Reconciling Covariances with Reliable Orbital Uncertainty

Zachary Folcik
MIT Lincoln Laboratory
Arthur Lue
MIT Lincoln Laboratory
Joshua Vatsky
MIT Lincoln Laboratory

ABSTRACT

There is a common suspicion that formal covariances do not represent a realistic measure of orbital uncertainties. By devising metrics for measuring the representations of orbit error, we assess under what circumstances such lore is justified as well as the root cause of the discrepancy between the mathematics of orbital uncertainty and its practical implementation. We offer a scheme by which formal covariances may be adapted to be an accurate measure of orbital uncertainties and show how that adaptation performs against both simulated and real space-object data. We also apply these covariance adaptation methods to the process of observation association using many simulated and real data test cases. We demonstrate that covariance-informed observation association can be reliable, even in the case when only two tracks are available. Satellite breakup and collision event catalog maintenance could benefit from the automation made possible with these association methods.

1. INTRODUCTION

Assessing the quality of orbital states estimated from observational data is a bedrock goal of astrodynamics. Tasks as fundamental to space situational awareness as resident space-object catalog maintenance and refinement, initial orbit determination (IOD), object correlation, space event detection and conjunction analysis all depend on a quantitative understanding of how well observations of a particular space object translate into an accurate picture of where that object is and where it will be in the future.

Conceptually, orbital uncertainty is straightforward. Measurements of a space object are made by a particular sensor with an understood confidence. The measurements are then translated into a state by means of a best-fit to a model of orbital dynamics. While the resulting state is the best fit to the measurements, it is not the only fit consistent with them. The "orbital uncertainty" is precisely the range of states that are consistent with the given measurements.

While the concept of orbital uncertainty is straightforward, the mathematical representation of orbital uncertainty in a compact and computationally efficient way is not, unless simplifying assumptions are made. Formal covariances are derived from such a set of simplifying assumptions. If the uncertainties are sufficiently small and if the uncertainty statistics are sufficiently gaussian, linear error theory is applicable and the state-uncertainty region may be represented by ellipsoidal volumes centered on the estimated state; the surfaces of these ellipsoidal volumes are constant-valued surfaces of the underlying multi-dimensional gaussian distributions. Covariances contain the parameters that define those multi-dimensional gaussian distributions and the corresponding ellipsoidal volumes.

As a practical matter, the computational simplification that a covariance represents rarely works under conditions of interest. But why is this the case? What is the origin of the discrepancy? Let us review some of the important insights that have already been made in the attempt to understand the problem.

2. THE EXISTING LITERATURE

Junkins et. al. [1] studied the limitations of applying linear error theory to orbital state propagation, describing how nonlinearities in the equations of motion of the chosen set of state variables affected error propagation. Monte-Carlo simulations were applied to compare the accuracy of the covariance to the true uncertainty for three sets of state variables: osculating position/velocity in cartesian space, osculating variables in polar coordinates and mean keplerian orbital elements. The results indicated that the covariance associated with a set of mean keplerian elements produced the best agreement with the Monte-Carlo point cloud. Hill et. al. [2] followed up the Junkins analysis by using a set of state variables in rectilinear coordinates.

Sabol et. al. [3] compared covariance and Monte-Carlo simulated point clouds for state variables in osculating position and velocity, osculating equinoctial elements and mean equinoctial elements. A realistic initial covariance and set of Monte-Carlo runs was initiated using simulated radar tracks. The analysis showed that the mean equinoctial elements provided the best covariance agreement with Monte-Carlo simulations, while the osculating equinoctial and osculating position/velocity coordinates providing the 2nd best and 3rd best agreement, respectively.

The results from these studies suggest there are two distinct reasons for errors in propagating the covariance. There are those due to imposing a rectangular coordinate system onto what is inherently a nonlinear trajectory, i.e. geometry induced errors, and those due to applying linear error theory to nonlinear equations of motion. Sabol, et. al. surmised that the improved performance of the osculating equinoctial variables over the osculating rectangular coordinates was due to geometry, and that the improved performance of

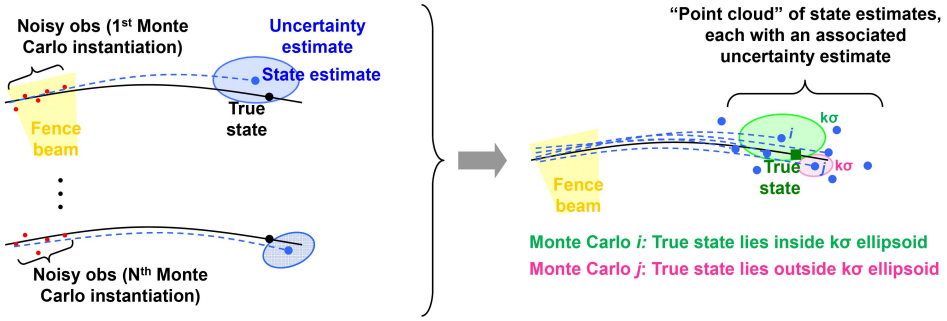

Figure 1. Schematic depiction of the algorithm behind the performance metric.

the mean equinoctial elements over the osculating equinoctial elements was due to the mean elements being more linear in time than the osculating set.

3. ROADMAP

This paper presents work to produce covariance and Monte Carlo simulation results that confirm findings in Junkins [1], and Sabol [3]. The covariance associated with a set of mean equinoctial elements outperforms the covariance of the osculating position and velocity vectors and the covariance of osculating equinoctial elements. In addition, a method to improve the accuracy of the initial and propagated mean equinoctial covariance is discussed. This method is shown to be effective both in Monte-Carlo simulations and with real observational data on satellites. A prototype observation-association and catalog-initialization/maintenance system which makes use of the accurate mean equinoctial covariance is also described. The results from the prototype indicate that accurate covariance can be effective even when it is based on short tracks of radar observations.

