Enhanced Process Noise Application in Estimation Filters For Cislunar Orbit Determination and Maneuver Detection

Charles J. Wetterer[1], Jason, Baldwin[2], Paul Billings[1], Christopher Craft[1], Juan Gutierrez[1], John Gaebler[1], Micah Dilley[1], and Jill Bruer[3]

[1]KBR

[2]Complex Futures

[3]Air Force Research Laboratory

ABSTRACT

Orbit determination and maneuver detection for objects in cislunar space using a sequential estimation filter require a robust dynamics model. At a minimum, this time-update state function needs to include the gravitational influence of the Earth and Moon, but oftentimes higher-order aspherical gravitational effects, as well as the influences of other bodies (e.g., Sun) and solar radiation pressure (SRP), must also be included. Due to their complexity and the lack of full knowledge of the underlying forces, however, even a high-fidelity dynamics model will always have an unavoidable mismatch when compared with reality, manifesting as both mismodeled and unmodeled effects. This mismatch is especially deleterious in regimes of chaotic dynamics, such as in the cislunar environment. To account for the dynamics mismatch, a small process noise is typically added to the state covariance during the time update step in the filtering sequence. Insufficient process noise will lead to filter overconfidence, unresponsiveness, and divergence, while excessive process noise will reduce the accuracy and precision of the state estimate. This technique is also common in maneuver detection using an interacting multiple model (IMM) approach, where an inflated process noise is used to account for statistically possible maneuver accelerations in a separate maneuver model, but excessive process noise (in the nominal model) greatly reduces sensitivity to maneuvers. A common process noise model assumes an additive random process with fixed variance applied at the time of each measurement. This fails to work adequately, however, when the observation cadence varies or when there are large observation gaps due to object observability, both of which are common occurrences in the cislunar regime. This model can be improved somewhat by implementing a time-dependent scaling factor on the variance, but the scaled noise approach still suffers from directional assumptions, usually resulting in a choice of process noise that overestimates the uncertainty in some directions in order to retain adequate uncertainty in other directions. Ideally, to capture the effect of process noise most accurately, it would be embedded in the propagation machinery, continuously integrated and allowed to evolve as the encountered orbital environment dictates. For a highly nonlinear environment such as cislunar, an adaptive sigma point filter is an effective approach for capturing the time evolution of the probability density function (PDF). Within this filtering approach, we apply process noise to the PDF within the numerical integrator by directly augmenting the sigma points in lieu of inflating the state covariance at the discrete measurement times. This paper compares the performance of this new process noise modeling to the previously developed simple and scaled process noise models for objects in cislunar halo orbits. The integrated process noise model performs as well as or better than the other process noise models with no need to tune the time interval between applications to avoid excessive information loss or misorientation of the uncertainty ellipse.

1. INTRODUCTION AND MOTIVATION

Our previous papers provide an overview of the expected photometric and astrometric measurements of objects in cislunar orbits [1], demonstrate the use of these observations to conduct initial orbit determination[2][3] and orbit maintenance using estimation filters [4][5], examine the debris environment and evolution of cislunar object breakups [6], and explore the detection and characterization of impulsive maneuvers for an object in a cislunar orbit using similar estimation filters[7]. Implicit in the use of estimation filters is the need to apply a process noise to account for unmodeled or mismodeled dynamics in the filter's state transition function. Insufficient process noise will lead to filter overconfidence, unresponsiveness, and divergence, while excessive process noise will reduce the accuracy and precision of the state estimate and complicate measurement association. The application of process noise is also common in maneuver detection using an interacting multiple model (IMM) approach, where an inflated process noise is used to account for statistically possible maneuver accelerations in a separate maneuver model, but excessive process noise (in the nominal model) greatly reduces sensitivity to maneuvers. In this paper, we first examine the typical way process noise is applied, namely the discrete addition of process noise to the state

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

covariance during the time update step in the estimation filter (e.g. an Unscented Kalman filter [UKF]). Even though a commonly described approach (as in, for example, [8])–which we term the "simple" process noise model in this paper–adjusts the amount of process noise applied based on the time since the previous application, we will show that the magnitude of the process noise needs to be further scaled to preserve the reference motion statistics when the time between applications varies.

In UKF filtering when considering process noise, two goals/operations are in tension: process noise addition and sigma point propagation. Every time the discrete noise is added, the distribution described by the sigma points must be "collapsed" to its Gaussian representation (because of the addition of the process noise covariance). This, by construction, removes the non-Gaussianity from the distribution's representation, and thus interferes with any ongoing measure of the non-Gaussianity (such as differential entropy). So, to maintain an accurate running measure of non-Gaussianity, we would prefer the time between noise additions Δt to be large; however, a large Δt widens the accuracy gap between the true discrete process noise and its approximation Q_D. This is strongly true if Q_D comes from the simple or scaled process noise models and is less strongly true (but still significant) if Q_D is a zero-order hold approximation using, e.g., the Van Loan model [9]. To achieve a fully accurate Q_D, continuous noise should be integrated numerically; however, this is computationally expensive, and still maintains the drawback of requiring the disposal of non-Gaussianity information when it is applied We will instead develop an integrated process noise model to enforce the noise power constraint onto the sigma point distribution itself without discretization, allowing the distribution to stay in sigma-point representation (with the inherent progressive nonlinearities) as long as possible, e.g., until a measurement update. This also enables more direct and well-behaved measures of non-Gaussianity, e.g., skewness. If desired, the continuous process noise power could be varied as a function of state and/or time (e.g., encoding a larger/smaller maneuver capability in specific places in an orbit or in Earth-moon system), though we do not address that here. We show this integrated process noise model works as well as, and often better than, the discrete process noise models through various examples.

2. BACKGROUND AND DEVELOPMENT

2.1 "Simple" and "Scaled" models: scaling discrete process noise with irregular measurement intervals

Typically, process noise is applied to the state covariance during the time update step in the estimation filter. In this paper, we use the Unscented Kalman Filter (UKF) [10][11] although its use can be applied in extensions of the UKF such as in the Adaptive Entropy-based Gaussian mixture Information Synthesis (AEGIS) filter [12] and IMM approach for maneuver detection [8].

The continuous-time equations of motion of a particle (assuming ballistic dynamics) typically take the form

$$\dot{x}(t) = f(x(t), t) + G(t)w(t) \tag{1}$$

where $f(\cdot)$ is the dynamics function, $\mathrm{G}(t)$ is a process noise mapping matrix (which is meant to distribute the process noise into the state derivatives), and $w(t)$ is a white, zero-mean Gaussian process noise vector such that

$$E[w(\tau)w(\tau)^T] = \tilde{Q}(\tau)\delta(t-\tau) \tag{2}$$

where $\tilde{Q}(t)$ is the process noise power and $\delta(\cdot)$ is the Dirac delta function.

For simplicity of explanation, if we assume a one-dimensional state, discretize and linearize the equations of motion, and use a discrete white noise acceleration model, Eq. (1) becomes

$$x(k+1) = Fx(k) + \Gamma\upsilon(k) \tag{3}$$

where the state vector $x = [position \quad velocity]^T$ for one-dimensional motion, and the process noise $\nu(k)$ is a white, scalar acceleration sequence with variance σ_a^2. The transition matrix is

$$F = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \tag{4}$$

and the vector gain on the process noise is

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

$$\Gamma = \begin{bmatrix} \frac{1}{2}T^2 \\ T \end{bmatrix} \tag{5}$$

The covariance of the process noise multiplied by the gain, $\Gamma v(k)$, is $Q = E[\Gamma v(k) v(k) \Gamma^T]$, with a velocity squared uncertainty or variance term after one reference time interval T of

$$\Delta U^2(T) = T^2 \sigma_a^2 \tag{6}$$

As this process is assumed cumulative and independent, over multiple time intervals the variance of the sum equals the sum of the variances. After n time intervals, the change in velocity variance becomes

$$\Delta U^2(nT) = nT^2 \sigma_a^2 \neq (nT)^2 \sigma_a^2. \tag{7}$$

From this, we see that a filter which does not add process noise at every period T – e.g., because measurements were not available at all those times – will have an incorrect uncertainty if Q (via Γ) is based on the elapsed time without any other adjustments. The "simple" approach is this erroneous case – a fixed value for σ_a^2 that is quadratically scaled by a kinematic factor in the form of (5) with T equal to the time since the previous observation.

Consider then, a change in the sampling time interval to be a (possibly non-integer) multiple of the original: $\tau = nT$. Rearranging to find n and substitution gives

$$\Delta U^2(\tau) = \frac{\tau}{T} T^2 \sigma_a^2 = \tau T \sigma_a^2 \tag{8}$$

Conceptually, we desire an adjusted process noise variance for a filter operating at the new sampling period to match the form of Eq. (6). In other words,

$$\Delta U^2(\tau) = \tau^2 \sigma_{adj}^2 \tag{9}$$

Comparing Eqs. (8) and (9), we find the required adjusted acceleration process uncertainty must be

$$\sigma_{adj} = \sqrt{\frac{1}{n}} \sigma_a \tag{10}$$

In a similar fashion, when $\tau < T$, it is the τ-based filter that accumulates, but this leads to the same scale factor as before where now n is a fraction less than one. The scaling factor will decrease the acceleration process uncertainty for time intervals longer than the reference time T ($n > 1$) and increase the acceleration process uncertainty for time intervals shorter than the reference time T ($n < 1$). The "scaled" process noise approach therefore uses this adjusted uncertainty when computing the additive term ($Q = \Gamma \sigma_{adj}^2 \Gamma^{\mathrm{T}}$) for the state covariance, and Γ is again based on the time since the previous measurement.

