AAS 23-386

BAYESIAN ANGLES-ONLY CISLUNAR SPACE OBJECT TRACKING

Keith A. LeGrand*, Aneesh V. Khilnani†, and John L. Iannamorelli†

Angles-only orbit determination involves the estimation of a satellite's unknown orbit without access to range or velocity measurements. Angles measurements, which are typically passive, are the primary means for tracking distant orbits, where active sensing is energy-prohibitive. This paper presents a Bayesian approach for tracking objects in the cislunar regime using noisy angles measurements. The proposed filter addresses challenges in the cislunar setting, which include dynamics-induced rapid uncertainty growth and long periods of non-observation. The filter is shown to be effective over a range of L_2 halo orbit determination problems and to significantly outperform a central-moment based filter.

INTRODUCTION

Space mission trends in recent years show increasing interest and activity in cislunar space. While the growing number of planned missions in this orbital regime reflects the great potential of cislunar space in commercial and military enterprises, it has also raised concerns about the inadequacy of existing infrastructure to support these operations. In particular, the well-established approaches to space domain awareness (SDA) for the low Earth orbit (LEO) and geostationary Earth orbit (GEO) regimes are ill-suited for the cislunar region due to the vast distances, chaotic dynamics, and extremely long periods of space object non-detectability.[1]

Recent work in cislunar space object tracking has considered the use of the deep space network (DSN) for tracking objects in halo orbits.[2,3] DSN based tracking is a highly attractive option for special uses, such as tracking crewed spacecraft, due to the extremely high accuracy measurements it provides and the availability of multiple ground stations across the globe. On the other hand, the use of active radio frequency (RF) measurements for general cislunar SDA would require prohibitive amounts of energy.[1] Space-based observers may play a key role in future cislunar SDA, as they can increase sensing coverage, geometric diversity of observations, and observation opportunities.[4–6]

Work in [7] explored the observation opportunities of ground-based telescopes and showed the complicated uncertainty structures that emerge during extended periods of non-observation. On the other hand, certain families of multi-body orbits such as 2:1 resonant orbits can result in nearly Gaussian uncertainty propagation, in which case central-moment filters, such as the unscented Kalman filter (UKF)[8] may suffice for tracking.[5]

*Assistant Professor, School of Aeronautics and Astronautics, Purdue University, 701 W Stadium Ave., West Lafayette, IN, United States

†Graduate Research Assistant, School of Aeronautics and Astronautics, Purdue University, 701 W Stadium Ave., West Lafayette, IN, United States

Little existing work has been dedicated to probabilistic–that is, Bayesian–orbit determination of halo orbits.[9] In principle, Bayesian estimation approaches, which operate on probability density function (pdf) representations of uncertainty, can maintain accurate belief representations even when the underlying dynamics and observations result in non-Gaussian distributions. For instance, Gaussian mixture (GM) filters have been successfully employed in traditional SDA applications, and can be used to approximate any arbitrary pdf with a finite number of discontinuities.[10–12] Despite these advancements, the cislunar angles-only SDA remains a challenge due to the underlying multi-body dynamics and observational sparsity.

This paper introduces a new adaptive Bayesian filter based on recent advances in GM filtering using recursive Gaussian splitting techniques.[13,14] By identifying GM mixture components that overlap discrete field-of-view (FoV) bounds and splitting them into a finer resolution mixture, the negative information content of non-detection can be leveraged to further refine the object state pdf. This paper presents the first application of the splitting approach to 3D FoVs and shows how the approach can be applied to the incorporation of nonlinear state constraints. Splitting is also employed in the processing of angles measurements to reduce linearization error and improve estimation performance. In order to capture the rapid nonlinear uncertainty growth between observations, the adaptive entropy-based Gaussian-mixture information synthesis (AEGIS) approach[15,16] is employed to automatically split GM components based on the predicted performance of their local linearization approximation. Additionally, this paper presents a new prediction-stage split criterion based on Jacobi constant variance. The resulting framework is shown to successfully maintain accurate uncertainty representations in a challenging series of L_2 halo orbit determination problems.

PROBLEM FORMULATION

This paper investigates the problem of tracking an object in cislunar space that is subjected to multi-body dynamics. In particular, we consider the problem of estimating the unknown state of a satellite in an Earth-Moon halo orbit using only sparse noisy angles measurements from an Earth-based observer. The space-object state $\mathbf{x} \in \mathbb{R}^n$ is defined as

$$\mathbf{x}^T(t) = \begin{bmatrix}\mathbf{r}^T(t) & \mathbf{v}^T(t)\end{bmatrix} = \begin{bmatrix}x & y & z & \dot{x} & \dot{y} & \dot{z}\end{bmatrix}^T \tag{1}$$

where $n = 6$, $\mathbf{r} \in \mathbb{R}^3$ denotes the relative position of the space object with respect to the Earth-Moon barycenter and expressed in the synodic frame, and $\mathbf{v} \in \mathbb{R}^3$ denotes the relative velocity of the space object as seen in the rotating synodic frame. Throughout this paper, vectors and vector-valued functions are represented by bold lower-case letters, and matrices are represented by bold upper-case letters. Although in reality, the Earth-Moon orbit is slightly eccentric, this paper adopts the circular restricted three-body problem (CR3BP) assumptions. Denote by $m_\oplus$ and $m_\leftmoon$ the mass of the Earth and Moon, respectively, and let

$$\mu \triangleq \frac{m_\leftmoon}{m_\oplus + m_\leftmoon} \tag{2}$$

Then, the CR3BP equations of motion (EOMs) are given by

$$\ddot{x} - 2\dot{y} = \frac{\partial U}{\partial x}, \qquad \ddot{y} + 2\dot{x} = \frac{\partial U}{\partial y} \qquad \ddot{z} = \frac{\partial U}{\partial z} \tag{3}$$

$$U = \frac{1-\mu}{r_{/\oplus}} + \frac{\mu}{r_{/\leftmoon}} + \frac{x^2+y^2}{2} \tag{4}$$

where U is the pseudo-potential, $r_{/\oplus}$ is the distance between the satellite and the Earth center of mass, and $r_{/\leftmoon}$ is the distance between the satellite and the Moon center of mass. The CR3BP relative dynamics admit an integral of motion known as the Jacobi constant, given as

$$C = C(\mathbf{x}) = 2U - (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) \tag{5}$$

The Jacobi constant remains constant in the absence of perturbations, and thus can be used to evaluate numerical stability in orbit propagation, among other uses.

This paper considers probabilistic orbit determination using angles measurements from ground-based sensors. The sensor state $\boldsymbol{\xi}$ consists of the sensor position and orientation and determines the bounded sensor FoV $\mathcal{S}(\boldsymbol{\xi})$, where $\mathcal{S}$ is modeled as a compact and bounded subset of $\mathbb{R}^3$. The satellite is detectable by the sensor if and only if the line of sight from the sensor to the satellite is not occluded, the satellite is properly illuminated, and the satellite position is within the sensor FoV. Accordingly, the probability of detection is a function of both the space object state and sensor state and defined as

$$p_{D,k}(\mathbf{x}, \boldsymbol{\xi}) = 1_{\mathcal{S}}(\mathbf{r}) \cdot \phi(\mathbf{x}, \boldsymbol{\xi}) \tag{6}$$

where the indicator function

$$1_B(\mathbf{a}) \triangleq \begin{cases} 1 & \mathbf{a} \in B \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

and where $\phi(\mathbf{x}, \boldsymbol{\xi}) = 1$ if the sensor-facing side of the satellite is illuminated and non-occluded, and zero otherwise. Note that more sophisticated models that consider visual magnitude[5,6] can be incorporated but are not considered in this paper.

It is assumed that prior knowledge of the object state is available in the form of a pdf $p_0(\mathbf{x})$ at some initial time t_0. At each observation time t_k, a sensor produces a measurement set Z_k that is either empty or contains a single vector measurement $\mathbf{z}_k$, depending on the probability of detection. This paper considers the case where $\mathbf{z}_k \in \mathbb{R}^2$ takes the form of a noisy right ascension/declination angle pair, which is related to the satellite and sensor state by the nonlinear relationship

$$\begin{bmatrix} \alpha_k & \delta_k \end{bmatrix}^T = \mathbf{z}_k = \mathbf{h}(\mathbf{x}_k, \boldsymbol{\xi}_k) + \boldsymbol{\nu}_k \tag{8}$$

where the additive measurement noise $\boldsymbol{\nu}_k \in \mathbb{R}^2$ is zero-mean, Gaussian, and white with covariance $\mathrm{E}[\boldsymbol{\nu}_k \boldsymbol{\nu}_k^T] = \mathbf{R}_k$. The problem considered in this paper is the formation of the prior and posterior distributions $p_{k|k-1}(\mathbf{x}_k | Z_{0:k-1})$ and $p_{k|k}(\mathbf{x}_k | Z_{0:k})$, respectively, at each observation step, so as to provide a probabilistic representation of a cislunar space object's orbit over time.

METHODOLOGY

This paper presents a novel Bayesian filtering algorithm to address the challenging angles-only cislunar space object tracking problem. The problem is characterized by strong nonlinearity in both the object dynamics and observation functions, which result in non-Gaussian state uncertainty and prohibit the direct adoption of central-moment based filters, such as the extended Kalman filter (EKF) or UKF. A key feature of the proposed approach is its ability to incorporate information from non-detections and nonlinear constraints and automatically adapt its approximation accuracy in key points of the pdf.

