

Cislunar Orbit Determination: Improvements in Uncertainty Realism and Data Fusion

**C. Channing Chow II¹, Jason, Baldwin², Charles J. Wetterer³, Micah Dilley³,
Keric Hill³, Paul Billings³, Christopher Craft³, and James Frith⁴**

¹Cloudstone Innovations, ²Complex Futures, ³KBR, and

⁴Air Force Research Laboratory

ABSTRACT

The high levels of nonlinearity evident in the cislunar domain present stressing challenges when estimating the trajectories for and maintaining custody of cislunar objects. In our previous paper, we demonstrated the use of the Unscented Kalman Filter (UKF) and Gaussian Mixture Models (GMM) in the orbit determination (OD) and tracking of objects in cislunar space using astrometric and photometric observations of non-Keplerian orbits that arise due to a more pronounced influence of the Moon's gravity field. It was also shown that the use of high-fidelity 3-body dynamics is essential in capturing accurate representations of the motion of objects in cislunar space. The UKF provided relatively good results for a variety of orbits and data cadences, while the GMM demonstrated promise in improving performance through large data gaps. In this paper, we expand our cislunar OD research in two areas: (1) we apply the Adaptive Entropy-based Gaussian-mixture Information Synthesis (AEGIS) approach to improve the uncertainty realism within the estimation filter, and (2) we compare filter performance using different observation types independently, namely electro-optical (EO) and passive radio frequency (PRF), and then show improvements with their data fusion. In the first OD area, performance using AEGIS is compared directly to use cases from our previous paper in which a single UKF and fixed GMMs were used. The orbits under consideration, here, are examples from three separate elemental periodic orbit families that have been differentially corrected in a high-fidelity dynamical system: H1, L1, and W4W5. The individual orbits sample a range of possibilities within each family, including orbits near theoretical bifurcation points between families (i.e., where a single orbit belongs to both families). In the second OD area, filter performance improvement is achieved primarily by reducing or eliminating large data gaps in the EO data by filling in with PRF data. The efficacy of the strategy is justified since PRF observations are still possible during times when EO observations are excluded due to EO constraints. This result is derived from processing synthetic observations from the four orbits from our previous paper that were the most challenging due to large data gaps. We model and simulate hypothetical, Earth ground-based, EO and PRF sensor networks to generate the expected right ascension (RA), declination (DEC), and visual magnitude (VMAG) measurements, and the time difference of arrival (TDOA) and frequency difference of arrival (FDOA) measurements, respectively. All filters are constructed from components using the modular Infinity Filter Framework (IFF).

1. INTRODUCTION

Our previous papers provided an overview of the expected astrometric and photometric measurements of objects in cislunar orbits [1] and a preliminary demonstration of the use of Earth ground-based observations to track objects in these cislunar orbits using estimation filters [2]. This previous work laid the foundation for expected observability for objects at these distances, the importance of using full ephemeris models when predicting trajectories, and the potential challenges that estimation filters may encounter when performing track updates, particularly when dealing with extended data gaps. Given our current understanding of the observability conditions for cislunar satellites, compounded by the dimness of objects at these distances, it is likely that extended data gaps will be an unavoidable challenge when observing the cislunar domain with ground-based electro-optical (EO) sensors.

This work expands upon the previous foundation to investigate three questions:

- 1) What are the characteristics of the time evolution of uncertainty in this domain over extended data gaps?
- 2) Can higher-fidelity uncertainty representations improve key operational considerations?
- 3) Might the use of additional observation modalities, specifically the fusion of EO and passive radio frequency (PRF) measurements, materially aid the object tracking problem for cislunar?

Distribution A: Approved for public release; distribution is unlimited. Public Affairs release approval AFRL-2022-3824.

To provide some background, it was demonstrated that the Unscented Kalman Filter (UKF) had success in processing EO observations of some cislunar objects, even in the presence of data gaps extending over multiple weeks [2]. Subset #3 in that work used a single observing site and a visual magnitude cutoff which led to most orbits experiencing upwards of 14-16 days without measurements. Despite the successes in Subset #3, there were a few orbits in the analysis which presented difficulty, wherein the UKF lost track of the object. A Gaussian Mixture Model (GMM), composed of multiple UKF components, was then employed to enhance the uncertainty representation. While it showed the ability to outperform the single UKF in many cases, there remained a few orbits for which the GMM and UKF-based filters still lost track.

Understanding the failure mode(s) experienced by the filters is an important task. The next logical step is to investigate the use of richer uncertainty representations, with a goal of achieving increased realism in the representation to help overcome these filter failures. Increased fidelity may not be the full solution to the problem, however, so developing some understanding of the time evolution of the uncertainty in this domain is a critical insight. Section 2 of this paper will take a brief look at this evolution. Given these results, a discussion of the new test parameters and the comparison of results across different filter types, to include an adaptive approach, will be covered in Section 3. The fusion of PRF with EO observations is shown to improve performance in this challenging domain for the use cases covered in Section 4. Finally, we return to the three questions posed above to frame the findings, draw conclusions, and discuss potential future work in Section 5.

2. EVOLUTION OF UNCERTAINTY IN CISLUNAR

In space domain awareness (SDA), an analyst's/operator's intuition is based largely on assumptions grounded on perturbed, restricted 2-body dynamics (e.g., orbits resemble simple conic sections, like ellipses). Expanding SDA to reach further into cislunar space requires new intuitions be built in the perturbed, restricted 3-body dynamical system. Specifically, it is important to understand the behavior or time evolution of objects and phenomena in this more complex domain, where the traditional simplifying assumptions are no longer valid (e.g., orbits are non-Keplerian).

For example, in LEO and GEO, the evolution of the position uncertainty trends towards deformations consistent with Keplerian motion (e.g., smears out largely in the velocity direction). Similar intuitions must be developed with respect to uncertainty evolution in cislunar space.

To seed the analysis, we return to the exemplar cislunar halo orbit of H1(100N) from the previous paper [2]. Observations from Subset #3 were processed for several days, producing a well-conditioned covariance to serve as the initial condition for propagating over multiple weeks. A Monte Carlo analysis was then performed to sample from this covariance initial condition, and these samples were propagated for 14 days. A quasi-random sampling technique based on Sobol sequences [5-6] was chosen to improve the space-filling characteristics for the Monte Carlo analysis and reduce the effective number of samples needed to achieve a fairly accurate representation of the underlying probability density function (PDF). Figure 1 shows the resulting 10,000-sample Monte Carlo distribution corresponding to the initial position covariance in the synodic frame. Monte Carlo samples are colored magenta, and the corresponding orthogonal projections are in light gray. The reference frame is the Earth-Moon synodic frame (+Z-axis is normal to the Moon's orbital plane, +X-axis is along the Earth-Moon line, and +Y-axis completes the orthonormal basis). Note that the extent of the uncertainty initial condition is predominately along the synodic X-axis, which effectively serves as a proxy for the radial direction.

Figure 2 shows propagated evolutions for this initial condition. The left image shows the Monte Carlo samples after being propagated in GMAT [7] for 12 days, while the right image shows the Monte Carlo samples two days later, or after being propagated for 14 days. The dynamical systems configuration in GMAT includes a full-ephemeris model (JPL's DE405) with the Earth as an aspherical gravitational body with low degree and order spherical harmonics, the Moon and Sun as point-source gravitational bodies, and solar radiation pressure effects. The blue line shows the full trajectory of the H1(100N) exemplar orbit from which the state was initially chosen, and the magenta points continue to signify the Monte Carlo samples with their orthogonal projections in light gray. In these figures, the distances are normalized by the Earth-Moon distance; hence, the instantaneous Moon distance is one synodic length unit.

The PDF of the position, as represented by the Monte Carlo samples, is being stretched in the directions of the gravitational bodies of the 3-body system. This intuition is important as we discuss the failure modes for previous filtering results.

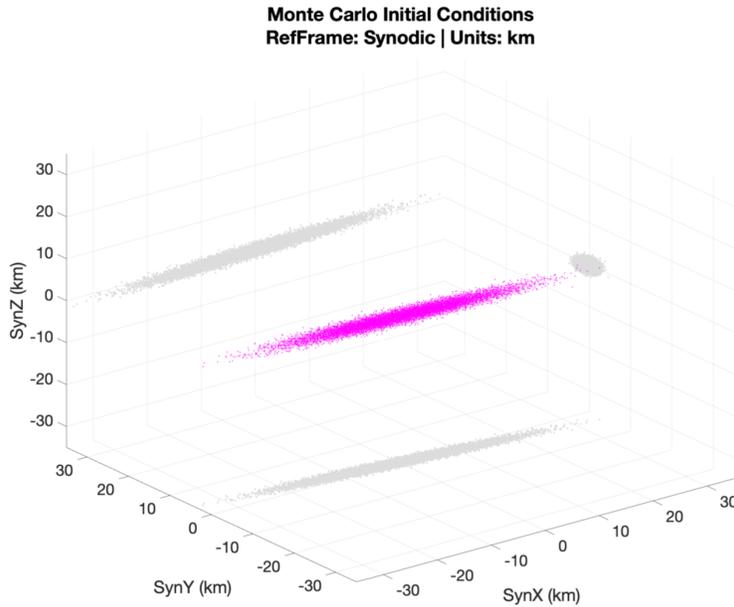


Figure 1. Initial Monte Carlo distribution for H1(100N) starts with a larger uncertainty in the radial direction. Monte Carlo particles are colored magenta, and their orthogonal projections are in light gray.