Several aspects of this study required high precision orbit determination and propagation tools. To satisfy this requirement we made use of the Goddard Trajectory Determination System (GTDS) [4]. Two of the propagation theories in GTDS were exercised, the Cowell propagator and the Draper Semianalytic Satellite Theory (DSST) propagator [5]. The perturbations modeled in GTDS for satellite propagation included 30×30 Earth gravity coefficients from the JGM–2 gravity model, atmospheric drag, lunar and solar point-mass gravity and solar radiation pressure. It should be noted that although drag was included in the modeling, the uncertainty in knowledge of the drag coefficient and uncertainty in the knowledge of atmospheric density were not addressed.

4. PERFORMANCE OF ORBIT-UNCERTAINTY MEASURES

For an uncertainty region to be considered meaningful, it should be the case that the true state lies somewhere within it. In practice, the uncertainty region is generally some form of P% confidence region. That being the case, we would expect to find the true state within the uncertainty region P% of the time; this directly points the way to our method of assessing uncertainty measures, which is to numerically estimate the probability that the true state is found within the predicted uncertainty region.

Given a sensor, whose configuration is characterized by a viewing geometry and noise characteristics, and an orbit, we generate simulated observations by adding zero-mean, gaussian measurement noise to true orbital positions. We then fit an orbit to the simulated observations, producing both a state and an estimate of the uncertainty in that state (the formal covariance, for example). Next, we propagate the state and uncertainty estimates forward in time. This process is repeated many times, with different instantiations of the measurement noise each time. Finally, we propagate the true state forward in time. Figure 1 depicts this methodology. Given this data, we can, at each time step, answer the question "What is the fraction of cases for which the true state lies within the predicted uncertainty region?" The cumulative measure of what fraction of these runs have truth that lies within the region of estimated state uncertainty at each time step is a measure of the quality of the orbit-uncertainty estimation scheme. This, therefore, lets us assess the quality of the uncertainty measure as a function of propagation time.

Covariances are matrix quantities that represent the statistics of uncertainty in the state space of an estimated state. If the uncertainties are sufficiently small and if the uncertainty statistics are sufficiently gaussian, the state-uncertainty region may be represented by ellipsoidal volumes centered on the estimated states. Covariances represent the parameters needed to fix those ellipsoidal volumes.

Covariances are thus correlated to uncertainty statistics via gaussian statistics. Specifically, if one constructs a series of nested ellipsoids of differing number of σ's, then, for an N–dimensional gaussian distribution, the expected fraction within each ellipsoid can be calculated and is presented in the table in Fig. 2. So, one may ask whether the truth lies within ellipsoidal volumes of differing number of sigmas, i.e., for different k–values where one is drawing ellipsoids of size $k\sigma$.

So, we can now use the performance metric to follow the quality of the uncertainty measure, but also to validate whether gaussian statistics is adhered to by using the values found in Fig. 2: for every k, we can

Probabilities of Lying Within the k σ Error Ellipsoid*				
N	1σ	2σ	3σ	4σ
1	0.683	0.954	0.997	0.9999
2	0.393	0.865	0.989	0.9997
3	0.199	0.739	0.971	0.9987
4	0.090	0.594	0.939	0.9970
5	0.037	0.451	0.891	0.9932
6	0.014	0.323	0.826	0.9862

* See, for example, "Modern Orbit Determination" by W. Wiesel

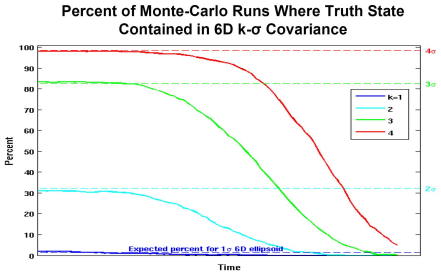

Figure 2. Left table shows containment probabilities of k–sigma shells given an N-dimensional gaussian distribution [6]. For the problem under consideration, state space is six-dimensional so we are particularly interested in the case where $N = 6$. The right graphic shows the containment probabilities over time for a notional case in which containment agrees with Gaussian statistics initially, but gradually degrades.

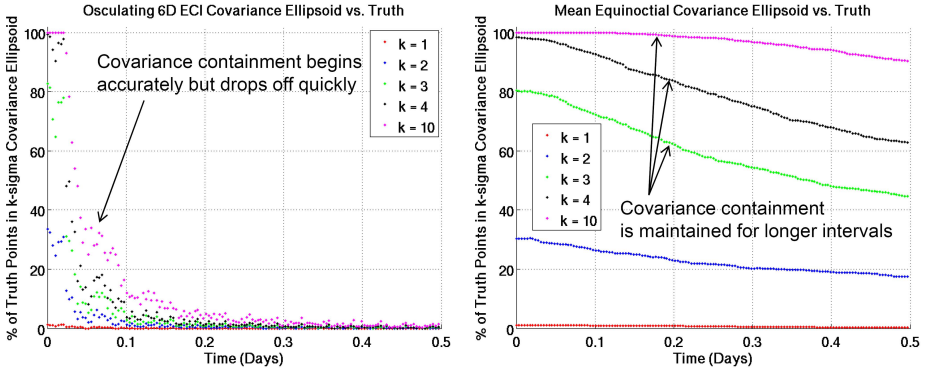

Figure 3. Osculating ECI vs. mean equinoctial 6D covariance containment; 10-20 degree rising pass scenario. Covariance containment is statistical agreement between the actual uncertainty cloud and the covariance The (k=4) covariance ellipsoids should contain truth about 98% of the time We confirm what Junkins, et al. [1] found: covariance of mean elements performs better than covariance of osculating ECI vectors.

