Probabilistic Initial Orbit Determination and Object Tracking in Cislunar Space Using Optical Sensors

Mark Bolden, Dr. Islam Hussein, Dr. Holly Borowski, Robert See, and Dr. Erin Griggs
Trusted Space, Inc.

1. ABSTRACT

In this work we use a very short-arc to probabilistically initialize and track an orbit for an object in the cislunar domain. Traditional techniques such as the Unscented Kalman Filter and the Particle Filter will not be able to handle the tracking of cislunar objects with such short arcs. Our approach delivers a new and powerful set of results that extends the advanced Probabilistic Admissible Region (PAR) from near earth orbit initialization to the cislunar domain with all its unique challenges. Furthermore, we integrate the PAR approach with an advanced new filter, the Particle Gaussian Mixture Filter (PGMF), to track the object after its orbit has been initialized. The PGMF can handle the highly nonlinear nature of cislunar dynamics that can undergo chaotic bifurcations, as well as the multi-modality (i.e., highly non-Gaussian nature) of the probability density describing the object's uncertainty.

2. INTRODUCTION

In this paper, we define the cislunar domain to encompass the full 4π steradian volume where three body dynamics are an accurate approximation of the Earth-Moon system. This domain is the ultimate high ground above Earth and securing and protecting that position is one of the five core competencies identified in the United States Space Force (USSF) Spacepower Doctrine [1]. Due to the vast distances and long duration solar/lunar exclusions, spacecraft in this domain can be unobservable for weeks at a time and then return to conventional orbital regimes, such as Geosynchronous (GEO) and Low Earth Orbit (LEO). To ensure safe operations in all orbit regimes, it is necessary to monitor and attribute spacecraft in the cislunar domain. Initializing and maintaining custody of spacecraft in this domain presents challenges that include initial orbit determination (IOD) for newly discovered objects or recovered lost objects, as well as for tracking objects after orbit initialization.

The topic of cislunar Orbit Determination (OD) has been discussed in multiple publications over the last several years. In 2020, Bolden pointed out that the Too Short Arc (TSA) problem, originally identified for GEO, was significantly more challenging in the cislunar domain, potentially requiring multiple observers to achieve range observability [2]. Chow pointed out that "since cislunar dynamics approach the chaos boundary, where small perturbations can lead to wildly disproportionate outcomes, traditional data processing techniques (e.g., IOD) become far more difficult [3]." In 2021, Chow discussed the superiority of an approach to processing periodic orbits using Gaussian Mixture Model Estimation Filters versus filters using two-body dynamics and the Unscented Kalman Filter [4]. Furfaro presented IOD and OD results using real data of the Chang'e 5 spacecraft. The approach demonstrated Gauss [5], Herget [6] and Vaisaala [7] methods for IOD and OD methods "typically employed in tracking Near Earth Objects (e.g., iterative batch least square, simplex methods) to iteratively reconstruct trajectories for both spacecraft and boosters" [8]. Greaves showed that observability is achievable with optical observations in the cislunar domain when the observer is space-based using a modified version of the Optimal Control Based Estimator (OCBE) with Unscented uncertainty propagation adapted to provide filtering solutions and an estimated control policy [9]. Thompson showed the performance of a current epoch state U-D factorized covariance filter with backwards smoothing to a reference epoch that is implemented in NASA/JPL's MONTE software [10]. Wishnek showed an IOD method that formulates a trial solution state as a cost function to optimize over the unobserved state space dimensions leveraging multiple angles measurements [11].

In this paper, on the other hand, we use a very short-arc (one angles and photometric measurement in the extreme case) to probabilistically initialize and track an orbit for an object in the cislunar domain. Observations of an object in the cislunar domain are usually very short-arced because of the colossal distances of the environment, causing multiple observations to effectively act as a single observation point. To rapidly generate an uncertain trajectory for object tracking, high-fidelity probabilistic IOD techniques, such as the Probabilistic Admissible Region (PAR) approach [12], are needed. Such techniques need to complement the single observation with other hypothesized information (e.g., target albedo-area product that correlates with range to the object from the sensor location) to generate a probabilistic representation of the object's orbit. A primary challenge to generating an initial track for an object is that cislunar trajectories cannot be easily described by parameters we currently use to describe near-earth orbits, such as classical orbital elements. For near-earth orbits, statistics for these parameters are often

inferred using available extraneous information such as the natural distribution of space objects in Earth orbit. Such information is currently unavailable for the cislunar domain.

Once an orbit is initialized, subsequent tracking of a cislunar object also presents us with additional unique challenges. We first note that, as is the case with IOD in Earth orbit, the resulting initialized orbit is typically multi-modal [12]. Any subsequent object tracking technique must therefore be able to handle multi-modal object state uncertainties. On top of that, the cislunar domain is subject to highly nonlinear, possibly chaotic dynamical behavior. It is therefore desired to develop and implement high-fidelity nonlinear filters that can accommodate multi-modal probability distributions, chaotic bifurcations, and that are robust to modeling errors. Such algorithms need to be light-weight for onboard implementation given that communications with on-orbit cislunar sensing assets might be limited. One advanced algorithm that is capable of handling all of these challenges is the Particle Gaussian Mixture Filter (PGMF) [13, 14]. The PGMF combines the best of the two worlds of particle filtering (specifically, its propagation step) and the Gaussian Mixture Model (GMM) form of the Unscented Kalman Filter (UKF) (specifically, its nonlinear update step). We will discuss the PGMF in greater detail in Section 4.

In summary, in this paper we (a) extend the use of PAR that was developed for near-Earth IOD to the cislunar domain using possibly only a single short arc optical measurement, and (b) implement the PGMF for the processing of subsequent optical observations of the target in an integrated framework. This integrated PAR-PGMF solution is a rigorous initial orbit determination and filtering framework that is scalable, robust to modeling errors and large IOD uncertainties, and can handle multi-modal uncertainties and highly nonlinear and chaotic dynamical systems. We describe in detail how the PAR and the PGMF will be used for cislunar object orbit initialization, and tracking. We will demonstrate the power of the PAR-PGMF framework on multiple use cases.

3. TECHNICAL CHALLENGES

In the cislunar domain, distances between observation platforms and objects are typically very large. This presents several challenges. First, objects are fainter and more difficult to detect due to large ranges and illumination conditions. Once detected, orbit initialization requires at least six independent parameters to be able to solve for the six-dimensional state of the target object. With the very long distances, the object, however, will appear to move slowly relative to the observing sensor. This implies that even for an extended period of time observations of the object contain too little independent information to initialize an orbit. For optical sensors, for example, this means that we effectively only have two parameters (e.g., azimuth and elevation from the sensor), which is four fewer than the needed six variables to determine the orbit of the object. This implies that other parameters, such as the object's orbital energy and albedo area product, need to be hypothesized, as opposed to being directly measured.

Another challenge is that both the observed and hypothesized quantities are uncertain in nature. This uncertainty translates into an uncertain orbit when the observed and hypothesized quantities are mathematically mapped during orbit initialization. When properly performed in a probabilistic fashion, the uncertainty in the initialized orbit is typically multi-modal. The resulting uncertain orbit will not be, for example, a simple single Gaussian distribution as is typically assumed in practice. Intuitively, this results from the fact that with short arc angles only data, it is typically unclear which direction the object is traveling. This results in two modalities in the initial PAR. Other similar ambiguities in the hypothesized parameters will likewise result in additional modalities. A Gaussian distribution model, while it simplifies the object tracking process, introduces assumptions that are easily violated in highly non-linear dynamical systems, such as in the cislunar domain. Therefore, it is important to retain the non-Gaussian, multi-modal nature of the initialized uncertain orbit.

This leads to the final challenge faced in the cislunar domain. Any object tracking technique needs to be able to process non-Gaussian, multi-modal orbital uncertainty. Techniques, such as the Extended Kalman Filter (EKF) or the UKF [15], that assume Gaussianity of the underlying uncertainty in the state of the object will fail at tracking the object due to the non-Gaussianity of the initialized orbit. Furthermore, even if the initial orbit's uncertainty was Gaussian, the physics of the cislunar domain is such that objects move according to highly nonlinear and chaotic dynamics. This means that even very small errors or mismodeling of the uncertainty can result in very large deviations between where the object really is and where the algorithm believes it is. This will result in losing track of the object. Traditional quasi-Gaussian nonlinear filtering approximations, such as UKF, cannot handle non-Gaussian filtering problems. They are, however, powerful for their computationally tractable uncertainty update step when a measurement of the target is obtained –an observation that we will return to when we discuss the PGMF later in the paper. The curse of dimensionality reveals itself in the form of particle depletion in the Particle Filter (PF) [16], which we will discuss more in section 4.