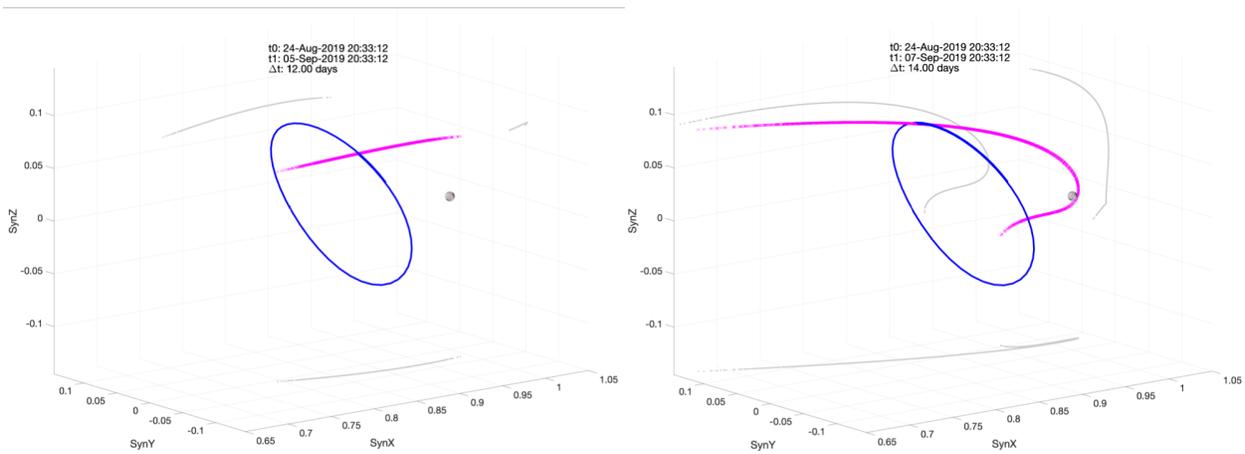


Figure 2. Monte Carlo propagated cloud (left) at 12 days versus (right) at 14 days, shows the rapid deformation of the distribution as the particles come closer to the Moon's gravity well. Monte Carlo particles are colored magenta, their orthogonal projections are in light gray, and the full reference orbit for H1(100N) is in blue.

It is noteworthy how rapidly the Monte Carlo samples transition from “well-behaved” after 12 days of propagation to “less well-behaved” after 14 days. In other words, the PDF quickly grows more nonlinear as its proximity to the Moon increases, causing significant distortions to portions of the underlying PDF.

It is precisely this lunar gravity effect that created the failure modes in previous filtering attempts. This effect was found to exist across all evaluated Lyapunov and Halo differentially corrected orbits. To highlight one example, Figure 3 shows the Monte Carlo results for H1(1N), where the initial condition was drawn from the UKF filter results for H1(1N) going into the first data gap while processing data from Subset #3. The length of the data gap in the observations for this test case is 15.7 days, and the snapshot shown in Fig. 3 is 13.5 days into the data gap. As the PDF sweeps through the lunar gravity well, some Monte Carlo samples “collide” with the Moon. Given this outcome it is

Distribution A: Approved for public release; distribution is unlimited. Public Affairs release approval AFRL-2022-3824.

expected that the UKF will lose track of the object. At the points where the PDF collides with the Moon, the representation of the PDF achieved with a limited number of sigma points (in the case of a UKF) breaks down to the point where even the inclusion of additional observations is unable to redeem a viable state estimate; thus, the track is lost.

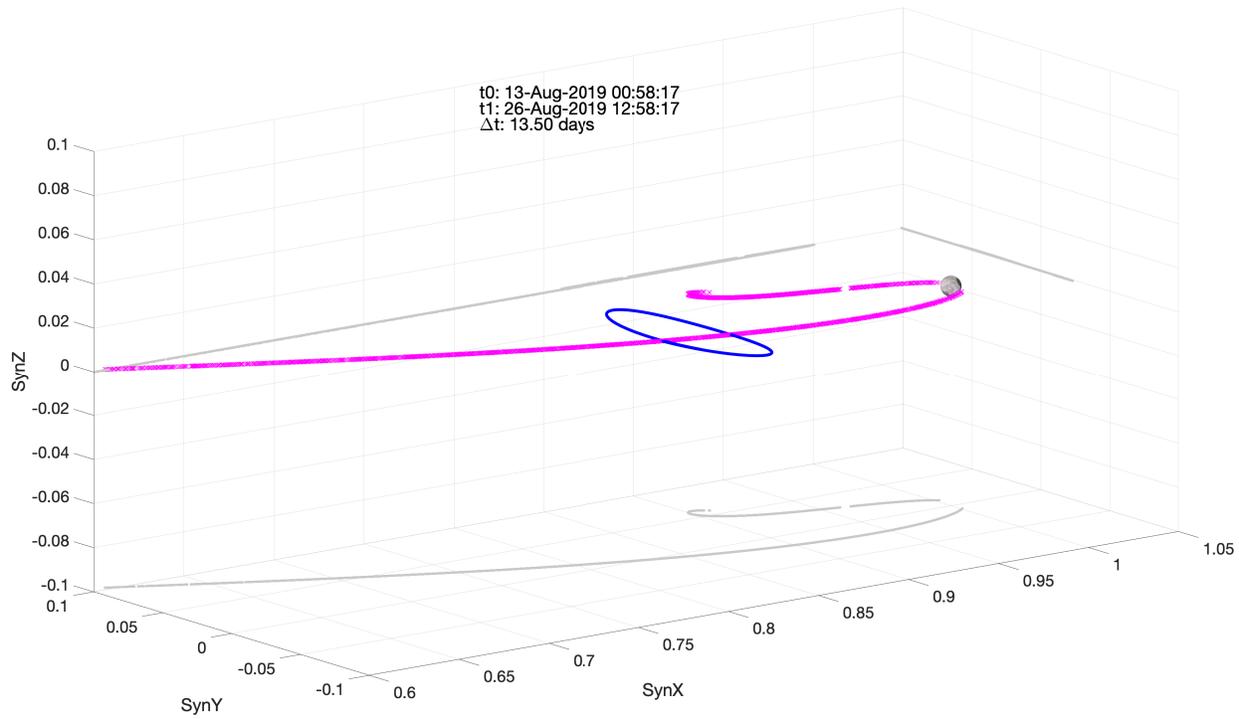


Figure 3. Monte Carlo propagated for 13.5 days shows portions of the PDF colliding with the Moon. Monte Carlo particles are colored magenta, their orthogonal projections are in light gray, and the full reference orbit for HI(1N) is in blue.

Different features of the initial conditions interplay to determine where, when, and how strongly/weakly perturbed the PDF gets as it encounters the Moon. Features of particular note include the structure and stability of the orbit family under consideration, the size and shape of the initial uncertainty, and the location of the object (i.e., timing) within an orbit. Again, due to the inherent instabilities of the cislunar regime, small changes have the potential to significantly impact the results.

Despite how powerful the UKF and GMM are in producing effective uncertainty representations in highly nonlinear environments, it will be necessary to achieve richer and more realistic representations to overcome some of the challenges discussed above. In some (perhaps many) cases, the UKF will provide ample capability. The key may be in pursuing adaptive filtering capability that can provide insight on the difficulty of the experienced dynamical conditions and to add complexity when needed to meet the demands of those conditions. Adaptive Entropy-based Gaussian-mixture Information Synthesis (AEGIS) is one such approach and is discussed in the next section.

The objective here is to highlight the strengths and weaknesses of all filter strategies under consideration. To create a baseline, we use a test condition under which all filtering strategies could successfully maintain track through an extended data gap. The tracks can then be assessed for quality and the extent to which the achieved uncertainty realism is impacting critical operational considerations. The details of this new use case are presented next.

3. AEGIS

The way we implement AEGIS is an adaptive extension to previous methods [4]. To keep the estimation problem confined to a near-Gaussian regime, AEGIS performs regular checks on the covariance during propagation and

conducts a Gaussian-mixture decomposition into independent Gaussian components in the presence of a deteriorating Gaussian condition. The metric used to monitor this Gaussian condition is the differential entropy, which is calculated for a Gaussian PDF as

$$H = \frac{1}{2} \log |2\pi e \mathbf{P}|, \quad (1)$$

where \mathbf{P} is the component state covariance. The differential entropy, H , is thus a scalar measure of the randomness or surprisal in the system (covariance), with an increasing differential entropy signifying a likely departure from linear (or Gaussian) behavior.

The original formulation for AEGIS suggested the use of two simultaneous filters, one linear and one nonlinear¹. The differential entropy for each of the two filters can be compared, and then a threshold can be set for the maximum desired divergence between the linear and nonlinear cases to signify a departure from linear behavior [4]. In cislunar, establishing a well-behaved linear filter is exceedingly difficult, so the implementation of AEGIS presented here sets a threshold on the allowable increase in differential entropy. In our analysis we set that threshold to 5 nats,² although different values for the entropy threshold were explored. (This choice satisfies the desire to perform enough Gaussian-mixture splits to keep all filter components in the vicinity of the Gaussian regime while not splitting so much that the computational load becomes burdensome.) In this light, the threshold for change in differential entropy becomes a filter tuning parameter.