count occurences and calculate the fraction of the Monte Carlo runs in which the true state is contained within the kσ error ellipsoid surrounding each individual estimated state. When calculating whether the truth state, $\mathbf{x}_{true}$, lies inside the k-sigma shell of the estimate, $\mathbf{x}_{est}$, and its covariance, C_{est}, the following k-distance, i.e. Mahalanobis distance, calculation is made:

$$k = \sqrt{(\mathbf{x}_{est} - \mathbf{x}_{true}) C_{est}^{-1} (\mathbf{x}_{est} - \mathbf{x}_{true})} \tag{1}$$

Let us examine examples of state variables and the associated determination and propagation of covariance in each. Although many orbits and tracking schedules were tested using the containment metric, Figures 3 and 4 show two examples of the containment results when using osculating ECI, osculating equinoctial and mean equinoctial coordinates. The states and covariance were determined from a single track of observations, and for these examples, a 474 km altitude circular orbit at 97.3 degrees inclination was used. The state and covariance were built and evolved according to high precision models in all of these coordinates. We see that while, initially, gaussian statistics is respected (i.e., the percentages found in Fig. 2 are followed fairly faithfully at the start of the evolution), covariance containment is extremely poor for osculating ECI coordinates; it fails almost immediately after one orbit. Containment is improved for osculating equinoctial coordinates and is even better still for mean equinoctial coordinates. However, in both cases, containment is not particularly satisfactory, except possibly in the most favorable of observation conditions and then only for a short period of time. Covariances, when produced using a short, single track of radar observations, form a poor measure of actual orbital uncertainties when measured for several days.

5. STATE SPACE AND DYNAMICS

In order to understand the structural flaws of covariance and how it fails to accurately represent orbital uncertainties, we need to recall the formalism underlying both orbit dynamics and the computation of covariances based on those dynamics. Each satellite orbit can be represented by a state (its kinematic properties such as position and velocity, as well as physical properties such as mass and cross-sectional area, etc.) that allows us to identify the object as well as predict its future evolution.[1] Then, the state space is an

[1] Alternatively, the kinematic properties may be elements of an element set (e.g., inclination, eccentricity, mean anomaly, etc.). Abstractly, exchanging positions and velocities for an element set is simply a one-to-one reparametrization of state space, and represents no fundamental difference. However, due to the nonlinearity of such a change in parameters, attention must be paid to the best choice of state-space components. A scenario that is properly gaussian in one representation may not be

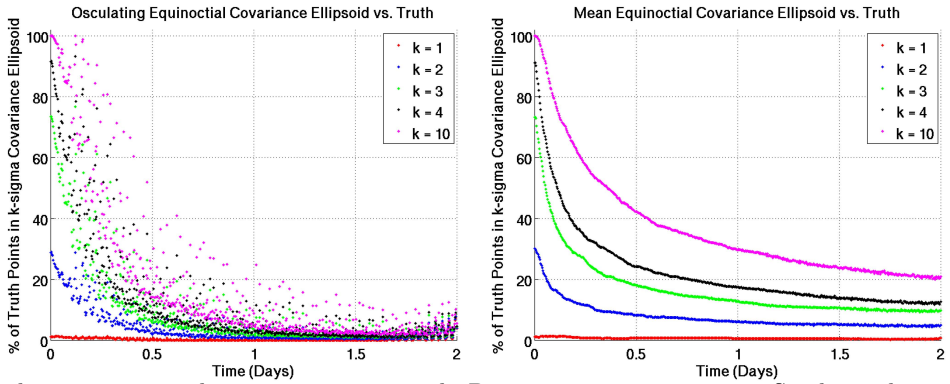

Figure 4. Osculating equinoctial vs. mean equinoctial 6D covariance containment. Single track, 10 degree zenith tracking pass of a 474 km altitude circular orbit at 97° inclination. The mean equinoctial element covariance performs better than osculating equinoctial covariance in containing the truth state. The osculating equinoctial containment also often includes oscillations (4 each rev in this case).

n–dimensional vector space, $\{\mathbf{x}\}$, where each satellite orbit in our scenario may be represented at any time, t, by a vector $\mathbf{x}(t)$ such that the components of that vector are given by its kinematic properties. Then, the state space will be filled with a cloud of points, each representing a different satellite orbit, where each point follows a specific trajectory through state space as time marches forward. An orbiting satellite located at $\mathbf{x}$ at time t is dictated by

$$\dot{\mathbf{x}} = \mathbf{F}(t, \mathbf{x}) \ . \tag{2}$$

where $\mathbf{F}$ is a vector-valued function given by familiar evolution equation, such as Newton's laws of gravitation, etc. So a trajectory in n-dimensional phase space has a one-to-one correspondence with each orbit or trajectory of a satellite in real space.

Each set of observations is processed into a likelihood distribution on the state space by some well-understood estimation scheme (e.g., least squares or Kalman filtering or even some IOD scheme). This likelihood distribution represents the probability distribution, Prob($\mathbf{x}$), that the observations taken correspond to a specific state $\mathbf{x}$ at a particular time. *Assuming* small uncertainties and well-behaved, linear statistics, that likelihood distribution may be treated as gaussian and can be represented as being located around a state $\mathbf{x}_i$ and characterized by a covariance $\mathbf{C}_i$ (associated with the i–th set of observations) such that

$$\mathrm{Prob_i}(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n \det \mathbf{C}_i}} \int d^n\mathbf{x} \ \exp\left[-\frac{1}{2}(\mathbf{x}-\mathbf{x}_i)^T \mathbf{C}_i^{-1}(\mathbf{x}-\mathbf{x}_i)\right] \ . \tag{3}$$

We may now define confidence ellipsoids, i.e., surfaces of constant $\mathrm{Prob}_i(\mathbf{x})$ exhibiting volumes of state space containing, for example, 90%, 95%, or 99% of the probability. These ellipsoids have sizes and shapes associated with the components of $\mathbf{C}_i$.