A similar derivation can be conducted for a discrete white noise velocity increment model; and interestingly, the scaling factor is the reciprocal of what was found for the acceleration model due to differences in the form of Γ.

$$\sigma_{adj} = \sqrt{n} \sigma_a \tag{11}$$

For the velocity increment model, the scaling factor will increase for time intervals longer than the reference time T ($n > 1$) and decrease for time intervals shorter than the reference time T ($n < 1$). For the typical case of longer intervals, the simple approach with a velocity increment model may be prone to overconfidence and measurement association difficulty due to insufficient noise addition.

While the need to adjust the process noise strength is corroborated as a footnote in [8], it is our opinion that this additional scaling may have been overlooked in some filter implementations in the industry and is therefore mentioned here. Most of the remainder of this paper will be a comparison of the "scaled" and the "integrated" approaches (described below).

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

2.2 "Integrated" model: process noise within numerical integrator

The method for process noise addition described in the previous section is based on inflation of the state covariance during the measurement update using the computed discrete noise covariance. This covariance increment could be anisotropic (e.g., increased noise in the in-track direction for LEO applications) or isotropic as desired. This approach constitutes a significant approximation and is appropriate in timescales where the dynamics may be regarded as linear. In a highly nonlinear astrodynamics scenario with long intervals between measurements, it seems necessary to account for the process noise evolution more robustly. For example, in a continuous formulation of a cislunar orbit, an incremental increase of uncertainty specifically along the in-track direction at the beginning of a period will be transformed by the dynamics into an uncertainty that is decidedly not along the in-track direction when a measurement arrives several days later.

Returning to Eqs. (1) and (2), we again discretize the continuous equations of motion, but now maintain the nonlinear dynamics mapping, and obtain

$$x(t_{k+1}) = \tilde{f}(x(t_k), t_k, t_{k+1}) + w_k \tag{12}$$

where $\tilde{f}(\cdot)$ is now the discrete propagation mapping and w_k is the discrete process noise random variable with

$$E[w_k w_k^T] = Q_{D,k} \tag{13}$$

where $Q_{D,k}$ is the discretized process noise covariance for time step k. To compute Q_D for a given time step, we must solve the equation

$$Q_{D,k} = \int_{t_{k-1}}^{t_k} \Phi(t_k, \tau) G(\tau) \tilde{Q}(\tau) G(\tau)^T \Phi(t_k, \tau)^T \, d\tau \tag{14}$$

where $\Phi(t_2, t_1)$ denotes the state transition matrix from time t_1 to time t_2. This integral is generally nontrivial to calculate, so one estimation method involves the Van Loan approximation [9]:

$$M = \begin{bmatrix} -\frac{\partial f}{\partial x}\Big|_{x(t_{k-1})} & G\tilde{Q}G^T \\ \mathbf{0} & \frac{\partial f}{\partial x}\Big|_{x(t_{k-1})}^T \end{bmatrix} (t_k - t_{k-1}) \tag{15}$$

$$B = e^M = \begin{bmatrix} \cdots & \mathrm{K}^{-1} Q_D \\ \mathbf{0} & K^T \end{bmatrix} = \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \tag{16}$$

$$Q_{D,k} = B_{22}^T B_{12} \tag{17}$$

where K is a state-transition-like matrix that does not require direct computation (because it factors out in the final expression), and, importantly, e^M involves the matrix exponential. From the above equations, it is directly evident that the Van Loan model's primary assumption is that of a zero-order hold with respect to the dynamics (i.e., assumes linear dynamics across the time interval), and so it is only accurate for "small enough" time intervals.

Alternatively, returning to Eqs. (1) and (2): to maintain the continuous formulation, we can enforce the noise power constraint onto the sigma point distribution itself without discretization. Consider the propagation of the ith sigma point within the distribution. We have

$$\dot{\chi}_i(t) = f(\chi_i(t), t) + f_{Q,i}(t) \tag{18}$$

where now the propagation of a sigma point $\chi_i(t)$ depends on the dynamics f *as well as* a fictitious force $f_{Q,i}(t)$ imposed collectively to ensure that the overall sigma-point distribution matches the desired noise statistics. For the evolution of the process noise covariance contribution to be zero-mean, Gaussian, and have the growth rate as described by its noise power, there are [at least] two conditions on the noise-driven evolution of the first two moments of the distribution:

$$\dot{P} = \frac{d}{dt} P = \tilde{Q}(t) \tag{19}$$

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

$$\frac{d}{dt}\bar{\chi} = \mathbf{0} \tag{20}$$

The first equation states that the sample covariance rate induced by $f_{Q,i}(t)$ should be equal to the noise power, while the second constraint enforces the zero-mean nature of the noise process by requiring the time rate of sample mean, $\bar{\chi}$, induced by $f_{Q,i}(t)$ to be zero.

The following development assumes a 2N+1 set of sigma points, though it is likely that comparable results hold for the 2N method. First, define $\delta\chi_i = \chi_i - \bar{\chi}$ where, again, $\bar{\chi}$ is the standard unscented transform's sample estimate of the component mean: $\bar{\chi} = \sum_{j=0}^{2N} w_j^0 \chi_j$, and w_j^0 is the mean weighting of the jth sigma point. Then the estimated component covariance is

$$P = \sum_{j=0}^{2N} w_j^C \delta\chi_j \left(\delta\chi_j\right)^T \tag{21}$$

where w_j^C is the standard covariance weight for the jth sigma point defined in the unscented transform. The time rate of change of this sample covariance is:

$$\dot{P} = \sum_{j=0}^{2N} w_j^C \left\{\delta\chi_j (v^j)^T + v^j \left(\delta\chi_j\right)^T\right\} \tag{22}$$

where $v^j = \frac{d}{dt}\left(\delta\chi_j\right) = \frac{d}{dt}\left(\chi_j - \bar{\chi}\right) = \frac{d\chi_j}{dt}$. When written elementwise, the covariance rate is

$$\dot{P}_{mn} = \sum_{j=0}^{2N} w_j^C \left\{\delta\chi_{j,m} v_n^j + v_m^j \delta\chi_{j,n}\right\} \tag{23}$$

where $\delta\chi_{j,m}$ is the mth element of the jth sigma point displacement. Now, since this expression for $\dot{P}_{mn}$ is linear in the various displacement rate components v_k^j, to enforce the constraint given in Eq. (19), it is convenient to define a set of equations in the form:

$$\sum_{j,k} \frac{\partial \dot{P}_{mn}}{\partial v_k^j} v_k^j - \tilde{Q}_{mn} = 0 \tag{24}$$

Assembling both the sought displacement rates and the noise power constraint values into vectors, we obtain the system

$$EV = W \tag{25}$$

where V is the stacked vector of displacement rate vectors,

$$V = \begin{bmatrix} v^0 \\ v^1 \\ \vdots \\ v^{2N} \end{bmatrix} \tag{26}$$

the constraint matrix $W = \left[\tilde{Q}_{11}, \tilde{Q}_{12}, \ldots, \tilde{Q}_{NN}\right]^T$, and the gradient matrix E is a matrix whose elements are defined such that

$$\frac{\partial \tilde{Q}_{mn}}{\partial v_k^j} = \frac{\partial \dot{P}_{mn}}{\partial v_k^j} = w_j^C \left\{\delta\chi_{j,m}\delta_{kn} + \delta\chi_{j,n}\delta_{km}\right\} \tag{27}$$

where δ_{ab} is the Kronecker delta; i.e.:

$$\delta_{\mathrm{ab}} = \begin{cases} 1 & a = b \\ 0 & a \neq b \end{cases} \tag{28}$$

Due to the symmetry of the covariance relationship and avoidance of redundant equations, only the upper triangular region of $\tilde{Q}(t)$, i.e., $n \geq m$, is used when forming the E and W matrices, resulting in fewer than N^2 rows.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

An additional modification is required to satisfy the mean rate constraint (20). We augment the matrix E and the right-hand side W as

$$\tilde{E} = \begin{bmatrix} E \\ B \end{bmatrix}, \tilde{W} = \begin{bmatrix} W \\ \mathbf{0}_{N\times 1} \end{bmatrix} \tag{29}$$

where $\mathbf{0}_{N\times 1}$ is a vector of zeros of the same dimension as the state, and B is a $[N \times (2N+1)N]$ matrix constructed as

$$B = [w_0^0 I_N \quad w_1^0 I_N \quad \cdots \quad w_{2N}^0 I_N] \tag{30}$$

Where I_N is the $[N \times N]$ identity matrix.

The augmented system is now

$$\tilde{E}V = \tilde{W} \tag{31}$$

and represents a linear, but underdetermined, system of equations in V (i.e., there is an infinite number of sets of velocities v^j of 2N+1 sigma points to achieve the desired statistics). So, we choose the least-norm solution

$$V_{LN} = \tilde{E}^T\left(\tilde{E}\tilde{E}^T\right)^{-1}\tilde{W} \tag{32}$$

which effectively specifies the minimum acceleration V that satisfies the constraints. The vectors constituting V_{LN} now represent the additional process noise accelerations $f_{Q,i}(t)$ to be added to the sigma points at each evaluation of the dynamics ODEs during integration. The application of process noise in this manner will be termed the integrated process noise model.

