation, the posteriori mean and covariance matrix, and mixture weights are obtained using the measurement update of the standard GSF.

Table 1: Three-component splitting library

i	ω_i	μ_i	σ_i
1	0.2252246249	-1.0575154615	0.6715662887
2	0.5495507502	0	0.6715662887
3	0.2252246249	1.0575154615	0.6715662887

When allowing the number of Gaussian components to grow unbounded, AEGIS is an accurate and consistent estimator. In this work we are interested not only in estimation accuracy, but also in computational efficiency to maintain custody of a very large number of SOs. The proposed solution to achieve this balance of performance versus accuracy is introduced next.

3.2. Modified Kernel-Based Ensemble Gaussian Mixture Filtering

As a recursive algorithm, the knowledge of the distribution $p(\boldsymbol{x}_{k-1}|\boldsymbol{y}_{k-1})$ at the prior time is assumed and approximated by N independent and identically distributed

(i.i.d.) samples $\boldsymbol{x}_{k-1}^{(i)}$ such that

$$p(\boldsymbol{x}_{k-1}|\boldsymbol{y}_{k-1}) \approx \sum_{i=1}^{N} \frac{1}{N} \delta(\boldsymbol{x}_{k-1} - \boldsymbol{x}_{k-1}^{(i)}) \tag{14}$$

where k is an integer that indicates the discrete time step, $\boldsymbol{y}$ is a measurement vector, and $\delta(\cdot)$ is the Dirac delta function. Following the same procedure as the BPF (Arulampalam et al., 2002), a set of samples at the next time step is obtained using the Markov transition kernel $p(\boldsymbol{x}_k|\boldsymbol{x}_{k-1})$. The Markov kernel indicates the dynamics of a system and all estimators use the true dynamic model without process noise in this paper.

The next step is to convert the samples into Gaussian mixtures using KDE. In other words, each particle is considered as a Gaussian component with non-zero covariance. The approximated GMM of the propagated samples is then expressed as follows:

$$p(\boldsymbol{x}_k) \approx \sum_{i=1}^{N} \frac{1}{N} n(\boldsymbol{x}_k; \boldsymbol{x}_{k|k-1}^{(i)}, B) \tag{15}$$

where the bandwidth matrix B is can be calculated by (Liu et al., 2016)

$$B = \beta \hat{P}_{k|k-1} \tag{16}$$

where β is the bandwidth parameter, $0 \leq \beta \leq 1$, and $\hat{P}_{k|k-1}$ is the sample covariance matrix calculated from the particles. The Gaussian components' means are the particles $\boldsymbol{x}_{k|k-1}^{(i)}$ and all GMM weights are equal to $1/N$. The covariance matrix of each Gaussian component is determined by the bandwidth parameter. The larger bandwidth parameter β, the smaller the probability assigned to the particle and vice versa.

Finally, we can incorporate the measurement information by updating the means, covariance matrices, and the weights of all N Gaussian components in the same way as the measurement update of the GSF. N i.i.d. samples are then drawn from the GMM approximation of the posterior distribution. These samples are used as a starting point for the next iteration. The details of the measurement update of the GSF and the method to draw N i.i.d. samples from a GMM are explained in Yun and Zanetti (2019).

In the EnGMF algorithm, it is crucial to choose the most appropriate bandwidth

which determines the performance of the filter. Bandwith selection is an accuracy vs. computational cost trade off, with the most accurate algorithms numerically solving an optimization problem. In this paper, we propose to use Silverman's rule of thumb (Silverman, 1986) to estimate the bandwidth (i.e., covariance) matrix B_S as follows:

$$B_S = \beta_S \hat{P} = \left(\frac{4}{n_{\boldsymbol{x}} + 2} \right)^{\frac{2}{n_{\boldsymbol{x}}+4}} N^{-\frac{2}{n_{\boldsymbol{x}}+4}} \hat{P}_{k|k-1} \tag{17}$$

where $n_{\boldsymbol{x}}$ is the dimension of the state. We can, therefore, obtain a near-optimal bandwidth parameter for orbit determination with sparse observation data without the need of performing any numerical optimization. If the sampling distribution were Gaussian, Silverman's rule of thumb would provide the optimal bandwidth parameter based on the mean integrated square error (MISE) as a performance criterion (Silverman, 1986). However, it may result in conservative (large) estimates when the distribution is not close to Gaussian. This is a very desirable feature, since inaccuracies results in conservatism rather than over-confidence and divergence. The flow chart of the modified EnGMF for orbit determination is described in Figure 1.

4. Adaptive EnGMF with Bi-Fidelity Propagation

This section presents the bi-fidelity propagation (Jones and Weisman, 2019) and an adaptive algorithm for selecting an appropriate number of particles being applied to the EnGMF to reduce the computational burden while guaranteeing accuracy and consistency.

4.1. Bi-Fidelity Orbit Uncertainty Propagation

This research focuses on the use of the bi-fidelity approaches (Narayan et al., 2014; Zhu et al., 2014) to reduce the computational cost of the EnGMF. The bi-fidelity approach leverages both high- and low-fidelity models to reduce the cost of orbit uncertainty propagation with an acceptable loss of accuracy in the prediction step of the proposed filter.

As in the EnGMF, this process starts from the knowledge of the prior distribution $p(\boldsymbol{x}_{k-1}|\boldsymbol{y}_{k-1})$ which, for the sake of simplicity, is considered as a multivariate Gaus-

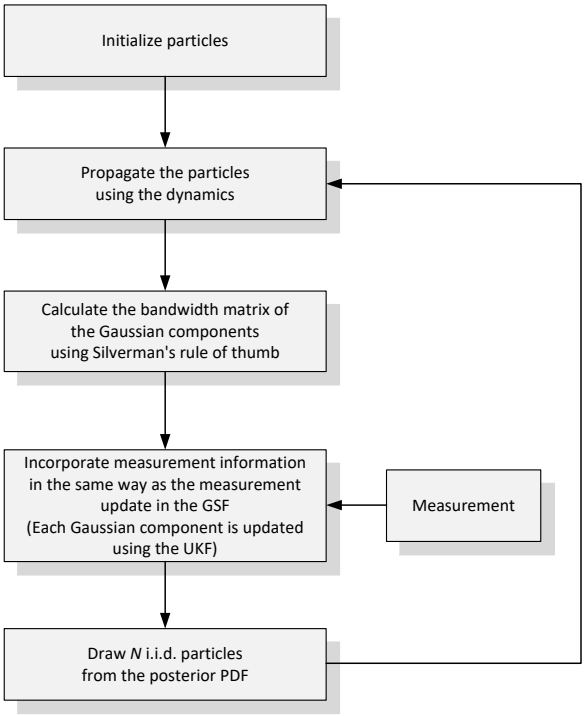

Figure 1: The flow chart of the modified EnGMF for orbit determination

sian distribution with mean $\hat{\boldsymbol{x}}_{k-1|k-1}$ and covariance $P^{xx}_{k-1|k-1}$. This prior distribution is approximated by N i.i.d. samples and each sample is drawn from the distribution as follows:

$$\boldsymbol{x}_{k-1}(\boldsymbol{\xi}^{(i)}) = \hat{\boldsymbol{x}}_{k-1|k-1} + \left(P^{xx}_{k-1|k-1}\right)^{1/2} \boldsymbol{\xi}^{(i)} \tag{18}$$

where $\boldsymbol{\xi}^{(i)}$ indicates the random inputs to the system which follows a standard multivariate normal distribution. A set of low-fidelity samples $\boldsymbol{x}^L(\boldsymbol{\xi}) \in \mathbb{X}$ in the state space $\mathbb{X}$ is then generated using a low-fidelity propagator. Given a set of random inputs $\Xi = \{\boldsymbol{\xi}^{(i)}\}_{i=1}^N$, the low-fidelity propagated samples are represented by the matrix

$$X^L(\Xi) \equiv \left[\boldsymbol{x}^L(\boldsymbol{\xi}^1) \ \cdots \ \boldsymbol{x}^L(\boldsymbol{\xi}^N)\right] \in \mathbb{R}^{n_a \times N} \tag{19}$$