It is critical that each of the above challenges are properly addressed by any cislunar IOD/OD tool combination. Current approaches to IOD/OD will, however, typically fail at addressing at least one of these challenges, resulting

in a failure to initialize an orbit and/or losing the object altogether. The main reason for this is that existing tools used for space object IOD/OD typically assume two-body dynamics and perturbations thereof that do not apply to the three-body dynamics nature in the cislunar domain. Existing IOD and OD techniques cannot cope with these peculiar aspects of the cislunar domain. For example, the Constrained Admissible Region (CAR) approach described in [17] requires specification of ranges on parameters such as semi-major axis and eccentricity and employ the physics of the unperturbed two-body (Earth-spacecraft) system. Orbits in the cislunar domain, however, have no notion of semi-major axis or eccentricity, and clearly do not evolve according to two-body dynamics. While a technique such as the PAR was also developed for two-body systems, it is extensible to the cislunar domain.

In the above paragraphs we noted the reasons why popular nonlinear methods such as the UKF and the PF will fail for tracking objects in the cislunar domain. We have also highlighted their strengths since the PGMF method that we propose in this effort combines the powers of both methods and yet avoids their flaws.

4. TECHNICAL SOLUTION

This paper will discuss our probabilistically consistent combined IOD/OD framework that manages the cislunar domain's complexity. The framework is robust to modeling errors and large IOD uncertainties and can handle multi-modal probability density functions and highly nonlinear and chaotic dynamical systems. The framework is algorithmically scalable, lightweight, and easy to implement. Our framework relies on the use of the PAR for IOD and the PGMF for OD. The combined IOD/OD framework is called PAR-PGMF and, as will be shown below, meets all desired qualities.

4.1 Astrometric-Radiometric PAR Initialization for IOD

The first component of the proposed PAR-PGMF solution is the PAR. The core idea of the PAR technique is to map the measurements and their uncertainties along with any information about complementary variables using a sampling approach to a probabilistic representation of the target's state using a particle representation. It is a probabilistic IOD technique that is in contrast with deterministic single trajectory IOD approaches. Deterministic approaches to IOD generate a single track for an object. Because they generate only a single trajectory, deterministic techniques are typically very sensitive, especially for nonlinear and chaotic regimes such as cislunar, to sensor noise and other assumptions made about the object's orbit. Unlike deterministic IOD approaches, methods such as the CAR and the PAR generate a family of orbits around the object's true orbit using the observed measurement from the sensor. When performed in a probabilistically rigorous way, the generated family of orbits are guaranteed to encompass the true orbit of the object. With every additional observation of the target after orbit initialization, the family of orbits, also called the "particle cloud" for simplicity, converge towards the true object orbit if the subsequent OD method is robust and unbiased.

The PAR approach starts with the single sensor observation. The PAR technique can be tailored to any sensor phenomenology. However, we will use an optical observation model for its relevance to the cislunar domain context and for its ease of understanding. Given an observation and the sensor noise standard deviation, PAR proceeds by sampling a large number, N, of simulated observations. The more the number of samples, the higher the fidelity of the PAR procedure in modeling the uncertainty of the object's orbit. The number of particles is only limited by the processing and memory of the available computational resources, and by any limits on meeting tactical mission timelines. For each sample, we need four additional independent parameters to generate a corresponding full orbit sample.

Our approach first takes advantage of the light intensity measured by the optical sensor. This provides a third direct observable parameter that provides information on the range to the target. We emphasize here that use of light intensity is only for the PAR, and not for the PGMF in subsequent orbit determination updates in this implementation. Using the light intensity along with the associated noise statistics, one can then generate N light intensity sample points. Since the first three parameters, two angles (i.e., azimuth and elevation) and a light intensity, are independent, we can then solve for the object's positional information uniquely, first using the equation of a diffuse/specular sphere to compute range:

$$M_v = -26.74 - 2.5\log(A\rho[\beta F_1(\varphi) + (1-\beta)F_2(\varphi)]) + 5\log R$$

where

$$F_1(\varphi) = \frac{2}{3\pi^2}[(\pi - \varphi)cos\ \varphi + \sin\varphi]$$

$$F_2(\varphi) = \frac{1}{4\pi}$$

and where M_v is the visual magnitude of the observed object, $A\rho$ is the albedo area product, φ is the solar phase angle, β is the sample mixing coefficient, and R is the range from the observer to the object [18]. Here we denote F_2 as a function of solar phase angle by convention. For other shapes there is a solar phase angle dependence to this term. The visual magnitude is sampled based on the uncertainty of the measurement. The albedo area product is sampled uniformly across all spacecraft possibilities. The range of solar phase angles is sampled uniformly and can be significantly constrained by combining the angles data with the maximum allowable range based on the cislunar domain definition. Given the line of sight, the solar phase angle is very insensitive to R within the domain. This process provides a range to the target for each sample. The approach enables the technique to work with no knowledge of object material or size. The result is large uncertainty in range that is constrained by independent measurements. Using angles and range we can convert to relative positions from the sensor to the target:

$$x_{ot} = R\cos(AZ)\cos(EL)$$
$$y_{ot} = R\sin(AZ)\cos(EL)$$
$$z_{ot} = R\sin(EL)$$

where AZ, EL are the observed angles, x_{ot}, y_{ot}, z_{ot} are the positional cartesian coordinates from the sensor to the target, and R is the range from the sensor to the target.

We then need three additional independent variables that give us information on the object's velocity. For these, we use known limits on the sensor's maximum angular rates (i.e., azimuth and elevation rates) to bound the maximum rate an object can move within the sensor's frame in order for the object to be detected by the sensor. If the object moves any faster across the sensor field of view than that, we know that the sensor cannot, by definition, detect it. This provides two hypothesized (since they are not directly measured) independent parameters. The final bound we can hypothesize is the object's speed. For cislunar, we take advantage of the equation governing the object's velocity and the Jacobi integral for the circular restricted three-body problem (CR3BP) [19]:

$$\dot{AZ} = \sum_{i=0}^{n} Rand[0, \frac{Max\ Streak}{Exposure\ Duration}]$$
$$\dot{EL} = \sum_{i=0}^{n} Rand[0, \frac{Max\ Streak}{Exposure\ Duration}]$$

Figure 1: Maximum Streak Length Detectable

$$\dot{x}^2 + \dot{y}^2 + \dot{z}^2 = (x^2 + y^2) + 2\left(\frac{1-\mu}{r_{13}}\right) + \frac{2\mu}{r_{23}} - C$$

where $\dot{x}, \dot{y}, \dot{z}$ are the cartesian velocity components of the target, x, y are the cartesian position components, μ is the non-dimensional mass of the moon, r_{13} is the range from earth to the object, r_{23} is the range from the moon to the object, and C is the Jacobi constant. The CR3BP approximation of the more complex three-body dynamics governing the cislunar domain introduces relatively small errors that will be corrected with a robust subsequent OD method such as the PGMF, which we will discuss later in this section. *This equation is the equivalent of using the energy equation in the CAR and PAR to hypothesize an object's orbital semi-major axis for the two-body problem.*

The Jacobi integral equation is only one equation with three unknowns (the cartesian components of velocity). We complete the system of three equations by taking the derivative of the relationship between angles and the observer and target positions to get the following two equations:

$$\dot{y}_t - \dot{x}_t tan\, AZ = \dot{y}_o - \dot{x}_o tan\, AZ + \frac{\dot{AZ}}{(cos\, AZ)^2}(x_t - x_o)$$

and

$$(\tan EL)^2 (x_t - x_o)\dot{x}_t + (tan\, EL)^2 (y_t - y_o)\dot{y}_t - (z_t - z_o)\dot{z}_t =$$
$$-\tan EL \frac{\dot{EL}}{(\cos EL)^2}[(x_t - x_o)^2 + (y_t - y_o)^2] + (\tan EL)^2 [(x_t - x_o)\dot{x}_o + (y_t - y_o)\dot{y}_o] - (z_t - z_o)\dot{z}_o$$

where x_o, y_o, z_o are the position terms of the observer, x_t, y_t, z_t are the position terms of the target, $\dot{x}_o, \dot{y}_o, \dot{z}_o$ are the velocity terms for the observer velocity, $\dot{x}_t, \dot{y}_t, \dot{z}_t$ are the velocity terms for the target velocity, AZ and EL are the measured angles, and $\dot{AZ}$, and $\dot{EL}$ are the constrained angle rates. The above three equations are independent and can thus be used to solve for the three velocity equations.

While the above equations seem to be nonlinear in the unknown velocity components, we note here that the angle rate equations are linear in the three unknown velocity variables. In these two equations, one can solve for two of the components, say the $\dot{x}_t$ and $\dot{y}_t$ velocity components, in terms of the third variable $\dot{z}_t$ velocity component, and then replace the occurrence of $\dot{x}_t$ and $\dot{y}_t$ in the Jacobi integral equation to get a single equation that is quadratic in $\dot{z}_t$. We then solve for $\dot{z}_t$ given the hypothesized values for the Jacobi constant and the angle rates. Given the value for $\dot{z}_t$, we can then obtain the solution for $\dot{x}_t$ and $\dot{y}_t$. With this we have a particle cloud with each particle having all six state variables (three position and three velocity) fully solved for.