A crucial component of a correlation algorithm is the evolution of the likelihood distribution, or in this case, the evolution of the covariance, $\mathbf{C}_i(t_i)$, associated with a set of observations taken near t_i, to a covariance at a future time, $\mathbf{C}_i(t)$. This evolution is dictated by variations of $\mathbf{F}(t, \mathbf{x})$ with respect to $\mathbf{x}$; moreover, if there is an uncertainty in $\mathbf{F}$, that uncertainty acts as a process noise. Assuming a gaussian process noise source with zero mean and covariance $\mathbf{Q}(t)$, the evolution of the covariance $\mathbf{C}$ is dictated by

$$\dot{\mathbf{C}} = \frac{\partial \mathbf{F}}{\partial \mathbf{x}}\mathbf{C} + \mathbf{C}\left(\frac{\partial \mathbf{F}}{\partial \mathbf{x}}\right)^T + \mathbf{Q} \ . \tag{4}$$

Here is a circumstance for which an inspired choice of state-space coordinates will be crucial. Gaussian process noise for one set of $\mathbf{x}$ may be extremely nongaussian if the state space is reparametrized by another set of coordinates $\mathbf{y}$.

Perturbations around circular orbits are of particular interest for analyzing orbit error dynamics.[2] For two-body dynamics, consider the reference orbit governed by the equations

$$r(t) = r_0 \ , \qquad \theta = 0, \qquad \phi(t) = \omega_0 t \ , \tag{5}$$

in standard polar coordinates for a given r_0 and ω_0. The reference angular frequency follows the expression $\omega_0 = \sqrt{GM/r_0^3}$. Then, second-order element set perturbations around this circular orbit are governed by

so in another representation.

[2] Moreover, for orbits that are not circular, the estimates of perturbations will still be well-approximated by the circular-orbit relationships to an order characterized by the orbit's eccentricity. I.e., the equations governing the perturbations will only be wrong to $\mathcal{O}(e)$, where e is the (small) eccentricity of the baseline truth orbit.

eight parameters: Δa, the semi-major axis perturbation; Δe_1 and Δe_2, the components of the eccentricity-vector perturbation; Δi_1 and Δi_2, the components of the inclination-vector perturbation; and $\Delta\phi_0$, the initial true longitude perturbation.

For IOD, and for short arcs of tracking in general, there is a hierarchy of perturbation scales. Generally, the initial velocity uncertainties, while small, are much larger than spatial position uncertainties. This requires the careful inclusion of certain second-order perturbations, particularly those second-order in velocity uncertainties.

After carefully establishing that to second-order, the eccentricity parameters are not perturbed by forces normal to the orbit plane, the in-plane circular orbit equations to second-order are as follows:

$$\begin{aligned} r(t) &= a\left[\left(1-\frac{\Delta e_1^2+\Delta e_2^2}{2}\right)-(\Delta e_1'\cos\omega t+\Delta e_2'\sin\omega t)\right. \\ &\quad \left.-\left(\frac{1}{2}(\Delta e_1^2-\Delta e_2^2)\cos 2\omega t+\Delta e_1'\Delta e_2\sin 2\omega t\right)\right] \qquad (6) \\ \phi(t) &= \Delta\phi_0+\omega t+2(\Delta e_1\sin\omega t-\Delta e_2(\cos\omega t-1)) \\ &\quad +\frac{5}{2}\left(\frac{1}{2}(\Delta e_1^2-\Delta e_2^2)\sin 2\omega t-\Delta e_1\Delta e_2\cos 2\omega t\right)\ , \qquad (7) \end{aligned}$$

with

$$\omega=\omega_0\left(\frac{r_0}{a}\right)^{3/2}\ \text{ and }\ \omega t=\omega_0 t\left(1-\frac{3}{2}\frac{\Delta a}{r_0}+\cdots\right)\text{ being the most relevant contribution.} \qquad (8)$$

Here, the secular growth in the anomaly is evident as the dominant long-term behavior. Using the definition of the element-set parameters, to second order we may establish the perturbed orbit parameter, $\Delta\omega$, in terms of initial position and velocities:

$$\Delta\omega = \omega_0\left[\left(1+\frac{\Delta a}{r_0}\right)^{-3/2}-1\right] \qquad (9)$$

Because of the simple nature of orbit-element evolution, one may reduce the complex dynamics of orbit-confidence evolution to that of a single position/velocity degree-of-freedom, $\{\Delta\phi,\Delta\omega\}$, governed by Eqs. (7) and (9). We see that long-term (i.e., longer than one orbit period) uncertainty dynamics in orbits are almost completely determined by along-track evolution as described by Eq. (7) to second-order. Focusing on the secular components of evolution, i.e., ignoring the oscillatory components, we see that one simply has

$$\begin{aligned} \Delta\phi(t) &= \Delta\phi_0+\Delta\omega t \qquad (10) \\ \Delta\omega(t) &= \text{constant}\ . \qquad (11) \end{aligned}$$

Let us recharacterize the variables as a single position/velocity degree of freedom such that $x=\Delta\phi$ and $v=\Delta\omega$ and the time variable $\tau=\omega_0 t$ is time-normalized using the period of the unperturbed orbit. If one wishes to assess the evolution of the along-track covariance, one may define the following relations and use simple evolution relations for ϕ and $\dot\phi$ to subsequently write expressions for evolving an initial covariance in time:

$$\begin{aligned} C_{xx} &= \langle xx\rangle-\langle x\rangle^2\rightarrow C_{xx}(\tau)=C_{xx}|_0+2\tau C_{xv}|_0+\tau^2 C_{vv}|_0 \qquad (12) \\ C_{vv} &= \langle vv\rangle-\langle v\rangle^2\rightarrow C_{vv}(\tau)=C_{vv}|_0 \qquad (13) \\ C_{xv} &= \langle xv\rangle-\langle x\rangle\langle v\rangle\ \rightarrow C_{xv}(\tau)=C_{xv}|_0+\tau C_{vv}|_0. \qquad (14) \end{aligned}$$