3. SAMPLE ORBITS AND OBSERVATIONS

A northern halo orbit about the L1 Lagrange point (HN1)—similar to the one used in our previous papers—and a southern near-rectilinear halo orbit (NRHO) about the L2 Lagrange point were used to compare and evaluate the various process noise models. The reference CR3BP orbits for the HN1 orbit and NRHO are HN1-0207 and NRHOS2-0184 respectively from the Air Force Research Laboratory (AFRL) cislunar periodic orbit library, an upgraded version of the orbit library described in our original cislunar paper[1]. HN1-0207 is unstable and has a period of approximately 12.1 days, with a Jacobi constant of 3.117, while the NRHOS2-0184 is a 6.53-day orbit at a Jacobi constant of 3.047, with a perilune radius of approximately 3140 km. These reference orbits were used to generate pseudo-periodic orbits in the full ephemeris model. Both full-ephemeris orbits are purely ballistic, generated using (all DE440 ephemeris) Earth 8×8, Moon 20×20, CrA/m=0.01092 m^2/kg, and E_solar_1AU = 1360.8 W/m^2 and constrained to stay as close to the reference orbit as possible. The initial epoch for both orbits is UTC 2024-01-01, 12:00:00.0 and both are propagated forward for 121 days. Fig. 1 shows the first month of both orbits in an Earth-Moon pulsating synodic frame (Moon fixed to [1, 0, 0]).

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

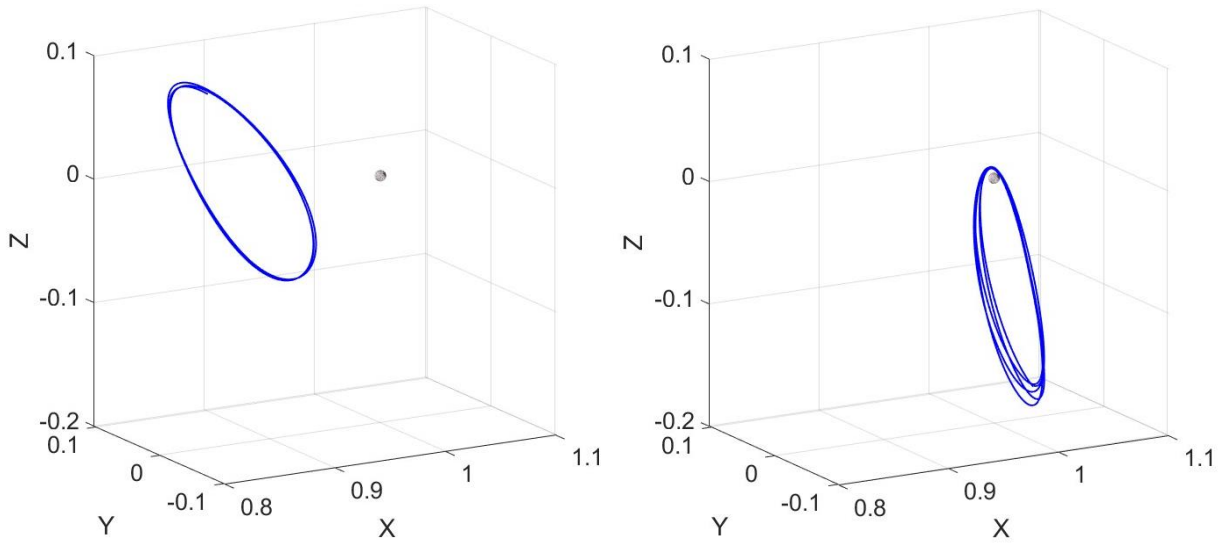

Fig. 1. HN1 (left) and NRHO (right) ballistic orbit (first month or about 3 periods (HN1) and 5 periods (NRHO)) in Earth-Moon pulsating synodic frame

A notional set of longitudinally separated Earth ground-based sensor sites are used in generating simulated observations to ensure the object is continuously visible from at least one of the sites: Maui (20.7˚ N, 156.4˚ W), Azores (37.8˚ N, 25.5˚ W), and Cocos Islands (12.2˚ S, 96.9˚ E). Five observation tracks with an obs-to-obs cadence of 60 seconds and a track-to-track cadence of 3600 seconds are generated for each site when the object is visible. Visibility thresholds will greatly affect the days and number of observations when the object is visible, as shown in Fig. 2, where three different sets with different thresholds are plotted for the NRHO ballistic orbit. The first set was designed to provide observations every day for use in later analysis, albeit with unrealistic thresholds: object elevation > 0 degrees above horizon, Sun elevation < 15 degrees above horizon and no limit on visual magnitude. The second set provides physically possible but still unrealistic observations to show the best-case limit for electro-optical observations: object elevation > 10 degrees, Sun elevation < -10 degrees (Sun exclusion), and visual magnitude < 28 (for a 1 m^2 cross sectional area diffuse sphere with 0.2 reflectance). Finally, the last set uses more realistic thresholds to show a more realistic picture on observation opportunities: object elevation > 20 degrees, Sun elevation < -20 degrees (Sun exclusion), and visual magnitude < 19. All sets require object > 1 degree from limb of Moon (Moon exclusion). Each observation has a measurement uncertainty (mostly set to 1 arcsecond in analyses to follow) in both right ascension and declination, where one arcsecond corresponds to about 1.86 km at the Moon's mean distance. These three sets produce 2588, 1687, and 589 tracks respectively over the 121-day simulation. Again, although the two dense sets of observations are unrealistic, they are easily decimated to both decrease the volume and increase the time gaps as needed for the analysis. Similarly, observations are generated for the HN1 orbit with 2700, 1761, and 937 tracks respectively for the three sets over the 121-day simulation. Table 1 summarized the exclusion limits and resulting tracks for the observation sets.

Table 1. Observation Sets

Set	Object elevation	Sun elevation	Visual magnitude	Angle from Moon	HN1 tracks	NRHO tracks
1	> 0˚	< 15˚	< Inf	> 1˚	2700	2588
2	> 10˚	< -10˚	< 28	> 1˚	1761	1687
3	> 20˚	< -20˚	< 19	> 1˚	937	589

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

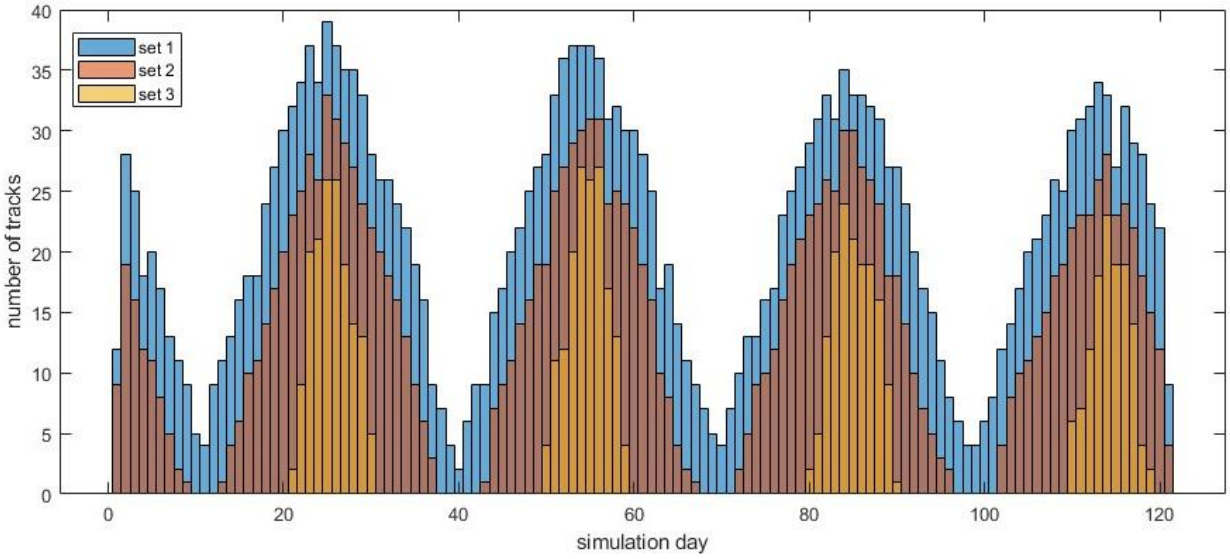

Fig. 2. Number of tracks versus simulation day for NRHO ballistic orbit – three different sets with different thresholds

4. ANALYSIS AND RESULTS

4.1 Tuning of the process noise using a ballistic orbit

To demonstrate the need for additional scaling of the process noise magnitude when applied at discrete time intervals (i.e., the adjusted noise variance in Eq. (10)), the ballistic orbit is processed with the estimation filter using two decimations of Set 2 for the ballistic NRHO. First, a maximum of one observation in two different tracks each day (for a total of 204 observations with an average observation cadence of ~12 hours), and second, all observations in every track (for a total of 8435 observations with an average observation cadence of ~15 minutes). A dynamics mismatch was imposed using Earth 6×6 and Moon 12×12 gravity models, and CrAmr = 0.01 m^2/kg which required an acceleration-based process noise of 3×10^{-11} km/s^2 with a reference time of 1 hr for satisfactory statistics, or equivalently 8.66×10^{-12} km/s^2 for a 12-hour reference time. (Note that one could account for these mismatches with consider parameters within the estimation state; a comparison between using consider parameters and process noise is a potential future analysis.) Fig. 3 displays the squared Mahalanobis distance against the expected chi-squared statistics at both cadences for the scaled process noise with consistent results. (In the titles, "PDFsum" is the displayed portion of the error PDF, with a value < 1 indicating extent beyond the graph limits. "AOC" is the area under the chi-squared curve.)