In the case of orbit-state uncertainty propagation, n_a should be increased by considering the state trajectory, i.e., $\boldsymbol{x}^L(\boldsymbol{\xi})$ should be the state vector $\boldsymbol{x}_0(\boldsymbol{\xi})$ propagated to the

time of interest t using the low-fidelity propagator (Jones and Weisman, 2019). The samples then define a subset of $\mathbb{X}$

$$\mathbb{X}^L(\Xi) \equiv \text{span}\big(X^L(\Xi)\big) = \text{span}\big[\boldsymbol{x}^L(\boldsymbol{\xi}^1) \ \cdots \ \boldsymbol{x}^L(\boldsymbol{\xi}^N)\big] \subseteq \mathbb{X}, \tag{20}$$

which is a function of Ξ. The matrix $X^H(\Xi)$ and space $\mathbb{X}^H(\Xi)$ for the high-fidelity samples also follows a similar definition. Based on the stochastic collocation method (e.g., see Chapter 20 of (Ghanem et al., 2017)), we approximated the propagated samples via the surrogate

$$\boldsymbol{x}^L(\boldsymbol{\xi}) \approx \hat{\boldsymbol{x}}^L(\boldsymbol{\xi}) = \sum_{l=1}^{r} c_l(\boldsymbol{\xi}) \boldsymbol{x}^L(\bar{\boldsymbol{\xi}}^l) \tag{21}$$

where c_l are expansion coefficients, $\bar{\boldsymbol{\xi}}^l$ are the random inputs for the collocation points (or called as important points) in the expansion, and r is the rank of the surrogate with $r \ll N$. The set of the random inputs corresponding to the collocation points, $\bar{\Xi} \equiv \{\bar{\boldsymbol{\xi}}^l\}_{l=1}^r$, is identified using the low-fidelity model. We then approximate the high-fidelity samples via

$$\boldsymbol{x}^H(\boldsymbol{\xi}) \approx \hat{\boldsymbol{x}}^H(\boldsymbol{\xi}) = \sum_{l=1}^{r} c_l(\boldsymbol{\xi}) \boldsymbol{x}^H(\bar{\boldsymbol{\xi}}^l) \tag{22}$$

To summarize, in the bi-fidelity approach, the low-fidelity samples are used to calculate the coefficients c_l and the important points $\bar{\boldsymbol{\xi}}^l$. This approach includes three assumptions (Zucchelli et al., 2021):

- $X^L(\Xi)$ allows for identifying the r samples required for Eq. (22),
- The r high-fidelity samples produce a sufficiently accurate basis for $\boldsymbol{x}^H(\boldsymbol{\xi}) \approx \hat{\boldsymbol{x}}^H(\boldsymbol{\xi})$, and
- Coefficients c_l of the expansion in Eq. (21) are sufficiently accurate to be leveraged in Eq. (22).

Jones and Weisman (2019) demonstrated through several Earth-orbit test cases that these three assumptions are reasonable in the context of orbit determination.

The coefficients c_l and important points $\bar{\boldsymbol{\xi}}^l$ are jointly computed and identified in a single algorithm using the following optimization problem:

$$\bar{\Xi} = \arg\min_{\Xi} \inf_{\boldsymbol{y} \in \mathbb{X}^L(\Xi)} \|\boldsymbol{x}^L(\boldsymbol{\xi}) - \boldsymbol{y}\| \tag{23}$$

By solving the optimization problem in Eq. (23), we build a space with basis vectors $\boldsymbol{x}^L(\bar{\boldsymbol{\xi}}^l)$ for $l = 1, \cdots, r$ which minimizes the distance between the points in $X^L(\Xi)$ and the space $\mathbb{X}^L(\bar{\Xi})$; however, it is generally not tractable (Zhu et al., 2014). A greedy algorithm is therefore used to generate $\bar{\Xi}$ by leveraging a solution to the pivoted Cholesky decomposition (Zhu et al., 2014)

$$\left[X^L\right]^{\mathrm{T}} G^L \left[X^L\right] = A^{\mathrm{T}} L L^{\mathrm{T}} A \tag{24}$$

where G^L is a Gramian matrix generated via $X^L(\Xi)$ and A is a pivot matrix that orders the samples based on the distance defined in Eq. (23). The details of the algorithm are explained in Zhu et al. (2014). The coefficients $\boldsymbol{c} = [c_1, \cdots, c_r]^{\mathrm{T}}$ are then generated via

$$L L^{\mathrm{T}} \boldsymbol{c} = \boldsymbol{g} \tag{25}$$

where

$$g_l = \left\langle \boldsymbol{x}^L(\boldsymbol{\xi}), \boldsymbol{x}^L(\bar{\boldsymbol{\xi}}^l) \right\rangle, l = 1, \cdots, r, \tag{26}$$

and $< \cdot, \cdot >$ indicates the discrete inner product. Note that a fixed value of r is required for solving Eqs. (24)-(26), and it is autonomously obtained using the procedure described in Jones and Weisman (2019).

4.2. Adapting the number of particles

Along with the bi-fidelity approach, an adaptive approach is also used to reduce the computational cost of the proposed algorithm. The EnGMF can extract as much information as possible from the sparse data when each Gaussian component has a small enough covariance matrix such that the nonlinearity of the measurement function can be accurately approximated to a linear function for the likely realizations of each Gaussian component. In the EnGMF, each Gaussian component has the same covariance

matrix and its magnitude decreases as the number of particles becomes larger, thus the more particles the more accurate the EnGMF performance. This study employs an algorithm to adaptively select an appropriate number of particles, the adaptation rule is based on the convergence assessment method (Elvira et al., 2017).

The adaptive algorithm is based on the predictive PDF of the observations of the EnGMF

$$\begin{aligned} p(\boldsymbol{y}_k|\boldsymbol{Y}_{k-1}) &= \int p(\boldsymbol{y}_k|\boldsymbol{x}_k)p(\boldsymbol{x}_k|\boldsymbol{Y}_{k-1})dx_k \\ &\approx \sum_{i=1}^{N} p(\boldsymbol{y}_k|\boldsymbol{x}_k^{(i)})p(\boldsymbol{x}_k^{(i)}|\boldsymbol{Y}_{k-1}) \\ &= \frac{1}{N}\sum_{i=1}^{N} n\big(\boldsymbol{y}_k; \boldsymbol{h}_k\big(\boldsymbol{x}_{k|k-1}^{(i)}\big), H_k^{(i)} B_S H_k^{(i)\mathrm{T}} + R_k\big) \end{aligned} \tag{27}$$

where $H_k^{(i)}$ is the Jacobian of the measurement evaluated at the prior mean $\boldsymbol{x}_{k|k-1}^{(i)}$ and $\boldsymbol{Y}_{k-1}$ is the collection of all measurement vectors up to and including the current time, $\boldsymbol{Y}_{k-1} = \boldsymbol{y}_1, \cdots, \boldsymbol{y}_{k-1}$. The approximated predictive observation PDF converges almost surely to the true one as the number of particles tends to infinity (Moral, 2004). Based on this convergence assessment, the algorithm for adapting the number of particles is described below.

The first step is the generation of C fictitious observations at each time step k, $\{\tilde{\boldsymbol{y}}_k^{(j)}\}_{j=1}^{C}$, using Eq. (27). To generate fictitious observations $\tilde{\boldsymbol{y}}_k^{(j)}$, we follow these two steps:

1. Draw C samples $u^{(j)}$ from the discrete uniform distribution $\{1, 2, \cdots, N\}$.

2. For each j, draw $\tilde{\boldsymbol{y}}_k^{(j)}$ from the GMM
   $\frac{1}{N}\sum_{j=1}^{N} n\big(\boldsymbol{y}_k; \boldsymbol{h}_k(\boldsymbol{x}_{k|k-1}^{(u^{(j)})}), H_k^{(u^{(j)})} B_S H_k^{(u^{(j)})\mathrm{T}} + R_k\big)$.

The next step is to compare the actual observation value y_k with the all fictitious observations at each time step. A 1-dimensional observation is herein assumed for the sake of simplicity and a discussion of how to use multi-dimensional observations is presented in the next subsection. The number of the fictitious observations smaller

than the actual one is calculated as follows:

$$A_{k,N,C} = \left| \{ y \in \{\tilde{\boldsymbol{y}}_k^{(j)}\}_{j=1}^{C} : y < y_k \} \right| \in \{0, 1, \cdots, C\} \tag{28}$$

where $A_{k,N,C}$ is a random variable and $|\cdot|$ represents the number of elements of a set. The value of the random variable $A_{k,N,C}$ indicates the relative position of the actual observation in the fictitious observations. Note that, when the number of particles tends to infinity, the probability mass function (PMF) of $A_{k,N,C}$ almost surely converges to the discrete uniform distribution on $\{0, 1, \cdots, C\}$ under mild assumptions (Elvira et al., 2017).