We note here the quadratic nature of the equation in $\dot{z}_t$. This implies that there are two solutions for the cartesian velocity components. This results in the bimodal nature of the uncertainty cloud mentioned in Section 3. This bimodality refers to the fact that when using angles and range information, we do not have any information on velocity (i.e., travel direction). This ambiguity results in the two modes in the initialized PAR. Thus, the resulting cloud of orbits for the object that represent its orbital uncertainty will not be Gaussian. Another contributing factor to the non-Gaussian nature of the resulting PAR is the fact that while the angles and range are typically Gaussian in nature, angle rates and the Jacobi constant uncertainties are assumed to be uniform. This mix of Gaussian and uniform uncertainties are mapped through the various nonlinear geometric and Jacobi integral equations to solve for the state are highly nonlinear algebraic equations. That will result in a complex, multi-modal and non-Gaussian distribution in the object's state uncertainty. We note here that the PAR process is embarrassingly parallelizable. For example, for $N = 10{,}000$ particles, on an industry-standard laptop, the particle cloud is generated from a single observation in about 70 milliseconds without optimizing the code that generates the cloud. With optimized code and advanced onboard processors, far many more particles can be generated in less time to generate a higher fidelity representation of the target's orbital uncertainty. We will seek to understand the trade space balancing between the number of particles against the available onboard computations in future research.

4.2 PGMF for Uncertainty Reduction and Object Custody

As discussed earlier in the paper, uncertainty is multi-modal and bifurcations can occur in the chaotic three-body system and by the nature of the IOD process that introduces highly non-Gaussian uncertainties. For orbit determination, the main challenge for current techniques is the highly nonlinear nature of the dynamics combined with the non-Gaussian nature of uncertainty generated from the IOD task. Traditional quasi-Gaussian nonlinear filtering approximations, such as the Kalman Filter (KF) and the Unscented Kalman Filter (UKF) [15], cannot handle non-Gaussian filtering problems. They are, however, powerful for their computationally tractable uncertainty *update* step when an observation of the target is obtained. The curse of dimensionality reveals itself in the form of particle depletion in the Particle Filter (PF) [16], where the initial uncertainty is in the form of a particle cloud generated from the IOD step that are then propagated and updated in a Bayesian framework whenever a new observation of the target is obtained. Particle depletion is the phenomenon where in the PF update step particles with very low likelihood of being close to the truth are removed, resulting in effectively a very small number of particles near the true state of the target. This, in turn, may result in a loss of the object's true track, especially for highly non-linear and chaotic dynamical systems where particle abundance and diversity become crucial. We note, however, that otherwise the PF is powerful for *propagation* of uncertainty for nonlinear, possibly chaotic, systems and is highly parallelizable.

While above we noted the reasons why popular nonlinear methods such as the UKF and the PF will fail for tracking objects in the cislunar domain, we have highlighted their strengths since the PGMF combines the powers of both methods and yet avoids their flaws. The PGMF is a robust approach that can handle the non-Gaussian nature of the particle cloud and that is robust to errors introduced into the PAR process such as the use of the CR3BP assumption, as well as other errors introduced in the subsequent uncertainty propagation.

Next, we describe how the PGMF approach to object tracking works [14, 13]. With the object's orbit probabilistically initialized using the PAR, presented as a non-Gaussian cloud of orbital particles, the PGMF performs OD on that cloud as subsequent observations of the target are received. Referring to Figure 2, the PGMF is kicked off by first propagating the particle cloud using the dynamical equations governing the cislunar domain. This propagation is performed to the time when the next observation of the target arrives, at which point we convert the propagated particle cloud into a Gaussian Mixture Model (GMM) by clustering the particles using a clustering algorithm such the k-means++ [20] or the Expectation

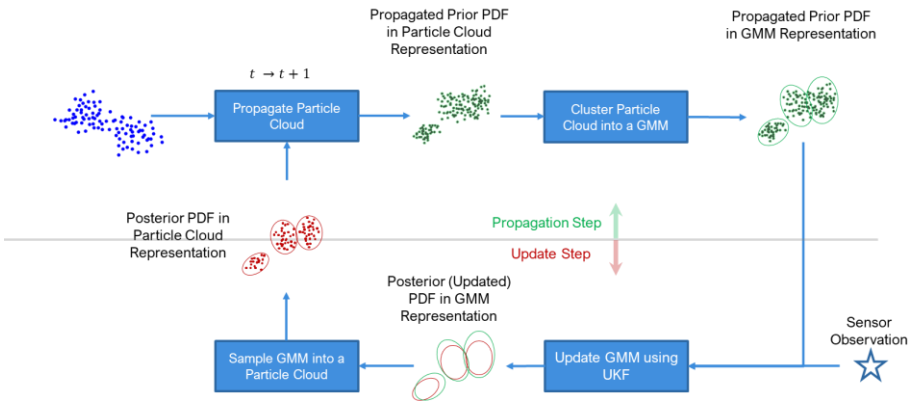

Figure 2: PGMF Approach

Maximization algorithm (EM) [21]. A UKF update step is then performed on the propagated prior GMM analytically to obtain the posterior GMM. The observable here is either just the optical angles-only observation and, possibly but not necessarily, also the light intensity of the target that relates to range from the observer to the target. In the use cases presented later in the paper, only the angles observations are used to update the track. The light intensity is only used to aid in the PAR initialization of the target's track as explained in previous section. With the GMM updated, one then samples this posterior GMM to obtain a new particle cloud that is then propagated to the time when the next observation is obtained. This process is repeated for all subsequent observations collected by the sensor. As can be seen, the PGMF is essentially based on the PF except for the updated step where an observation updated GMM is used to resample the cloud.

A key factor that makes the PGMF attractive is that every operation described in Figure 2 is scalable, rendering the entire algorithm scalable. This is because the PGMF uses the UKF GMM update step, which is *computationally tractable*. In our simulation studies of the PGMF, it was found that the update step is far less computationally expensive than the propagation step. For propagation, the PGMF takes advantage of the *embarrassingly parallelizable* particle filter's propagation of uncertainty for nonlinear, possibly chaotic, systems step. Furthermore, the PGMF is immune to the "curse of dimensionality" due to the clustering-based UKF GMM update step that results in resampling all particles so that they are all located near the true orbit as measured using the sensor observation. As such, the PGMF allows us to get the best of the two worlds of the GMM-based UKF (analytical, nonlinear, and scalable update step) and the PF (nonlinear and scalable propagation step), while avoiding the shortcomings of both. Finally, the PGMF is robust to modeling errors (see [14, 13] for details) due to the particle nature at the heart of the approach. This is important for nonlinear and chaotic dynamical systems where slight modeling errors can result in loss of custody of the target. Robustness also implies that we can use simpler physical models for the dynamics and yet be able to maintain custody of the target. This is key to being able to keep the software lightweight with small computing and power requirements, especially when used for on-orbit IOD/OD. In the use cases shown later in the paper, this robustness quality enabled us to use simplified propagation that are crude approximations of the ground truth models used to generate the data ingested by the PGMF. The PGMF (propagation and update) was able to process data at rates as high as 1000 times faster than real-time.

4.3 Additional Considerations

During this development effort, we prioritized considering potential risks to this approach. Two major risks were considered and mitigated, and that generated additional insights.

First, we considered the use of a photometric measurement for quantifying the range uncertainty. Measured brightness for a single spacecraft can vary substantially based on unknown parameters such as pose, material composition, and shape resulting in large range uncertainties. To accommodate this, we considered both diffuse and specular contributions and large ranges of albedo area products. Though not presented in this paper, we also considered shape diversity for future applications of this technique. Additionally, the PAR/PGMF is robust to any unintentional biases introduced by the photometric range estimate. A key insight on this risk is that due to the vast distances of the cislunar domain, large uncertainties in range still provide meaningful information to sample the possible ranges.

Second, we considered the use of CR3BP assumptions for the generation of velocity terms for the PAR. The real Earth Moon System is not consistent with the CR3BP representation resulting in biases in the PAR velocity components. Similar to potential range biases, the PAR/PGMF is robust to these biases as subsequent measurements are utilized to update the particle cloud. A key insight on this risk is that this approach is the exact same concept utilized in two-body dynamics when using the classical orbital elements for CAR/PAR techniques.

We are confident that these risks have been fully considered and that the combined PAR/PGMF technique effectively handles these large uncertainties and biases in the initial PAR for cislunar domain.

5. RESULTS & ANALYSIS

As proof of concept, we have implemented this approach in software code and processed simulated data generated based on our extensive experience with simulated and real data on cislunar objects. Based on previous experience developing cislunar trajectories, we generated a truth trajectory specifically designed 1) to produce one or two observations before entering the lunar exclusion zone for seventeen days, and 2) to result in complex bifurcations of uncertainty due to a swing-by of the Moon. Figure 3 shows the rendering of the truth trajectory over a 30-day propagation period in the Earth Moon Body Barycentric Rotating (EMBBR) frame. For all observations we modeled a single ground-based observer operating on Haleakala in Maui, HI.