For simplicity, it is assumed that the object-to-measurement data association is known or the rate of false alarms is zero. In that case, the measurement random finite set (RFS) Z_k is characterized by the likelihood function

$$g_k(Z_k|\mathbf{x}_k) = \begin{cases} 1 - p_{D,k}(\mathbf{x}_k; \mathcal{S}) & Z_k = \varnothing \\ p_{D,k}(\mathbf{x}_k; \mathcal{S}) g_k(\mathbf{z}_k|\mathbf{x}_k) & Z_k = \{\mathbf{z}_k\} \\ 0 & |Z_k| > 1 \end{cases} \tag{9}$$

where $g_k(\mathbf{z}_k|\mathbf{x}_k)$ denotes the single-object measurement likelihood and is distinguished from the multi-object likelihood function by the nature of its vector measurement argument. With this, the Bayes filter recursion is given by the following prediction and measurement update steps:

$$p_{k|k-1}(\mathbf{x}_k|Z_{0:k-1}) = \int p_{k|k-1}(\mathbf{x}_k|\mathbf{x}_{k-1}) p_{k-1}(\mathbf{x}_{k-1}|Z_{0:k-1}) \mathrm{d}\mathbf{x}_{k-1} \tag{10}$$

$$p_k(\mathbf{x}_k|Z_{0:k}) = \frac{g_k(Z_k|\mathbf{x}_k) p_{k|k-1}(\mathbf{x}_k|Z_{0:k-1})}{\int g_k(Z_k|\mathbf{x}) p_{k|k-1}(\mathbf{x}|Z_{0:k-1}) \mathrm{d}\mathbf{x}} \tag{11}$$

where $p_{k|k-1}(\mathbf{x}_k|\mathbf{x}_{k-1})$ is the state transition density. Unfortunately, closed-form exact solutions of (10) and (11) are available only for special cases: namely, when $p_0(\mathbf{x})$ is Gaussian, the state transition and measurement likelihood are linear-Gaussian, and the probability of detection $p_{D,k}(\mathbf{x}_k; \mathcal{S}) = p_{D,k}$ is constant. Notably, the problem of angles-only cislunar space object tracking involves nonlinear measurements, nonlinear dynamics, and a state-dependent probability of detection, and therefore requires an approximation of the Bayes filter recursion.

Gaussian Mixture Bayes Filter

While central moment filtering approaches (e.g., the EKF and UKF) have shown to be effective in quasi-linear problems, they are severely diminished or otherwise inapplicable in truly nonlinear settings. Alternatively, probabilistic approaches maintain full descriptions of the filtering densities (as opposed to the first two central moments) and are applicable to a wider range of nonlinear estimation problems, such as the cislunar space object tracking problem. In particular, the GM is a highly flexible representation of a pdf and a universal function approximator for the class of pdfs with a finite number of discontinuities (which are of practical interest).

Consider the discrete-time nonlinear dynamics model given by

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}) + \mathbf{w}_{k-1}, \qquad \mathrm{E}[\mathbf{w}_{k-1}\mathbf{w}_{k-1}^T] = \mathbf{Q}_{k-1} \tag{12}$$

where the discrete-time dynamics model is obtained through integration of the continuous-time dynamics. The process noise $\mathbf{w}_k = \mathbf{w}(t_k)$ is taken to be zero-mean, Gaussian, and white. Let the posterior density $p_{k-1}(\mathbf{x}_{k-1}|Z_{0:k-1})$ be a GM

$$p_{k-1}(\mathbf{x}_{k-1}|Z_{0:k-1}) = \sum_{\ell=1}^{L_{k-1}} w_{k-1}^{(\ell)} \mathcal{N}(\mathbf{x}_{k-1}; \mathbf{m}_{k-1}^{(\ell)}, \mathbf{P}_{k-1}^{(\ell)}) \tag{13}$$

Then, by the Chapman-Kolmogorov equation (10), the prior pdf at t_k is the weighted sum

$$p_{k|k-1}(\mathbf{x}_k|Z_{0:k-1}) = \sum_{\ell=1}^{L_{k|k-1}} w_{k|k-1}^{(\ell)} \int \mathcal{N}(\mathbf{x}_k;\, \mathbf{f}(\mathbf{x}_{k-1}), \mathbf{Q}_{k-1}) \mathcal{N}(\mathbf{x}_{k-1}; \mathbf{m}_{k-1}^{(\ell)},\, \mathbf{P}_{k-1}^{(\ell)}) \mathrm{d}\mathbf{x}_{k-1} \tag{14}$$

For sufficiently small component covariances $\mathbf{P}_{k-1}^{(\ell)}$, the prior pdf can be accurately approximated by the GM[12]

$$p_{k|k-1}(\mathbf{x}_k|Z_{0:k-1}) \approx \sum_{\ell=1}^{L_{k|k-1}} w_{k|k-1}^{(\ell)} \mathcal{N}(\mathbf{x}_k; \mathbf{m}_{k|k-1}^{(\ell)}, \mathbf{P}_{k|k-1}^{(\ell)}) \tag{15}$$

where

$$w_{k|k-1}^{(\ell)} = w_{k-1}^{(\ell)} \tag{16}$$

$$\mathbf{m}_{k|k-1}^{(\ell)} = \int \mathbf{f}(\mathbf{x}_{k-1}) \mathcal{N}(\mathbf{x}_{k-1}; \mathbf{m}_{k-1}^{(\ell)}, \mathbf{P}_{k-1}^{(\ell)}) \mathrm{d}\mathbf{x}_{k-1} \tag{17}$$

$$\mathbf{P}_{k|k-1}^{(\ell)} = \int (\mathbf{f}(\mathbf{x}_{k-1}) - \mathbf{m}_{k|k-1}^{(\ell)})(\mathbf{f}(\mathbf{x}_{k-1}) - \mathbf{m}_{k-1}^{(\ell)})^T \mathcal{N}(\mathbf{x}_{k-1}; \mathbf{m}_{k-1}^{(\ell)}, \mathbf{P}_{k-1}^{(\ell)}) \mathrm{d}\mathbf{x}_{k-1} \tag{18}$$

Closed-form analytic solutions of the integrals in (17) and (18) are not available in general but are readily approximated by linearization or sigma point approximations, as discussed in the following subsection.

A question of practical importance, especially as it relates to the propagation of uncertainty under CR3BP dynamics is, what constitutes a sufficiently small component covariance $\mathbf{P}_{k-1}^{(\ell)}$? Indeed, as $\mathbf{P}_{k-1}^{(\ell)} \to \mathbf{0}$, the right-hand side (RHS) of (15) approaches the true pdf $p_{k|k-1}(\mathbf{x}_k|Z_{0:k-1})$ uniformly in $\mathbf{x}_k$ and $\mathbf{z}_k$ [17, Ch. 8]. Another case in which (15) is exact is when the dynamics model $\mathbf{f}(\mathbf{x}_{k-1})$ is linear. Thus, for the purpose of approximating the prior pdf as a GM, the covariances $\mathbf{P}_{k-1}^{(\ell)}$ should be small enough that state deviations from $\mathbf{m}_{k-1}^{(\ell)}$ behave linearly over the propagation period $[t_{k-1}, t_k]$. This consideration is central to adaptive GM filtering, as discussed later in the paper.

Let the prior pdf $p_{k|k-1}(\mathbf{x}_k|Z_{0:k-1})$ be a GM. Then, given the measurement RFS Z_k, the posterior density is determined using Bayes' rule as

$$\begin{aligned} p_k(\mathbf{x}_k|Z_{0:k}) \propto & \delta_{\varnothing}(Z_k) \sum_{\ell=1}^{L_{k|k-1}} (1 - p_{D,k}(\mathbf{x}_k; \mathcal{S})) w_{k|k-1}^{(\ell)} \mathcal{N}(\mathbf{x}_k; \mathbf{m}_{k|k-1}^{(\ell)}, \mathbf{P}_{k|k-1}^{(\ell)}) \\ & + (1 - \delta_{\varnothing}(Z_k)) \sum_{\ell=1}^{L_{k|k-1}} p_{D,k}(\mathbf{x}_k; \mathcal{S}) w_{k|k-1}^{(\ell)} g_k(\mathbf{z}_k|\mathbf{x}_k) \mathcal{N}(\mathbf{x}_k; \mathbf{m}_{k|k-1}^{(\ell)}, \mathbf{P}_{k|k-1}^{(\ell)}) \end{aligned} \tag{19}$$