6. THE ROLE OF NONGAUSSIANITIES

Now, we may start to address why containment fails. Simple along-track dynamics are sufficient until one needs to consider second-order contributions to evolution. This is where the underlying assumption of linearity girding covariances breaks down. The source of covariance failure is non-gaussianities. These nongaussianities appear because of irreducible nonlinearities in the evolution equations of the orbital state. The assumption of linearity is a generally good assumption so long as errors are small. However, as errors grow large, the straight ellipsoids that covariances represent do not well-characterize the curved banana-like shapes that the true orbital uncertainties follow in Earth centered inertial (ECI) rectangular coordinates. We

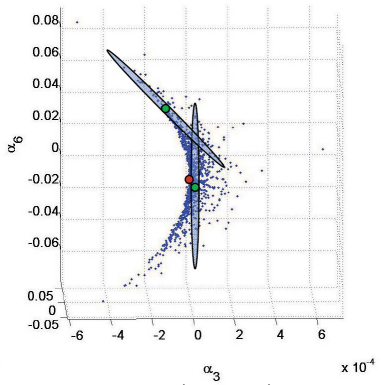

Figure 5. Type 3 nongaussianity. For a given truth (red dot), one may establish an uncertainty volume (blue points) given initial observations. State estimates (green dots) lie in the uncertainty volume, but their corresponding covariance estimates are straight, thin ellipsoids that follow the shape of the blue points. If the true uncertainty volume were flat, the estimated covariances would contain the truth; however, because it is curved, the truth state (red) will not fall inside the estimated covariances most of the time

see three places where important nonlinearities arise: (1) ECI vs. Element Set coordinates, (2) osculating versus mean, and (3) IOD position vs. velocity.

Type 1 nongaussianities are the most straightforward to understand. Ellipsoids in ECI are initially aligned with the linear geometry, however, as uncertainty grows with time, we know the true state will follow the curved orbit. That discrepancy represents a nonlinearity that may be addressed by taking element set coordinates rather than raw position-velocity ones. This type of nongaussianity is discussed in Junkins, et al. [1].

Type 2 nongaussianities involve per-orbit oscillations that introduce nonlinearities on the scale of the orbit period. Mean element sets average out those contributions making the evolution more monotonic and secular, those being more conducive to linearizations that are valid over long periods of time. Linear error theory as shown by equation 4, after all, assumes linearity.

The third type of nongaussianity that appears in short-track IOD is a subtle one. When long-track (e.g. horizon-to-horizon) IODs are made, this nongaussianity disappears and using mean-equinoctial covariances is valid and successful. However, when IOD tracks are short (e.g. from a radar fence sensor), initial position uncertainties (dx) are much smaller than initial velocity uncertainties (dv) in a quantifiable fashion. The covariances appear more pancake-shaped. Nonlinearities become important as dx is smaller than dv^2 and one cannot ignore dv^2 type terms in the evolution equations. What look like pancake shapes get turned into saddle shapes, and as one can see from Fig. 5, covariance does not cover the uncertainty volume well.

7. COVARIANCE INFLATION

How do we solve these issues and get covariances that work, i.e., that contain properly the evolving uncertainty of an initially determined orbit state? The first two types of nongaussianities can be cured by careful choice of orbital state coordinates (Type 1: use elset coordinates versus ECI coordinates, Type 2: use mean element sets versus osculating element sets). To resolve the final problem we take a different action. As we said, the problem is in containing the saddle shape with a pancake. The solution is to fatten the pancake in a strategic way. The scheme is to inflate (in effect) initial position uncertainty to the scale of the initial velocity uncertainty (or one can be more aggressive and rather than inflate $dx \rightarrow dv$, choose $dx \rightarrow 10 * dv^2$). With a proper choice of initial inflation of the covariance, the governing *spatial* size of subsequent covariance volume is not much larger than before, but still contains the entire saddle-shaped uncertainty, even as time evolves. Figures 6 and 7 schematically depict the scheme.

More precisely the following scheme is to be followed:

1. Observation with given noise statistics are taken
2. Initial state is estimated from observations as well as corresponding initial covariance
3. Eigenstates, $\{\alpha_i\}$ and eigenvalues $\{\lambda_i\}$ of covariance are identified in natural units[3]
4. Eigenstates whose eigenvalues are $< 10\lambda_{\max}^2$ have their eigenvalues re-scaled up to the value $\lambda_{\max}$ (or to $10\lambda_{\max}^2$)

[3] Natural units are dictated by the orbit regime under consideration. For low-earth orbit objects, we take distances to be in units of Earth radius, $R_\oplus$, and velocities to be in units of $\sqrt{GM_\oplus/R_\oplus}$, the velocity of an object in circular orbit of Earth's radius. The corresponding natural time unit is inverse of such an orbit's angular frequency: $\sqrt{R_\oplus^3/GM_\oplus}$. For objects in geosynchronous orbit, natural units are the same, except that $R_\oplus$ are replaced by the geosynchronous radius, $\sim 6.61 R_\oplus$.

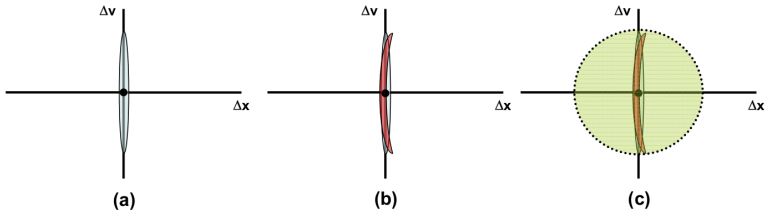

Figure 6. (a) Schematic uncertainty volume of an IOD state appears in blue, where the uncertainty in velocity is much larger than the uncertainty in position, estimated using the linear truncation of the evolution equations. (b) the red uncertainty volume of the same IOD if the the full evolution equations are used. Note that when there is a large hierarchy of uncertainty scales, the linear (gaussian) approximation of uncertainty does not represent the actual uncertainty well. (c) However, the linear (gaussian) uncertainty in green does formally contain the true nongaussian IOD uncertainty and does not have a large hierarchy between the position and velocity uncertainties.