Finally, after a set of W consecutive statistics, $S_k = \{a_{k-W+1,N,C}, a_{k-W+2,N,C}, \cdots, a_{k,N,C}\}$, has been obtained from the random variable $A_{k,N,C}$, a chi-squared test (Plackett, 1983) is conducted to check whether S_k is a sequence of samples from the uniform distribution on $\{0, 1, \cdots, C\}$. The chi-squared statistic is calculated as follows:

$$\chi_k^2 = \sum_{j=1}^{C} \frac{(O_j - E_j)^2}{E_j} \tag{29}$$

where O_j is the frequency of the actual observations being in the j-th relative position in the window size W and E_j is the expected frequency under the null hypothesis, i.e., $E_c = W/(C+1)$. The number of particles of the EnGMF is then adaptively selected by comparing the p-value associated with the chi-squared statistic $p_{k,C}$ with a preassigned threshold p_l. If $p_{k,C}$ is less than or equal to the significance level, we can conclude that the sequence S_k is not sampled from the uniform distribution on $\{0, 1, \cdots, C\}$. Therefore, if $p_{k,C} < p_l$, we increase the number of particles to increase the accuracy.

4.3. Algorithm Summary

The adaptive EnGMF with bi-fidelity propagation algorithm is summarized in Algorithm 1, and the details of the algorithm are described below.

1. N_0 and N_{max} are the preassigned initial and maximum number of particles, respectively, $N_0 < N_{max}$.

2. For orbit determination with sparse data problem, observations passes last only a few minutes after a long time propagation. The chi-squared test is therefore performed for every observation pass to efficiently assess the convergence of the proposed algorithm, i.e., the convergence assessment includes only one long time propagation. In addition, when the number of particles is increased, only the newly generated particles $\{\boldsymbol{x}_{k-1}^{\prime(i)}\}_{i=1}^{N/2}$ are propagated using the bi-fidelity propagation to reduce the computational burden. The added high-fidelity particles $\{\boldsymbol{x}_{k}^{\prime(i)}\}_{i=1}^{N/2}$ are then appended to the existing high-fidelity particles $\{\boldsymbol{x}_{k}^{(i)}\}_{i=1}^{N/2}$.

3. In the previous section, a 1-dimensional observation is used to assess the convergence of the algorithm. The same scheme can be applied over each observation element for a multi-dimensional case. As a result, multiple p-values, $p_{k,C,t}$ with $t = 1, \cdots, n_{\boldsymbol{y}}$ where $n_{\boldsymbol{y}}$ is the dimension of an observation, are available for the convergence assessment. As a conservative approach, in the proposed algorithm, the number of particles is doubled when any of the p-values is less than or equal to the significance level.

5. Numerical Results

5.1. Simulation Conditions

To evaluate the performance of the UKF, AEGIS, and EnGMF, one numerical example is considered. The system dynamic equations are numerically integrated with an embedded Runge-Kutta 8(7) method (Dormand and Prince, 1980) and the setup of the model is summarized in Table 2. Range, range-rate, and angle measurements are simulated using a ground station located at the North Pole (latitude = 90°, longitude = 0°, altitude = 0 km). Using a fictitious tracking station at the north pole allows for a LEO SO in near-polar orbit to be visible every orbit. Hence a sensitivity study of the accuracy of the tracking algorithms as a function of the frequency of the measurement passes (once per orbit, once every two orbits, etc.) is easily performed without the need of scheduling sensors at different geographic locations to insure the desired measurement frequency. In this simulation, tracking passes are short and sparse. The

Algorithm 1 Adaptive EnGMF with Bi-Fidelity Propagation

// Initialization //
Set $N = N_0$
Initialize the particles $\boldsymbol{x}_0^{(i)} \sim p(\boldsymbol{x}_0), \quad i = 1, \cdots, N$
Initialize the weights $\omega_0^{(i)} = 1/N, \quad i = 1, \cdots, N$
Set $a = 0$
Set $z = 1$

for $k = 1 : T$ **do** % T: Final simulation time
  // Bi-Fidelity Propagation //
  if $a = 0$ **then**
    Propagate $\boldsymbol{x}_{k-1}^{(i)}$ using the low-fidelity model
    Compute $\bar{\Xi} \equiv \{\bar{\boldsymbol{\xi}}^l\}_{l=1}^r$ through Eqs. (24)-(26)
    Propagate $\boldsymbol{x}_{k-1}^{(i)}$ using the bi-fidelity stochastic collocation approach in Section 4.1
  else
    Propagate $\boldsymbol{x}_{k-1}^{\prime(i)}$ using the low-fidelity model
    Reuse the $\bar{\Xi} \equiv \{\bar{\boldsymbol{\xi}}^l\}_{l=1}^r$ obtained from $a = 0$ case
    Propagate $\boldsymbol{x}_{k-1}^{\prime(i)}$ using the bi-fidelity stochastic collocation approach in Section 4.1
    Append $\{\boldsymbol{x}_k^{\prime(i)}\}_{i=1}^{N/2}$ to $\{\boldsymbol{x}_k^{(i)}\}_{i=1}^{N/2}$, i.e., $\{\boldsymbol{x}_k^{(i)}\}_{i=1}^{N} = \{\boldsymbol{x}_k^{(i)}, \boldsymbol{x}_k^{\prime(i)}\}_{i=1}^{N/2}$
    Set $a = 0$
  end if

  // Gaussian Mixture Model //
  Compute B_S through Eq. (17)
  Set $\omega_k^{(i)} = 1/N, \quad i = 1, \cdots, N$

  // Fictitious Observations //
  Draw $\tilde{\boldsymbol{y}}_k^{(j)} \sim n\big(\boldsymbol{y}_k; \boldsymbol{h}_k(\boldsymbol{x}_{k|k-1}^{(u^{(j)})}), H_k^{(u^{(j)})} B_S H_k^{(u^{(j)})\mathrm{T}} + R_k\big), \quad j = 1, \cdots, C$
  Compute $a_{k,N,C}$ through Eq. (28)

  if $k = Wz$ **then** *// Convergence Assessment //*
    For each observation element, compute χ_k^2 through Eq. (29) and the associated $p_{k,C,t}, \quad t = 1, \cdots, n_{\boldsymbol{y}}$
    if $p_{k,C,1} \le p_l$ **or** $\cdots$ **or** $p_{k,C,n_{\boldsymbol{y}}} \le p_l$ **and** $N < N_{max}$ **then**
      Draw $\boldsymbol{x}_{k-W}^{\prime(i)} \sim \sum_{i=1}^N \omega_{k-W}^{(i)} n\big(\boldsymbol{x}_{k-W}; \hat{\boldsymbol{x}}_{k-W}^{(i)}, \hat{P}_{k-W}^{xx(i)}\big), \quad i = 1, \cdots, N$
      Increase $N = 2N$
      Set $k = k - W + 1$
      Set $a = 1$
    else
      // Measurement Update //
      Convert $\boldsymbol{x}_k^{(i)}$ from Cartesian coordinates to the equinoctial orbital elements through Eq. (8)
      Incorporate measurement information in the same way as the measurement update of the GSF
      Set $N = N_0$
      Set $z = z + 1$
      Draw $\boldsymbol{x}_k^{(i)} \sim \sum_{i-1}^N \omega_k^{(i)} n\big(\boldsymbol{x}_k; \hat{\boldsymbol{x}}_k, \hat{P}_k^{xx}\big), \quad i = 1, \cdots, N$
      Convert $\boldsymbol{x}_k^{(i)}$ from the equinoctial orbital elements to Cartesian
    end if
  else
    // Measurement Update //
    Convert $\boldsymbol{x}_k^{(i)}$ from Cartesian coordinates to the equinoctial orbital elements through Eq. (8)
    Incorporate measurement information in the same way as the measurement update of the GSF
    Draw $\boldsymbol{x}_k^{(i)} \sim \sum_{i=1}^N \omega_k^{(i)} n\big(\boldsymbol{x}_k; \hat{\boldsymbol{x}}_k^{(i)}, \hat{P}_k^{xx(i)}\big), \quad i = 1, \cdots, N$
    Convert $\boldsymbol{x}_k^{(i)}$ from the equinoctial orbital elements to Cartesian coordinates
  end if
end for

measurements are available every 10 seconds with a pass lasting only 2 minutes, i.e., 12 measurements per pass. Each observation consists of range, range-rate, right ascension, and declination and the measurements are corrupted by additive zero-mean Gaussian white noise with standard deviation of 30 m and 0.3 m/s for the range and range-rate, respectively, and 100 arc-seconds on the right ascension and declination observation.