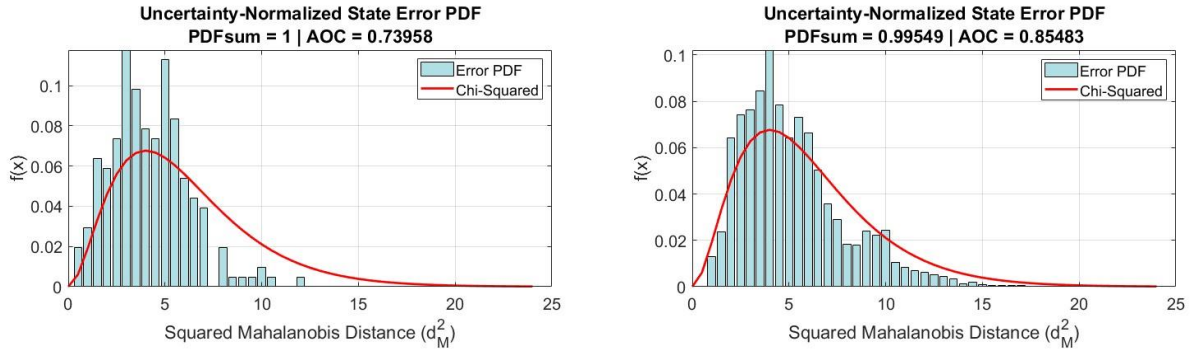

Fig. 3. Histogram of squared Mahalanobis distance compared against expected chi-squared statistics for ~12-hour cadence run (left) and ~15-minute cadence (right) for scaled process noise model

Whereas the scaled model's noise variance is scaled at each observation per Eq. (10), the simple model's noise variance is fixed for every observation in both sets of data. (Both are subsequently scaled by the same time-dependent factor Γ to account for the kinematics.) The simple process noise variance was set to 8.66×10^{-12} km/s^2 to target the ~12-hour cadence and Fig. 4 shows the results. The ~12-hr cadence run yielded the expected statistics

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

(similar to the scaled runs) while the ~15-min cadence statistics illustrate insufficient process noise (more instances of higher squared Mahalanobis distance).

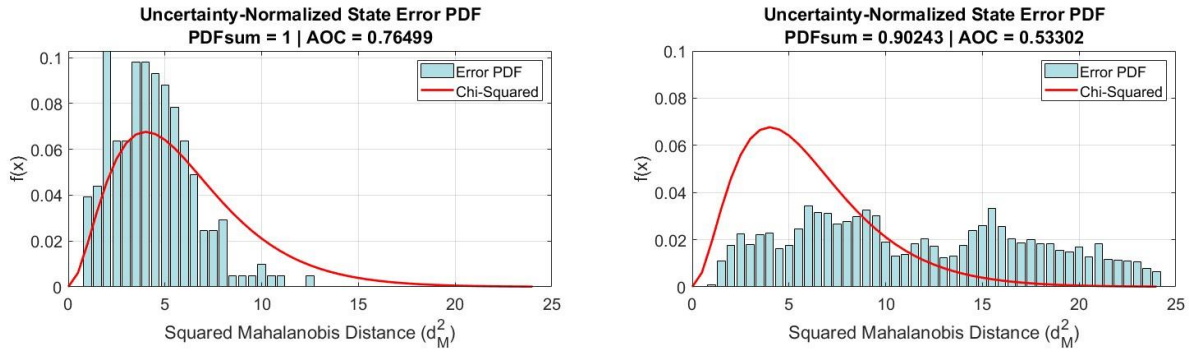

Fig. 4. Histogram of squared Mahalanobis distance compared against expected chi-squared statistics for ~12-hour cadence run (left) and ~15-minute cadence run (right) for simple process noise model

The integrated process noise model requires a different base factor, namely the continuous process noise power, and this was tuned to 2×10^{-9} km^2/s^3 using the 8435-observation dataset (~15 minute cadence). Fig. 5 displays the squared Mahalanobis distance against the expected chi-squared statistics for both runs with the integrated process noise model. Both runs yield very similar results indicating, as desired, that the continuous process noise power is independent of observation cadence, similar to the scaled process noise model.

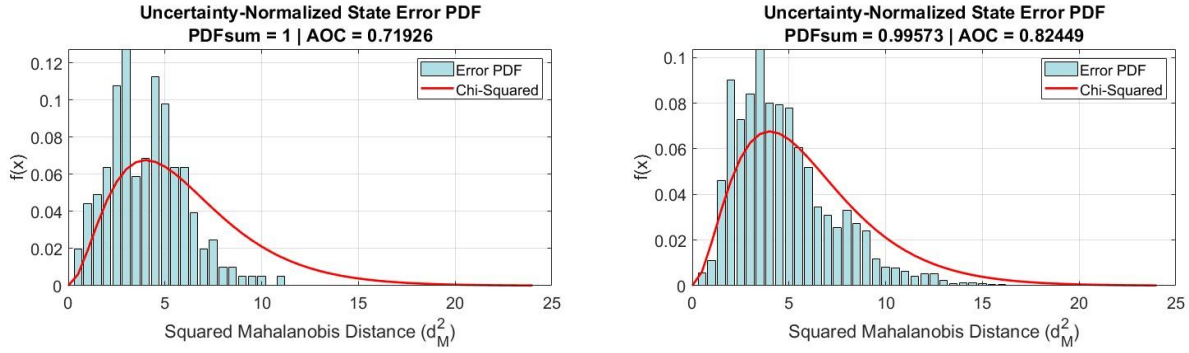

Fig. 5. Histogram of squared Mahalanobis distance compared against expected chi-squared statistics for ~12-hour cadence run (left) and ~15-minute cadence run (right) for integrated process noise model

4.2 Comparison of uncertainty using various models

To give some intuition for the behavior of the three methods (scaled process noise model, Van Loan model, and integrated process noise model), numerical simulations were conducted at different application cadences (only relevant for the scaled and Van Loan models) and compared against "truth." The "truth" is a Monte Carlo simulation with 90,000 particles (generated using a Sobol quasi-random distribution [13]) via a sequential resampling process for each particle from its Van Loan uncertainty approximation, using a dense and adaptive time interval Δt. Each interval is lengthened (if change in the largest eigenvalue of the dynamics Jacobian is less than 0.8%) or shortened (if the change in the largest eigenvalue of the dynamics Jacobian is greater than 1%) accordingly.

The scaled process noise model is assumed to have an isotropic acceleration of 1.65×10^{-7} km/s^2 for the HN1 orbit and 3.4×10^{-8} km/s^2 for the NRHO, and the Van Loan and integrated process noise models are assumed to have an isotropic acceleration noise power of 1×10^{-5} km^2/s^3 for the HN1 orbit and 2×10^{-6} km^2/s^3 for the NRHO. Note that the acceleration and acceleration power values for each orbit represent a relatively large amount of process noise and are chosen to visualize the differences over a short time span. Initial uncertainties were sub-millimeter 1-sigma in position and sub-micron per second 1-sigma in velocity. The HN1 orbit is propagated 5 days forward and the NRHO is propagated 3.5 days forward. Fig. 6 shows the portion of the orbits traversed during these times and the final Monte Carlo distributions in a pulsating Earth-Moon synodic frame.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