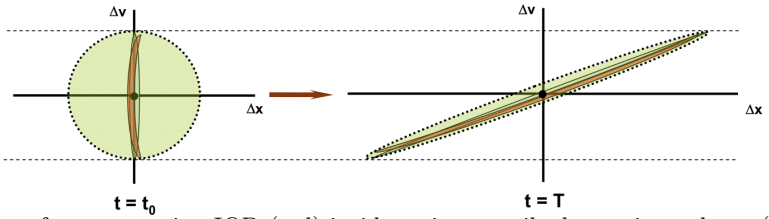

Figure 7. The evolution of a nongaussian IOD (red) inside a circumscribed gaussian volume (green) remains in that volume until large-scale nongaussianities unrelated to IOD become important (i.e., when $\Delta v t \sim 1$). The virtue of that choice of inflation is two-fold: (1) The containment of the true uncertainty persists with time (2) After a brief period of time during which the covariance stretches out in space, the spatial extent of the inflated covariance volume is not substantially larger than the spatial extent of the true uncertainty.

5. New covariance is taken as the initial covariance and the state and covariance are propagated as normal

We have truncated the evolution equations to terms linear, or at most quadratic, in uncertainties. An argument needs to be made for why this is a valid procedure to undertake. In fact, in the linear truncation, a case can be made that this is in fact *not* sufficient. Figure 6 depicts the scenario. When there is a large hierarchy of uncertainty scales, i.e. while $\Delta v \sim \epsilon \ll 1$ in some normalized units, $\Delta x \ll \Delta v$, is even smaller such that there are terms in the governing equations such that $\Delta v^2 \sim \Delta x$, a linear (gaussian) treatment of uncertainties is insufficient. In order to proceed with a gaussian treatment of the uncertainty evolution, one needs to contain the true uncertainty in a volume that is know to obey linear (gaussian) evolution.[4]

So, we surrender having proper gaussian statistics of the one, two, and three–sigma ellipsoids in favor of ensuring overall containment by some fixed n–sigma modified ellipsoid. The difference between this approach and merely naively increasing the initial volume of the original covariance ellipsoid (by taking a sufficiently large number of sigmas) is that the inflation is asymmetric in position versus velocity (thus maintaining a relatively manageable size in real space), and the containment resulting from this inflation is *persistent in time*.

To see how well covariance inflation works to address the shortcomings of covariances as measures of orbital uncertainty, our reference test orbit is again taken as truth. Figure 8 shows one example of the results. It has been shown that raw mean equinoctial covariance does not provide sufficient containment when based on short initial tracks and when propagated 12 hours. Covariance inflation as prescribed improves containment. Many additional cases have been attempted using this inflation technique. Several types of radar passes with varying track lengths and observation noise levels were tested. Also, several types of orbits were tested including manned spaceflight low-altitude, sun synchronous, NASA 2030 catalog peak density, Iridium, 2000 km altitude, and Molniya orbits. In each of these cases, containment improved sufficiently to make short-track estimates and covariance useful for representing actual uncertainty. More specifically, the 4σ containment remained above 70% while propagating the estimates and covariance for 12 hours. This criteria is important

[4] One needs to be careful, however, to ensure that the circumscribing volume continues to be gaussian even as it evolves in time. In the first case (i.e., no hierarchy of small scales), treating just the $\mathcal{O}(\epsilon)$–terms should be sufficient. The most perilous terms ignored are $\mathcal{O}(\epsilon^2 t)$–terms. This implies that things should be fine until $\epsilon t \sim 1$. In the second case, taking $\Delta x \sim \epsilon_1$ and $\Delta v \sim \epsilon_2$ such that $\epsilon_2^2 \sim \epsilon_1$ where $\mathcal{O}(\epsilon^2)$–terms are kept, the most perilous terms ignored are $\mathcal{O}(\epsilon^3 t^2)$–terms. In all cases where there are $\epsilon_1 \sim \epsilon_2^2$, there are ϵ_2^2–terms accounted for that dominate. Thus again, one is safe until $\epsilon_2 t \sim 1$.

Under circumstances where one is neglecting second-order terms, one needs to make sure terms that look like $\epsilon^2 t$ don't get too large. If there is a hierarchy of scales such that $\epsilon_1 \gg \epsilon_2^2$ so that it is safe to neglect second-order terms initially, one still needs to worry about neglecting $\epsilon_2^2 t$–terms. One may show that if one uses the containing gaussian volume of an initially nongaussian pancake-like volume (as depicted in Fig. 6), then the evolution is safely gaussian until such time as $\epsilon_2 t \sim 1$ (see Fig. 7).

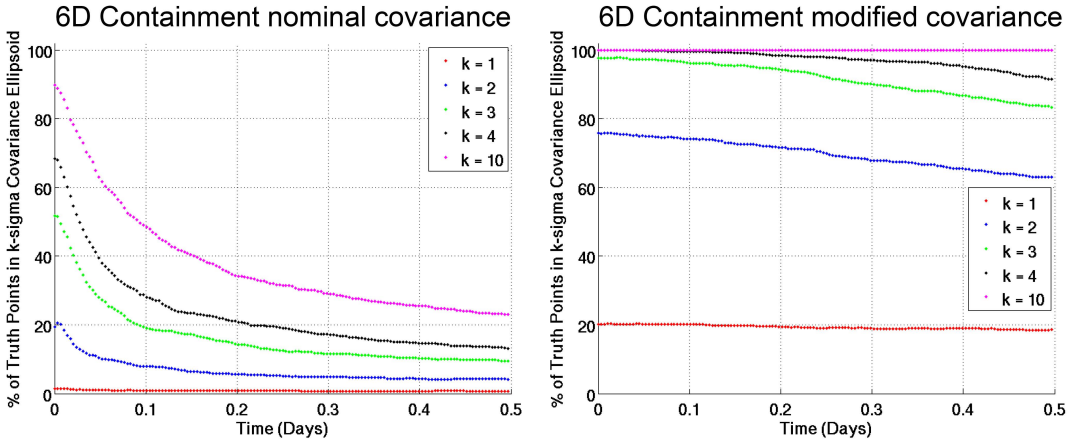

Figure 8. Inflation results (10 degree zenith). 474 km altitude circular orbit. 97.3° inclination. 1σ range noise: 40 m. 1σ azimuth noise: 10 mdeg. 1σ elevation noise: 10 mdeg.