To communicate the complexity of the uncertainty we present a rendering of 20 trajectories with relatively tight initial deviations from the true object orbit that was used in the PAR-PGMF analysis and propagated over a 17-day period. On the left in Figure 4, the first 4 days of propagation are shown. Many of the trajectories perform lunar swing-bys at different times that send them in wildly different directions. On the right, the 20 trajectories are shown after 17-days, where they scatter in multiple bifurcation groupings around the domain. This trajectory breaks conventional approaches, such as a UKF when it attempts to update the initial state 17-days later since, say, the sigma points generated at the initial time step can diverge along various bifurcation paths, resulting in an unrealistically large Gaussian at an update step 17-days later. This use case is an extreme case specifically chosen to challenge the PAR/PGMF approach.

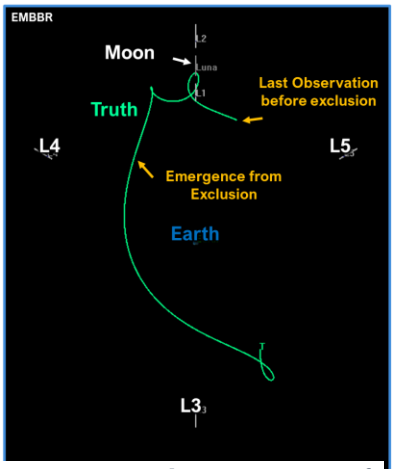

Figure 3: 30-day Propagation of Truth Trajectory

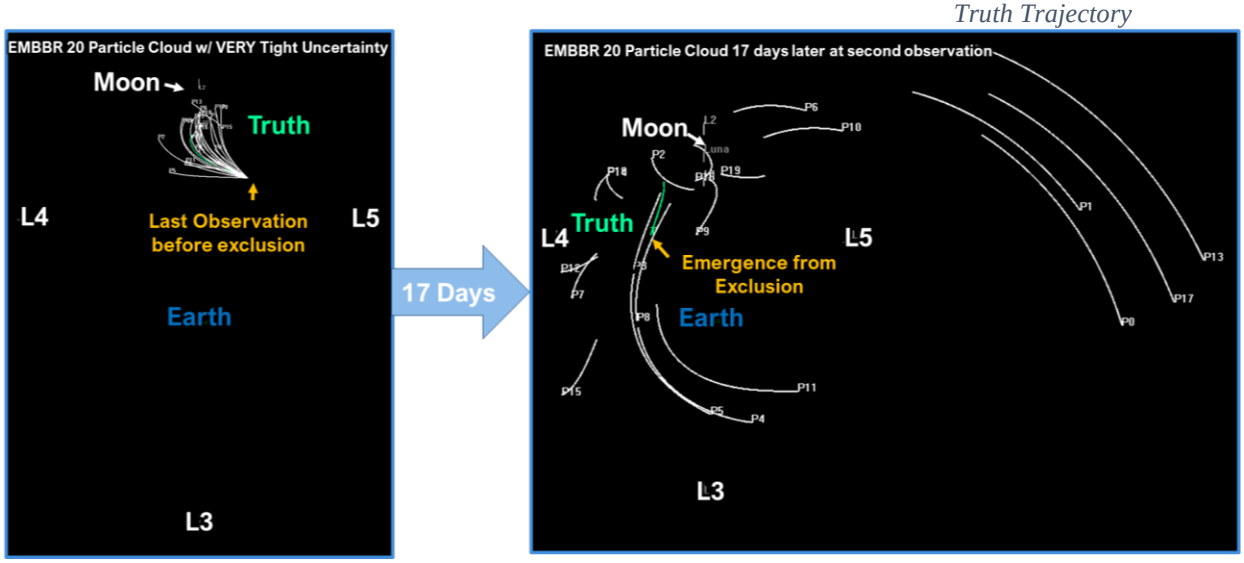

Figure 4: Extremely Tight Initial Uncertainty Propagation Visualization

5.1 PAR-PGMF Implementation Results

We now shift to the PAR-PGMF implementation. The core idea of the PAR technique is to map the measurements and their uncertainties along with any information about complementary variables using a sampling approach to a probabilistic representation of the target's state using a particle representation. The PAR's particle cloud is then propagated to the time when the next observation arrives, at which point we convert the propagated particle cloud into a Gaussian Mixture Model (GMM) by clustering the particles using a clustering algorithm. One then samples the posterior GMM probability distribution function (PDF) to obtain a new particle cloud that is then propagated to the time when the next observation is obtained. This process is repeated for all subsequent observations collected by the sensor.

In this next example we show the propagated particle clouds between each observation update with small uncertainties in the sampled parameters. Figure 5 plots each particle with a cyclic colormap to indicate the passage of time between observation updates. For this example, two observations were collected one hour apart before entering lunar exclusion. The propagated particles after the initial PAR can be seen in red in the negative x-direction. After 17-days the propagated particles can be seen in blue. The uncertainties have resulted in multiple bifurcations that are very non-Gaussian. As additional observations are collected every hour after the initial 17-day propagation, the PGMF is robust to these bifurcations and converges to the truth.

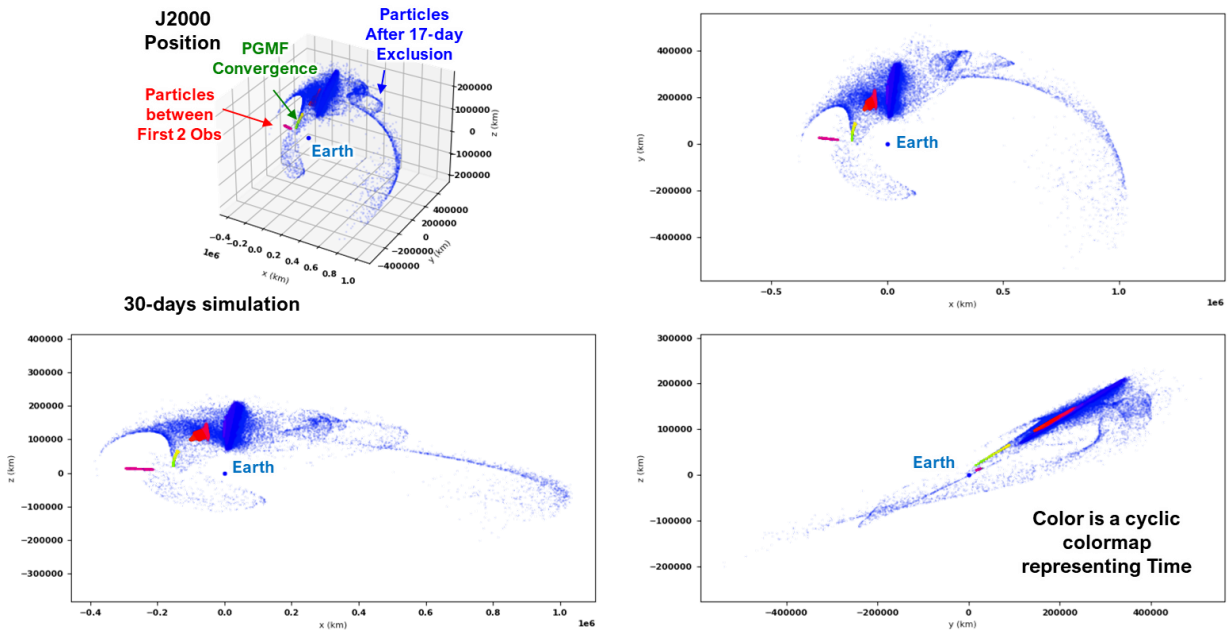

Figure 5: Propagated Particles in J2000 for 2 Observations Before 17 Day Exclusion w/ Small Uncertainties

In Figure 6 we plot a single propagated particle cloud time step and overlay the Gaussian mixture that k-means++ fit to the data. The time step is after 2 observations separated by 1 hour and a 17 day period with no observations. Visually we observe that the uncertainty of the particle cloud is significantly better captured than it would be by a single Gaussian. A single Gaussian for a UKF update would span the entire domain and cause the updated step of the UKF to fail.