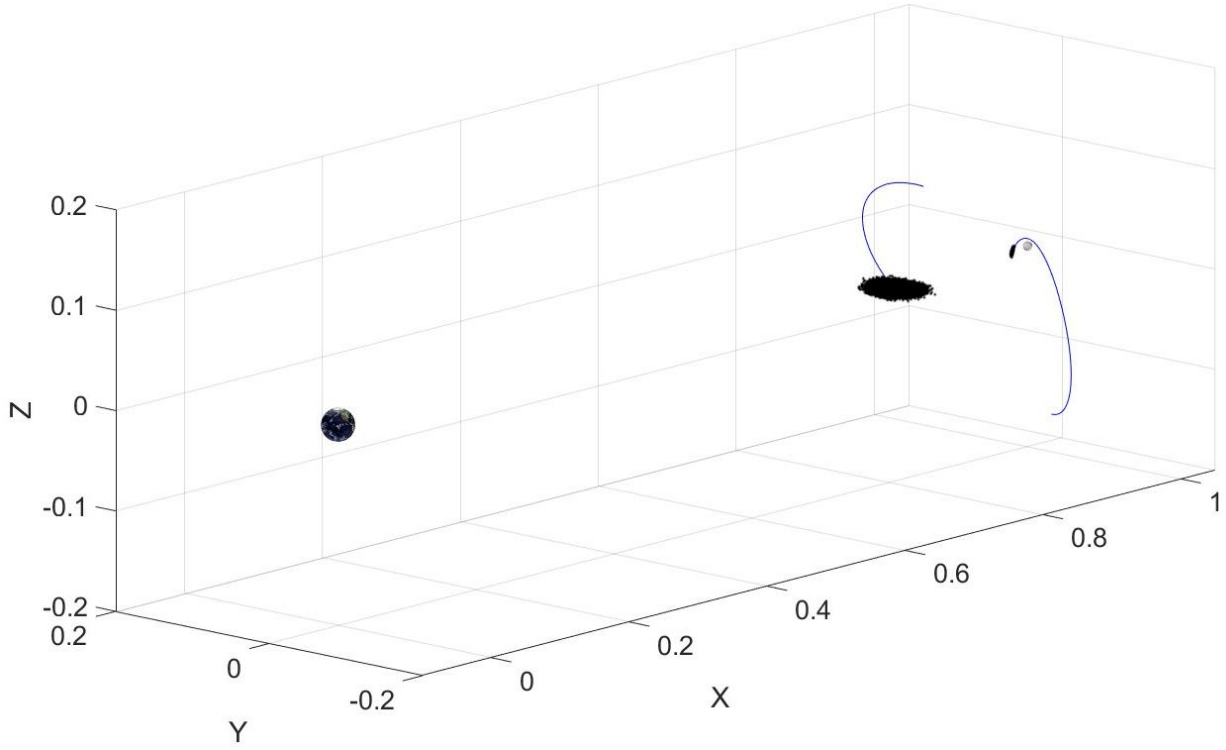

Fig. 6. Portions of HN1 orbit (propagated 5 days) and NRHO (propagated 3.5 days) in pulsating Earth-Moon synodic frame showing ending Monte Carlo distribution starting from near zero initial uncertainties

Fig. 7 shows the 3-σ position covariance ellipsoids (the 3x3 projection in the x,y,z space of the 6x6 covariance) of the discrete process noise covariance matrices Q_D at the end of a 5-day simulation period for the HN1 orbit. The left-hand plot applies all the process noise at the end of the simulation period (so-called "zero order hold") representing one extreme of sampling cadence while the right-hand plot applies the process noise using the same adaptive time scale as the Monte Carlo particles (Δt = 30 minutes at start of simulation and 25 minutes (minimum) at end) representing the other extreme of sampling cadence. In the left-hand plot, it is evident that any dynamical effects are completely absent from the scaled process noise model, due to the lack of any propagation after the noise addition. The Van Loan model tracks well with the Monte Carlo distribution for a while but then begins to diverge as the dynamics depart from the zero-order dynamics hold assumption. (This is similar to the difference between an Extended Kalman filter and a UKF.) Interestingly, it stays remarkably close to truth in hypervolume but the orientation is increasingly degraded. Lastly, the integrated process noise model yields a final state covariance that is the most well aligned and properly scaled to the Monte Carlo distribution. Also of note is that the right-hand plot shows that all three process noise models behave essentially identically when the process noise is repeatedly applied using the adaptive time scale for this orbit and simulation period. However, while the frequent application allows the dynamics to shape the noise after application, there are several downsides beyond the additional computational demand, as will be shown with the NRHO and discussed subsequently.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

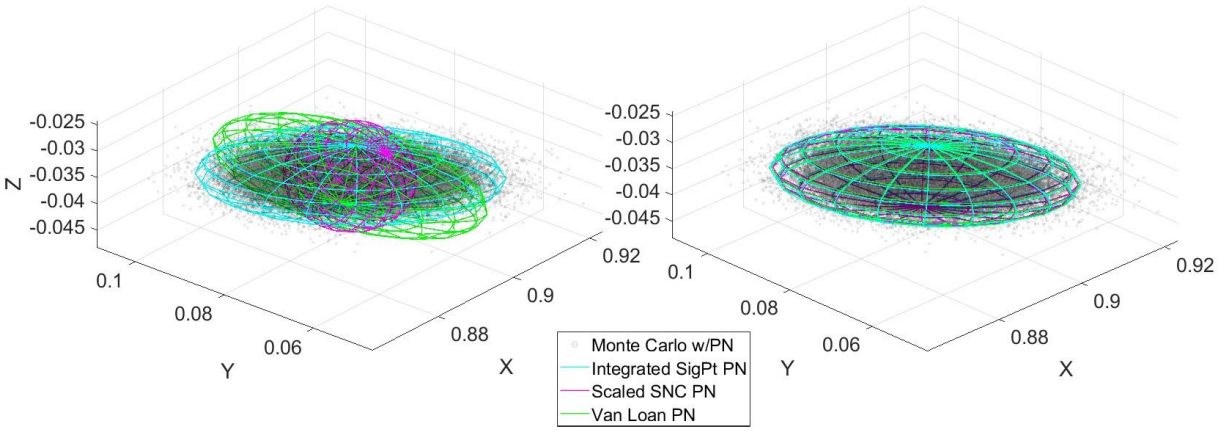

Fig. 7. Comparison of 3-σ covariance ellipsoids of the integrated, scaled, and Van Loan process noise models against Monte Carlo particle distribution over 5-days for HN1 orbit. The left-hand plot is zero-order hold (Δt = 5 days), and right-hand plot is adaptive time step based on Jacobian (Δt = 25-30 minutes).

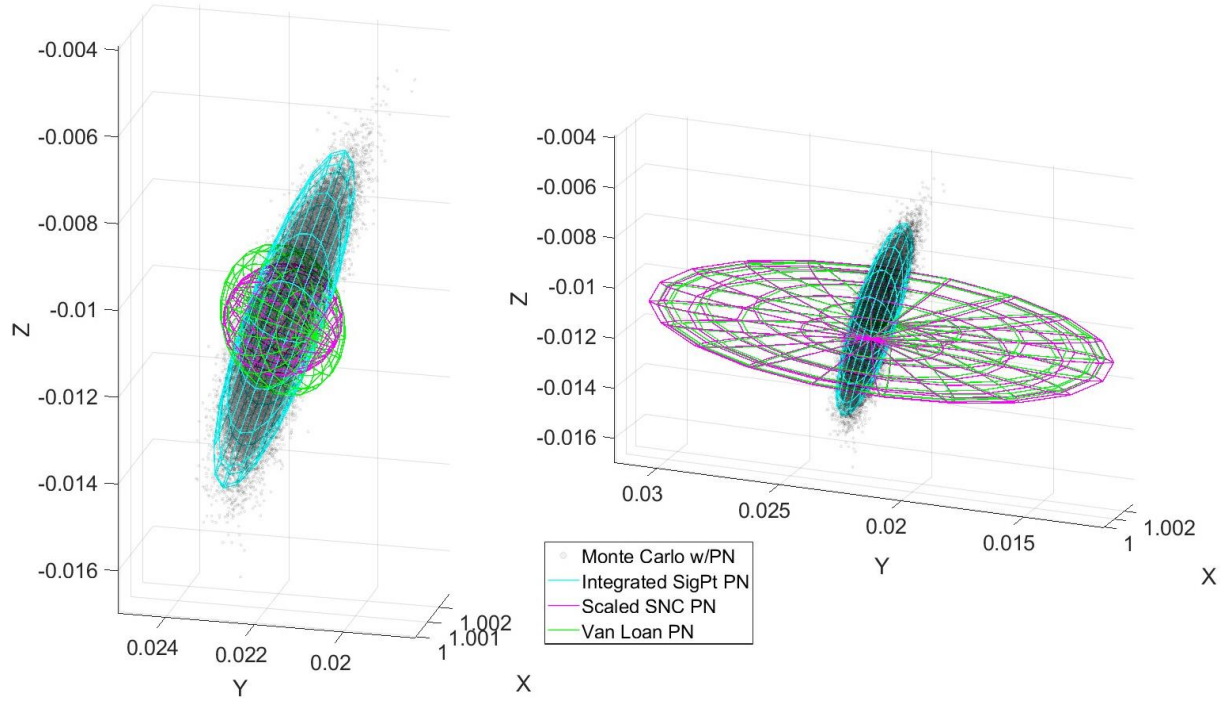

Fig. 8. Comparison of 3-σ covariance ellipsoids of the integrated, scaled, and Van Loan process noise models against Monte Carlo particle distribution over 3.5 days for NRHO. The left-hand plot is zero-order hold (Δt = 3.5 days), and right-hand plot is adaptive time step based on Jacobian (Δt = 18 seconds – 30 minutes).

The NRHO presents a different challenge. This orbit is marginally stable in the Lyapunov sense; but with its high eccentricity relative to the Moon, it experiences sharp nonlinearities near perilune. This yields strong non-Gaussianity near perilune, but a near-return to Gaussian behavior back out at apolune. Fig. 8 shows the 3-σ position covariance ellipsoids of the discrete process noise covariance matrices at the end of 3.5 days (+2.5 hours after perilune) for the NRHO. Again, the left-hand plot is the zero-order hold while the right-hand plot applies the adaptive time methods. The left-hand plot suffers the same deficiencies as with the HN1 orbit, but of particular note here is that unlike with the HN1 orbit, the adaptive time scale results (right-hand plot) for the scaled and Van Loan are now significantly worse than the integrated model (Δt = 30 minutes at start of simulation and 18 seconds (minimum) at perilune). This is due to the necessity of collapsing to Gaussian and resampling the sigma points after every process noise addition occurring in the highly nonlinear region near perilune. Not only does this destroy all measures of nonlinearity from the representation (which affects the overall filtering process), it introduces small

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

representation errors which are magnified by the dynamics at perilune; it thus rapidly becomes inferior in representing the actual PDF. The resulting distributions do not recover, even after the true state distribution returns to near-Gaussian behavior as the NRHO object approaches apolune. The integrated process noise model, however, only resamples the sigma points at the end time of the simulation period, maintaining a far better representation of the actual PDF.