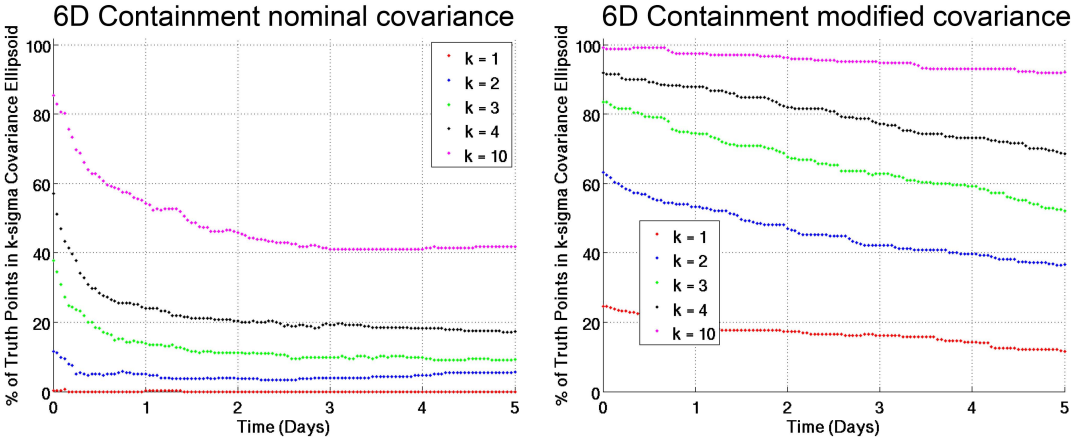

Figure 9. Inflation results (real data). Object 7646 800 - 1100km Object 16908 1500 - 1500km. Eglin, PARCS, ALTAIR, and HAX. Based on 272 independent 60-sec tracks at varying elevation angles.

for short-tracks taken with radar because revisit times for a fence could be on the order of 12 hours, and correlation methods relying on covariance require agreement between the covariance and actual uncertainty.

The inflation technique has also been corroborated using real data. These real data tests are constructed on calibration satellites (EGP (SSN 16908) and STARLETTE (SSN 7646)) for which we have truth orbits with better than 10 meter errors. The truth orbits are based on ILRS laser-range measurements; the same truth orbits are used by the MIT Lincoln Laboratory Space Surveillance Center (LSSC) to calibrate the Millstone, Haystack and HAX radars. Radar observations were gathered from the Eglin, Parcs, HAX and Altair radars on EGP and STARLETTE. Because these are calibration satellites, most of the tracks were taken horizon-to-horizon. Four weeks of calibration orbits and four weeks of observations were gathered in order to collect a significant sample. Each horizon-to-horizon radar track was divided up into 60-second sub-tracks yielding 272. Orbit estimates and covariances were generated from each sub-track and then predicted for the next five days. The k–values were then calculated using the state and covariance estimates and the coincident truth states.

Figure 9 shows that the simulations of short tracks are verified by the tracks of real data. The nominal covariance performs poorly over the 5 day span. The containment of the 4-sigma covariance drops to under 20% after 5 days of propagation. The modified, i.e. inflated, mean equinoctial covariance performs reasonably well in that the 4-sigma covariance only drops to about 70% after 5 days.

8. IOD AND OBSERVATION CORRELATION

The orbital dynamics play directly into the evolution of the volume of the confidence ellipsoids. Approaching correlations in terms of state space volumes allows us to see that the growth in those volumes should occur almost exclusively along one direction in state space. That direction can be parametrized by a single value: either by the uncertain along-track velocity value or the uncertain drag coefficient. So, while the size of the confidence ellipsoid may grow in that direction, the other five directions are relatively well-contained, thus preventing the volume of the confidence region from becoming uselessly large.

Using the evolution equations Eqs. (2) and (4), one may conceive of a scenario consisting of a sequence of observations tagged by their times of occurrence, $\{t_i\}$. Each set of observations will be processed by some estimation scheme into likelihood distributions characterized by $\{\mathbf{x}_i(t_i), \mathbf{C}_i(t_i)\}$, each corresponding to some unknown target or object in the catalog. Each of these objects (as represented in state space by the likelihood distributions) must be evolved forward in time by Eqs. (2) and (4).

Consider the initial object represented by $\{\mathbf{x}_0(t), \mathbf{C}_0(t)\}$. Eventually, a new set of observations, corresponding to $\{\mathbf{x}_i(t_i), \mathbf{C}_i(t_i)\}$, is such that $\mathbf{x}_i(t_i)$ falls within the confidence ellipsoid given by $\{\mathbf{x}_0(t), \mathbf{C}_0(t)\}$ where $t = t_i$. Then there is a significant probability that this new set of observations is associated with the original object.