The RFS measurement representation and the corresponding state-dependent probability of detection function enable the incorporation of negative information into the posterior pdf, as seen in (19). In other words, states that were detectable at time t_k but not detected are "downweighted" according to their associated probability of detection. Let the measurement noise be zero-mean, Gaussian, white, and independent of the process noise such that the measurement likelihood function $g_k(\mathbf{z}_k|\mathbf{x}_k) = \mathcal{N}(\mathbf{z}_k; \mathbf{h}(\mathbf{x}_k), \mathbf{R}_k)$, where $\mathbf{R}_k$ is the measurement error covariance. Then, the posterior pdf is approximated as the GM

$$p_k(\mathbf{x}_k|Z_{0:k}) \approx \sum_{\ell=1}^{L_k} w_k^{(\ell)} \mathcal{N}(\mathbf{x}_k; \mathbf{m}_k^{(\ell)}, \mathbf{P}_k^{(\ell)}) \tag{20}$$

where

$$w_k^{(\ell)} \propto \begin{cases} w_{k|k-1}^{(\ell)}(1 - p_{D,k}(\mathbf{m}_{k|k-1}^{(\ell)}; \mathcal{S})) & Z_k = \varnothing \\ w_{k|k-1}^{(\ell)} p_{D,k}(\mathbf{m}_{k|k-1}^{(\ell)}; \mathcal{S}) q_k^{(\ell)}(\mathbf{z}_k) & Z_k = \{\mathbf{z}_k\} \end{cases} \tag{21}$$

$$\mathbf{m}_k^{(\ell)} = \begin{cases} \mathbf{m}_{k|k-1}^{(\ell)} & Z_k = \varnothing \\ \mathbf{m}_{k|k-1}^{(\ell)} + \mathbf{K}_k(\mathbf{z}_k - \hat{\mathbf{z}}_k^{(\ell)}) & Z_k = \{\mathbf{z}_k\} \end{cases} \tag{22}$$

$$\mathbf{P}_k^{(\ell)} = \begin{cases} \mathbf{P}_{k|k-1}^{(\ell)} & Z_k = \varnothing \\ \mathbf{P}_{k|k-1}^{(\ell)} - \mathbf{C}_k^{(\ell)}\mathbf{K}_k^{(\ell)T} - \mathbf{K}_k^{(\ell)}\mathbf{C}_k^{(\ell)T} + \mathbf{K}_k^{(\ell)}\mathbf{P}_z^{(\ell)}\mathbf{K}_k^{(\ell)T} & Z_k = \{\mathbf{z}_k\} \end{cases} \tag{23}$$

$$q_k^{(\ell)}(\mathbf{z}_k) = \mathcal{N}(\mathbf{z}_k;\, \hat{\mathbf{z}}_k^{(\ell)},\, \mathbf{P}_{z,k}^{(\ell)}) \tag{24}$$

$$\hat{\mathbf{z}}_k^{(\ell)} = \int \mathbf{h}(\mathbf{x}_k)\mathcal{N}(\mathbf{x}_k;\, \mathbf{m}_{k|k-1}^{(\ell)},\, \mathbf{P}_{k|k-1}^{(\ell)})\mathrm{d}\mathbf{x}_k \tag{25}$$

$$\mathbf{C}_k^{(\ell)} = \int (\mathbf{x}_k - \mathbf{m}_{k|k-1}^{(\ell)})(\mathbf{h}(\mathbf{x}_k) - \hat{\mathbf{z}}_k^{(\ell)})^T \mathcal{N}(\mathbf{x}_k;\, \mathbf{m}_{k|k-1}^{(\ell)},\, \mathbf{P}_{k|k-1}^{(\ell)})\mathrm{d}\mathbf{x}_k \tag{26}$$

$$\mathbf{P}_{z,k}^{(\ell)} = \int (\mathbf{h}(\mathbf{x}_k) - \hat{\mathbf{z}}_k^{(\ell)})(\mathbf{h}(\mathbf{x}_k) - \hat{\mathbf{z}}_k^{(\ell)})^T \mathcal{N}(\mathbf{x}_k;\, \mathbf{m}_{k|k-1}^{(\ell)},\, \mathbf{P}_{k|k-1}^{(\ell)})\mathrm{d}\mathbf{x}_k + \mathbf{R}_k \tag{27}$$

$$\mathbf{K}_k^{(\ell)} = \mathbf{C}_k^{(\ell)} \left(\mathbf{P}_{z,k}^{(\ell)}\right)^{-1} \tag{28}$$

The GM approximation of the posterior pdf is the result of several underlying approximations. The first of these approximations is the assumed Gaussian form of the measurement likelihood products $g_k(\mathbf{z}_k|\mathbf{x}_k)\mathcal{N}(\mathbf{x}_k;\, \mathbf{m}_{k|k-1}^{(\ell)},\, \mathbf{P}_{k|k-1}^{(\ell)})$, which are generally non-Gaussian under a nonlinear measurement function. Second, the predicted measurements, cross covariances, and innovations covariances (25)-(27) must be approximated via linearization, sigma point methods, or Monte Carlo integration. Third, the state-dependent probability of detection $p_{D,k}(\mathbf{x}_k; \mathcal{S})$ is replaced by a zeroth-order Taylor expansion about the GM means.[14]

In the proposed filter, the integrals in (17)-(18) and (25)-(27) are approximated using the square-root scaled unscented transform,[18] as discussed in the following subsection. A subtle but important observation is that exact knowledge of the aforementioned integrals does not guarantee that the GM approximations (15) and (19) are exact. This is due to the assumed Gaussian form of transformed components under time- and measurement-updates, which will be non-Gaussian in general. To further reduce the approximation error, the GM resolution can be increased in key regions to better capture non-Gaussian effects, as discussed later in the paper.

Gaussian Mixture Square-Root Sigma-Point Approximation

The nonlinear CR3BP dynamics and nonlinear measurement model prevent exact closed-form solution of the integrals in (17)-(18) and (25)-(27). Therefore, this paper employs sigma point approximations,[18,19] where a set of weighted samples are deterministically computed from each component and propagated through the nonlinear dynamics and measurement equations. In particular, this paper adopts the square-root scaled unscented transform.[18] In the square-root formulation, covariance matrices are factorized as $\mathbf{P}_k^{(\ell)} = \mathbf{S}_k^{(\ell)}\mathbf{S}_k^{(\ell)^T}$, where $\mathbf{S}_k^{(\ell)}$ is the lower-triangular Cholesky square-root factor of the full covariance matrix $\mathbf{P}_k^{(\ell)}$. The Cholesky factor formulation offers many

advantages, including reduced computational cost, improved numerical stability, and guaranteed positive semi-definiteness of the state covariance.[18]

The predicted means (17) and covariances (18) are approximated as follows:

$$\boldsymbol{\mathcal{X}}^{(\ell)}_{0:2n,k-1} = \begin{bmatrix} \mathbf{m}^{(\ell)}_{k-1} & \mathbf{m}^{(\ell)}_{k-1} + \gamma \mathbf{S}^{(\ell)}_{k-1} & \mathbf{m}^{(\ell)}_{k-1} - \gamma \mathbf{S}^{(\ell)}_{k-1} \end{bmatrix} \tag{29}$$

$$\dot{\boldsymbol{\mathcal{X}}}_i{}^{(\ell)}(t) = \boldsymbol{f}(\boldsymbol{\mathcal{X}}^{(\ell)}_i(t)), \qquad \boldsymbol{\mathcal{X}}^{(\ell)}_{i,k|k-1} = \boldsymbol{\mathcal{X}}^{(\ell)}_i(t_k) \tag{30}$$

$$\mathbf{m}^{(\ell)}_{k|k-1} = \sum_{i=0}^{2n} \omega^{(m)}_i \boldsymbol{\mathcal{X}}^{(\ell)}_{i,k|k-1} \tag{31}$$

$$\bar{\mathbf{S}}^{(\ell)}_{k|k-1} \leftarrow \mathtt{qr}\left\{ \begin{bmatrix} \sqrt{\omega^{(c)}_1} \left(\boldsymbol{\mathcal{X}}^{(\ell)}_{1:2n,k|k-1} - \mathbf{m}^{(\ell)}_{k|k-1} \right) & \sqrt{\mathbf{Q}} \end{bmatrix} \right\} \tag{32}$$

$$\mathbf{S}^{(\ell)}_{k|k-1} \leftarrow \mathtt{choldowndate}\left\{ \bar{\mathbf{S}}^{(\ell)}_{1:2n,k|k-1}, \sqrt{-\omega^{(c)}_0} \left(\boldsymbol{\mathcal{X}}^{(\ell)}_{0,k|k-1} - \mathbf{m}^{(\ell)}_{k|k-1} \right) \right\} \tag{33}$$

where γ denotes the composite scaling parameter $\gamma = \sqrt{n + \lambda}$ and $\lambda = \alpha^2(n + \kappa) - n$. The scaling parameters α, β, and κ control the spread of the $2n + 1$ sigma points and determine the sigma point weights,

$$\begin{aligned} \omega^{(m)}_0 &= \frac{\lambda}{n + \lambda} \\ \omega^{(c)}_0 &= \frac{\lambda}{n + \lambda} + (1 - \alpha^2 + \beta) \\ \omega^{(m)}_i &= \omega^{(c)}_i = \frac{1}{2(n + \lambda)}, \qquad i = 1, \ldots, 2n \end{aligned} \tag{34}$$

Here, $\mathtt{qr}(\boldsymbol{A})$ is shorthand for the QR decomposition $\boldsymbol{A}^T = \boldsymbol{Q}\boldsymbol{R}$, where only the transpose of the upper-triangular part of the matrix $\boldsymbol{R}$ is returned. The Cholesky factor of the rank-1 downdate $\tilde{\boldsymbol{S}}\tilde{\boldsymbol{S}}^T = \boldsymbol{S}\boldsymbol{S}^T - \boldsymbol{x}\boldsymbol{x}^T$ is abbreviated by $\mathtt{choldowndate}(\boldsymbol{S}, \boldsymbol{x})$. In the case that the second argument $\boldsymbol{x}$ is a matrix, the downdate is applied successively for each column of $\boldsymbol{x}$.