It is possible to improve the scaled process noise model and Van Loan results on a case-by-case basis (i.e., for a particular orbit and for a particular ending time) by tuning the noise application cadence. Noise application more often than once (as with the zero-order hold extreme) improves the performance by subjecting noise accruals to the dynamics more frequently, allowing propagation of the incremental noise additions. Avoiding too-frequent application (as with the adaptive time scale extreme) mitigates the issue of discarding nonlinear information when resampling sigma points in very nonlinear regions of the orbit. Fig. 9 plots the Bhattacharyya coefficient of each process noise model's Gaussian position PDF to the Gaussian approximation of the Monte Carlo distribution (i.e., using the sample mean and covariance) as a function of time for the HN1 orbit (top) and NRHO (bottom) for three different sampling cadences: zero-order hold, the adaptive time scale, and Δt = 6 hours ("tuned" to provide good results for the scaled and Van Loan +2.5 hours after NRHO perilune or ~4 days since epoch). A Bhattacharyya coefficient of 1 indicates a perfect match. In the top plots for the HN1 orbit, the scaled process noise model for the zero-order hold cadence is the first to degrade (at ~2-days) followed by the Van Loan for the zero-order hold (at ~4-days), while the adaptive time scale for both and the integrated process noise model stay undegraded until ~6-days. In the bottom plots for the NRHO, the zero-order hold results for both the scaled process noise model and Van Loan model degrade significantly at perilune and recover somewhat afterward; the adaptive time scale version for both methods also degrades dramatically at perilune but now continue to degrade afterward; and the integrated process noise model again stays relatively undegraded until ~6-days. The hand-tuned 6-hr cadence plot does indeed provide better results than the zero-order hold and adaptive time scale for the NRHO just after perilune (~4-days). As is evident, the tuning of the sampling cadence for the two discrete methods varies considerably based on which orbit is considered and even the subset of interest within a given orbit. In stark contrast, the integrated process noise model completely avoids the vagaries of such tuning.

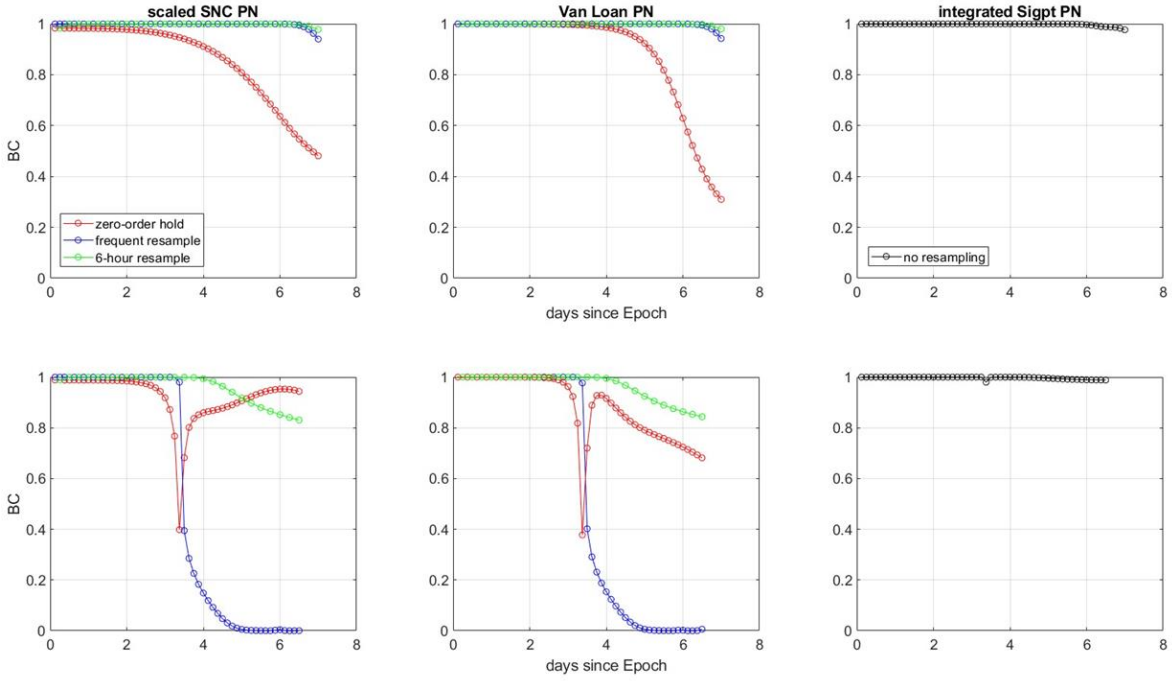

Fig. 9. Bhattacharyya distance for scaled (left), Van Loan (middle) and integrated (right) process noise models for zero-order hold (red), adaptive time scale (blue), and Δt = 6 hours (green) sampling cadences (only relevant for scaled and Van Loan) for the HN1 orbit (top) and NRHO (bottom)

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

4.3 Scaled vs. Integrated over observation gaps using a UKF

In the previous section, the initial covariance is nearly zero and the process noise magnitude is set to a very large value to emphasize the differing performance between process noise models. To demonstrate and compare these differences in actual estimation filter performance when using more realistic initial conditions and process noise magnitudes for the scaled process noise model and the integrated process noise model, observations of the ballistic NRHO were processed using several decimations of the Set 1 data with varying observation gap sizes. This was achieved by first processing only one randomly selected observation per day, and second only processing certain days to adjust the sizes of observation gaps. Observations were chosen to create gaps with sizes between ~5 days and ~30 days. The same dynamics mismatch and process noise magnitudes from Section 4.1 were used (significantly smaller than those from Section 4.2), and the scaled approach added noise prior to each measurement (essentially zero-order hold cadence of previous section). Additionally, only a single UKF component was processed throughout the filtering, and the initial uncertainties were 100 km in position and 10 m/s in velocity.

All estimation runs with both process noise models are successful in that all observations associate, and the object's track is well maintained. Fig. 10 plots the radial, in-track, and cross-track errors from truth along with the 1-σ uncertainty as a function of time for a sample integrated process noise model run (using the Set 3 "realistic thresholds" and one observation per day) showing small errors and reasonable uncertainty estimates for even very sparce observations.

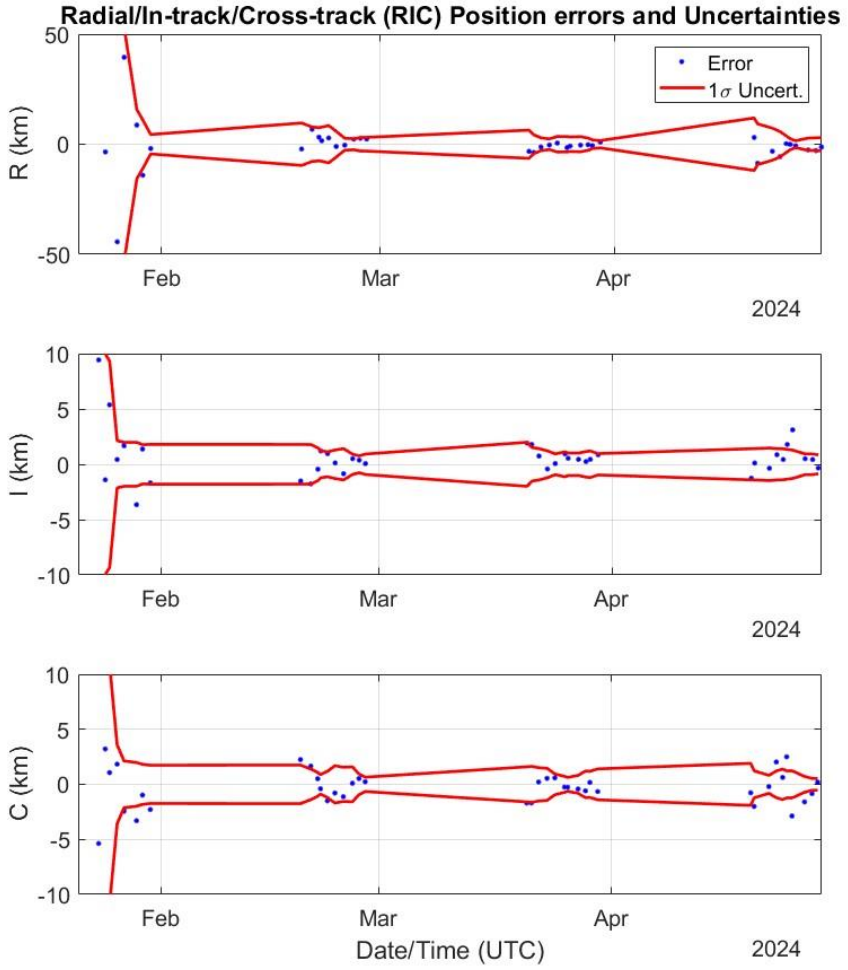

Fig. 10. Radial, in-track, and cross-track errors from truth and 1-σ estimated uncertainties as a function of time for sample filter run with 20 total observations and integrated process noise model

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

As a means of comparing the filter process utilizing the two process noise models quantitatively, the innovations of the first measurement after each observation gap are recorded and compared for runs using the same observations (as determined by adjusting which days are processed and the initial random number seed). Gaps ranged from ~5 days to ~30 days with 4 observations processed at the start and between every gap. Many factors affect the variations observed in these innovations: the state at the start of the observation gap as determined by all the observations/gaps before the gap, actual error and measurement uncertainty of the first observation after the gap, start and end points in the orbit during the gap, length of gap, and the process noise model itself. To provide a wide variety of possibilities, nineteen different runs are completed, varying which observations are used, with various gaps for both the scaled and integrated process noise models, for a total of 525 innovation comparisons. Fig. 11 is a box-and-whisker plot of these innovations for the scaled process noise model as a function of the length of the observation gap as well as the difference between the innovations for the scaled and integrated process noise models. (The central mark is the median value, the bottom and top edges of the box indicate the 25th and 75th percentiles, the whiskers extend to the most extreme data points not considering outliers, and the outliers are plotted individually using the '+' marker symbol).