Table 2: High-fidelity dynamic model for the numerical simulations

Dynamic Model	Description
Primary Body Gravity	70×70
Third-Body Perturbations	Sun and Moon
Atmospheric Drag	Cannonball
Solar Radiation Pressure	Cannonball

The SO is in a near polar orbit with the following Keplerian orbital elements: a = 7,078.0068 km, e = 0.01, i = 85°, and $\omega = \Omega = \nu = 0$. The orbit of the SO is shown in Figure 2. The simulation epoch is 4-January-2010 at 00:00:00 UTC. The shape of the SO is assumed to be a sphere with a cross-sectional area of 1 m^2 and a mass of 500 kg. The drag coefficient and the coefficient of reflectivity of the SO are set to be 2 and 1.5, respectively. The initial distribution is defined in Cartesian coordinates as

$$\boldsymbol{x}_0 \sim n\left(\boldsymbol{x}_0;\ \boldsymbol{\mu}_0,\ P_0\right) \tag{30}$$

where

$$\boldsymbol{\mu}_0 = \begin{bmatrix} 7007.2175 & (\text{km}) \\ 0 & (\text{km}) \\ 0 & (\text{km}) \\ 0 & (\text{km/s}) \\ 0.6606 & (\text{km/s}) \\ 7.5509 & (\text{km/s}) \end{bmatrix} \tag{31}$$

$$P_0 = \begin{bmatrix} 1.481e{+}2 & 0 & 0 & 0 & -9.237e{-}2 & -5.333e{-}2 \\ 0 & 2.885e{+}1 & 9.994 & -3.232e{-}2 & 0 & 0 \\ 0 & 9.994 & 5.770 & -1.242e{-}2 & 0 & 0 \\ 0 & -3.232e{-}2 & -1.242e{-}2 & 3.687e{-}5 & 0 & 0 \\ -9.237e{-}2 & 0 & 0 & 0 & 6.798e{-}5 & 3.145e{-}5 \\ -5.333e{-}2 & 0 & 0 & 0 & 3.145e{-}5 & 3.166e{-}5 \end{bmatrix} \tag{32}$$

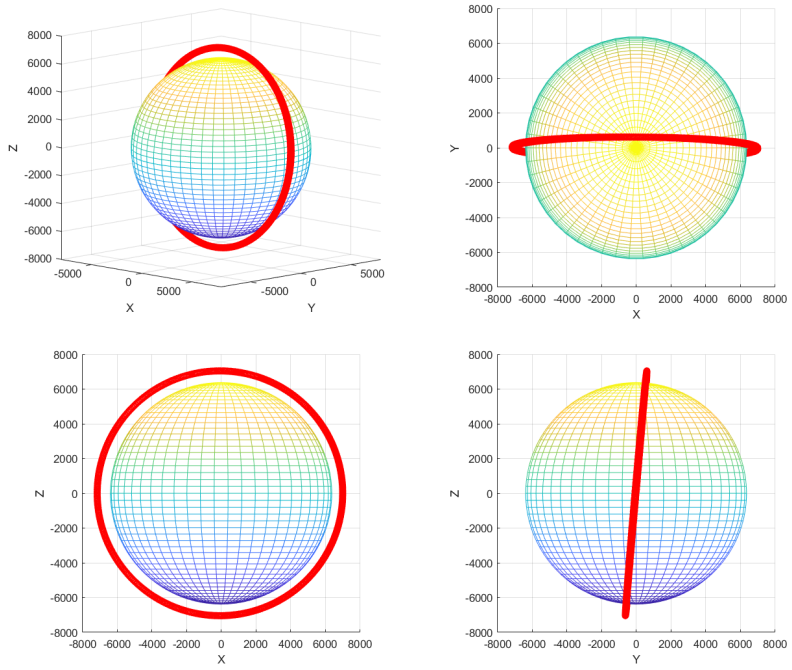

Figure 2: The orbit of the SO in the ECI frame, km

First, a Monte Carlo analysis is performed with 100 simulations, and each simulation has one measurement pass every orbital period (5926 seconds). Note that, throughout this paper, the starting time of each measurement pass is randomly selected in close proximity of a multiple of the orbital period. The UKF uses the following tuning parameters: α = 1, β = 2, κ = 3 - $n_{\boldsymbol{x}}$ = -3, for its sigma points spread. The parameter α tunes the spread of the sigma points, the parameter β is used to include prior knowledge of the distribution of state, and κ is a secondary scaling parameter

(Wan and Merwe, 2000). For the AEGIS method, the three-component splitting library is used (AEGIS-3), and the threshold on the allowed deviation of the differential entropy is set as $\Delta H = 0.001 H_0$ (DeMars et al., 2013). The value of H_0 is unique for each mixture component and based on the covariance at the latter of the last posterior estimate or the output of a splitting operation. After each measurement pass, the AEGIS algorithm is forced to have only one Gaussian component with the posterior mean and covariance matrix. This simple merging algorithm reduces AEGIS computational burden and adds conservatism that cannot cause divergence (a Gaussian distribution is the most uncertain given any finite covariance matrix). Unless otherwise specified, the EnGMF method uses 1000 particles. Both the AEGIS and the EnGMF use the UKF measurement update equations for incorporating measurement information in each GMM component. For the UKF and EnGMF, two implementations with Cartesian coordinates and the equinoctial orbital elements are compared. AEGIS is only implemented in Cartesian coordinates. An AEGIS implementation in equinoctial coordinates will result in very few component splits as the splits occur due to nonlinearity in the propagation and the SO dynamics expressed in equinoctial elements is nearly linear, making an equinoctial AEGIS implementation very similar to the equinoctial UKF.

These three algorithms are compared based on accuracy, computational complexity, and consistency. The accuracy of the filters is represented by their root-mean-square error (RMSE), which is computed from the true and estimated states at each measurement update time for all Monte Carlo simulations. The filters' complexity is represented by their average execution time per filtering run in a C++ implementation on a 3.2 GHz single-core Ubuntu operating system. The filters' consistency is examined using the scaled normalized estimation error squared (SNEES) β^R which is defined as follows:

$$\beta_k^R = \frac{1}{Md} \sum_{j=1}^{M} (\boldsymbol{x}_k^{(j)} - \hat{\boldsymbol{x}}_k^{(j)})^{\mathrm{T}} (\boldsymbol{P}_k^{(j)})^{-1} (\boldsymbol{x}_k^{(j)} - \hat{\boldsymbol{x}}_k^{(j)}) \tag{33}$$

where M is the number of Monte Carlo simulations, $\boldsymbol{x}_k^{(j)}$ are the true states, $\hat{\boldsymbol{x}}_k^{(j)}$ are the estimated states, $\boldsymbol{P}_k^{(j)}$ are the filter's estimated error covariance matrix of the j-

th Monte Carlo run at the time step k. The size of the state space $d = 6$ is used to scale the NEES value (Bar-Shalom et al., 2001) such that a consistent filter will result in a SNEES of one rather than a NEES of d. If the SNEES value is much greater than 1, it means the estimator is divergent; however, if the value is much smaller than 1, it indicates the estimator is overly conservative. When the estimator is consistent, SNEES should be nearly one at all times.

5.2. Simulation Results: 1-orbit interval period

The time history of the RMS position errors of the 100 simulations is depicted in Figure 3 and the position's RMSE values of each filter are listed in Table 3. Due to their nonlinear nature, AEGIS and EnGMF provide better performance than the UKF at the very first measurement update. However, in this measurement-rich environment, equinoctial UKF performs near the top in accuracy after the initial transient, and it is the most consistent at a small fraction of the computational cost of nonlinear filters. From the results, it is also shown that the UKF and EnGMF with equinoctial orbital elements outperforms the corresponding filter with Cartesian coordinates. Nevertheless, the best performance in terms of estimation accuracy is obtained with the AEGIS, closely followed by the equinoctial UKF.

Table 3: Monte Carlo averaged RMSE, SNEES, and computation time for 100 simulations

	Position's RMSE (km)	SNEES	Computation time (sec)
UKF (Cartesian)	0.2212	507.4479	3.36
UKF (Equinoctial)	0.1839	1.2425	3.55
AEGIS-3 (Cartesian)	0.1810	1.5058	460.34
EnGMF (Cartesian)	0.3320	0.4986	189.47
EnGMF (Equinoctial)	0.3284	0.4920	190.49

Figure 4 shows the SNEES value for 100 Monte Carlo simulations and the average computation time per filtering run for all the filters. In Figure 4(a), the SNEES value of the EnGMF is smaller than 1, which means the EnGMF is conservative. For the EnGMF, the covariance matrix calculated by Silverman's rule for each Gaussian component is over-smoothed since the density is not truly Gaussian. The value of the