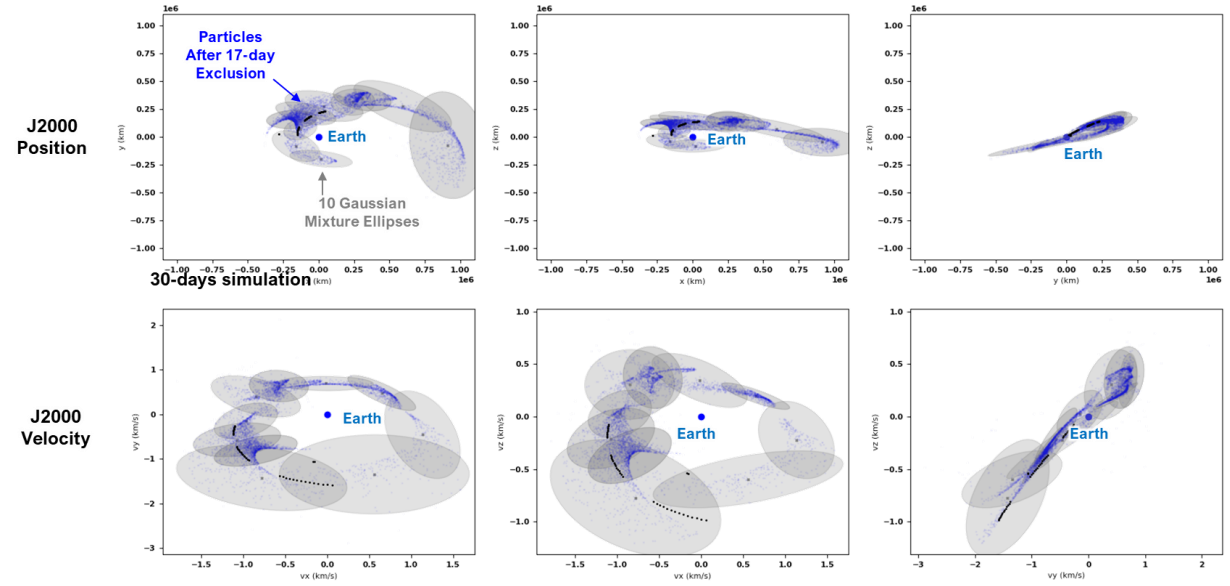

Figure 6: Propagated Particle Cloud after 17-day Observation Gap at a Single Time Step with GMM Overlay

We can also view this information in the Radial, In-track, and Cross-track (RIC) frame with respect to the truth trajectory as shown in Figure 7. In this reference frame we can visualize the complex bifurcations of uncertainty more clearly. As described in Figure 5 we can see that the PGMF converges to near zero error.

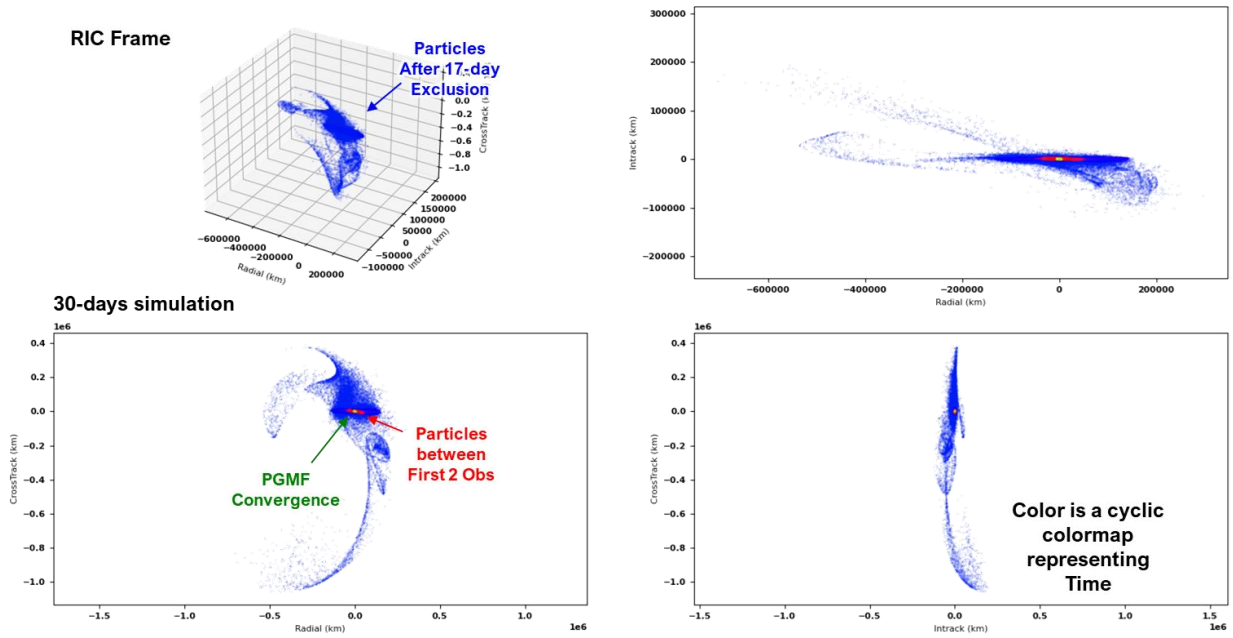

Figure 7: Propagated Particles in RIC for 2 Observations Before 17 Day Exclusion w/ Small Uncertainties

Finally, we can also view the propagated particle cloud in velocity space, as shown in Figure 8. In this coordinate system the bimodal structure of the initial propagated cloud becomes apparent and is expected. This is due to the lack of direction of travel knowledge from the initial observation.

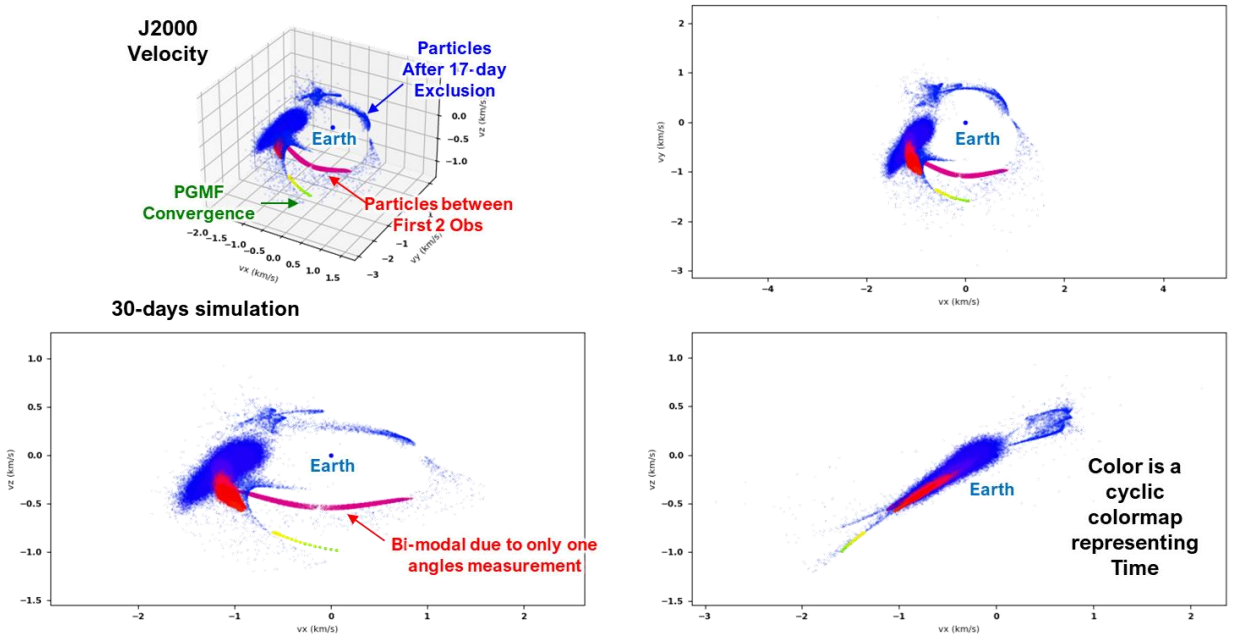

Figure 8: Propagated Particles in J2000 for 2 Observations Before 17 Day Exclusion w/ Small Uncertainties

In Figure 9 we observe that the PGMF successfully updates and converges after the 17-day gap. This is shown by plotting the RIC frame particle error over time and the absolute particle error over time with a logarithmic y-axis. As expected when using angles only data, the range uncertainty requires more observations over time to converge. It also remains the largest source of uncertainty during convergence.