Sigma point approximations are also employed in the Bayes measurement update in the case where the space object is detected and a right-ascension/declination pair is produced. In that case, the measurement RFS $Z_k = \{\mathbf{z}_k\}$, and the posterior means (22) and covariances (23) are approxi-

mated as follows:

$$\boldsymbol{\mathcal{Z}}_{i,k}^{(\ell)} = \boldsymbol{h}\left(\boldsymbol{\mathcal{X}}_{i,k|k-1}^{(\ell)}\right) \tag{35}$$

$$\hat{\mathbf{z}}_k^{(\ell)} = \sum_{i=0}^{2n} \omega_i^{(m)} \boldsymbol{\mathcal{Z}}_{i,k}^{(\ell)} \tag{36}$$

$$\bar{\mathbf{S}}_{z,k}^{(\ell)} \leftarrow \texttt{qr}\left\{\left[\sqrt{\omega_1^{(c)}}\left(\boldsymbol{\mathcal{Z}}_{1:2n,k}^{(\ell)} - \hat{\mathbf{z}}_k^{(\ell)}\right) \quad \sqrt{\mathbf{R}}_k\right]\right\} \tag{37}$$

$$\mathbf{S}_{z,k}^{(\ell)} \leftarrow \texttt{choldowndate}\left\{\bar{\mathbf{S}}_{z,k}^{(\ell)},\, \sqrt{-\omega_0^{(c)}}\left(\boldsymbol{\mathcal{Z}}_{0,k}^{(\ell)} - \hat{\mathbf{z}}_k^{(\ell)}\right)\right\} \tag{38}$$

$$\mathbf{C}_k^{(\ell)} = \sum_{i=0}^{2n} \omega_i^{(c)}\left[\boldsymbol{\mathcal{X}}_{i,k|k-1}^{(\ell)} - \mathbf{m}_{k|k-1}^{(\ell)}\right]\left[\boldsymbol{\mathcal{Z}}_{i,k}^{(\ell)} - \hat{\mathbf{z}}_k^{(\ell)}\right]^T \tag{39}$$

$$\mathbf{K}_k^{(\ell)} = \left[\mathbf{C}_k^{(\ell)}(\mathbf{S}_{z,k}^{(\ell)})^{-T}\right](\mathbf{S}_{z,k}^{(\ell)})^{-1} \tag{40}$$

$$\mathbf{m}_k^{(\ell)} = \mathbf{m}_{k|k-1}^{(\ell)} + \mathbf{K}_k^{(\ell)}\left(\mathbf{z}_k - \hat{\mathbf{z}}_k^{(\ell)}\right) \tag{41}$$

$$\mathbf{S}_k^{(\ell)} \leftarrow \texttt{choldowndate}\left\{\mathbf{S}_{x,k|k-1}\,, \mathbf{K}_k^{(\ell)}\mathbf{S}_{z,k}^{(\ell)}\right\} \tag{42}$$

where $\sqrt{\mathbf{R}_k}$ represents the square-root factor of $\mathbf{R}_k$ such that $\sqrt{\mathbf{R}_k}\sqrt{\mathbf{R}_k}^T = \mathbf{R}_k$. Note that the "$-T$" operation represents the transpose of the inverse of a matrix.

Gaussian Mixture Splitting

Key challenges of the cislunar tracking problem include chaotic dynamics, long observation gaps, and the massive volume of the cislunar region of interest. The culmination of these challenges is a nonlinear non-Gaussian estimation problem for which central-moment filters (e.g., the EKF, and UKF) are inadequate. Standard GM filters address many of the pitfalls of central-moment filters by approximating the true non-Gaussian state density as a sum of Gaussian kernels. Successful GM filtering relies on constructing GMs with small enough component covariances such that, on a component-level, the measurement and dynamics models can be accurately linearized. In the cislunar tracking problem, fixed-sized (non-adaptive) GM filters degrade over long observation gaps due to the growth of the component covariances and, consequentially, the violation of the underlying assumption that time- and measurement-updated components remain Gaussian.

Adaptive GM filtering addresses the shortcomings of fixed-size GM approaches by dynamically increasing the GM resolution in regions of the state-space where nonlinearity-induced effects are the strongest. This subsection builds on previous work in adaptive GM filtering and presents a new algorithm suitable for angles-only cislunar tracking. Unlike existing approaches, the proposed algorithm uses adaptation to accommodate four sources of nonlinearity: nonlinear dynamics, nonlinear measurements, nonlinear probability of detection, and nonlinear constraints. This subsection also presents new Gaussian splitting equations based on Cholesky square root factors.

A key mechanism of adaptive GM filtering is Gaussian splitting, where a Gaussian component is replaced by GM approximation of itself. Splitting is performed efficiently by utilizing a pre-generated library of optimal split parameters for the univariate standard Gaussian,

$$q(x) = \mathcal{N}(x;\, 0,\, 1) \approx \tilde{q}(x) = \sum_{j=1}^{R} \tilde{w}^{(j)} \mathcal{N}(x;\, \tilde{m}^{(j)},\, \tilde{\sigma}^2) \tag{43}$$

as first proposed in [20] and refined in [15, 21]. By this approach, the univariate split parameters $\{\tilde{w}^{(j)}, \tilde{m}^{(j)}, \tilde{\sigma}^{(j)}\}_{j=1}^{R}$ can be referenced at runtime and applied to arbitrary multivariate Gaussian distributions via scaling, shifting, and covariance diagonalization.

To achieve an end-to-end estimation algorithm that is entirely square-root factor-based, this paper introduces a new splitting method based on Cholesky factors. Omitting the time index subscripts for brevity, let $\lambda_i^{(\ell)}$ and $\boldsymbol{v}_i^{(\ell)}$ be the i^{th} eigenvalue and corresponding normalized eigenvector, respectively, of the full covariance matrix $\mathbf{S}^{(\ell)}\mathbf{S}^{(\ell)^T}$. Let $\mathbf{V}^{(\ell)} = [\boldsymbol{v}_1 \cdots \boldsymbol{v}_n]$ and $\mathbf{\Lambda}^{(\ell)} = \text{diag}([\lambda_1^{(\ell)} \cdots \lambda_n^{(\ell)}])$. Then, the multivariate GM component can be split along its i^{th} eigenvector $\boldsymbol{v}_i^{(\ell)}$ as follows:

$$w^{(\ell)}\mathcal{N}(\mathbf{x};\, \mathbf{m}^{(\ell)},\, \mathbf{S}^{(\ell)}(\mathbf{S}^{(\ell)})^T) \approx \sum_{j=1}^{R} w^{(\ell,j)}\mathcal{N}(\mathbf{x};\, \mathbf{m}^{(\ell,j)},\, \mathbf{S}^{(\ell,j)}(\mathbf{S}^{(\ell,j)})^T) \tag{44}$$

where

$$w^{(\ell,j)} = \tilde{w}^{(j)} w^{(\ell)} \tag{45}$$

$$\mathbf{m}^{(\ell,j)} = \mathbf{m}^{(\ell)} + \sqrt{\lambda_i^{(\ell)}}\tilde{m}^{(j)}\boldsymbol{v}_i^{(\ell)} \tag{46}$$

$$\mathbf{S}^{(\ell,j)} \leftarrow \texttt{choldowndate}(\mathbf{S}^{(\ell)}, \mathbf{a}_i^{(\ell)}) \tag{47}$$

$$\mathbf{a}_i^{(\ell)} = \mathbf{V}^{(\ell)}(\sqrt{\lambda_i^{(\ell)}(1-\tilde{\sigma}^2)}\mathbf{e}_i) \tag{48}$$

where the i^{th} element of the canonical basis vector $\mathbf{e}_i$ is unity and all other elements are zero. The split Cholesky factor is obtained using the standard Cholesky rank-1 downdate, which computes the Cholesky factor of $\mathbf{S}^{(\ell)}(\mathbf{S}^{(\ell)})^T - \mathbf{a}_i^{(\ell)}(\mathbf{a}_i^{(\ell)})^T$. While it is possible to split along an arbitrary non-eigenvector direction, the covariance of the resulting GM will be rotated with respect to the original component, and thus should be avoided. Note that the computation of $\mathbf{V}^{(\ell)}$ and $\mathbf{\Lambda}^{(\ell)}$ does require an intermediate computation of the full covariance matrix. Fortunately, when applying the split approximation recursively, this computation is required only once, as the spectral factorization of the child components is simply

$$\mathbf{V}^{(\ell,j)} = \mathbf{V}^{(\ell)}\,, \qquad \mathbf{\Lambda}^{(\ell,j)} = \text{diag}([\lambda_1^{(\ell)} \cdots \tilde{\sigma}^2\lambda_i^{(\ell)} \cdots \lambda_n^{(\ell)}]) \tag{49}$$

Generally speaking, splitting can be applied to any component in any spectral direction at any step in GM filtering. Splitting inherently increases the GM size and filter's computational cost, and thus must be performed judiciously. Tractable and accurate GM splitting filtering therefore requires thoughtful component selection criteria and split direction determination schemes. One approach is to select components whose linearized and sigma point-based mappings result in significantly different Shannon entropy values.[15] Another simple approach considers only the component weights and splits components with the large posterior weights.[22] Weight-based selection criteria may fail to split important high covariance components that are lower weight individually but nonetheless contribute significant probability mass in combination with other low-weight components. A better approach is to choose components based on both their weight and associated linearization error,[23] as is discussed later in this subsection.

Splitting in Prediction Due to the chaotic nature of CR3BP dynamics, space object state uncertainty evolves to form increasingly complex non-Gaussian distributions over time, particularly for periodic orbits about collinear libration points, which are unstable. Thus, GM tracking solutions must adapt by increasing the mixture resolution in regions where the assumed Gaussian form of the nonlinear transformation is most invalid. This paper adopts the AEGIS filter[15] structure for prediction, which employs mid-propagation splitting, thereby mitigating repeated solution of the CR3BP differential equations over the full propagation period.