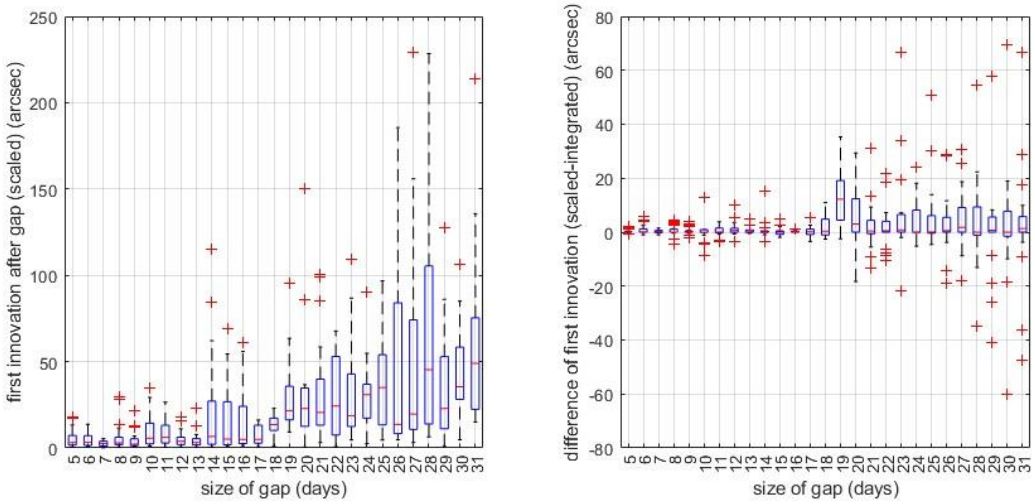

Fig. 11. Bar and whisker plot of post-gap innovations as the observation gap varies for the scaled process noise model (left – 6 outliers above 250" not shown) and the difference between the scaled and integrated process noise models (right – 2 outliers above 100" difference not shown)

As expected, as the size of the observation gap increases, the median and maximum size of the resulting innovation increases. For all gap sizes, both models perform about the same (near zero values in difference plot) with a slight advantage to the integrated process noise model (the median value and 25-75 percentile box with positive values in difference plot).

The 1-σ volume after each gap ($(4/3)\pi\sqrt{det(P_x)}$ where P_x is the 3×3 submatrix of the position portion of the pre-update state covariance for that first observation after a gap), can also be compared, and this shows a stark contrast between the 1-σ volumes. Fig. 12 are scatter plots of the scaled process noise model 1-σ volume (left) and the scaled/integrated 1-σ volume ratio (right) as a function of gap size. The phase of the gap within the orbit does influence the volume, which is indicated by red points when the object at the start of the gap is within 35000 km of the Moon's center. As expected, the 1-σ volume increases as the gap size increases. Most significant is the 1-σ volume using the scaled process noise model is nearly always greater than the 1-σ volume when using the integrated process noise model (ratio > 1 in right-hand plot). (There was one case where they were essentially equal.) Also evident is the increased volatility (i.e. significantly larger volumes) when the gap starts near perilune (red points).

Not surprisingly, when the 1-σ "innovation area" after each gap is examined (i.e., $4\pi\sqrt{det(S)}$, where S is the 2×2 innovation covariance for the first RA/DEC observation), the scaled process noise model's area is also greater than that of the integrated process noise model. This is shown in Fig. 13 which plots the scaled process noise model 1-σ

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

innovation area (left) and the scaled/integrated 1-σ innovation area ratio (right) with red points again indicating gaps when the object at the start of the gap is within 35000 km of the Moon's center.

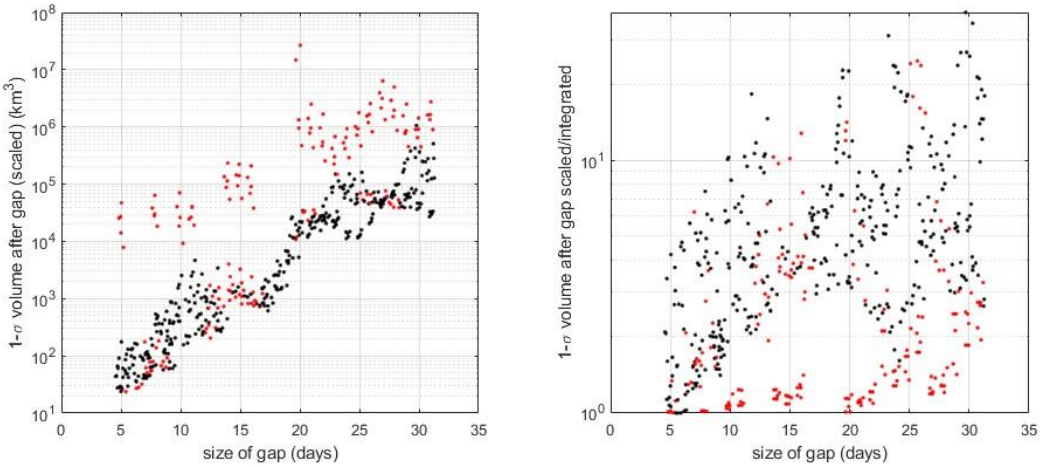

Fig. 12. Post-gap 1-σ volumes as the observation gap varies for the scaled process noise model (left) and the ratio of the scaled to the integrated process noise models (right). Red points indicate the object at the start of the gap is within 35000 km of the Moon's center.

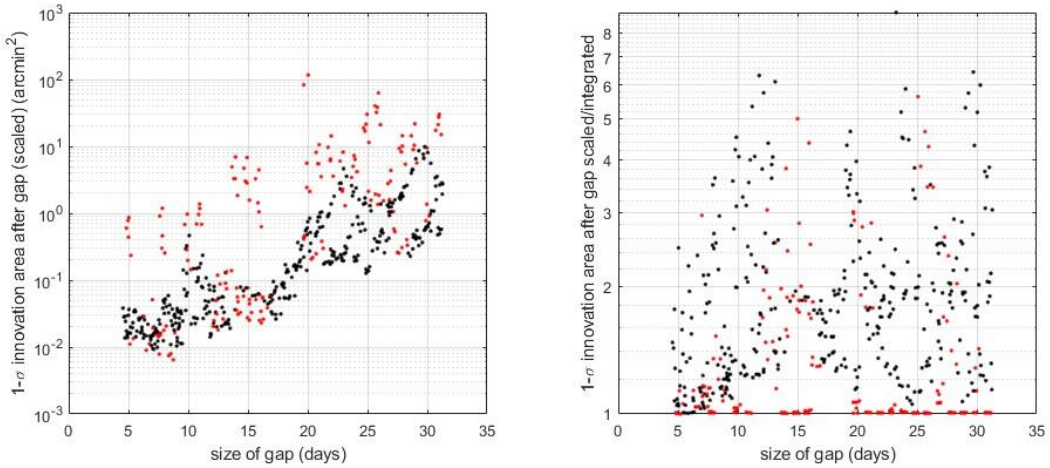

Fig. 13. Post-gap 1-σ innovation areas as the observation gap varies for the scaled process noise model (left) and the ratio of the scaled to the integrated process noise models (right). Red points indicate the object at the start of the gap is within 35000 km of the Moon's center.

Variations to the above analysis can be made. For example, a smaller or larger dynamics mismatch with a correspondingly smaller or larger process noise could be applied, or more observations could be processed per day to improve the state and state covariance going into each gap. Fig. 14 compares the general trends (ignoring outliers) of the 1-σ volume when using the scaled process noise model for a variety of cases: NRHO with original dynamics mismatch/process noise (shown above), NRHO with higher dynamics mismatch/process noise (4×10^{-10} km/s^2 for scaled and 2×10^{-8} km^2/s^3 for integrated), NRHO with no dynamics mismatch and lower process noise (5×10^{-13} km/s^2

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

for scaled and 2×10^{-11} km^2/s^3 for integrated), HN1 orbit with original dynamics mismatch/process noise and 1, 3, and 5 observations processed per day. The trends observed are as expected:

- The more Lyapunov-unstable HN1 orbit causes a larger "slope" in this semi-log plot (corresponding to a greater exponential growth in the volume) than NRHO
- More observations per day yield better state/covariance going into gaps and correspondingly lower volumes (HN1 results compared to each other)
- Larger dynamics mismatch (and correspondingly larger process noise) leads to larger volumes (NRHO results compared to each other).