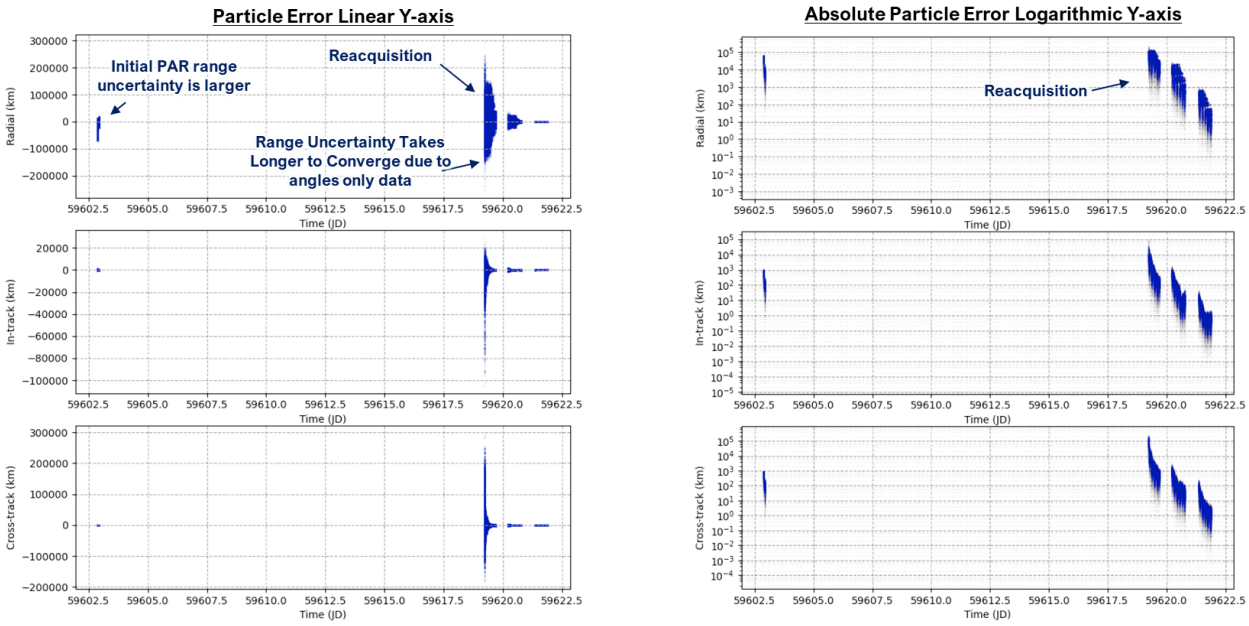

Figure 9: Particle Error and Absolute Particle Error Overtime

Another way to quantify uncertainty throughout the PAR-PGMF simulation is to use entropy to evaluate the information gain over time to ensure the PGMF is effectively reducing uncertainty. In Figure 10 we see that the entropy continues to decrease as more observation are fed to the filter. We note here that there is no analytical expression for entropy, measured in nats in the figure, for particle cloud representations of uncertainty. To circumvent this issue, we use the GMM representation of the particle cloud uncertainty. While there is a closed form representation for entropy for Gaussian distributions, there is no analytical, closed-form expression for the entropy of the GMM. We evaluated entropy by performing Monte Carlo integration of the GMM. This results in a conservative estimate of entropy for the particle cloud since the GMM is generally an over-conservative approximation of the uncertainty contained in the cloud.

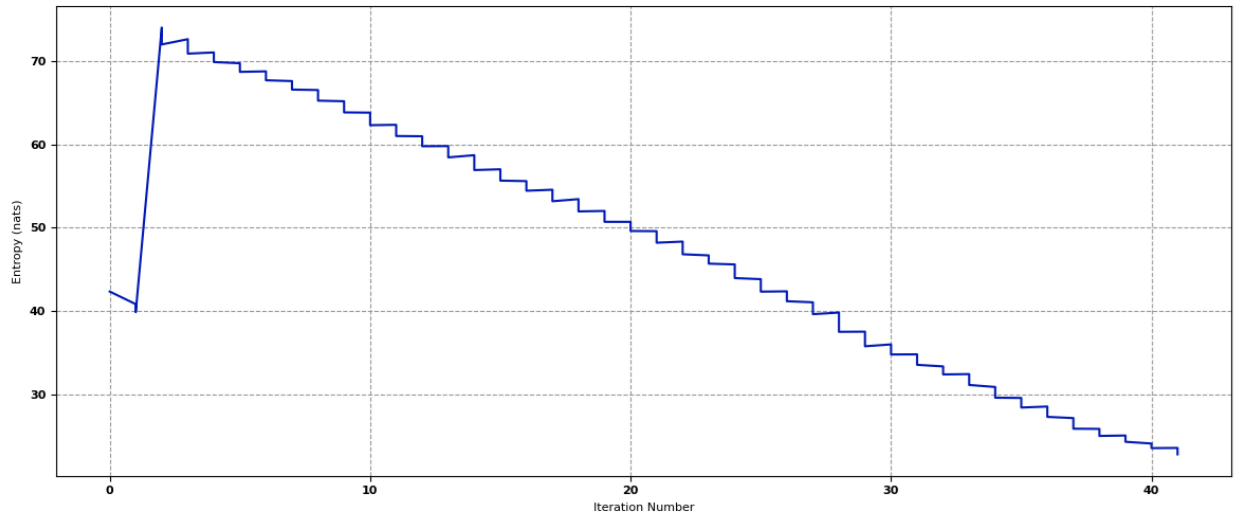

Figure 10: Entropy as a Function of Iteration Number

While this is an interesting case for validating the PAR/PGMF approach and visualizing the results, it is insightful to analyze a more difficult case with only one observation prior to the 17-day exclusion and realistic uncertainties in the sampled parameters. Additionally, we implemented lower fidelity force models for the PGMF

propagator than for the simulated truth to create a model mismatch to evaluate the PGMF's robustness to modeling errors. Both utilized Orekit for propagation. See Table 1 for more details.

Table 1: Forces Modeled

Truth Simulator Forces:	**PGMF Propagator Forces:**

• Earth Holmes Featherstone
  ○ Gravity Order: 36 Deg
  ○ Min Step: 0.00001
  ○ Max Step: 300.0
• Jupiter, Moon, Sun point sources
• SRP
• Atmospheric drag: HarrisPriester
	• Earth Holmes Featherstone
	  ○ Gravity Order: 12 Deg
	  ○ Min Step: 0.00001
	  ○ Max Step: 300.0
	• Jupiter, Moon, Sun point sources
	• SRP: Not modeled
	• Atmospheric drag: Not modeled

For the next case we again show the propagated particles in J2000 in Figure 11. We learn three things from this plot; 1) unlike the previous case this result shows that after the 17-day exclusion zone period the uncertainty in the cloud extends significantly beyond lunar distance, 2) it is non-Gaussian, and though tough to see all the patterns, there is structure to the uncertainty that derives from the true measurements (angles and visual magnitude) taken by the sensor. In some sense, structure in the particle cloud that is hard to observe with the naked eye is since the measurement combined with the Jacobi integral equation bias the otherwise uniformly distributed particle cloud towards internal cloud structures that can be detected by the clustering step of the PGMF, and 3) the PGMF is robust to this very large complex uncertainty where a UKF would simply fail due to the wide-spread locations of the propagated sigma points across multiple trajectory bifurcations.

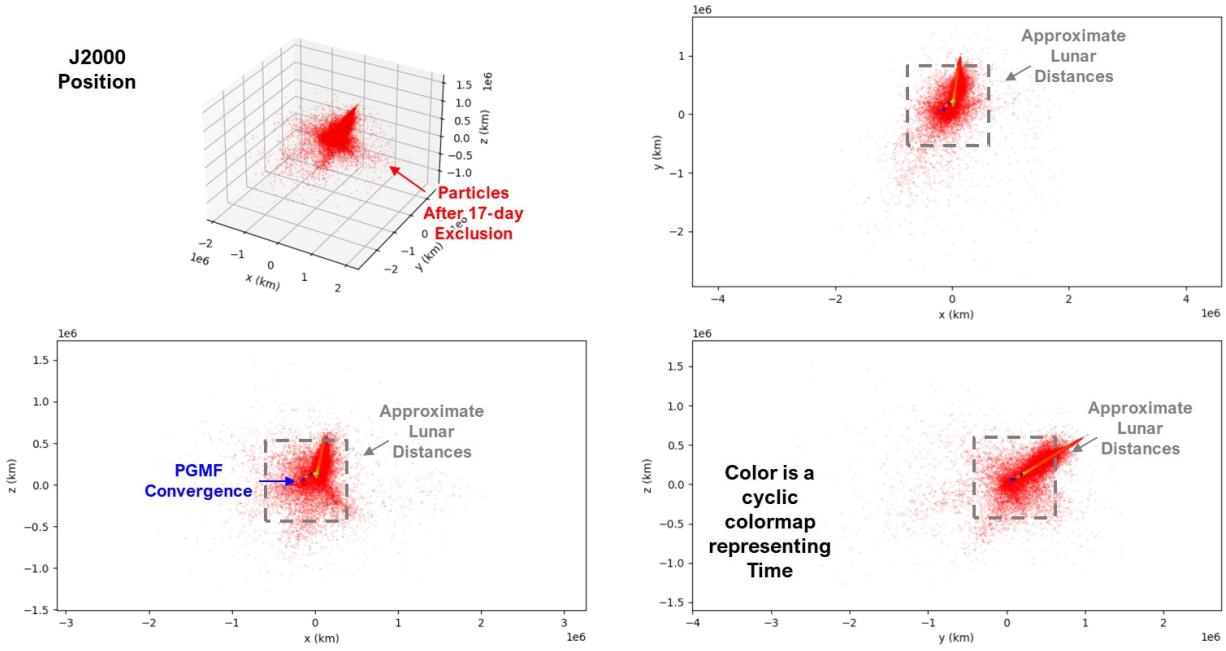

Figure 11: Propagated Particles in J2000 for 1 Observation Before 17 Day Exclusion w/ Realistic Uncertainties

Figure 12 conveys the same message by viewing the particles in velocity space. While these plots are informative, we really want to see how the PAR/PGMF performs with respect to error. Figure 13 shows the particle error as a function of time in the RIC frame relative to the truth along with the initial PAR. The left side shows the 20-day period, and the right side shows the PGMF converging after reacquisition. There are two major takeaways from this plot; 1) the PGMF converges to the truth, and 2) the range uncertainty takes longer to converge. This is expected since the initial range information is highly uncertain, and it is only utilized for the initial PAR.