In addition to the AEGIS entropy-based split criterion, this paper proposes a new split criterion tailored to cislunar space object tracking. In particular, a component is split if the variance of its Jacobi constant distribution exceeds a user-specified threshold. In the absence of process noise, the Jacobi constant variance can be computed from the sigma points as

$$\mathrm{var}(C^{(\ell)}(t)) = \sum_{i=0}^{2n} \omega_i^{(c)} \left(C(\boldsymbol{\mathcal{X}}^{(\ell,i)}(t)) - \mathrm{E}[C^{(\ell)}(t)] \right)^2 \tag{50}$$

$$\mathrm{E}[C^{(\ell)}(t)] = \sum_{i=0}^{2n} \omega_i^{(m)} C(\boldsymbol{\mathcal{X}}^{(\ell,i)}(t)) \tag{51}$$

If the Jacobi constant variance exceeds the user-specified threshold $\sigma^2_{C,\max}$, the component is split along the eigenvector corresponding to the highest state covariance eigenvalue. In principle, a component with high Jacobi constant variance represents large deviations in orbit geometry (as opposed to different positions along the same orbit) and will inevitably produce a large covariance non-Gaussian distribution when propagated forward in time.

Aside from this additional criterion, the prediction step is equivalent to the original AEGIS algorithm and is summarized as follows. Consider a single GM component with index ℓ. Let $t_s \in [t_{k-1}, t_k]$ denote the time at which nonlinearity is detected via the Shannon entropy measure, the time at which the Jacobi constant covariance exceeds the user-specified threshold, or the end of the propagation interval t_k, whichever comes first. Let t_{s-1} denote the time at which the last split occurred. If no split has occurred in the interval, t_{s-1} is initialized as t_k. The component is propagated over $t \in [t_{s-1}, t_s]$ using the square-root scaled unscented transform (29)-(33). If $t_s = t_k$, then the prediction step is complete for the component. Otherwise, the component is split into an R-component GM along its highest-variance spectral direction according to (44).

Splitting in Update Adaptive splitting is also applied in the processing of right ascension and declination measurements to reduce linearization errors in the square-root UKF equations. The unscented transform, which belongs to a broader category of sigma point approximations, is a form of *statistical linearization*, which approximates a general nonlinear transformation as $\mathbf{y} = \mathbf{g}(\mathbf{x}) \approx \mathbf{G}\mathbf{x} + \mathbf{b}$.[23] In view of the nonlinear measurement transformation $\mathbf{h}(\mathbf{x})$, the optimal $\mathbf{G}_k^{(\ell)}$ and $\mathbf{b}_k^{(\ell)}$ are given by[23]

$$\mathbf{G}_k^{(\ell)} = \mathbf{C}_k^{(\ell)^T} (\mathbf{S}_k^{(\ell)})^{-T} (\mathbf{S}_k^{(\ell)})^{-1} \qquad \text{and} \qquad \mathbf{b}_k^{(\ell)} = \hat{\mathbf{z}}_k^{(\ell)} - \mathbf{G}_k^{(\ell)} \mathbf{m}_{k|k-1}^{(\ell)} \tag{52}$$

and the associated linearization error is

$$\boldsymbol{e}_k^{(\ell)} = \mathbf{h}(\mathbf{x}_k) - \mathbf{G}_k^{(\ell)} \mathbf{x}_k - \mathbf{b}_k^{(\ell)} \tag{53}$$

The linearization error in (53) is zero-mean and can be characterized by its covariance matrix

$$\mathrm{E}\left[\boldsymbol{e}_k^{(\ell)} \boldsymbol{e}_k^{(\ell)^T} \right] = \mathbf{P}_{e,k}^{(\ell)} = \mathbf{P}_{\tilde{z},k}^{(\ell)} - \mathbf{G}_k^{(\ell)} \mathbf{S}_k^{(\ell)} \mathbf{S}_k^{(\ell)^T} \mathbf{G}_k^{(\ell)^T} \tag{54}$$

where $\mathbf{P}_{\tilde{z},k}$ is the predicted measurement covariance for the ℓ^{th} component in the absence of measurement noise and is given by

$$\mathbf{P}_{\tilde{z},k}^{(\ell)} = \sum_{i=0}^{2n} \omega_i^{(c)} \left(\boldsymbol{\mathcal{Z}}_{i,k|k-1}^{(\ell)} - \hat{\mathbf{z}}_k^{(\ell)} \right) \left(\boldsymbol{\mathcal{Z}}_{i,k|k-1}^{(\ell)} - \hat{\mathbf{z}}_k^{(\ell)} \right)^T \tag{55}$$

From this readily computable covariance matrix, the linearization error incurred by the sigma point approximation can be quantified for each component. Components with high linearization error are identified and split into smaller components to reduce the approximation error. To consider both the linearization error and the overall statistical significance of a given component, the selection criteria

$$s_k^{(\ell)} = (w_{k|k-1}^{(\ell)})^{\Gamma} \cdot (1 - \exp(-\epsilon^{(\ell)}))^{1-\Gamma} \in [0,1] \qquad \text{where} \qquad \epsilon^{(\ell)} = \text{tr}(\mathbf{P}_{e,k}^{(\ell)}) \tag{56}$$

is adopted, which represents a geometric interpolation between the weight and linearization error.[23] The tuning parameter $\Gamma \in [0,1]$ determines the relative importance of component weight and linearization error, with $\Gamma = 0$ resulting in only the consideration of linearization error (with no regard to weight), and $\Gamma = 1$ resulting in only the consideration of the component weight. If $s^{(\ell)}$ exceeds the user-specified threshold $\epsilon_{\max}$, the component is selected for splitting.

Splitting for Negative Information In cislunar space object tracking, it is important to make use of all available information, including evidence of where an object is not observed. This form of evidence is referred to as *negative information* and is incorporated into the posterior pdf via the state-dependent probability of detection function $p_{D,k}(\mathbf{x}_k; \mathcal{S})$. Recall that the GM approximation (20) of the posterior involves the expansion of $p_{D,k}(\mathbf{x}_k; \mathcal{S})$ about the component means. In effect, by this approximation, the probability of detection is assumed to be constant over a given component's local support and equal to the probability of detection at that component's mean.

In some cases, the probability of detection may be different over the local support of a component, resulting in approximation error and in extreme cases, filter divergence. This is particularly true for components whose position-marginal density overlaps the sensor FoV boundaries, resulting in a sharp change in the probability of detection within the component support. To address these challenges, this paper employs a recently developed recursive splitting algorithm that automatically identifies components that overlap FoV bounds and splits them until certain stopping criteria are satisfied,[14] as illustrated in Figure 1. Details of the FoV splitting algorithm are omitted for brevity but can be found in [14].

Nonlinear Constraints In the proposed filter, it is straightforward to incorporate nonlinear constraints or soft data from human sources in the same way that negative information is incorporated. For instance, with knowledge of a satellite's ΔV capacity and its state at some time, bounds can be established for the Jacobi constant as

$$C_{\min}(\Delta V) \leqslant C(\mathbf{x}) \leqslant C_{\max}(\Delta V) \tag{57}$$

This constraint can then be incorporated to refine the probabilistic belief by eliminating unreachable orbits from the probabilistic solution. Defining the set $\mathcal{C} = \{\mathbf{x} : C(\mathbf{x}) \in [C_{\min}, C_{\max}]\}$, the GM weights are updated as

$$w_{k|k-1}^{(\ell)} \propto \begin{cases} w_{k|k-1}^{(\ell)} & \mathbf{m}_k^{(\ell)} \in \mathcal{C} \\ 0 & \text{otherwise} \end{cases} \tag{58}$$

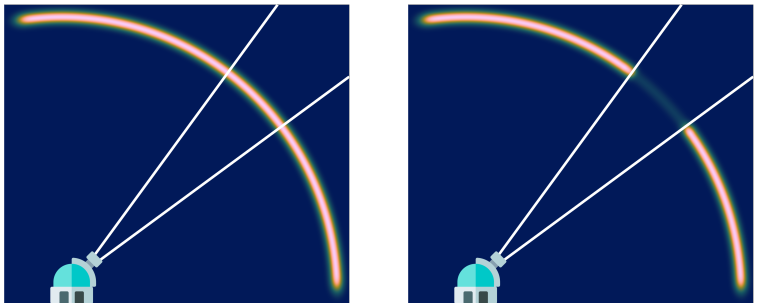

Figure 1. Example space object state probability density function before (left) and after (right) incorporating negative information (non-detection), obtained using recursive Gaussian mixture splitting.

Note that (58) is found using the same zeroth-order Taylor expansion applied in (21); see [14] for details. If a component overlaps the 6D Jacobi constant set boundaries $\partial\mathcal{C}$, it is split to improve the means-based expansion approximation. A simpler but partial incorporation of the Jacobi constant constraint can be achieved through splitting about the zero-velocity curves and pruning components whose positions fall in forbidden regions.