An adaptive, discretized (component-wise) approach like AEGIS is an effective blend of adding complexity when necessary, managing computation as the PDF becomes richer in representation, and facilitating more capacity to query the structure of the evolving PDF for improved decision making. All this added capability comes while maintaining the benefits of exploiting Gaussian processes for linear correlations, data association, Kalman state update, and many other features enabled by remaining sufficiently near the confines of a Gaussian process for each filter component. A suitably tuned AEGIS serves as an effective “sweet spot” in the trade space between the efficiency, nonlinear capacity, and limited representational richness of the UKF versus the realism, richness, yet extreme computational burden of particle filters and Monte Carlo methods. It’s an adaptive shape shifter, transforming itself across this spectrum as the complexity of the dynamics model demands.

3.1 Infinity Filter Framework

Numerous modifications were required to the Infinity Filter Framework (IFF) modules to adapt it from the original UKF and GMM implementations and make it possible to implement AEGIS. To name a few, the new initialization parameters include the cadence of entropy checking, process noise application options, the entropy splitting threshold, the number of components per Gaussian-mixture split, the component pruning and merging tolerances.

¹ Here, linear and nonlinear refer to how the estimation engine approximates and thus represents the true dynamics.

² A “nat” is the natural unit for information entropy. It is similar to a “bit” but is based on natural logarithms and powers of e rather than base 2 logarithms and powers of 2.

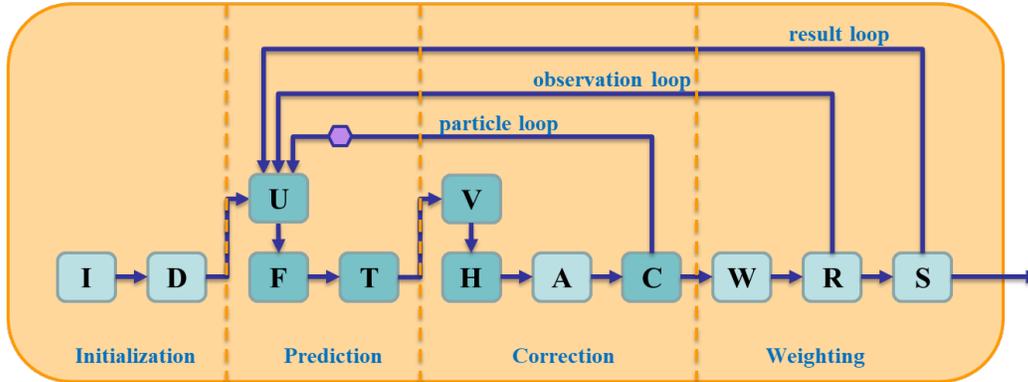


Figure 4. IFF module schematic showing how any estimation engine can be implemented as a composition of these building blocks. On color shading: darker = parallel processes, lighter = serial processes.

Figure 4 displays the IFF module schematic. The functionality of the sigma-point generation, propagation, and time update (modules U, F, T, and V) need to be executed in a repetitive manner to allow for the monitoring of the differential entropy multiple times throughout propagation between observations. To accommodate this requirement, a new F-module was created that absorbed and iterated over the functionality of the original U, F, T, and V modules at shorter timescales. Furthermore, the R-module, which handles resampling, was modified to handle component pruning and merging logic [3]. This example serves to illustrate the power and flexibility of the IFF, as AEGIS was systematically implementable without modification to the underlying framework or wholesale development of a separate filter altogether. Creating AEGIS from existing UKF and GMM components in the IFF only required the intuition of the developer to creatively design one or two modules, as other modules such as data association, measurement update, and weight update were reusable in their current forms.

Two additional noteworthy modifications are the handling of adaptive component creation and process noise. For component handling, two options were entertained: (1) create new IFF processing “particles” for every component, or (2) stack/contain all new components within a single processing particle. Since numerous advantages were found by employing a component stacking approach within a single particle, we selected the latter option as the strategy for our AEGIS implementation. The full treatment and application of process noise is a deep topic beyond the scope of this paper. Suffice it to say that for the purposes of this analysis, the traditional approach was used where process noise was accumulated and applied just before a measurement update; the same approach was applied across all filter types for consistency.

3.2 Orbits, Observations, and Test Conditions

Several cislunar orbit families from our previous work [2] are considered again here as the target spaces for our analysis: the planar L1 (long-period Lyapunov) family emanating from the $\mathcal{L}1$ Lagrange point, the H1 (halo) family which bifurcates from the L1 family at the L11 branch point, and the W4W5 (continuous east-west axial) family connecting the $\mathcal{L}4$ and $\mathcal{L}5$ Lagrange points and intersecting the H1 family at the H11 branch point [2]. As before, we down-select from these orbital families to focus on 15 individual orbits, each of which is then differentially corrected to find the nearest nearly-periodic orbit in the full ephemeris model. The propagator used to generate these differentially corrected orbits is TurboProp.

Some test conditions are modified from earlier work in this analysis. For example, the observations from the limited-EO data set (Subset #3) from the previous study is modified to have an increased visual magnitude threshold less than 20.5. This adjustment serves to shorten the gaps in the data to ensure the successful tracking of all objects across all filter types (UKF, GMM, AEGIS). In addition, all data prior to the first multi-day data gap is removed, leaving most of the 15 orbits with an observation profile containing a single large data gap similar to the one depicted in Figure 5. The W4W5 orbits are the exceptions that span much longer time frames and thus lead to numerous data gaps.

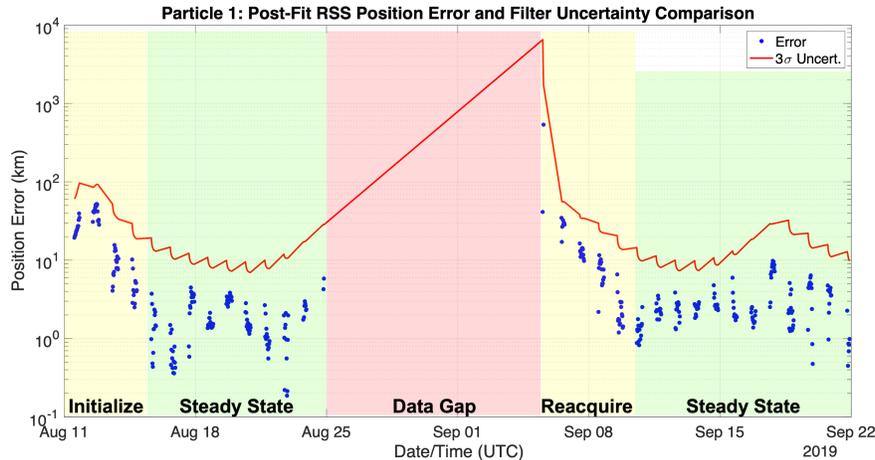


Figure 5. Post-fit RSS position error for H1(100N) using UKF; colored track segments notional.

Figure 5 shows the post-filter update root sum square (RSS) position errors for the H1(100N) orbit for the UKF run. This example highlights the notional segments of the track: initialize, pre-gap steady state, data gap, reacquire, and post-gap steady state. For the error distributions presented in Section 3.3, the initialize segment is discarded to eliminate potentially high error values that are the result of the track initialization process. The error distributions across the three filter types focus solely on the two steady state segments and the reacquire segment. The data gaps in question, given the new visual magnitude cutoff, now range from between 11.5 to 13.5 days, effectively reducing the gaps from the previous analysis by an average of two days. All filters across all orbits now successfully maintain track to provide a common baseline for comparison.

3.3 Error Comparison

Figure 6 shows a swarm plot of the RSS position error distributions across all filter types and orbits. The swarm plot is similar to a box plot, but it shows all the data points in the distribution instead. This swarm plot, specifically, is a composition of only the steady state and reacquire segments of all the RSS position error data (with the time dimension suppressed). Additionally, we add visual artifacts, such as the median lines (green) and 95th quantile lines (red) for each distribution to better mark key points. The median serves as a suitable proxy for the steady state error. Note that almost all medians reside near each other and fall anywhere from one to five kilometers of error. This result is expected, as the steady state segment for all three filter types is being handled by a UKF. This statement will be made clearer in Section 3.5. The steady state error is essentially a function of the measurement noise in the observations (e.g., one arcsecond) and the density/sparsity of the measurements.

The maximum error and the 95th quantile are reasonable indicators for how well the filters are performing during the data gap and reacquisition segments. Particularly for the difficult L1 and H1 family orbits found in last year's paper [2], AEGIS dramatically reduces the maximum error, in some cases by an order of magnitude. Also, the 95th quantile highlights how the reacquisition segment for these same orbits improves over the UKF and GMM, indicating a reacquisition phase of the tracks encumbered by less error (hence, likely a better and more realistic representation of the uncertainty condition). However, not every case where AEGIS is used shows improvement. AEGIS performs noticeably poorer during reacquisition of H1(400N) and performs essentially the same on orbits such as H1(300N) (see Figure 6). These results underscore the importance of the pruning and merging logic of the components.

Note that the state estimate – and therefore the error values – for a component-driven filter strategy (such as GMM and AEGIS) can be derived in multiple ways. To name a few, consider the weighted average across existing components, the weighted average across only “high-weighted” components above a given threshold, or simply the highest-weighted component. For this analysis, the highest-weighted component is utilized. An in-depth analysis, from a decision intelligence perspective, to understand the differences between these methods for assessing uncertainty and error for the PDF is warranted.