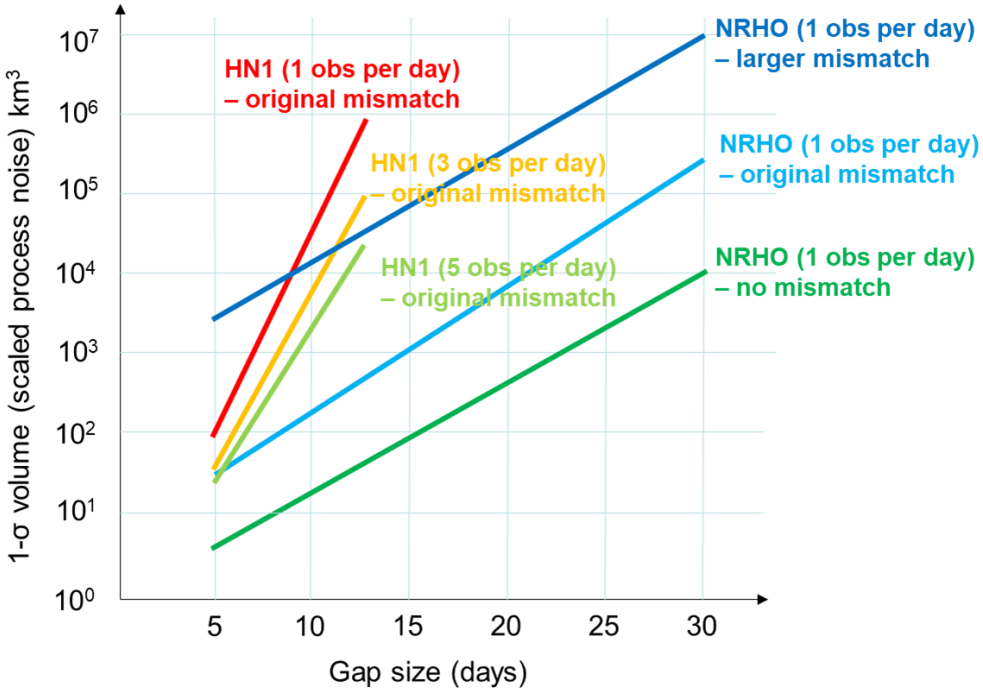

Fig. 14. General trends of post-gap 1-σ volumes as the observation gap varies for the scaled process noise model for various instances using the HN1 and NRHO ballistic orbits

The exact same processing and analysis can also be done allowing splitting into multiple Gaussian components via the AEGIS filter. This is achieved by monitoring the differential entropy for each component and splitting when this value exceeds a preset threshold. It was noticed, however, that allowing splitting during the most nonlinear portion of this NRHO (at perilune) greatly worsened the overall performance of the filter because of the extreme nonlinearity, which motivated the Monte Carlo study discussed in Section 4.2. Optimizing the splitting and merging of Gaussian mixture components is the subject of our companion paper [14] and combining with the integrated process noise modeling will be saved for future work. Also saved for future work is exploring the use of the integrated process noise model in an IMM for maneuver detection.

5. DISCUSSION

The results above demonstrate that the integrated process noise model provides an evolution in the noise-driven state that is, in most cases, superior to that achieved with the simpler methods commonly in use, and, at worst, performs similarly to those methods when they are specially tuned for a given orbit/trajectory. While this is an encouraging result, the greater advantage of this method lies in the ability of a filter component (here, a set of sigma points) to stay in its native representation for longer, avoiding the sometimes-catastrophic collapse to an approximate Gaussian representation.

Staying in the native [sigma-point] representation opens the door to more robust measures of non-Gaussianity, as (depending on the choice of the unscented transform's sample points) some higher-order information, like skewness, is available. Therefore, we see this method as an enabler for robust splitting criteria for Gaussian Mixture Model (GMM)-based filters (e.g. AEGIS), as each component will contain more pristine nonlinearity information.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

The integrated process noise model does come at some cost in terms of computational effort and algorithmic complexity. The requirement to solve a system of equations at each evaluation of the ODEs, as in Equation (32), clearly imposes a nominal computational cost. (It should be noted, however, that this cost is likely negligible when compared to other methods, like the Van Loan model, which require integration of a state transition matrix and, therefore, frequent evaluations of the dynamics Jacobian.) Perhaps more impactful is the algorithmic/implementation complexity cost. Specifically, state propagation is no longer performed separately for each sigma point but instead must be accomplished for an entire distribution (or component, in the case of a GMM) at once. This requires propagation routines which are able to handle this level of abstraction.

Finally, as mentioned above, the integrated process noise model involves constraints on the mean and covariance rates of the sigma point group; computation of these constraints requires computing the sample mean and covariance. These sample moments become less accurate as the true distribution becomes less Gaussian (e.g., more skewed), potentially causing "drift" in the propagated states over time, so this method is best for distributions which are still nearly Gaussian. This issue is, of course, obviated when using a GMM in which a component is split before becoming too strongly non-Gaussian. Additionally, in the case of zero-mean Gaussian continuous noise, an additional constraint on the third-order central moment (skewness) should ideally be imposed; however, this imposes additional computation burden and, in practice, we have not noticed measurable skewness induced by the constraints as stated above.

6. SUMMARY AND FUTURE WORK

We developed an integrated process noise model enforcing the noise power constraint onto the sigma point distribution itself without discretization or resampling and showed it performs equally well and often better than standard scaled discrete process noise models when used in an estimation filter. This was demonstrated with a single component UKF. The use of the integrated process noise model also enables a more direct and well-behaved measure of non-Gaussianity which can be exploited in the splitting and merging of Gaussian components, as in the AEGIS or other GMM filters. Future work will involve developing and using the integrated process noise model and skewness measure in combination with the merging and splitting techniques developed in our companion paper, as well as exploring its use in an IMM for maneuver detection. Additionally, using the integrated process noise model to account for a known dynamics mismatch, as done in this paper, will be compared to employing consider parameters for the appropriate dynamics parameters.

7. REFERENCES

[1] C. Chow, C. Wetterer, K. Hill, C. Gilbert, D. Buehler and J. Frith, "Cislunar Periodic Orbit Families and Expected Observational Features.," in *The Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, Maui, HI, September 15-18*, 2020.

[2] J. Gaebler, K. Hill and C. Wetterer, "Cislunar Initial Orbit Determination with Optical Tracklets," in *The Advanced Maui Optical and Space Survellance Technologies (AMOS) Conference, Maui, HI, September 19-22*, 2023.

[3] P. Billings, J. Baldwin, C. Wetterer, J. Gaebler, C. Craft, M. Dilley, K. Hill and J. Bruer, "Cislunar orbit determination using CAR-MHF," in *The Advanced Maui Optical and Space Surveilance Technologies (AMOS) Conference, Maui, HI, September 19-22*, 2023.

[4] C. Chow, C. 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 *The Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, Maui, HI, September 14-17*, 2021.

[5] C. Chow, J. Baldwin, C. Wetterer, M. Dilley, K. Hill, P. Billings, C. Craft and J. Frith, "Cislunar Orbit Determination: Improvements in Uncertainty Realism and Data Fusion," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, Maui, HI, September 27-30*, 2022.

[6] C. Wetterer, C. Craft, J. Baldwin, M. Dilley, P. Billings, J. Gaebler and J. Bruer, "Cislunar Debris from Halo Object Breakups," in *The Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, Maui, HI, September 19-22*, 2023.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com

[7] C. Wetterer, C. Chow, J. Baldwin, M. Dilley, K. Hill, H. Montell-Weiland, P. Billings, C. Craft and J. Frith, "Cislunar Maneuver Detection and Classification," in *Advanced Maui Optical and Space Surveillance Technologies (AMOS) Conference, Maui, HI, September 27-30*, 2022.

[8] Y. Bar-Shalom, P. Willett and X. Tian, Tracking and data fusion, Storrs, CT, USA: YBS publishing, 2011.

[9] C. Van Loan, "Computing integrals involving the matrix exponential," *IEEE Transactions on Automatic Control,* vol. 23, no. 3, pp. 395-404, June 1978.

[10] S. Julier and J. Uhlmann, "A new extnsion of the Kalman filter to nonlinear systems," in *Proc. AeroSense: 11th Int. Symp. Aerospace/Defense Sensing, Simulation and Controls*, 1997.

[11] J. Crassidis and F. Markley, "Unscented Filtering for Spacecraft Attitude Estimation," in *AIAA 2003-5484, AAIA Guidance, Navigation, and Control Conference*, Austin, Texas, Aug. 11-14, 2003.

[12] K. DeMars, J. Bishop and M. Jah, "An Entropy-based Approach for Inference of Uncertain Dynamical Systems," *Journal of Guidance, Control and Dynamics,* vol. 36, no. 4, 2013.

[13] I. Sobol, *USSR Comput. Maths. Math. Phys.,* vol. 7, pp. 86-112, 1967.

[14] J. Gaebler, J. Gutierrez, P. Billings, C. Craft, C. Wetterer, J. Baldwin, M. Dilley and J. Bruer, "Application of Uncertainty Propagation with Adaptive Gaussian Mixture Models for Cislunar Orbits," in *The Advanced Maui Optical and Space Survellance Technologies (AMOS) Conference, Maui, HI, September 16-19*, 2025.

Copyright © 2025 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS) – www.amostech.com