The complete GM Bayes filter for angles-only cislunar tracking is shown in block diagram form in Figure 2. Between observation opportunities, the GM pdf is propagated according to CR3BP dynamics. An AEGIS-inspired prediction step evaluates entropy and Jacobi constant variance at intermediate time steps to detect non-Gaussianity and splits components accordingly. Any nonlinear constraints, such as bounds on the Jacobi constant, are applied to prior pdf and the weights are re-normalized. Splitting is also employed in the measurement update stage. If no space object is detected at a given observation step ($Z_k = \varnothing$), the proposed filter splits GM components along FoV bounds to incorporate the negative information content of the empty measurement. If the space object is detected, the nonlinear measurement is incorporated into the posterior pdf via Bayes' rule and statistical linearization. If the resulting linearization error is too large, the relevant components are split until the stopping criterion is satisfied.

RESULTS

To assess the performance of the proposed adaptive GM filter, a tracking problem involving a satellite in an L_2 halo orbit is considered. Synthetic topocentric right ascension and declination measurements are generated for an Earth-based observer when the satellite is within the FoV and lighting conditions are satisfied. When lighting conditions are satisfied, but the satellite is not within the FoV, the resulting observation is the empty set $Z_k = \varnothing$. The sensor FoV is taken to be a rectangular prism with $3\,[\text{deg}]$ vertical and horizontal half angles. Measurements, when available, are corrupted by white Gaussian noise with standard deviation $10''$ in each angle, which corresponds to a $1[\text{pixel}]$ standard deviation error in a 2160×2160 focal plane array. The parameters of the adaptive GM filter are provided in Table 1. The FoV orientation is assumed to be such that the boresight remains fixed in the rotating synodic frame and maintains a fixed angular separation from the Moon to avoid image saturation. For simplicity, a global network of sensors is assumed such that an (potentially empty) observation can be made once per hour given suitable lighting conditions.

The initial state uncertainty is assumed to be Gaussian with $20\,[\text{km}]$ and $1\,[\text{m/s}]$ standard devia-

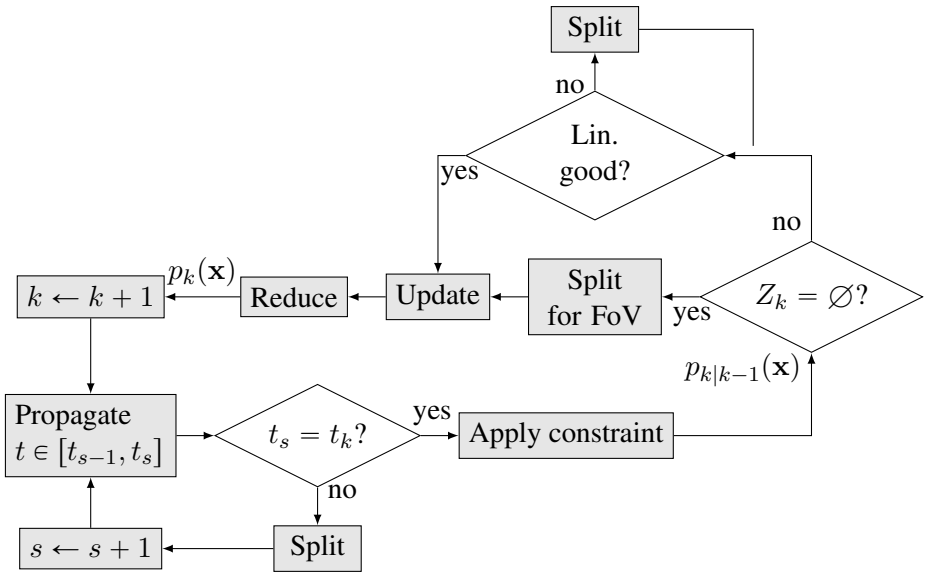

Figure 2. Block diagram of Gaussian mixture Bayes filter for cislunar angles-only tracking.

Table 1. Adaptive Gaussian Mixture Filter Parameters

	Parameter	Value
Sigma Point Scaling	α	0.001
	β	2
	κ	0
Noise	$\sqrt{\mathbf{Q}_k}$	$\mathbf{0}$
	$\sqrt{\mathbf{R}_k}$	diag([10″ 10″])
Gaussian Splitting	R	5
	Γ	0.5
	$\sigma^2_{C,\max}$	0.0001
	$\epsilon_{\max}$	0.01
	$L_{k,\max}$	500

tion in position and relative velocity, respectively. The initial mean is

$$\mathbf{m}_0^{(1)} = \begin{bmatrix} 1.0697 & 0 & 0.2015 & 0 & -0.1855 & 0 \end{bmatrix}^T \tag{59}$$

which corresponds to a L_2 Northern halo orbit with Jacobi constant 3.0169 and stability index 1.0000000000297. For numerical stability, regularization is applied when integrating the CR3BP EOMs, as described in the Appendix. The performance of the adaptive GM Bayes filter is analyzed over 500 Monte Carlo trials, in which the initial true state and measurement noise sequence is randomly varied. Note that randomly generated initial conditions generally result in non-periodic orbits and diverging visibility profiles. Thus, for the sake of analysis, these samples are differentially corrected, after which outliers are rejected according to a χ^2 hypothesis test. While the differentially corrected sample covariance is generally smaller than the assumed initial covariance, limiting a more complete consistency analysis, the sample population still serves to compare performance between filters.

An example of true satellite halo orbit and the sensor FoV are shown in Figure 3. The total simulation duration is 30 days, where due to the orbital geometry and solar phase, the satellite is undetectable for the first two weeks. During this long period of non-observation, the satellite state

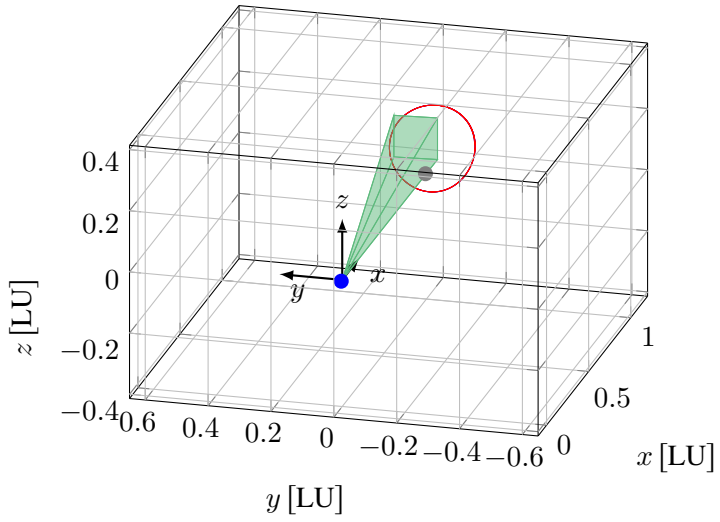

Figure 3. True satellite orbit and sensor FoV as seen in the rotating synodic frame and expressed in non-dimensional length units. The Earth and Moon positions are denoted by the blue and grey circles, respectively.

uncertainty grows considerably due to the chaotic dynamics, and the resulting pdf is highly non-Gaussian by the 15^{th} day, as shown in Figure 4. Despite the fact that the satellite does not enter the FoV until later in the 15^{th} day, the filter incorporates negative information from the absence of detections to significantly reduce the state uncertainty, as shown in Figure 5.

Filter performance is measured by the maximum a posteriori (MAP) estimation error over time. As a measure of the filter's reported uncertainty, the conditional covariance of the GM pdf is computed and the root-sum of the diagonal elements is taken and referred to as the covariance root sum square (RSS). The 2-σ RSS values are reported while noting that σ bounds generally do not correspond to curves of constant probability in a non-Gaussian distribution. The estimation error and 2-σ RSS histories are overlaid for all 500 Monte Carlo trials in Figure 6. The estimation error remains below the 2-σ RSS curves in all cases, suggesting that the proposed filter is consistent. After

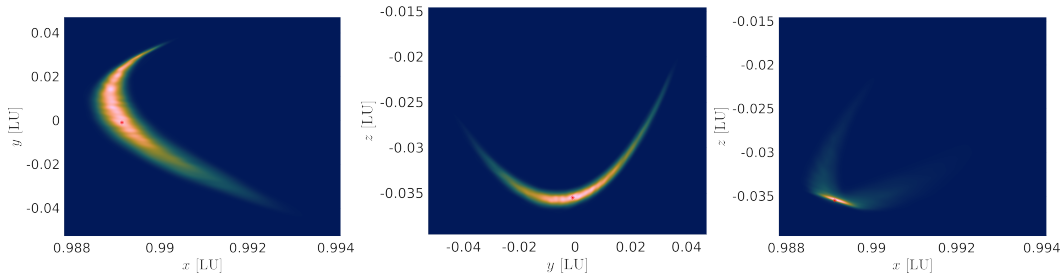

Figure 4. Satellite state pdf marginals at $t = 14.2083\,[\text{day}]$, where the true position coordinates are denoted by red asterisks.

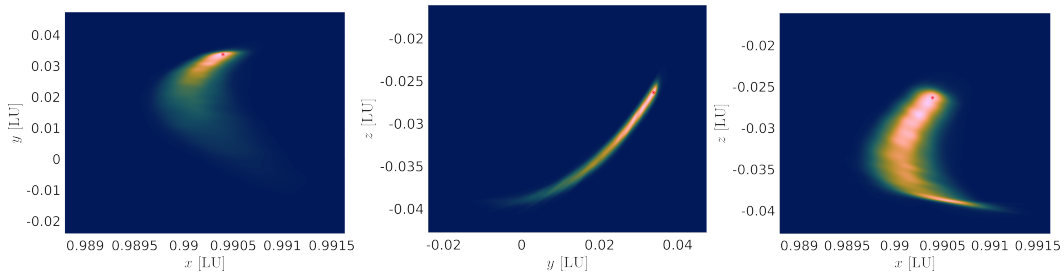

Figure 5. Satellite state pdf marginals at $t = 14.4167\,[\text{day}]$ after incorporating negative information, where the true position coordinates are denoted by red asterisks.

the first observation window, there is an approximately ten-day observation gap, during which the uncertainty grows by two orders of magnitude. During visibility periods (where one observation is made every hour), the filter achieves position and velocity estimates that are accurate to within $10\,[\text{km}]$ and $1\,[\text{m/s}]$, respectively (and lower in many cases). Of course, these accuracy metrics must be interpreted within the context of the assumed sensor model specifications. Higher accuracy estimates can undoubtedly be achieved with higher frequency observations, space-based observers, and higher accuracy measurements, among other variables.