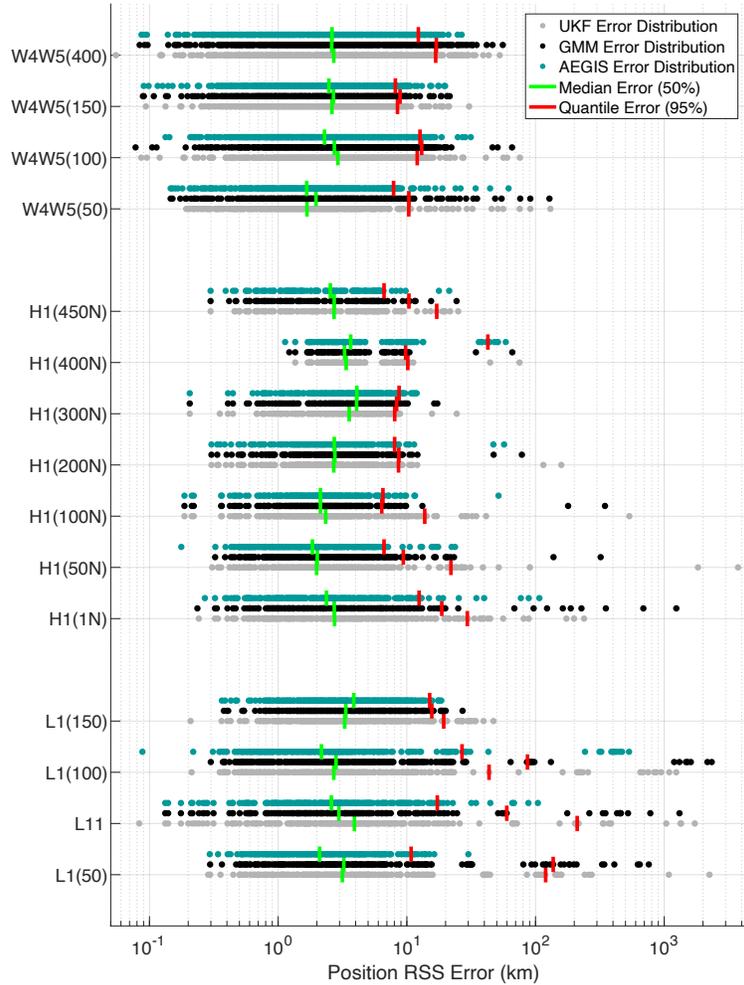


Figure 6. Swarm plot of the RSS position error distributions for the UKF, GMM, and AEGIS for all evaluated orbits showing how AEGIS (teal dots) dramatically reduces the maximum error and 95th-percentile (red lines) for H1 and L1 families, in some cases by an order of magnitude. The comparable median values (green lines) indicate similar steady-state performance (post-reacquisition).

3.4 Reacquisition and Filter Consistency

This section looks closer at the reacquisition phase and filter consistency by inspecting the error and error covariance behavior for selected examples. Figure 7 presents the RSS position errors and the three-sigma uncertainty for orbit L1(50). This orbit highlights a case where there is significant AEGIS improvement (as indicated by the error PDF in Figure 6). It is noteworthy that the AEGIS errors and uncertainty are lower than the UKF by over an order of magnitude during reacquisition. The AEGIS result on the right is for the highest weighted component at each time step. Also note that the effective reacquisition timeline for the AEGIS filter is merely a couple nights, where the UKF filter is still effectively in reacquisition after five nights of observations post-data gap.

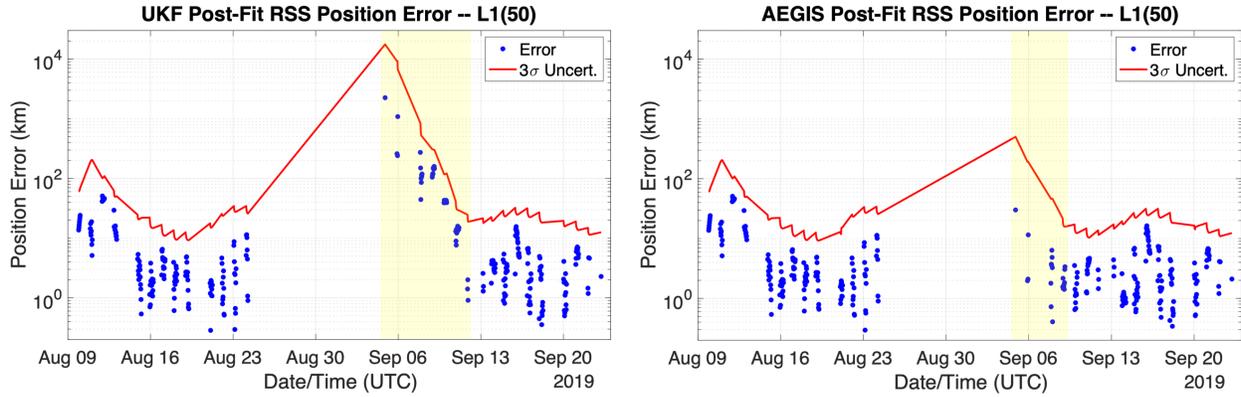


Figure 7. Post-fit RSS position error comparisons for L1(50) using UKF (left) and AEGIS (right) over the evaluation period. Reacquisition periods following the data gap are highlighted in yellow; based on the errors in this period, AEGIS is seen to reacquire faster than the UKF.

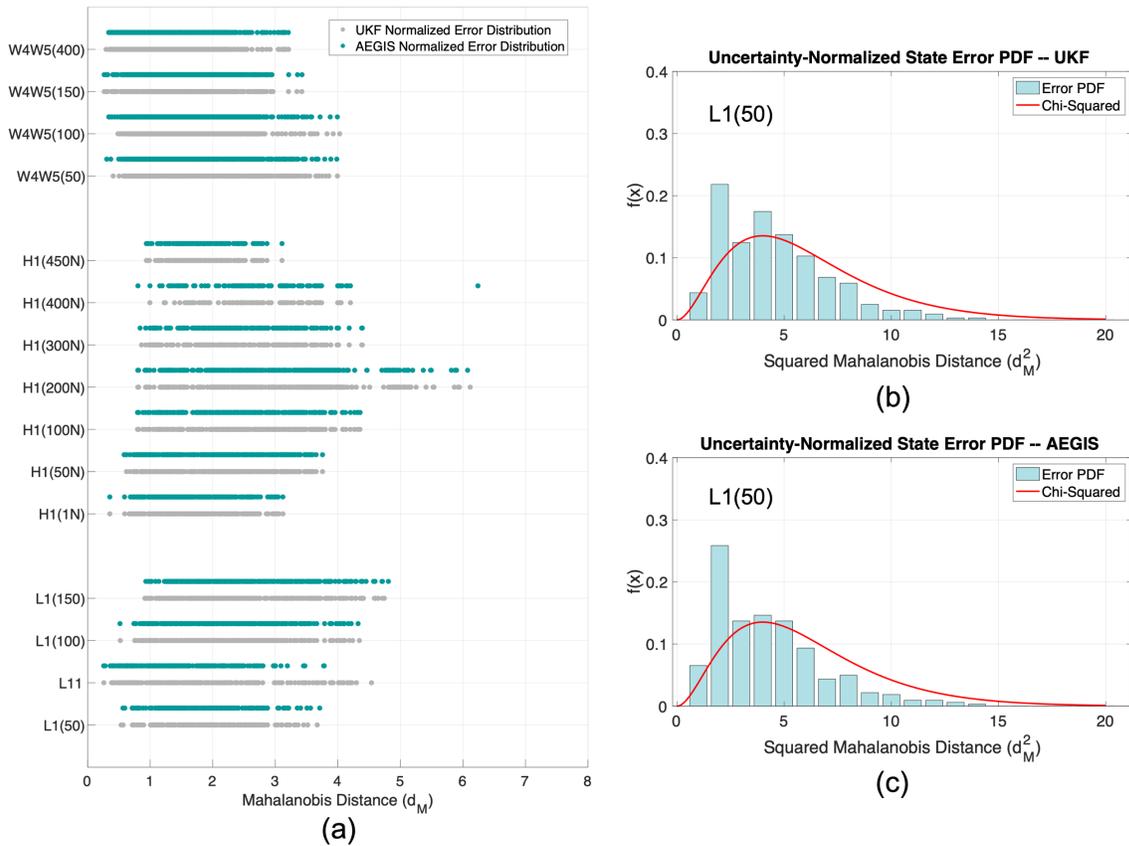


Figure 8. Reasonable consistency is seen between the UKF and AEGIS as evidenced by (a) Mahalanobis distance distributions for all orbits showing minimal variation between filter types, squared Mahalanobis distance and chi-squared comparisons for L1(50) for (b) UKF and (c) AEGIS.

A useful metric for assessing consistency is the covariance-normalized error behavior – i.e., the Mahalanobis distance. Figure 8 shows an analysis of the Mahalanobis distance, a multi-dimensional generalization of the idea of measuring standard deviations from a mean. The Mahalanobis distance normalizes the squared error by the covariance according to

$$d_M = \sqrt{(X - X^*)^T P^{-1} (X - X^*)} \quad (2)$$

where X is the estimated state vector, X^* is the “truth” state vector, and P the state covariance. Figure 8a shows the Mahalanobis distance distributions for both the UKF and AEGIS for all orbits. Similar to before, the errors from the initialization phase of the filter are removed from this analysis. Minimal error variation between filter types implies that the filters are similar with respect to consistency. For a process dominated by Gaussian noise (which is true for this simulation), the square of the Mahalanobis distance should follow a chi-squared distribution. Figure 8b and Figure 8c further show consistency from both filters, meaning that the estimated covariance is commensurate with the errors. Considering the (non-normalized) RSS errors, however, the nod must be given to AEGIS in this stressing case due to its improved realism. One additional take-away from the chi-squared analysis is that both filters may be a bit conservative for this orbit given the departure from chi-squared behavior at low values for the squared Mahalanobis distance.