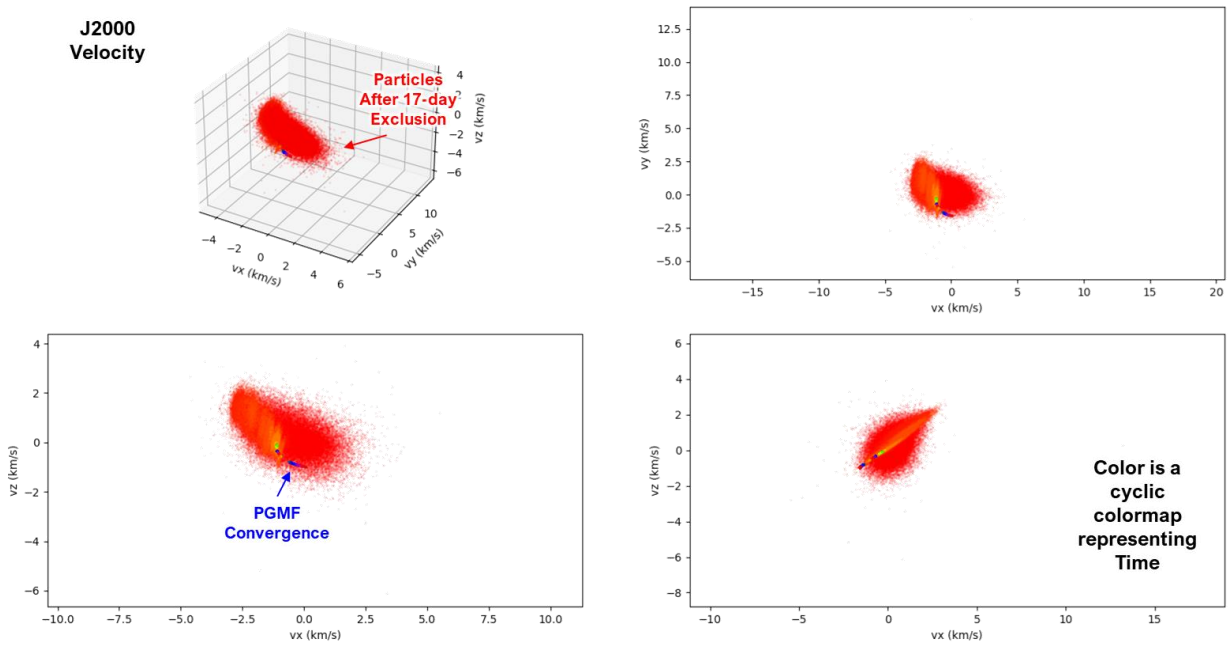

Figure 12: Propagated Particles in J2000 for 1 Observation Before 17 Day Exclusion w/ Realistic Uncertainties

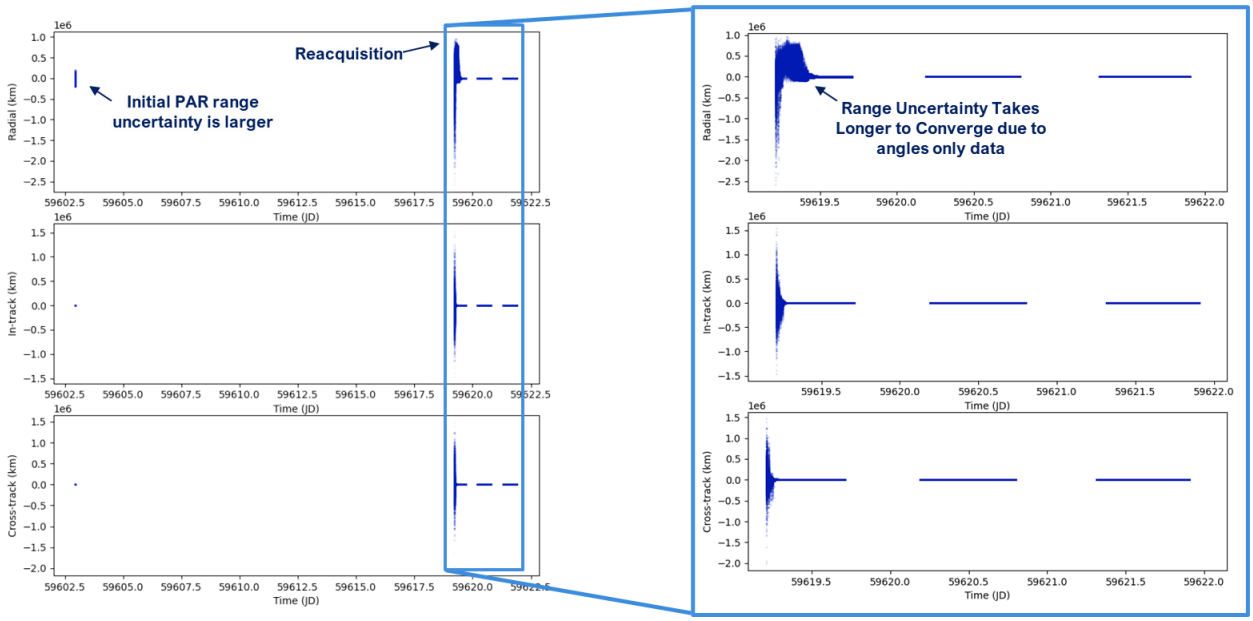

Figure 13: Particle Error for 1 Observation Before 17 Day Exclusion w/ Realistic Uncertainties

To better understand the accuracy achieved over time, we plot the absolute error of the particle cloud with logarithmic scaling on the *y*-axis in Figure 14. We have three observations from this figure; 1) again, we see that the range takes longer to converge, 2) we see that radial uncertainty reduces to 10 km and in-track/cross-track uncertainties reduce to approximately 1 km of uncertainty, and 3) we can visualize evidence of bifurcations in the uncertainty as the PGMF converges.

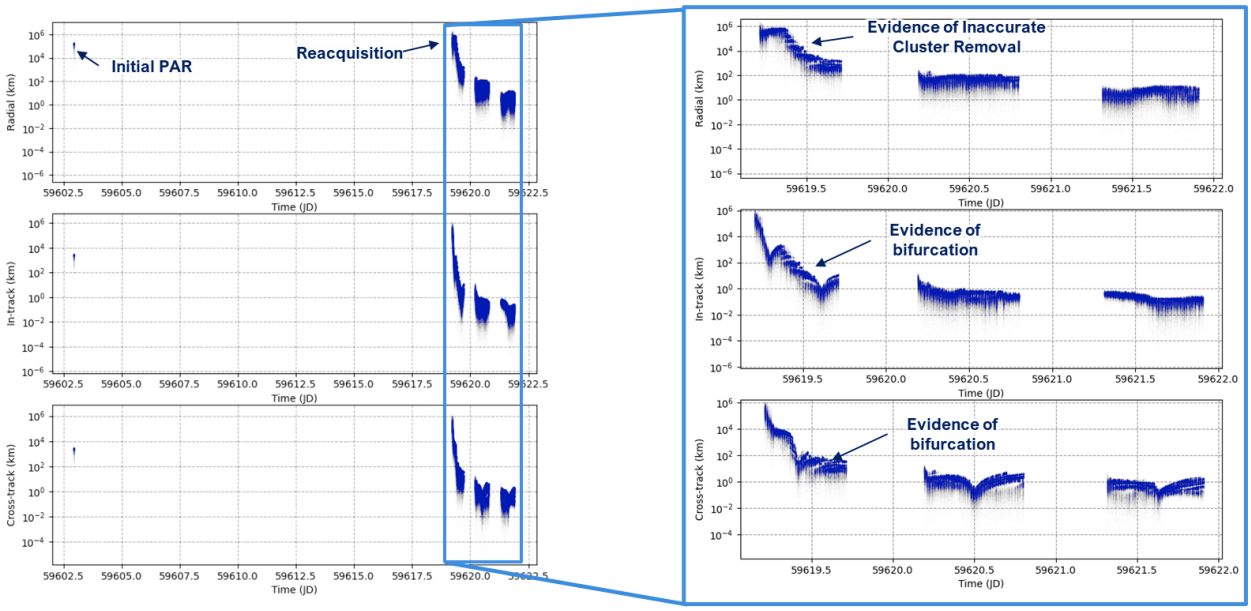

Figure 14:Particle Absolute Error for 1 Observation Before 17 Day Exclusion w/ Realistic Uncertainties

We can zoom in on the radial plot from Figure 14 in Figure 15 to more clearly observe an effect that looks like horizontal stripes in the range uncertainty. This effect is the gradual extinction of the likelihood of GMM clusters with higher error over time. This effect is visible in these plots due to plotting each particle with transparency to emphasize the structure in the cloud.