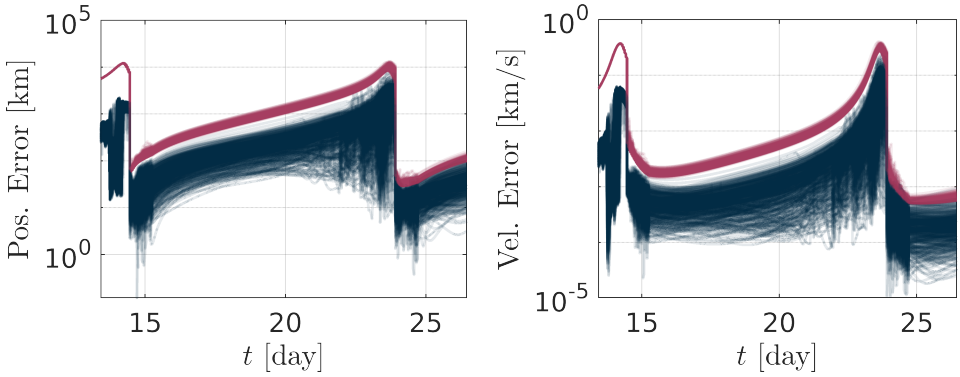

Figure 6. Estimation error over time (navy) of proposed adaptive GM Bayes filter in 500 Monte Carlo trials of a L_2 halo orbit determination problem, where the uncertainty is represented by the 2-σ RSS values shown in purple, and where the horizontal axis limits reflect the first and last instances of (potential) visibility within the FoV.

As a means of comparison, the same simulation is conducted in 250 Monte Carlo trials using the square-root unscented Kalman filter (SRUKF). The SRUKF estimation error and covariance RSS

are shown in Figure 7. In almost all of the 250 trials, the measurement likelihood agreement fell below the machine precision, resulting in a total filter failure and the termination of the trial. Of the remaining trials that did complete, all produced estimation errors that significantly exceeded the 2-σ RSS bounds, indicating filter divergence. The inadequacy of the SRUKF for this estimation problem is further evidenced by the magnitude of the position estimation error, which at times exceeds the Earth-Moon semi-major axis.

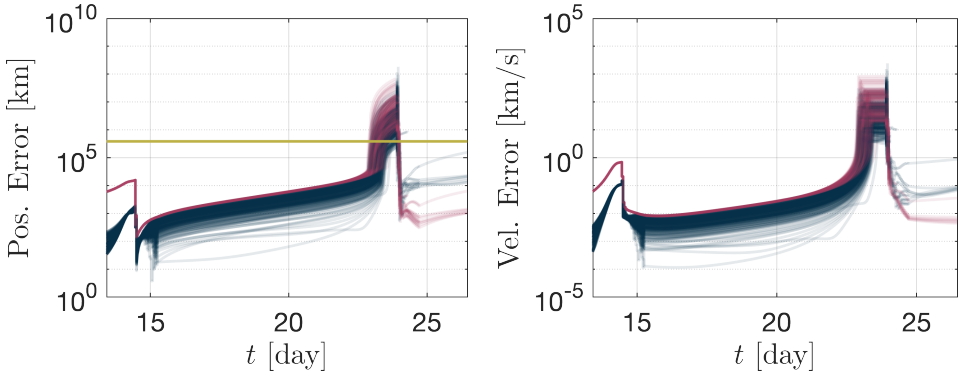

Figure 7. Estimation error over time (navy) of SRUKF in 250 Monte Carlo trials of L_2 halo orbit determination problem, where the uncertainty is represented by the 2(RSS) values shown in purple, and where the horizontal axis limits reflect the first and last instances of (potential) visibility within the FoV. The Earth-Moon semi-major axis is represented by the horizontal gold line for reference.

The stability of the proposed adaptive GM filter is also investigated. To demonstrate the filter performance in different settings, the filter is applied in 31 different L_2 Northern halo orbit determination problems with Jacobi constants varying between 3.0152 and 3.1584 and stability indices varying between 1.5801 and 583.7903. Filter convergence was observed in all 31 cases, suggesting that the proposed adaptive GM filter is applicable to a broad range of angles-only cislunar orbit determination problems. The 31 orbits are shown in Figure 8, where the color of the orbit reflects the minimum estimation error achieved during that particular orbit simulation. It should be noted that, unlike the previous analysis, only one trial is conducted per orbit, and thus the reported estimation error is expected to vary under different measurement noise sequences.

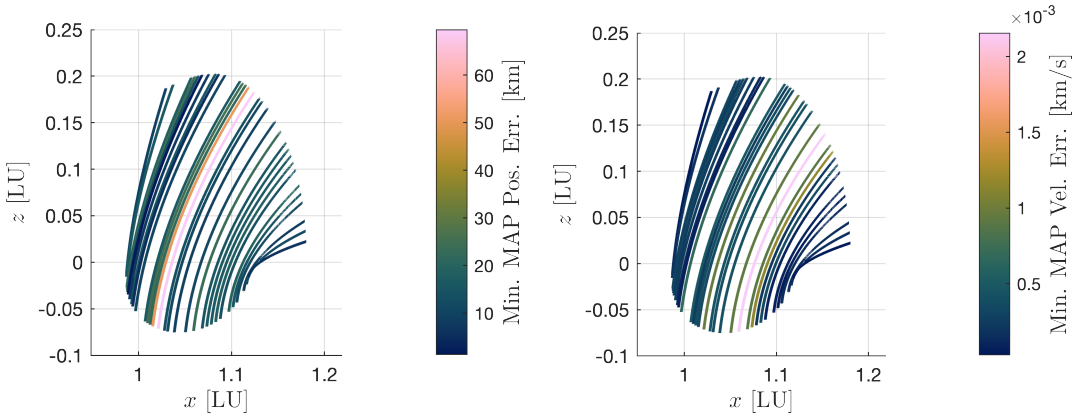

Figure 8. L_2 Northern halo orbit family considered in evaluation, where the color of the orbit reflects the minimum estimation error achieved during simulation.

CONCLUSIONS

This paper presents a novel Bayesian adaptive Gaussian mixture (GM) filter for angles-only cislunar orbit determination. The filter models the satellite state probability density function as a GM to appropriately capture the non-Gaussian structure of the evolving state uncertainty. Recursive Gaussian splitting is applied to improve the linearized time- and measurement-update approximations, as well as to incorporate negative information and nonlinear constraints. The efficacy and robustness of the filter is demonstrated in 31 representative simulated orbit determination problems across the L_2 Northern halo family, as well as in a 500-trial Monte Carlo analysis.

REFERENCES

[1] M. J. Holzinger, C. C. Chow, and P. Garretson, "A Primer on Cislunar Space," 2021.

[2] C. P. Newman, D. C. Davis, R. J. Whitley, J. R. Guinn, and M. S. Ryne, "Stationkeeping, Orbit Determination, and Attitude Control for Spacecraft in Near Rectilinear Halo Orbits," Astrodynamics Specialists Conference, 2018.

[3] N. L. Parrish, M. J. Bolliger, E. Kayser, M. R. Thompson, J. S. Parker, B. W. Cheetham, D. C. Davis, and D. J. Sweeney, "Near Rectilinear Halo Orbit Determination with Simulated DSN Observations," *2020 AIAA Scitech Forum*, American Institute of Aeronautics and Astronautics, 2020, 10.2514/6.2020-1700.

[4] M. Gupta, K. C. Howell, and C. Frueh, "Earth-Moon Multi-Body Orbit to Facilitate Cislunar Surveillance Activities," AAS/AIAA Astrodynamics Specialist Conference, 2021.

[5] C. Frueh, K. C. Howell, K. J. DeMars, S. Bhadauria, and M. Gupta, "Cislunar Space Traffic Management: Surveillance Through Earth-Moon Resonance Orbits," *8th European Conference on Space Debris*, Vol. 8, European Space Agency, 2021.

[6] J. A. Dahlke, A. P. Wilmer, and R. A. Bettinger, "Preliminary Comparative Assessment of L2 and L3 Surveillance Using Select Cislunar Periodic Orbits," AAS/AIAA Astrodynamics Specialist Conference, 2022.

[7] C. Frueh, K. C. Howell, K. J. DeMars, and S. Bhadauria, "Cislunar Space Situational Awareness," *31st AIAA/AAS Space Flight Mechanics Meeting*, Space Flight Mechanics Committee, 2021.

[8] S. Julier and J. Uhlmann, "Unscented Filtering and Nonlinear Estimation," *Proceedings of the IEEE*, Vol. 92, No. 3, 2004, pp. 401–422, 10.1109/JPROC.2003.823141.