Another approach for exploring uncertainty realism is mapping the uncertainty into measurement space to build intuitions for how the implementation of AEGIS could impact Earth-based sensor tasking for reacquisition following large data gaps. Any improvement in uncertainty realism should yield reduced sky collections to reacquire a satellite. We shift to orbit H1(50N) to inspect the results against a stressing Halo family orbit. Referring back to Section 2, we use a quasi-random Monte Carlo sampling of the uncertainty for H1(50N) that achieves a highly realistic uncertainty representation that can be propagated for the duration of the data gap, which in this case is approximately 11.72 days. At the end of the data gap, marginal PDF evaluations for the UKF and the components of AEGIS are used to generate contours for the filter uncertainty. The Monte Carlo samples and the marginal PDFs are then mapped into measurement space to portray the uncertainty condition for the Earth-based sensor tasked with collecting on the satellite to end the data drought.

Figure 9 showcases the improved realism achieved by the AEGIS filter’s representation of uncertainty. The contours represent the UKF uncertainty (left) and AEGIS uncertainty (right), the Moon is in the center, the red asterisk is the true satellite position in measurement space, and the black markers are the Monte Carlo samples, showing the condition that the uncertainties find themselves for this H1(50N) data gap.

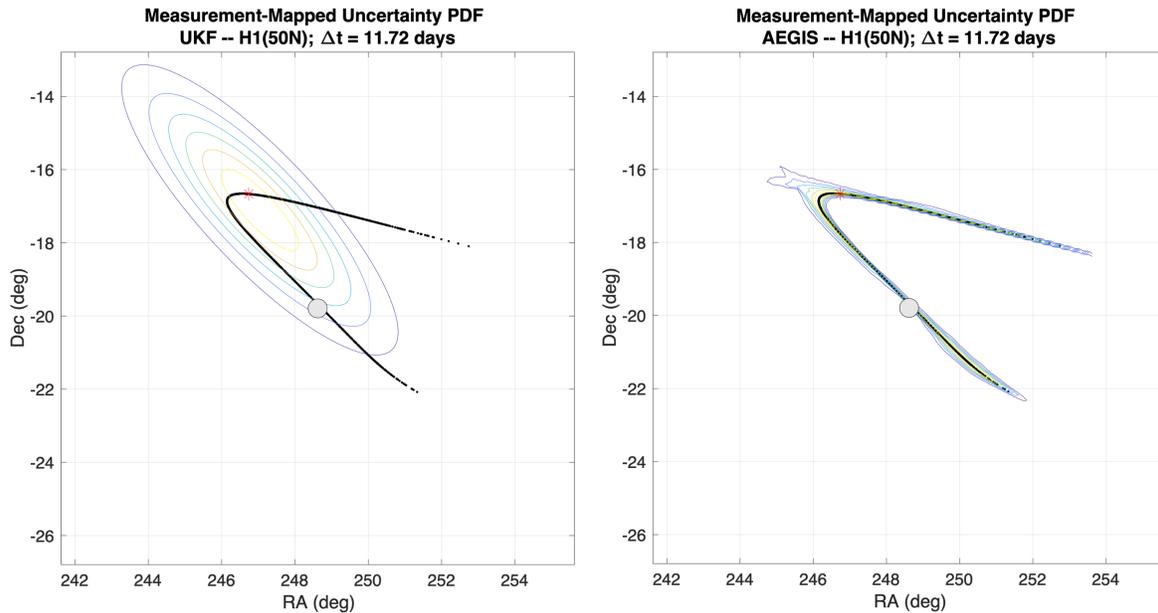


Figure 9. Measurement-space mapping of uncertainty PDFs for H1(50N) at the end of the data gap shows significantly less area to consider due to the improved covariance representation realism achieved by the AEGIS filter; UKF (left) and AEGIS (right); propagated Monte Carlo sampling in black and Moon in center.

3.5 Component Entropy Behavior

Perhaps the most powerful aspect of a discretized PDF representation is the ability to monitor the health of the individual components. Sticking with H1(50N), Figure 10 shows the entropy values for all 1800+ components at the end of the data gap. This dramatic increase in components comes with an increased computational load, with certain AEGIS runs taking hours compared to minutes for the UKF counterpart. The blue levels for the bars signify proximity to the Moon, with darker blues indicating the component is nearer the Moon. Note that the vast majority of the components are near the Moon with higher entropy levels. The increase in nonlinearity experienced as the components approach the Moon is leading to entropy accumulation, more frequent splitting, and runtime cost.

Figure 11 shows the time history of the number of components (top) and the time history of the entropy value ($H(t)$) and entropy split baseline ($H(0)$) for a select component (bottom). AEGIS is performing a split every few days for the first 9-10 days of the data gap, but then the number of components goes exponential as significant numbers of components approach the Moon. An operator could lean on these insights to help determine priority for sensor tasking as the entropy conditions for the components begin to show stress. Also note that on either side of the gap during the “steady state” segments, the number of AEGIS components is one, and thus the filter is operating as a single UKF.

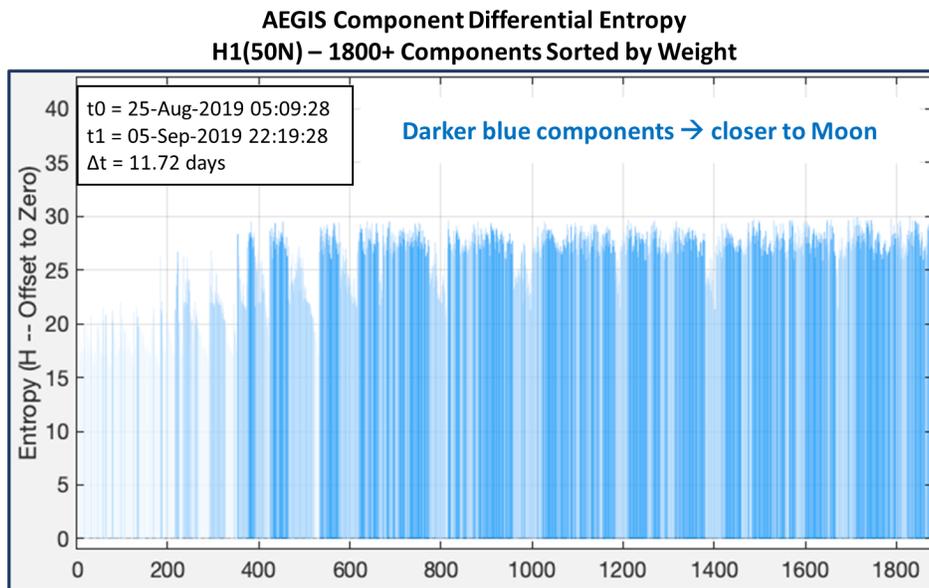


Figure 10. Component-level entropy values across the entire PDF at the end of the H1(50N) data gap; darker blue levels signify proximity to the Moon showing how a significant number of components are generated because of their proximity to the Moon.

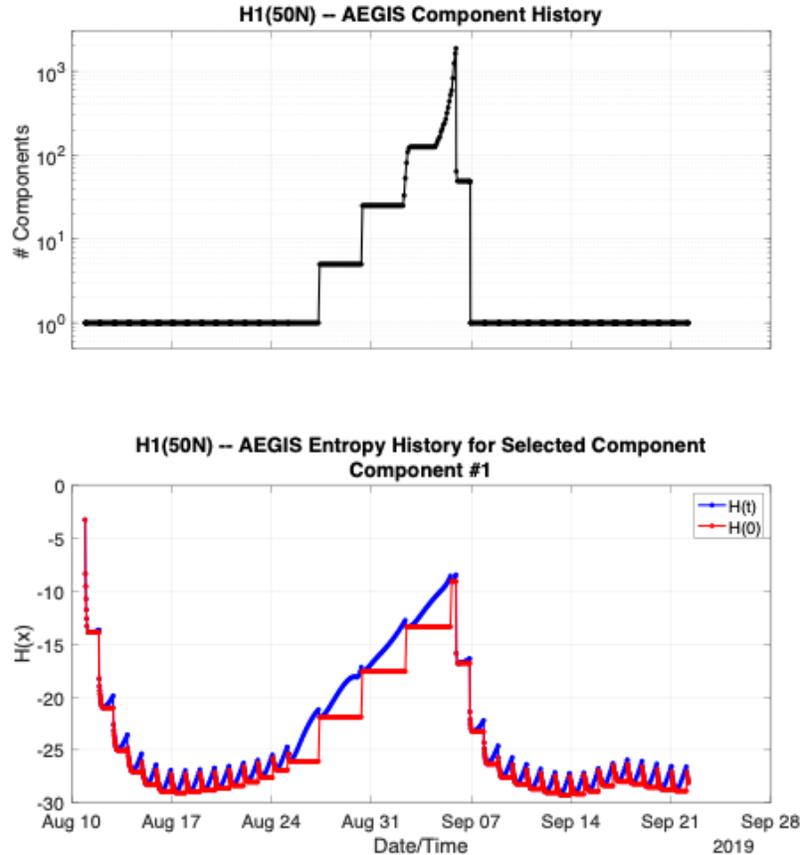


Figure 11. AEGIS component history may help to determine priority for sensor tasking as the entropy conditions for the components begin to show stress. Total components (top) and entropy values for component #1 (bottom). Number of components increases as entropy increases, by design, and both before and after the data gap, the number of components is one, indicating AEGIS is operating identically to a UKF.