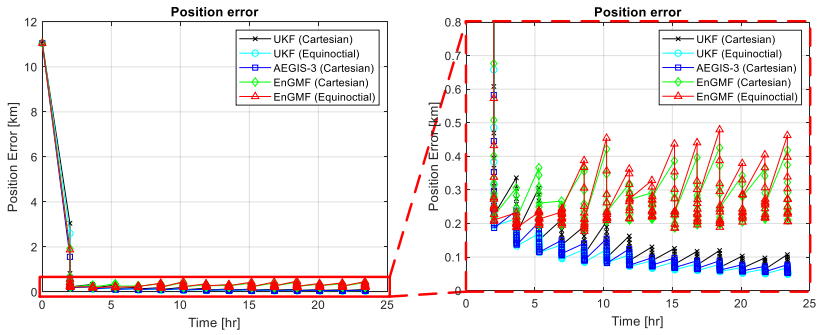

Figure 3: The average RMSE for 100 Monte Carlo simulations, 1 pass per orbit

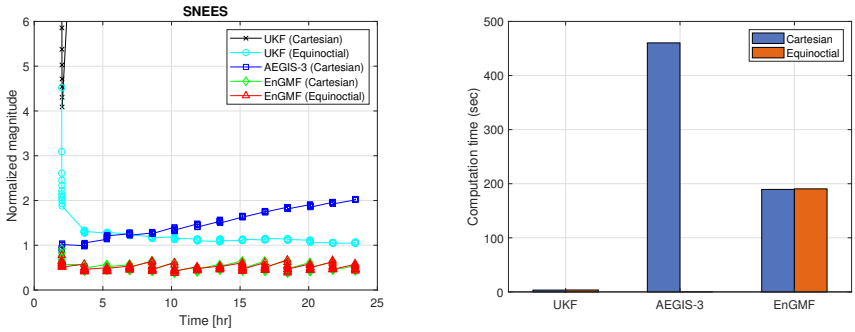

(a) The SNEES value

(b) The average computation time per filtering run

Figure 4: The SNEES value and the average computation time per filtering run for 100 Monte Carlo simulations

AEGIS filter is gradually increased starting from the value 1. The UKF works better when using equinoctial orbital elements than when using Cartesian coordinates. When the UKF uses Cartesian coordinates, it diverges in two out of 100 simulations, which means the estimation error widely exceeds the ± 3 sigma predicted standard deviations of the posterior covariance matrix. As is typical for linear estimators processing nonlinear measurements without underweighting (Zanetti et al., 2010), the equinoctial UKF is overly optimistic in processing the very first batch of measurements, but due to the measurement-rich scenario, it recovers nicely and achieved very good consistency.

The time-averaged SNEES value for 100 cases is listed in Table 3. The average computation time is also presented in Table 3. In terms of computation time, the best performance is obtained with the UKF by a wide margin (as expected from a simple

linear filter), and the EnGMF reduces the mean computation time by 58.73% in comparison with the AEGIS. Note that resetting the GMM in AEGIS to a single component after each measurement pass greatly reduces its computational cost when compared to other merging/pruning schemes.

5.3. Simulation Results: 1- to 6-orbit interval periods

Having established the baseline performance of the estimators with one measurement pass per orbital period, we focus on the real challenge addressed by this paper: scarcity of measurements. Additional simulations are performed when the gap between measurement passes is increased to 2, 3, 4, 5, and 6 orbital periods. As in the previous case, a Monte Carlo analysis is performed with 100 simulations. As we are concerned with computational speed, we set a maximum allowable number of GMM components for the AEGIS to be 1000 to contain its overall run time; note that relaxing this constraint will result in a more accurate, but slower filter.

Figure 5(a) displays the position RMSE of the UKF, AEGIS, and EnGMF in all the six cases and the Monte Carlo averaged RMS position errors for all cases are listed in Table 4, where n/a indicates the filter has diverged. The time-averaged SNEES value of each estimator for all the six cases is displayed in Figure 5(b). The AEGIS outperforms the EnGMF with Cartesian coordinates in terms of RMS accuracy for all the six cases. However, the EnGMF with the equinoctial orbital elements provides better estimation accuracy than the AEGIS when the interval time between measurement passes is 6 orbital periods. Also note that the RMS position error of the AEGIS increases more rapidly with the orbital periods than the EnGMF as shown in Figure 5(a). While the equinoctial UKF provides excellent performance for the one-orbit interval period, its performance is severely degraded in terms of accuracy and consistency for the two- and three-orbits case, and is completely diverging for 4–6 orbital periods between measurement passes. The UKF uses the Unscented Transformation (UT) to approximate mean and covariance across nonlinear transformations: the dynamics model and the measurement model. In the presence of very large prior uncertainties, the nonlinearities of range, range rate and bearing angles are sufficient to make the UT inaccurate and the UKF diverge.

The choice of using Silverman's rule in the EnGMF rather than performing bandwidth optimization is a trade between speed and accuracy/consistency. However, since the choice results in a conservative filter (estimated covariance larger than actual one) this trade off is deemed worthy when the goal is to maintain custody of a very high number of SOs. The EnGMF implementation in the equinoctial elements provides increasingly better performance than the EnGMF with Cartesian coordinates as the interval between measurement passes is increased. The UKF with Cartesian coordinates and equinoctial orbital elements diverge when the interval time is more than 3 and 4 orbital periods, respectively.

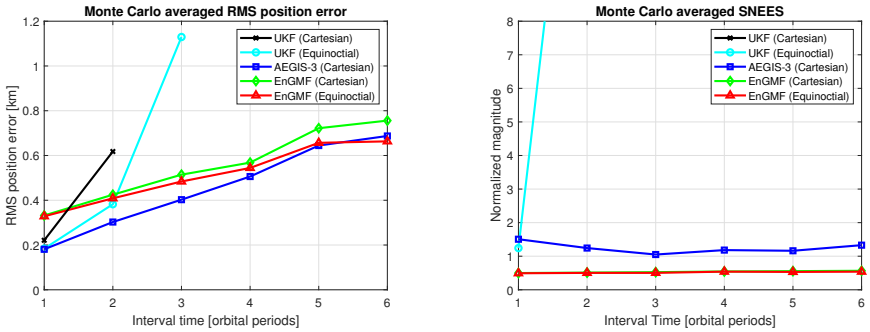

(a) Monte Carlo averaged RMSE (b) Monte Carlo averaged SNEES value

Figure 5: Monte Carlo averaged RMSE and SNEES value for all the six cases

Table 4: Monte Carlo averaged RMS position errors for all the six cases

	1	2	3	4	5	6
UKF (Cartesian)	0.2212	0.6178	n/a	n/a	n/a	n/a
UKF (Equinoctial)	0.1839	0.3823	1.1294	n/a	n/a	n/a
AEGIS-3 (Cartesian)	0.1810	0.3027	0.4026	0.5060	0.6445	0.6872
EnGMF (Cartesian)	0.3320	0.4248	0.5144	0.5682	0.7219	0.7559
EnGMF (Equinoctial)	0.3284	0.4086	0.4838	0.5444	0.6566	0.6632

Figure 6 shows the computation time of each filter, which is normalized by the value for the EnGMF with Cartesian coordinates. Compared to the AEGIS filter, the EnGMF reduces the computation time by 59.91% on average.

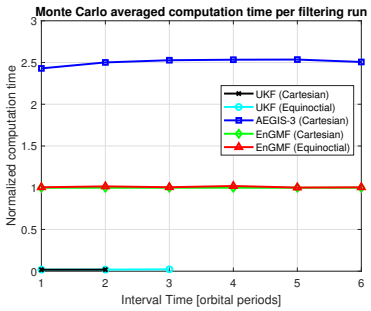

Figure 6: Monte Carlo averaged computation time per filtering run

5.4. Simulation Results: 6-orbit interval periods

In this section the performance of the proposed algorithm is evaluated for the case when measurement passes are taken every six orbits but under more challenging circumstances than the prior analysis. Firstly, to demonstrate that the proposed algorithm is able to handle parametric uncertainty in the dynamics, error is added to the drag and reflectivity coefficients. The drag coefficient C_D is typically a value between 2 and 4, depending on shape, altitude, and molecular content (Vallado and Finkleman, 2014) and the reflectivity coefficient C_R between 0.0 and 2.0 (Vallado, 2013). In the following analysis, the true values of the drag and reflectivity coefficients are corrupted with three levels of zero-mean errors having standard deviations set at 0.1, 0.2, and 0.3. The resulting time history of the RMS position errors of the 100 simulations is shown in Figure 7 and the position's RMSE values of the filters with different uncertain level of the drag and reflectivity coefficients are listed in Table 5. The EnGMFs implementation used is that with equinoctial orbital elements. The results show that the performance of the EnGMF degrades as the noise level of the parameters C_D and C_R increases (as expected); however, the EnGMFs is able to handle this additional error source as shown in Figure 8.