[9] C. C. Chow, C. J. Wetterer, J. Baldwin, M. Dilley, K. Hill, P. Billings, and J. Frith, "Cislunar Orbit Determination Behavior: Processing Observations of Periodic Orbits with Gaussian Mixture Model Estimation Filters," *The Journal of the Astronautical Sciences*, Vol. 69, Oct. 2022, pp. 1477–1492, 10.1007/s40295-022-00347-7.

[10] K. J. DeMars and M. K. Jah, "Initial Orbit Determination via Gaussian Mixture Approximation of the Admissible Region," *Proceedings of the 22nd AAS/AIAA Space Flight Mechanics Meeting*, No. AAS 12-260, 2012.

[11] K. A. LeGrand, K. J. DeMars, and H. J. Pernicka, "Bearing-Only Initial Relative Orbit Determination," *Journal of Guidance, Control, and Dynamics*, Vol. 38, 2015, pp. 1699–1713, 10.2514/1.G001003.

[12] H. W. Sorenson and D. L. Alspach, "Recursive Bayesian Estimation using Gaussian Sums," *Automatica*, Vol. 7, No. 4, 1971, pp. 465–479, https://doi.org/10.1016/0005-1098(71)90097-5.

[13] K. A. LeGrand and S. Ferrari, "The Role of Bounded Fields-of-View and Negative Information in Finite Set Statistics (FISST)," *2020 ISIF 23rd International Conference on Information Fusion (FUSION)*, International Society of Information Fusion, 2020, pp. 1–9, 10.23919/FUSION45008.2020.9190174.

[14] K. A. LeGrand and S. Ferrari, "Split Happens! Imprecise and Negative Information in Gaussian Mixture Random Finite Set Filtering," *(in press) Journal of Advances in Information Fusion*, 2022, 10.48550/ARXIV.2207.11356.

[15] K. J. DeMars, R. H. Bishop, and M. K. Jah, "Entropy-Based Approach for Uncertainty Propagation of Nonlinear Dynamical Systems," *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 4, 2013, pp. 1047–1057, 10.2514/1.58987.

[16] K. J. DeMars, *Nonlinear Orbit Uncertainty Prediction and Rectification for Space Situational Awareness*. Dissertation, University of Texas at Austin, 2010.

[17] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*, Vol. 1 of *Information and System Sciences Series*. Prentice-Hall, Inc.

[18] R. V. D. Merwe, "Sigma-Point Kalman Filters for Probabilistic Inference in Dynamic State-Space Models," 2004.

[19] S. Julier, J. Uhlmann, and H. Durrant-Whyte, "A New Approach for Filtering Nonlinear Systems," *Proceedings of 1995 American Control Conference - ACC'95*, Vol. 3, pp. 1628–1632 vol.3, 10.1109/ACC.1995.529783.

[20] U. D. Hanebeck, K. Briechle, and A. Rauh, "Progressive Bayes : A New Framework for Nonlinear State Estimation," No. April 2003, 2003, 10.1117/12.487806.

[21] M. F. Huber, T. Bailey, H. Durrant-Whyte, and U. D. Hanebeck, "On Entropy Approximation for Gaussian Mixture Random Vectors," *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 2008, pp. 181–188, 10.1109/MFI.2008.4648062.

[22] F. Faubel, J. McDonough, and D. Klakow, "The Split and Merge Unscented Gaussian Mixture Filter," Vol. 16, No. 9, pp. 786–789, 10.1109/LSP.2009.2024859.

[23] M. F. Huber, "Adaptive Gaussian Mixture Filter Based on Statistical Linearization," *14th International Conference on Information Fusion*, 2011, pp. 1–8.

[24] E. L. Stiefel and G. Scheifele, *Linear and regular celestial mechanics: perturbed two-body motion, numerical methods, canonical theory*, Vol. 174. Springer, 1971.

[25] K. Howell and J. Breakwell, "Almost rectilinear halo orbits," *Celestial mechanics*, Vol. 32, No. 1, 1984, pp. 29–52.

[26] D. Bettis and V. Szebehely, "Treatment of close approaches in the numerical integration of the gravitational problem of N bodies," *International Astronomical Union Colloquium*, Vol. 10, Cambridge University Press, 1971, pp. 133–150.

APPENDIX

The presence of the inverse of the target's distance between the primaries, $r_{/\oplus}$ and $r_{/\mathbb{C}}$, in (4) results in numerical instability during close approaches to either primary. Regularization can be employed to alleviate the instability, and applying it to a three-dimensional position space occurs in two steps: the introduction of a fictitious time and a coordinate transform to a four-dimensional position space.[24] The fictitious time, τ, is defined such that $\frac{dt}{d\tau} = r_{/\mathbb{C}}$ which has the effect of "slowing down" time when approaching the Moon. Following a similar derivation of regularized CR3BP EOMs,[25] (3) and (4) can be rewritten in vector form with $(\cdot)'$ denoting differentiation with respect to τ as

$$\bar{\boldsymbol{r}}'' = \frac{\bar{\boldsymbol{r}} \cdot \bar{\boldsymbol{r}}'}{\bar{\boldsymbol{r}} \cdot \bar{\boldsymbol{r}}} \bar{\boldsymbol{r}}' + \|\bar{\boldsymbol{r}}\| \mathbf{B} \bar{\boldsymbol{r}}' - \mu \frac{\bar{\boldsymbol{r}}}{\|\bar{\boldsymbol{r}}\|} + (\bar{\boldsymbol{r}} \cdot \bar{\boldsymbol{r}}) \bar{\boldsymbol{\zeta}} \tag{60}$$

where $\bar{\boldsymbol{r}} = \begin{bmatrix} x - 1 + \mu & y & z & 0 \end{bmatrix}^T$, $\bar{\boldsymbol{r}}' = \begin{bmatrix} x' & y' & z' & 0 \end{bmatrix}^T$, $\bar{\boldsymbol{\zeta}} = \begin{bmatrix} x - \frac{(1-\mu)(x+\mu)}{\|r_{/\oplus}\|^3} & y - \frac{(1-\mu)y}{\|r_{/\oplus}\|^3} & -\frac{(1-\mu)z}{\|r_{/\oplus}\|^3} & 0 \end{bmatrix}^T$, and

$$\mathbf{B} = \begin{bmatrix} 0 & 2 & 0 & 0 \\ -2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

It can easily be seen that (60) still contains singularities. Thus, a coordinate transform is defined such that $\bar{\boldsymbol{r}} = \mathcal{L}(\bar{\mathbf{u}}) \bar{\mathbf{u}}$ where $\bar{\mathbf{u}} = \begin{bmatrix} u_1 & u_2 & u_3 & u_4 \end{bmatrix}^T$ and

$$\mathcal{L}(\bar{\mathbf{u}}) = \begin{bmatrix} u_1 & -u_2 & -u_3 & u_4 \\ u_2 & u_1 & -u_4 & -u_3 \\ u_3 & u_4 & u_1 & u_2 \\ u_4 & -u_3 & u_2 & -u_1 \end{bmatrix}$$

Using this coordinate transform and its associated properties allows (60) to be rewritten without singularities as

$$\bar{\mathbf{u}}'' - \frac{1}{2}h\bar{\mathbf{u}} = \mathcal{L}^T(\bar{\mathbf{u}})\mathbf{B}\mathcal{L}(\bar{\mathbf{u}})\bar{\mathbf{u}}' + \frac{1}{2}(\bar{\mathbf{u}}\cdot\bar{\mathbf{u}})\mathcal{L}^T(\bar{\mathbf{u}})\bar{\boldsymbol{\zeta}} \tag{61}$$

where $h = \frac{1}{2}(\dot{\bar{\boldsymbol{r}}}_0\cdot\dot{\bar{\boldsymbol{r}}}_0) - \frac{\mu}{\|\bar{\boldsymbol{r}}_0\|}$ is a constant of integration, and subscript 0 denotes initial values. Initial conditions in the new four dimensional position space are defined based on the value of $x-1+\mu$.[26] If $x-1+\mu \geqslant 0$, then $u_{1,0} = \sqrt{\frac{1}{2}((x_0-1+\mu)+\|\bar{\boldsymbol{r}}_0\|)}$ and $\bar{\mathbf{u}}_0 = \begin{bmatrix} u_{1,0} & \frac{y_0}{2u_{1,0}} & \frac{z_0}{2u_{1,0}} & 0 \end{bmatrix}^T$. Otherwise, $u_{2,0} = \sqrt{\frac{1}{2}(\|\bar{\boldsymbol{r}}_0\| - (x_0-1+\mu))}$ and $\bar{\mathbf{u}}_0 = \begin{bmatrix} \frac{y_0}{2u_{2,0}} & u_{2,0} & 0 & \frac{z_0}{2u_{2,0}} \end{bmatrix}^T$. The derivatives can be related through $\bar{\mathbf{u}}_0' = \frac{1}{2}\mathcal{L}^T(\bar{\mathbf{u}}_0)\dot{\bar{\boldsymbol{r}}}_0$. With initial conditions properly defined, (61) can be numerically integrated to yield a time history of $\bar{\mathbf{u}}$ and $\bar{\mathbf{u}}'$. Final values in the original coordinate system are then recovered by $\bar{\boldsymbol{r}} = \mathcal{L}(\bar{\mathbf{u}})\bar{\mathbf{u}}$ and $\dot{\bar{\boldsymbol{r}}} = \frac{2}{\|\bar{\boldsymbol{r}}\|}\mathcal{L}(\bar{\mathbf{u}})\bar{\mathbf{u}}'$.