4. DATA FUSION WITH PASSIVE RF

The extended data gaps present in the previous analysis are the result of periods of time when the cislunar object is close to the Sun in the sky for an Earth-based sensor. With this phase angle, the spacecraft is not only backlit, and thus fainter and more difficult to detect optically, but it also has a shorter interval during the day when it is visible to the sensor due to observability restrictions such as (1) minimum altitude for the object and (2) minimum angles associated with the Sun and Moon (i.e., exclusion zones). By utilizing a different sensing modality that is not as constrained by the locations of the Sun and Moon, the presence of extended data gaps can be reduced or avoided altogether. Additionally, a different sensing modality offers a different mapping of uncertainties between measurement space and state space that could yield better results in various situations and enhance results when fused with the other sensing modality. In this section we explore these topics and consider PRF observations both separately and in combination with EO observations.

4.1 Infinity Filter Framework

In the following data fusion analyses, a single UKF – as used previously – is employed. The processing differs only in the number and type of input measurements used per observation time step. Changes to the IFF’s H-module are necessary to process the new measurement type, in this case PRF’s time difference of arrival (TDOA) and frequency difference of arrival (FDOA) values. Other minor changes are made to the T, A, C and R modules to allow for and to distinguish between the different measurement types within the filter. For example, the code needs to identify all

measurement types in real time during processing (as opposed to assuming, say RA, as the first measurement) in order to correctly apply any regularization needed for each of the processed measurement types.

4.2 Observations and Test Conditions

For our data set, we use the observations generated for last year's study (single EO sensors) [2] along with new PRF observations generated from a notional Maui (20.7° N, 156.4° W) - Midway Island (28.2° N, 177.4° W) - Palmyra Island (5.9° N, 162.1° W) sensor network. The forward light-curve modeling software, Forge, is again used to generate the light curves for a diffuse sphere (e.g., 1 m² cross sectional area and 0.2 reflectance). The resulting PRF measurements are the TDOA and the FDOA values from each of the three different baselines (constructed as pairs of sensors), thus yielding six measurements per time step³. The simulated TDOAs are measured in seconds and are the result of the different distances between the sensors and the object given the finite speed of light. The simulated FDOAs are measured in Hertz and are the result of the different relative velocities between the sensors and the object. The Ku-band (12.5 GHz) is selected as the receiving frequency and the simulated Gaussian uncertainty is set to 10⁻⁸ seconds for TDOA and 0.03 Hz for FDOA with no timing bias and an assumed known tropospheric delay. As before, the simulated uncertainty for the EO observations is set to 1 arcsecond.

EO observations from a single site are used with a 1-hour cadence and associated thresholds (Subset #3 from [2]): object elevation is greater than 10 degrees, Sun elevation is less than -10 degrees (Sun exclusion), the object is greater than one degree from limb of Moon (Moon exclusion), and the object's visual magnitude is less than 20. New PRF observations are produced for the three-site network with a 1-hour cadence and associated thresholds: object elevation above zero degrees for all three sites, object greater than five degrees from Sun (Sun exclusion). Additionally, five different combinations of measurements are produced using identical observation times determined by the same thresholds as with the single EO sensor but applied to all three sensor sites of the PRF network. This time alignment is contrived to allow for the probing of information content per observation for the different sensing modalities. These five combinations are: (1) EO1 = EO from a single site, (2) EO3 = EO from the three sites simultaneously, (3) PRF = PRF from the 3-site array, (4) PRF+EO1 = PRF combined with EO from a single site, and (5) PRF+EO3 = PRF combined with EO from all three sites. Wider data gaps are present here, versus previous work, because of the additional constraint imposed to require all sites be able to observe the cislunar object simultaneously.

4.3 Direct Comparison

Figure 12 shows a direct comparison of the tracking performance between a single EO site and single three-sensor PRF network. A side-by-side comparison of the EO versus PRF position errors over time is presented to highlight the stark contrast of where the UKF loses/maintains track of the object due to the presence/absence of data gaps, respectively.

³ Assuming the object is observable from the three baseline pairs simultaneously.

EO1

PRF

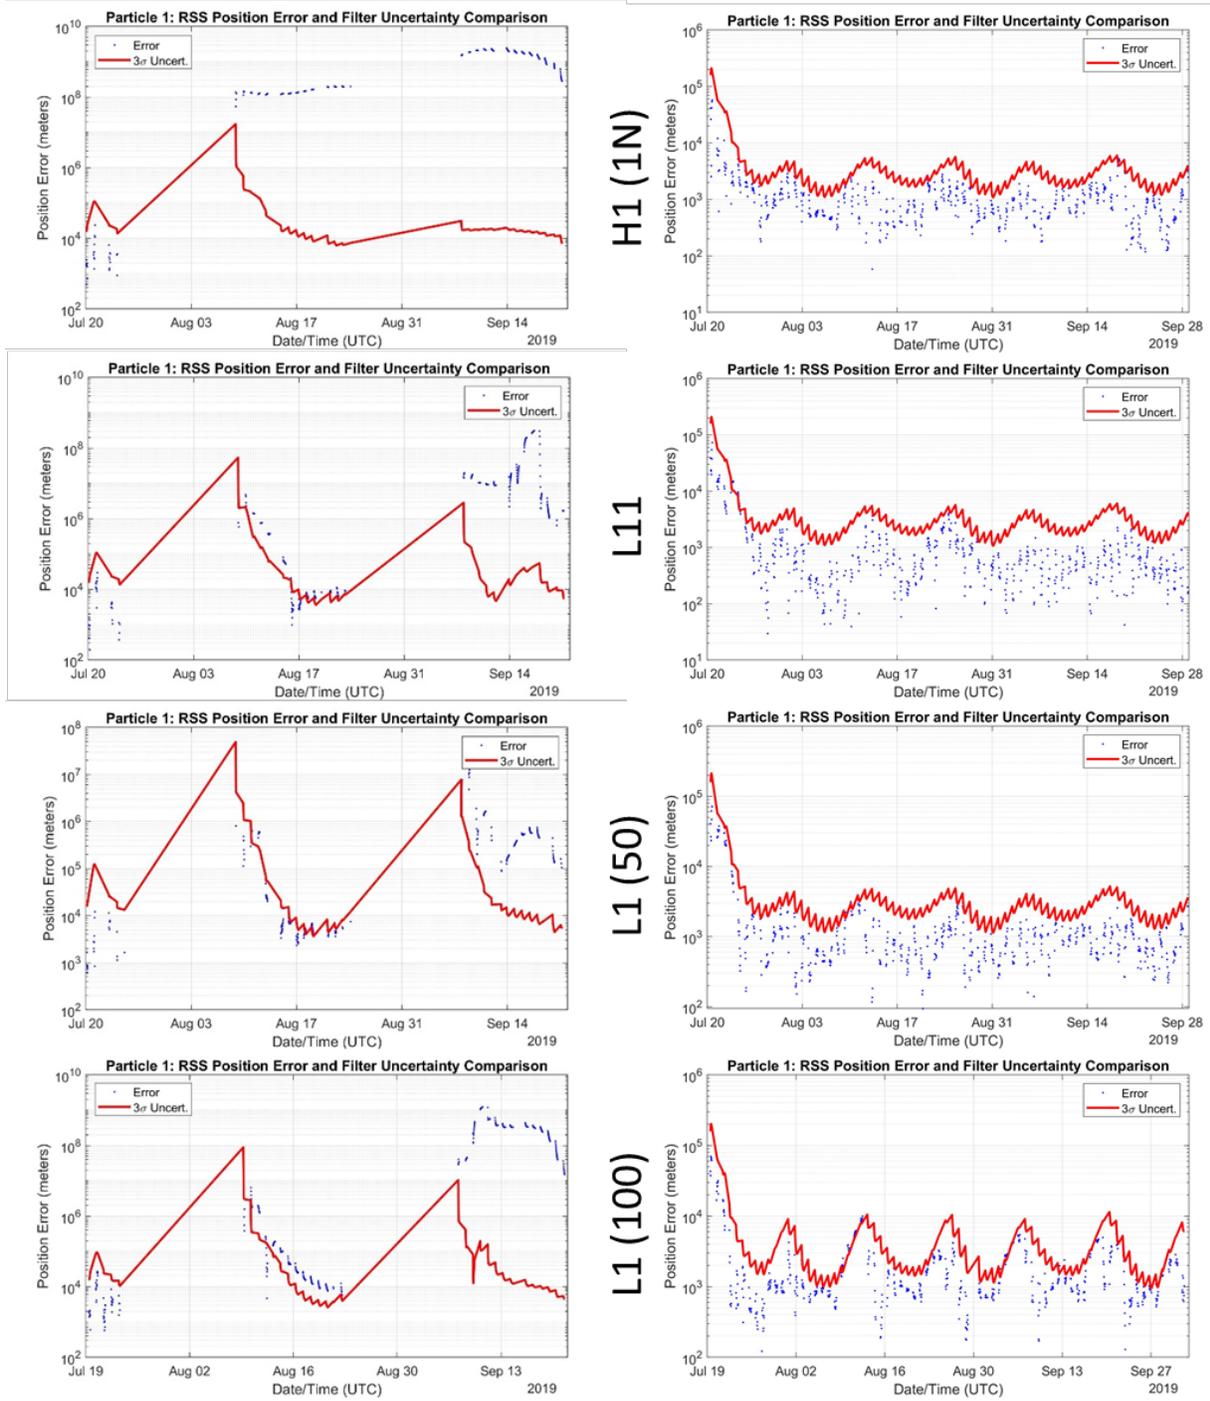


Figure 12. Direct comparison between EO-only and PRF-only, with associated exclusions, showing how elimination of data gap using PRF observations allows the filter to maintain better tracking.