Table 5: Monte Carlo averaged RMS position errors for the EnGMF with equinoctial orbital elements and C_D and C_R uncertainty

StdDev	0	0.1	0.2	0.3
RMSE	0.6632	0.6678	0.6734	0.6850

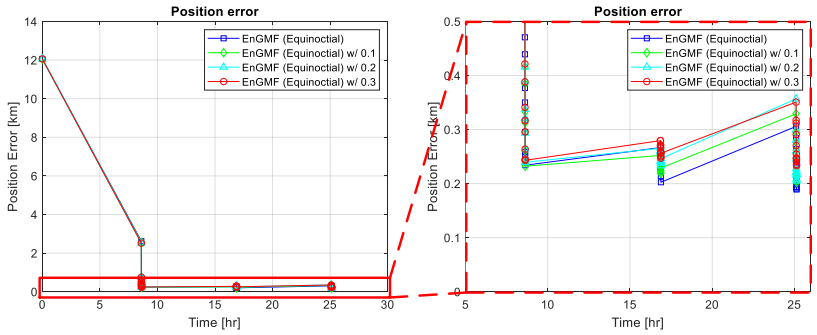

Figure 7: The average RMSE for 100 Monte Carlo simulations, six-orbits case with C_D and C_R uncertainty

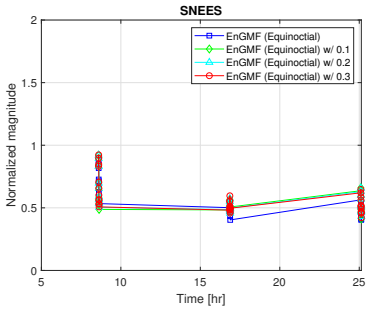

Figure 8: The SNEES value for 100 Monte Carlo simulations, the six-orbits case with C_D and C_R uncertainty

Secondly, the performance of the equinoctial EnGMF is evaluated with an angles-only measurements (i.e., the only measurement available are right ascension and declination while range and range rate are not available in this scenario). The interval between measurement passes is again set to six-orbits. Figure 9 compares the time history of the averaged RMSE and SNEES values of the equinoctial EnGMFs with all four types of measurement (i.e., range, rage-rate, right ascension, and declination) versus using only angle measurement. Figure 9(a) shows that with only angle measurement the performance of the EnGMF degrades compared to using all four measurement types (as expected); however, Figure 9(b).SNEES remains almost unchanged and hence the EnGMF remains a slightly conservative filter also in the angles-only case.

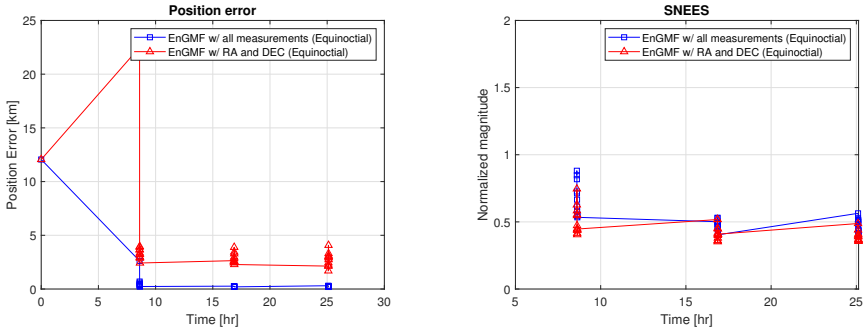

(a) The averaged RMSE

(b) The SNEES value

Figure 9: The averaged RMSE and SNEES value for 100 Monte Carlo simulations, six-orbit case with only angle measurements

5.5. Simulation Results: 10-orbit interval periods

A more in-depth comparison of the performances of various filters is shown for the challenging scenario of one measurement pass every ten orbits. Since the UKF fails under this challenging scenario, the EnGMF and AEGIS are compared to the EnKF instead. Figures 10 and 11 present the time history of the RMS position errors, SNEES values, and average computation time per filtering run of the EnGMF, AEGIS, and EnKF, and each value is also shown in Table 6. In terms of accuracy, the EnGMF with Cartesian coordinates or the equinoctial orbital elements outperforms the EnKF and AEGIS method. This is because the EnKF (Cartesian or equinoctial orbital elements) and AEGIS filters diverge in two and seven out of 100 simulations respectively, as shown in Figure 11(a), whereas the EnGMF is conservative. As expected, Figure 11(b) shows that the computation time of the EnGMF is comparable with that of the EnKF, since they share the same time propagation algorithm which requires the majority of the computations. Moreover, the EnGMF reduces the computation time by 60.57% compared to the AEGIS filter. In Table 6, we can see that the EnGMF using the equinoctial orbital elements obtains the best performance in terms of accuracy and mean computation time.

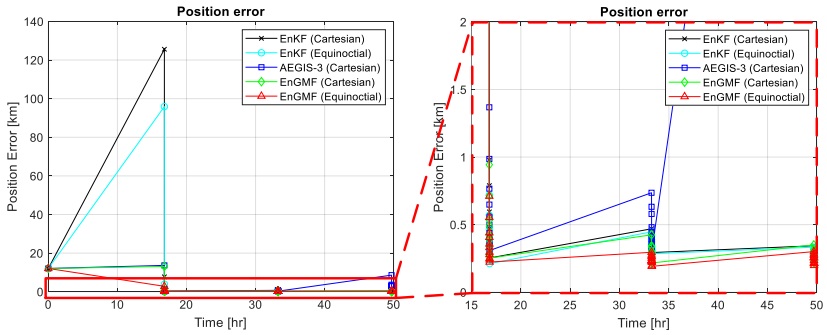

Figure 10: The average RMSE for 100 Monte Carlo simulations, the ten-orbits case

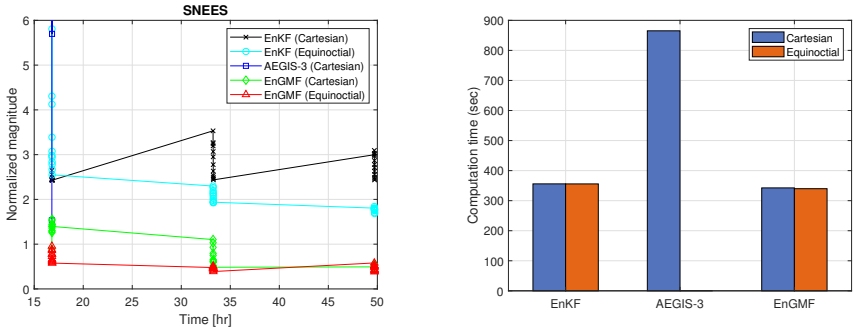

(a) The SNEES value

(b) The average computation time per filtering run

Figure 11: The SNEES value and the average computation time per filtering run for 100 Monte Carlo simulations, the ten-orbits case

Table 6: Monte Carlo averaged RMSE, SNEES, and computation time for the ten-orbits case

	Position's RMSE (km)	SNEES	Computation time (sec)
EnKF (Cartesian)	4.2699	2.7008	355.85
EnKF (Equinoctial)	3.3328	2.6713	355.74
AEGIS-3 (Cartesian)	2.1578	1.6459e+06	865.17
EnGMF (Cartesian)	0.9930	0.8595	342.46
EnGMF (Equinoctial)	0.6688	0.5504	339.86

5.6. Simulation Results: 20-orbit interval periods

The performance of the estimators worsen as the gap between measurement passes increases until they eventually diverges. To evaluate the performance of the adaptive EnGMF with bi-fidelity propagation under an even sparser measurement data condi-

tion, a Monte Carlo analysis is performed with 200 simulations when the gap between measurement passes is increased to 20 orbital periods. The value of 20 orbital periods is chosen because it causes two divergence cases out of 200 runs when 1000 particles are used. The initial number of particles of the adaptation is set to $N_0 = 250$ and the maximum number of particles is $N_{max} = 2000$. We use 5 fictitious observations $C = 5$ and 12 window size $W = 12$. The low-fidelity model is summarized in Table 7. The EnGMF with 2000 particles, the EnGMF with 2000 particles with bi-fidelity propagation, and the adaptive EnGMF with bi-fidelity propagation with different significance level of p-values, $p_l = [0.1, 0.15, 0.2, 0.25]$, are compared based on the accuracy, consistency, and complexity. The EnGMFs implementation is the one with equinoctial orbital elements in this 20-orbits scenario.