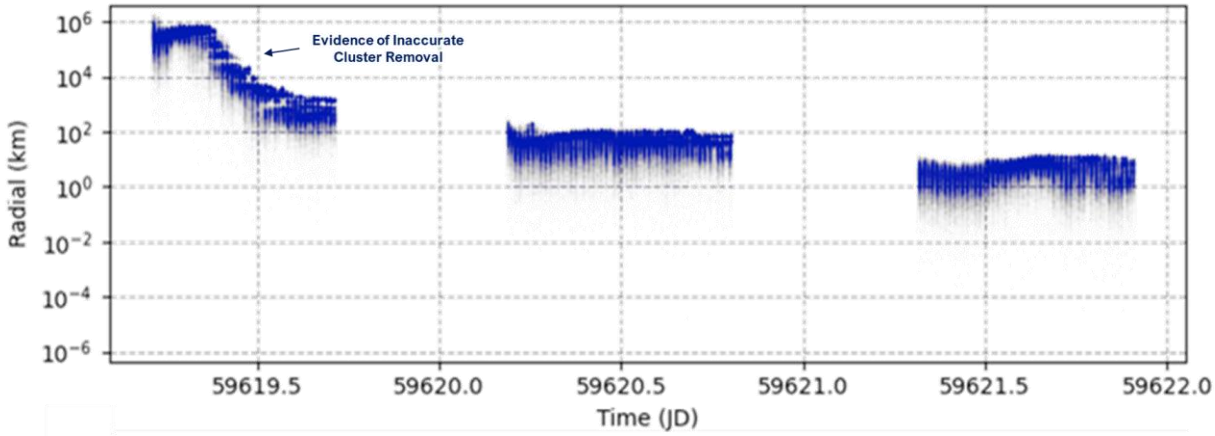

Figure 15: Particle Absolute Error for 1 Observation Before 17 Day Exclusion w/ Realistic Uncertainties

In Figure 16 we see that the entropy continues to decrease as more observation are fed to the filter. Intermittent large reductions in entropy are due to the removal of clusters that are not possible based on new collected observations that have low likelihood with the disappearing GMM components.

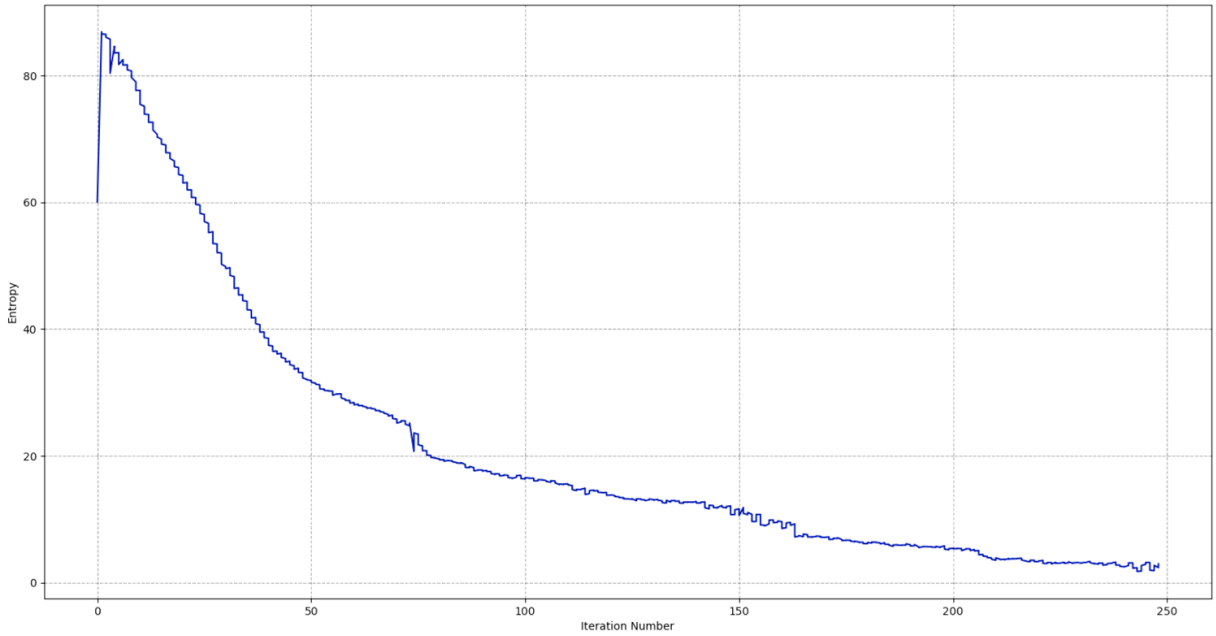

Figure 16: Entropy in nats as a Function of Iteration Number

6. CONCLUSIONS

The PAR-PGMF technique presents a closed form solution that; 1) can initialize a PAR on as few as a single optical observation, 2) accounts for complex uncertainty bifurcations, and 3) can process multi-modal uncertainty distributions. The discussed approach can generate a PAR for cislunar objects by leveraging optical data and performance knowledge of the sensor utilized. The PGMF is robust to very large and complex uncertainties to converge to an accurate object tracking solution. The combined PAR/PGMF approach is a robust probabilistic solution to the optical cislunar IOD/OD challenge. This has been demonstrated on simulated data using full force models for the cislunar domain. Future work will continue to stress test this approach leveraging both simulated and real data from both ground based and space based systems. The work will also expand the range estimate to be agnostic to object shape.

7. REFERENCES

[1] United States Space Force, Spacepower: Doctrine for Space Forces. Nimble Books, 2020.

[2] M. Bolden, T. Craychee and E. Griggs, "An Evaluation of Observing Constellation Orbit Stability, Low Signal-to-Noise, and the Too-Short-Arc Challenges in the Cislunar Domain," Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, Hawaii, 2020.

[3] C. C. Chow, C. J. Wetterer, K. Hill, C. Gilbert, D. Buehler and J. Frith, "Cislunar Periodic Orbit Families and Expected Observational Features," in *AMOS*, Wailea, HI, 2020.

[4] 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," in *AMOS*, Wailea, HI, 2021.

[5] H. D. Curtis, Orbital Mechanics for Engineering Students, Oxford: Elsevier Butterworth-Heinemann, 2005.

[6] P. Herget, The Computation of Orbits, Ann Arbor: Private, 1948.

[7] B. Marsden, "Initial orbit determination- The pragmatist's point of view," *The Astronomical Journal*, vol. 90, pp. 1541-1547, 1985.

[8] R. Furfaro, V. Reddy, T. Campbell and B. Gray, "Tracking Objects in Cislunar Space: The Chang'e 5 Case," in *AMOS*, Wailea, HI, 2021.

[9] J. A. Greaves and D. J. Scheeres, "Relative Estimation in the Cislunar Regime using Optical Sensors," in *AMOS*, Wailea, HI, 2021.

[10] R. M. Thompson, N. P. Re, C. Meek and B. Cheetham, "Cislunar Orbit Determination and Tracking via Simulated Space-Based Measurements," in *AMOS*, Wailea, HI, 2021.

[11] S. Wishnek, M. J. Holzinger and P. Handley, "Robust Cislunar Initial Orbit Determination," in *AMOS*, Wailea, HI, 2021.

[12] I. Hussein, C. W. T. Roscoe, M. Mercurio, M. P. Wilkins and P. W. Schumacher, Jr., "Probabilistic Admissible Region for Multi-Hypothesis Filter Initialization," *J. Guidance, Control and Dynamics,*, vol. 41, no. 3, 2017.

[13] D. Raihan and S. Chakravorty, "Particle Gaussian Mixture Filters II," *Automatica*, vol. 98, pp. 341-348, 2018.

[14] D. Raihan and S. Chakravorty, "Particle Gaussian mixture filters-I," *Automatica*, vol. 98, pp. 331-340, 2018.

[15] A. Gelb, Applied Optimal Estimation., MIT Press, 1974.

[16] T. Bengtsson, P. Bickel and B. Li, "Curse-of-dimensionality revisited: Collapse of particle filter in very large-scale systems," *Probability and Statistics: Essays in Honor of David A. Freedman*, vol. 2, p. 316–334, 2008.

[17] T. Kelecy, M. Jah and K. DeMars, "Application of a Multiple Hypothesis Filter to near GEO high area-to-mass ratio space objects state estimation," *Acta Astronautica*, vol. 81, pp. 435-444, 2012.

[18] M. D. Hejduk, "Specular and Diffuse Components in Spherical Satellite Photometric Modeling," in *AMOS*, Wailea, HI, 2011.

[19] V. Szebehely, Theory of Orbits: The Restricted Problem of Three Bodies, New York : Academic Press, 1967.

[20] S. P. Lloyd, "Least squares quantization in pcm," *IEEE Transactions on Information Theory*, vol. 28, no. 1, p. 129–137, 1982.

[21] A. K. Jain, M. N. Murty and P. J. Flynn, "Data clustering: a review," *ACM Computing Surveys*, vol. 31, no. 3, p. 264–323, 1999.