4.4 Data Fusion with Identical Observation Times

Table 1 arranges the results from comparing measurement types with identical observation times (defined by the data gaps created in the EO3 case). The numerical value after each measurement type in the table column headings is the number of individual measurements at each observation time step. For example, for EO1, one RA and one DEC value (for a total of two measurements) are provided, while for PRF, a TDOA and FDOA for each of the three different baselines (for a total of six) are available. The numerical value after the descriptor in each box is a “goodness of tracking” measure as determined by examining plots of the position error as a function of time for each of the orbits and measurement types (see Figure 13 and Figure 14). After each extended gap, the goodness of tracking has the following scores: 0.0 = lost ($>10^6$ km position error), 0.5 = converging but mostly still $>10^6$ km position error and well outside 3- σ error bound; 1.0 = converging but mostly outside error bound ($>75\%$ of time); 1.5 = converging but sometimes outside error bound; 2.0 = converging and mostly inside error bound ($>75\%$ of time). One score is calculated per data gap. With two data gaps in this particular data sets, two scores are calculated and added together, and so the best possible goodness of tracking score, here, is 4. The final row of the table gives the average goodness of tracking score for the four orbits analyzed.

Table 1. Data fusion with identical observation times unsurprisingly shows better performance with smaller data gaps and/or more data; results color-coded according to scoring criteria.

orbit	EO1 (2)	PRF (6)	EO1+PRF (8)	EO3 (6)	EO3+PRF (12)
H1 (1N)	Lost after 1 st gap (0)	Lost after 1 st gap (0)	Lost after 1 st gap (0)	Successful (3)	Successful (2.5)
L11	Lost after 1 st gap (0)	Lost after 2 nd gap (1)	Successful (3)	Successful (4)	Successful (3)
L1 (50)	Lost after 2 nd gap (1.5)	Successful (3)	Successful (4)	Lost after 2 nd gap (1.5)	Successful (4)
L1 (100)	Lost after 2 nd gap (1)	Successful (3)	Successful (4)	Successful (3)	Successful (4)
Average	0.625	1.750	2.750	2.875	3.375

From Table 1, it appears that the fusion of data (i.e., combining different measurement types per observation) generally leads to improved performance; in other words, the trend in the average “goodness of tracking score” as a function of measurements per observation has a positive slope. Figure 13 and Figure 14 highlight the improved error performance in the case of processing L1(100) and H1(1N) data.

Amongst varied performance in EO-only and PRF-only cases, the results show that the data fusion construct with the most EO+PRF measurements succeeds in all cases, leading to the intuition that it will perform better overall. In some instances, however, using more measurements is not strictly always beneficial. For example, although the EO3+PRF succeeds with the L1(50) orbit where EO3 alone loses tracking after the second gap, in two of the other three cases, EO3+PRF performs slightly worse than EO3 alone. This phenomenon is suspected to be driven by filter overconfidence. Connecting this topic and the previous uncertainty representation topic, this case seems to be a prime candidate to be enhanced with a more realistic uncertainty. Here, the application of AEGIS could shed more light into the nuanced behavior of using a multiplicity of measurements in context with different measurement types.

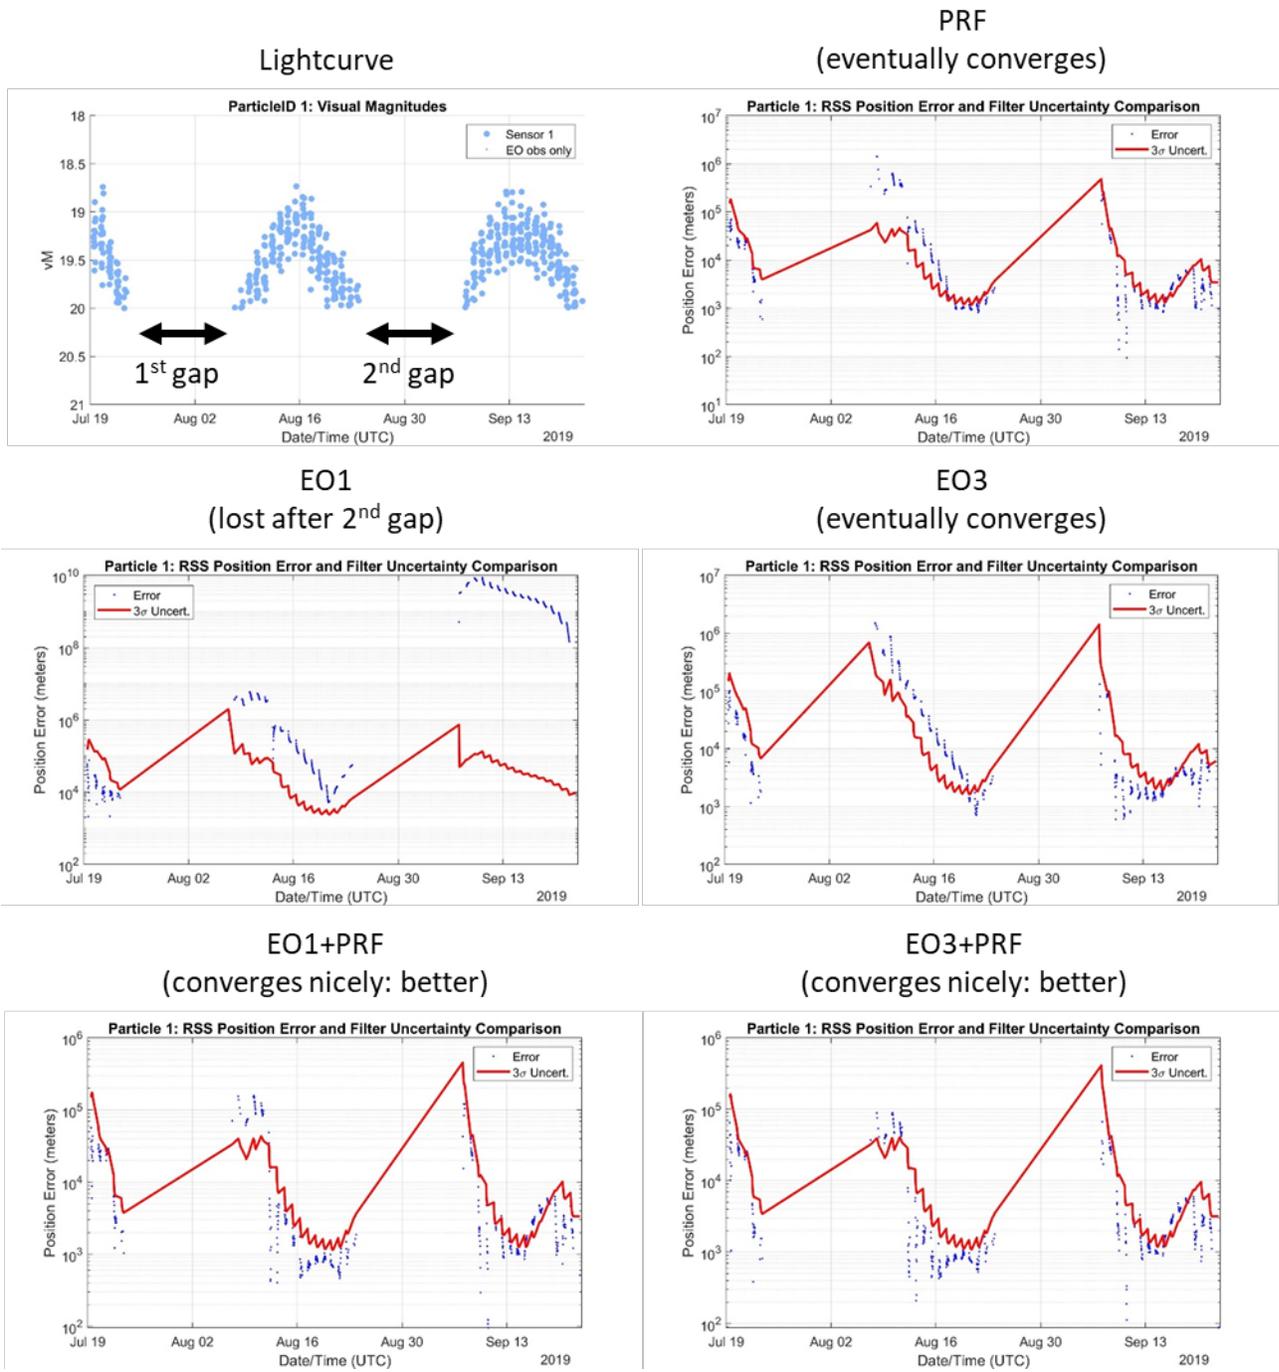


Figure 13. The L1(100) position errors over time show the successive improvements in convergence as more measurements and different data types are added. The light curve plot (upper left) is included to emphasize the data gaps due to the constraints on the EO data.