Table 7: Low-fidelity dynamic model for the numerical simulations

Model	Low-Fidelity
Central Body Gravity	Two-Body and J_2
Third-Body Perturbations	None
Atmospheric Drag	None
Solar Radiation Pressure	None

Figures 12 and 13 display the time history of and the time average of the RMS position errors, respectively, and the histogram of the number of particles in each window size used for the adaptive EnGMF with bi-fidelity propagation is shown in Figure 14. Figure 15 shows the time history of the SNEES values and average computation time of the 200 simulations. The result shows that, in the adaptive EnGMF with bi-fidelity propagation, as the significance level of p-values becomes higher, it provides improved accuracy with a larger average number of particles. Moreover, the algorithms are conservative as shown in Figure 15(a). As a result, it is demonstrated that the predictive observation PDF is effectively used as a convergence assessment of the EnGMF. The EnGMF with 2000 particles with or without bi-fidelity propagation outperforms any of the adaptive EnGMF with bi-fidelity propagation and the EnGMF with 2000 particles obtains the best performance in terms of accuracy. The EnGMF with 2000 particles, however, requires almost fourteen times the computation time of the EnGMF with 2000

particles with bi-fidelity propagation. The computational time can be reduced further by using the adaptive algorithm as shown in Figure 15(b). The time-averaged position's RMSE values, average number of particles, SNEES values, and computational times of each filter are listed in Table 8.

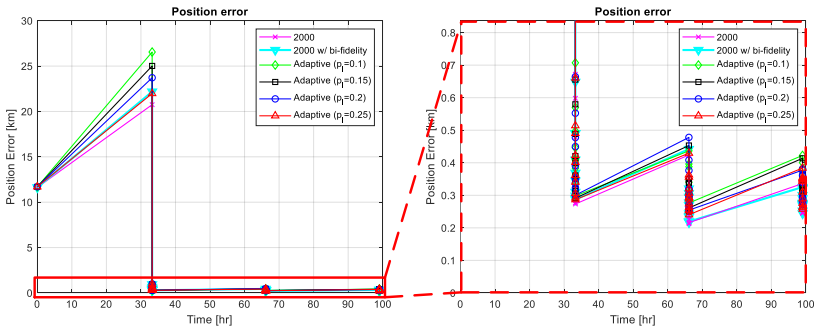

Figure 12: The average RMSE for 200 Monte Carlo simulations, the twenty-orbits case

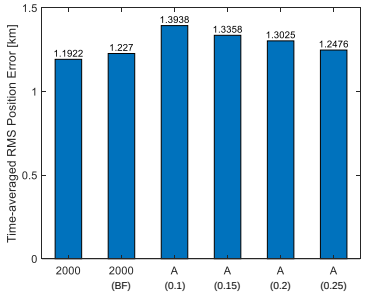

Figure 13: The time-averaged RMSE for 200 Monte Carlo simulations, the twenty-orbits case

Table 8: Monte Carlo averaged RMSE, SNEES, and computation time for the twenty-orbits case

	Position's RMSE (km)	Average number of particles	SNEES	Computation time (sec)
EnGMF-2000	1.1922	2000	0.6658	1430.5
EnGMF-2000 (bi-fidelity)	1.2270	2000	0.6308	105.5
Adaptive ($p_l = 0.1$)	1.3938	385.83	0.7797	41.76
Adaptive ($p_l = 0.15$)	1.3358	428.33	0.7409	45.29
Adaptive ($p_l = 0.2$)	1.3025	638.75	0.7227	56.33
Adaptive ($p_l = 0.25$)	1.2476	980.83	0.6796	80.33

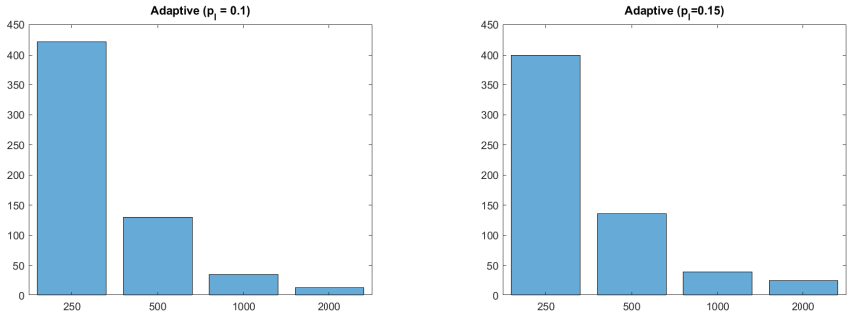

(a) $p_l = 0.1$ (b) $p_l = 0.15$

(c) $p_l = 0.2$ (d) $p_l = 0.25$

Figure 14: Histogram of the number of particles in each window size used for the adaptive EnGMF with bi-fidelity propagation with different significance level of p-values, $p_l = [0.1, 0.15, 0.2, 0.25]$

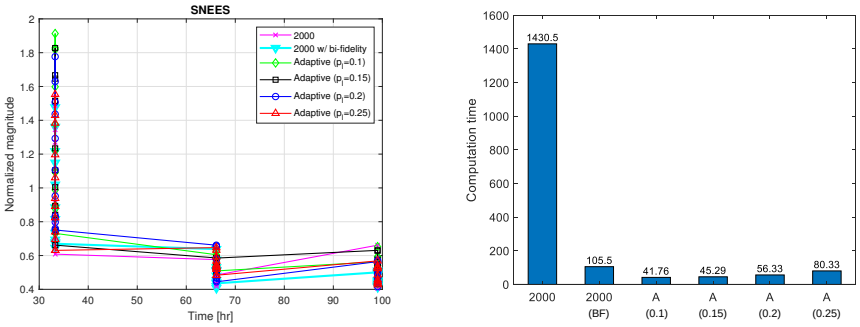

(a) The SNEES value (b) The average computation time per filtering run

Figure 15: The SNEES value and the average computation time per filtering run for 100 Monte Carlo Simulations, the twenty-orbits case

6. Conclusions

This paper studies a software-only solution to the orbit determination problem with sparse observation data. The motivation behind the study is the ability to maintain custody of a very large number of LEO objects. As such, it is of outmost importance in this study to strike a balance between estimation accuracy/consistency and computational burden of the methodology employed. Two linear filter implementations (unscented and ensemble Kalman filters) are shown to either be inadequate for very scarce measurement scenarios (UKF with measurement passes every three orbits or more regardless of the choice of Cartesian coordinates or equinoctial orbital elements) or to be less accurate at the same computational cost (EnKF). A state-of-the-art Gaussian sum filter named AEGIS is shown to perform well at a very high computational cost, but it fails when the number of Gaussian components is artificially capped in order to contain its total execution time.

A newly proposed approach is a modification of the kernel-based ensemble Gaussian mixture filter. Each propagated sample of the prior distribution is treated as a Gaussian component with a non-zero covariance matrix. The covariance matrix of a Gaussian component is calculated with Silverman's rule of thumb to reduce the computational cost of numerically optimizing a bandwidth parameter. The rule produces the optimal bandwidth when the samples are drawn from a Gaussian distribution, and results in a conservative estimate for non-Gaussian distributions. Numerical simulations show that the modified algorithm is more accurate and/or faster than the other approaches for sparse measurement scenarios. The conservatism inherent from using Silverman's rule cannot cause filter divergence but can result in slight loss of accuracy. This slight loss of accuracy is deemed an acceptable trade off to computational efficiency for the ultimate purpose of this work: tracking a very large number of space objects. While this conservatism can potentially trigger false collision alarms, an efficient strategy to maintaining a very large catalog is using the proposed lower complexity and conservative estimates for the population at-large and to only focus high precision and computationally expensive orbit determination solutions for the very small subset of objects that are deemed at risk for collision.

Finally, the adaptive kernel-based ensemble Gaussian mixture filter with bi-fidelity propagation was designed and shown to provide good performance while reducing the computational time. The bi-fidelity approach leverages both high- and low-fidelity models to reduce computational burden with an acceptable loss in accuracy for long time propagation. The adaptive algorithm determines the appropriate number of particles based on the convergence assessment using a predictive observation probability density function.

Acknowledgment

This work was sponsored in part by DARPA (Defense Advanced Research Projects Agency) under STTR contract number 140D0420C0062 and partially supported by the Air Force Office of Scientific Research under award number FA9550-19-1-0404.

References

Alspach, D., Sorenson, H., 1972. Nonlinear bayesian estimation using gaussian sum approximations. IEEE Transactions on Automatic Control 17, 439–448. doi: 10.1109/SAP.1970.270017.

Anderson, J.L., Anderson, S.L., 1999. A monte carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. Monthly Weather Review 127, 2741–2758. doi:10.1175/1520-0493.

Arulampalam, M., Maskell, S., Gordon, N., Clapp, T., 2002. A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking. IEEE Transactions on Signal Processing 50, 174–188. doi:10.1109/78.978374.

Arulampalam, S., Gordon, N., Ristic, B., 2004. Beyond the Kalman Filter: Particle Filters for Tracking Applications Partical Filters for Tracking Applications. Artech House.

Bar-Shalom, Y., Li, X.R., Kirubarajan, T., 2001. Estimation with Applications to Tracking and Navigation: Theory Algorithms and Software. Wiley.