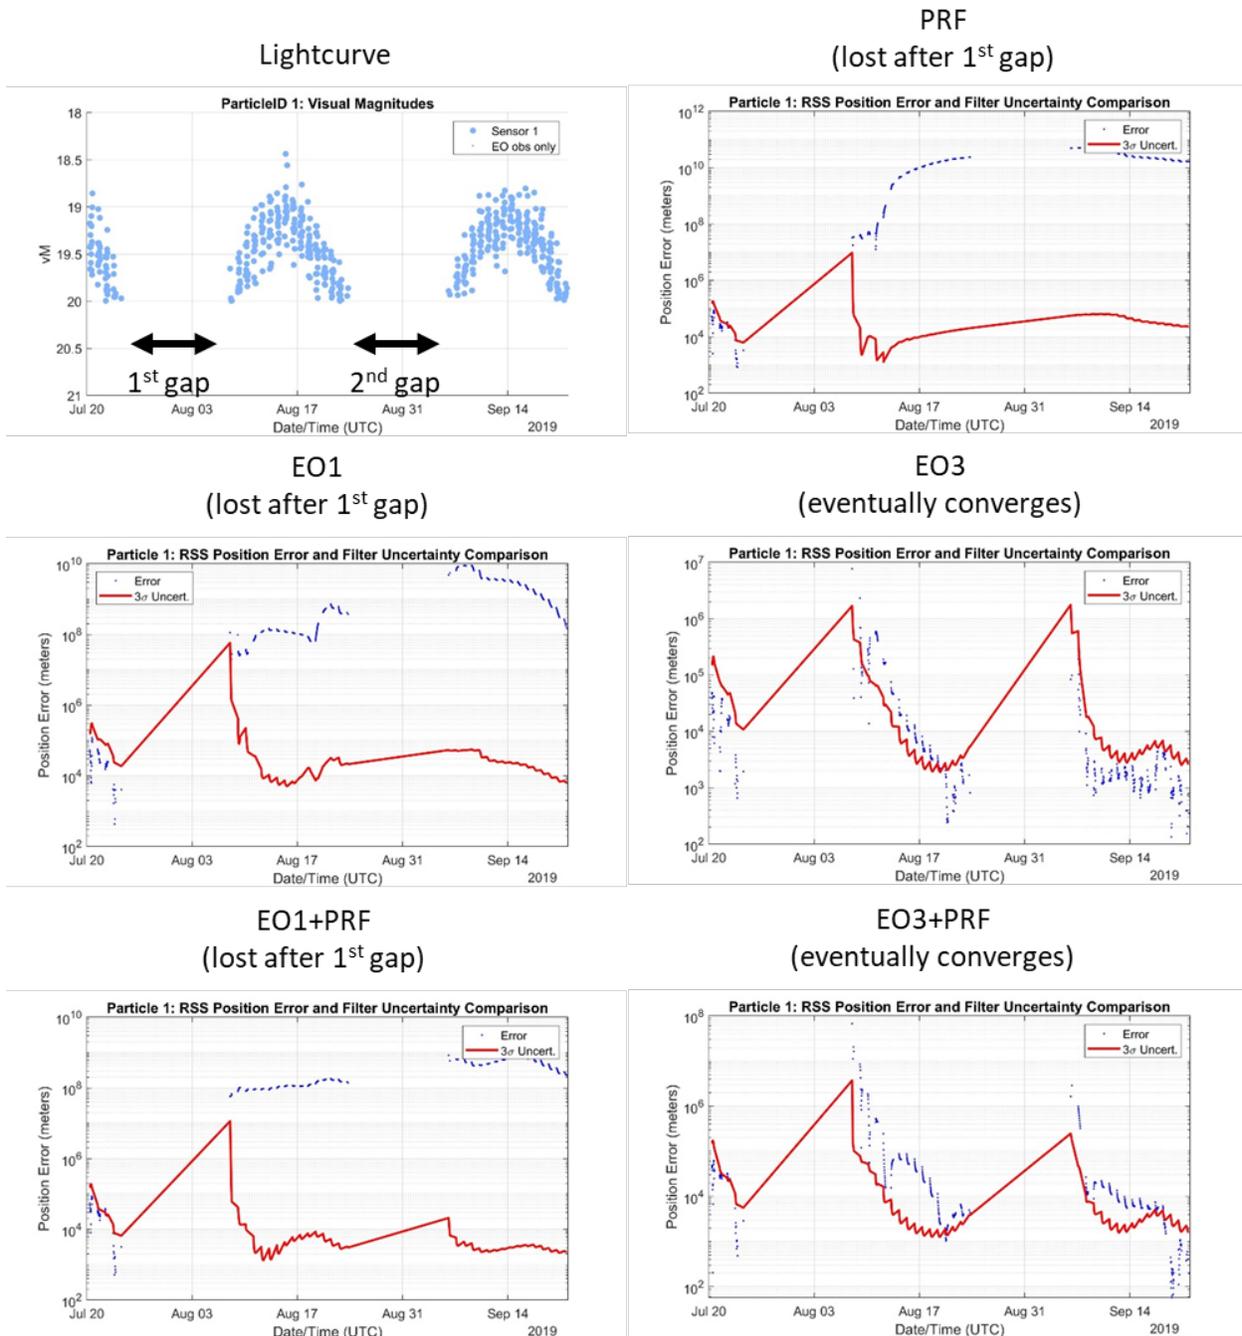


Figure 14. The $H1(1N)$ position errors over time show the successive improvements in convergence as more measurements and different data types are added. The light curve plot (upper left) is included to emphasize the data gaps due to the constraints on the EO data.

5. DISCUSSION

As discussed in Section 3, the AEGIS filter configuration used in the current work utilizes a basic threshold in differential entropy to determine when to split Gaussian parent components into child components. While differential entropy is a simple and easy-to-implement measure of growing non-Gaussianity, there may be metrics which can be used in place of (or in addition to) it to allow more efficient use of AEGIS components. Ideally, an improved technique

would have the properties of limiting the exponential growth of Gaussian component cardinality while retaining realism and explanatory power of the combined uncertainty estimates. In future work, investigation will be conducted into the use of adaptive schemes based on dynamic sensitivity and local orbit stability and incorporation of other metrics, such as the Tensor Eigenvalue Measure of Nonlinearity (TEMoN) [8].

A related area of investigation is the interplay between process noise inclusion and entropy evolution. When process noise is taken into account in the usual way – that is, integrated over the timespan between entropy checks, with Gaussian assumptions, into a discrete process noise covariance that is added to each component’s UKF covariance estimate – it requires a “collapsing” of each UKF component from its sigma points into the Gaussian approximation. This collapsing results in artificial modifications/discontinuities in the evolution of the entropy, which translates into additional complexity in component splitting heuristics during the overall AEGIS filtering process. Further investigation is required to mitigate these effects.

Applying uncertainty models with increasing realism (from UKF to GMM to AEGIS) improves filter performance through extended data gaps and reacquisition of the track. The breakdown of the PDF in lunar proximity, however, is inevitable given a long enough data gap, regardless of the representation approach. AEGIS and the monitoring of the associated differential entropy provides insight as to when this breakdown is occurring, thus informing operators and decision makers to the need for additional observations of an object to avoid loss of track. A potential mitigating strategy for this effect could be the curation of components (e.g., rejecting extreme components strategically from the AEGIS representation of the PDF), under assumptions of enforced mission viability on the part of the tracked RSOs, to push deeper into data gaps at the cost of robust uncertainty representation.

In summary, we show how the error growth is mitigated by improving uncertainty realism when processing through extended data gaps. We also demonstrate the greatest benefit of PRF observations over EO observations is the ability to “fill in the gaps” and thus extend the contiguous period of time the object can be observed, allowing the possibility of near continuous tracking. Data fusion at particular observation times is also shown to lead to better filter performance, although this result is largely due to having more measurements per time step rather than different measurement types because the types chosen for this work, EO and PRF, do not provide enough linearly independent information content to resolve ambiguities in the range dimension.

6. REFERENCES

- [1] Chow, C., Wetterer, C., Hill, K., Gilbert, C., Buehler, D., and J. Frith. Cislunar Periodic Orbit Families and Expected Observational Features. *The Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, September 15-18, 2020.
- [2] Chow, C., Wetterer, C., Baldwin, J., Dilley, M., Hill, K., Billings, P., and J. Frith. Cislunar Orbit Determination Behavior: Processing Observations of Periodic Orbits with Gaussian Mixture Model Estimation Filters. *The Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference*, Maui, HI, September 14-17, 2021.
- [3] DeMars, K. J., Nonlinear Orbit Uncertainty Prediction and Rectification for Space Situational Awareness, *Ph.D. Thesis*, Univ. of Texas at Austin, Austin, TX, 2010.
- [4] DeMars, K., Bishop, J., and Jah, M., An Entropy-based Approach for Inference of Uncertain Dynamical Systems. *Journal of Guidance, Control and Dynamics*, 36:4, July 2013.
- [5] Joe, S., and F.Y. Kuo. Remark on Algorithm 659: Implementing Sobol’s Quasirandom Sequence Generator, *ACM Transactions on Mathematical Software*. 29 49-57 (2003).
- [6] Joe, S., and F.Y. Kuo. Constructing Sobol Sequences with Better Two-Dimensional Projections, *SIAM Journal on Scientific Computing*. 30, 2635-2654 (2008).
- [7] GMAT Development Team, 2020. *User’s Guide*.
- [8] Jenson, E.L. and D.J. Scheeres. Semianalytical Measures of Nonlinearity Based on Tensor Eigenpairs. *AAS/AIAA Astrodynamics Specialist Conference*, 2021.