Broucke, R.A., Cefola, P.J., 1972. On the equinoctial orbit elements. Celestial Mechanics 5, 303–310. doi:10.1007/BF01228432.

Castaings, T., Pannetier, B., Muller, F., Rombaut, M., 2015. Track initiation in low-earth-orbit objects using statistical modeling of sparse observations. IEEE Transactions on Aerospace and Electronic Systems 51, 258–269. doi:10.1109/TAES.2014.130484.

DeMars, K.J., Bishop, R.H., Jah, M.K., 2013. Entropy-based approach for uncertainty propagation of nonlinear dynamical systems. Journal of Guidance, Control, and Dynamics 36, 1047–1057. doi:10.2514/1.58987.

DeMars, K.J., Jah, M.K., Schumacher, P.W., 2012. Initial orbit determination using short-arc angle and angle rate data. IEEE Transactions on Aerospace and Electronic Systems 48, 2628–2637. doi:10.1109/TAES.2012.6237613.

Dormand, J.R., Prince, R.J., 1980. A family of embedded Runge-Kutta formulae. Journal of Computational and Applied Mathematics 6, 19–26.

Doucet, A., de Freitas, N., Gordon, N., 2001. Sequential Monte Carlo Methods in Practice. Information Science and Statistics, Springer–Verlag, New York.

Elvira, V., Miguez, J., Djuric, P.M., 2017. Adapting the number of particles in sequential monte carlo methods through an online scheme for convergence assessment. IEEE Transactions on Signal Processing 65, 1781–1794. doi:10.1109/TSP.2016.2637324.

Folkner, W.M., Williams, J.G., Boggs, D.H., *at al.*, 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.

Gelb, A. (Ed.), 1974. Applied Optimal Estimation. The MIT press, Cambridge, MA.

Ghanem, R.G., Higdon, D., Owhadi, H., 2017. Handbook of Uncertainty Quantification. Springer International Publishing, Switzerland.

Horwood, J.T., Aragon, N.D., Poore, A.B., 2012. Gaussian sum filters for space surveillance: Theory and simulations. Journal of Guidance, Control, and Dynamics 34, 1839–1851. doi:10.2514/1.53793.

Horwood, J.T., Poore, A.B., 2011. Adaptive gaussian sum filters for space surveillance. IEEE Transactions on Automatic Control 56, 1777–1790. doi:10.1109/TAC.2011.2142610.

Hughes, P.J., Alfriend, K.T., 2020. Covariance based track association with modified equinoctial elements, in: Proceedings of the 2020 AAS/AIAA Astrodynamics Specialist Conference. Paper No. AAS 20-681.

Jones, B.A., Weisman, R., 2019. Multi-fidelity orbit uncertainty propagation. Acta Astronautica 155, 406–417. doi:10.1016/j.actaastro.2018.10.023.

Jones, M.C., Marron, J.S., Sheather, S.J., 1996. A brief survey of bandwidth selection for density estimation. Journal of the American Statistical Association 91, 401–407. doi:10.2307/2291420.

Julier, S.J., Uhlmann, J.K., 2004. Unscented filtering and nonlinear estimation. Proceedings of the IEEE 92, 401–422. doi:10.1109/JPROC.2003.823141.

Junkins, J.L., Akella, M.R., Alfriend, K.T., 1996. Non-gaussian error propagation in orbital mechanics. Journal of the Astronautical Sciences 44, 541–563.

Kalman, R.E., 1960. A new approach to linear filtering and prediction problems. Journal of Basic Engineering 82, 35–45. doi:10.1115/1.3662552.

Katzfuss, M., Stroud, J.R., Wikle, C.K., 2016. Understanding the ensemble kalman filter. The American Statistician 70, 350–357. doi:10.1080/00031305.2016.1141709.

Liu, B., Ait-El-Fquih, B., Hoteit, I., 2016. Efficient kernel-based ensemble gaussian mixture filtering. Monthly Weather Review 144, 781–800. doi:10.1175/MWR-D-14-00292.1.

Moral, P.D., 2004. Feynman-Kac Formulae: Genealogical and Interacting Particle Systems With Applications. Springer, New York, NY, USA. doi:10.1007/978-1-4684-9393-1.

Narayan, A., Gittelson, C., Xiu, D., 2014. A stochastic collocation algorithm with multifidelity models. SIAM Journal on Scientific Computing 36, A495–A521. doi:10.1137/130929461.

National Research Council, 2012. Continuing Kepler's Quest: Assessing Air Force Space Command's Astrodynamics Standards. The National Academies Press.

Park, B.U., Marron, J.S., 1990. Comparison of data-driven bandwidth selectors. Journal of the American Statistical Association 85, 66–72. doi:10.1080/01621459.1990.10475307.

Pate-Cornell, E., Sachon, M., 2001. Risks of particle hits during space walks in low earth orbit. IEEE Transactions on Aerospace and Electronic Systems 37, 134–146. doi:10.1109/7.913673.

Pavlis, N.K., Holmes, S.A., Kenyon, S.C., Factor, J.K., 2008. An Earth gravitational model to degree 2160: EGM2008, in: Proceedings of the European Geosciences Union General Assembly, Vienna, Austria.

Plackett, R.L., 1983. Karl pearson and the chi-squared test. International Statistical Review 51, 59–72. doi:10.2307/1402731.

Raihan, D., Chakravorty, S., 2018. Particle gaussian mixture filters-i. Automatica 98, 331–340. doi:10.1016/j.automatica.2018.07.023.

Sabol, C., Sukut, T., Hill, K., *at al.*, 2010. Linearized orbit covariance generation and propagation analysis via simple monte carlo simulations, in: 20th Annual AAS/AIAA Space Flight Mechanics Meeting, San Diego, California.

Scott, D.W., 1992. Multivariate Density Estimation: Theory, Practice, and Visualization. John Wiley and Sons, New Jersey, NJ. doi:10.1002/0470045345.

Silverman, B.W., 1986. Density Estimation for Statistics and Data Analysis. Chapman and Hall, London.

Sorenson, H.W., Alspach, D.L., 1971. Recursive Bayesian Estimation Using Gaussian Sums. Automatica 7, 465–479. doi:10.1016/0005-1098(71)90097-5.

Tapley, B.D., Schutz, B.E., Born, G.H., 2004. Statistical Orbit Determination. Elsevier Academic Press.

Terejanu, G., Singla, P., Singh, T., Scott, P.D., 2011. Adaptive gaussian sum filter for nonlinear bayesian estimation. IEEE Transactions on Automatic Control 56, 2151–2156. doi:10.1109/TAC.2011.2141550.

Tuggle, K., Zanetti, R., 2018. Automated splitting gaussian mixture nonlinear measurement update. Journal of Guidance, Control, and Dynamics 41, 725–734. doi:10.2514/1.G003109.

Vallado, D.A., 1997. Fundamentals of Astrodynamics and Applications, 3rd Edition. McGraw-Hill, New Jersey, NJ.

Vallado, D.A., 2013. Fundamentals of Astrodynamics and Applications, 4th ed. Microcosm Press.

Vallado, D.A., Finkleman, D., 2014. A critical assessment of satellite drag and atmospheric density modeling. Acta Astronautica 95, 141–165. doi:10.1016/j.actaastro.2013.10.005.

Wan, E.A., Merwe, R.V.D., 2000. The unscented kalman filter for nonlinear estimation. Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium doi:10.1109/ASSPCC.2000.882463.

Weedon, B.C., Cefola, P.J., 2010. Computer Systems and Algorithms for Space Situational Awareness: History and Future Development', in: Advances in the Astronautical Sciences, pp. 205–226.

Yun, S., Zanetti, R., 2019. Sequential monte carlo filtering with gaussian mixture sampling. Journal of Guidance, Control, and Dynamics 42, 2069–2077. doi:10.2514/1.G004403.

Yun, S., Zanetti, R., 2020. Nonlinear filtering of light-curve data. Advances in Space Research 66, 1672–1688. doi:10.1016/j.asr.2020.06.024.

Zanetti, R., DeMars, K.J., Bishop, R.H., 2010. Underweighting nonlinear measurements. Journal of Guidance, Control, and Dynamics 33, 1670–1675. doi:10.2514/1.50596.

Zhu, X., Narayan, A., Xiu, D., 2014. Computational aspects of stochastic collocation with multifidelity models. SIAM/ASA Journal on Uncertainty Quantification 2, 444–463. doi:10.1137/130949154.

Zucchelli, E., Delande, E., Jones, B., Jah, M., 2021. Multi-fidelity orbit uncertainty propagation. The Journal of the Astronautical Sciences 68, 695–727. doi:10.1007/s40295-021-00267-y.