To assess the closeness of two overlapping thin ellipsoids that are oriented in different directions, one needs to construct a more complex metric of association. Let us call this the "joint k–distance." This new distance metric provides a single number for automated correlation decisions that accounts for uncertainty. Consider the following: find the minimum combined k–distance

$$k_{\text{joint}}(\mathbf{x})^2 = k_1^2 + k_2^2 \text{ where } k_i^2(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_i)C_i^{-1}(\mathbf{x} - \mathbf{x}_i) \ , \tag{15}$$

and $\mathbf{x}_i$ is the estimated state of the i–th set of observations (propagated or freshly acquired), and C_i is its corresponding inflated covariance. By finding $\mathbf{x}$ that minimizes k_{joint}, we find

$$\begin{aligned} k_{\text{joint}}^2 = \mathbf{x}_1 C_1^{-1}\mathbf{x}_1 &+ \mathbf{x}_2 C_2^{-1}\mathbf{x}_2 \\ &- \left(\mathbf{x}_1 C_1^{-1} + \mathbf{x}_2 C_2^{-1}\right)\left(C_1^{-1} + C_2^{-1}\right)^{-1}\left(C_1^{-1}\mathbf{x}_1 + C_2^{-1}\mathbf{x}_2\right) \ . \end{aligned} \tag{16}$$

The k–distance represents the likelihood that both objects are the same state, assuming that the combined probability is just the product of the individual probabilities. The joint k–distance is guaranteed to be positive and represents a normalized distance scale that characterizes whether two states with given uncertainties are likely to be the same.

We can now look at how this set of schemes, i.e. covariance inflation, joint k–distance test, works for the association of observations of unknown objects. We begin by simulating a fence-type radar sensor with a specific beam shape, as well as range, azimuth, and elevation noise that observes objects in orbits found in the current AFSSN catalog. The orbital states were regarded as truth for these simulations. Gaussian noise was added to the observations of the truth trajectory and were generated with a spacing of between 2 to 8 seconds while the satellites were in the field of view of the radar. The resulting noisy observations were then processed using the 6D state and covariance correlation algorithm. Finally, correlation decisions were compared to truth to determine what percentage of observations were correctly correlated.

To manage all the tasks necessary to test these covariance-based correlation methods, a prototype system was developed. The correlation and catalog-maintenance prototype reads in track files that contain observations of satellites as would be collected from a radar sensor. The observations from the incoming track file are processed using IOD, specifically the range and angles method [4], and then least squares orbit determination [4] to create a new state and covariance representing the incoming track. The states and covariance from the catalog are propagated to the time of validity or epoch of the state and covariance of the incoming track. The joint k–distance is calculated by pairing each state and covariance from the propagated catalog with the state and covariance produced from the incoming track. Any pair with a joint k–distance less than eight is considered a match and is saved for further consideration. If there are multiple matches or matches involving long propagations, a differential correction is used to resolve the correct match or verify the match.

Once the prototype has processed all given track files, the observation tags determined by the prototype are compared with the original tags supplied with the observations. The results of tests using simulated track files indicate a high degree of success in using the estimates, modified covariance and the k-distance metric for correlation. Including all trials, 84192, the algorithm prototype successfully correlated 74112 of the 1st and 2nd track opportunities for a success rate of 88%. The percentage of 2nd tracks matched to the wrong 1st track was less than 1%. These trials consisted of all SSN and NASA 2030 catalogued satellites in which the track simulations provided 10 degree track extents from the radar fence locations. The observations were simulated with 18 millidegree angle noise and 30 meter range noise. The simulated radar fence sites were located in either Harold Holt, Australia, Ascension Island in the Atlantic Ocean or the Kwajalein Atoll in the Pacific Ocean. The satellite catalogs used were the AFSSN catalog as of 3 March 2010, the NASA2030 catalog (a catalog of $>$ 2cm objects simulated for the year 2030 as evolved from 2006, provided by the NASA Orbital Debris Program Office) and a synthetic catalog which contains satellite orbits distributed uniformly in apogee, perigee and inclination.

In addition to the simulated observation association cases mentioned already, real data cases were tested using the correlation prototype. These eight cases are summarized in Table 1. Radar observations from the Eglin, FPS–85, radar from March 2010 were gathered and fed through the prototype system. The tracks of observations chosen were selected so that Eglin tracked the observed satellite twice in a three day interval. For example, in case 1, there were 147 satellites for which this occurred and in case 2, there were 108 satellites. The interval for each case is identified in the "Catalog, Observation span" column of Table 1.

TABLE 1
Real observation correlation results from the correlation prototype.

Case #	Radar	Catalog, Observation span	Satellites with two passes	Satellites correctly correlated after two passes (with check DC)	Average matches per track
1	Eglin	SSN, 25-28 Mar 2010	147	119 (81%)	1.1
2	Eglin	SSN, 28-31 Mar 2010	108	91 (84%)	0.95
3	Eglin	SSN, 1-3 July 2010	142	91 (64%)	1.3
4	Eglin	SSN, 1-3 April 2010	89	76 (85%)	0.97
5	Eglin	SSN, 4-7 June 2010	59	45 (76%)	1.2
6	Eglin	SSN, 14-17 June 2010	49	40 (82%)	1.1
7	Eglin	SSN, 14-18 April 2010	135	104 (77%)	1.1
8	Eglin	SSN, 1-4 March 2010	112	83 (74%)	0.96

As one would expect when using real data instead of simulated data, the ratio of satellites correctly correlated after two passes over the satellites with two passes is less for these cases than for the simulated cases. The average ratio is 77%. This ratio is about 11% less than the average success rate of the simulated cases. Features of real data that probably contribute to the reduced correlation success rate are range, azimuth and element measurement biases, drag coefficient and atmospheric density uncertainties and longer gap times between 1st and 2nd tracks. Indeed, the correlation decisions are ultimately based on the k–distance metric which relies on accurate covariance.

9. CONCLUDING REMARKS

We have highlighted how standard measures of orbital uncertainty break down in important domains of interest. For long-track, i.e. horizon-to-horizon tracks of initial orbits, results show that mean equinoctial state variables provide for covariance propagation while osculating ECI and equinoctial state variables do not. For short-track initial orbits, mean equinoctial covariances are also inaccurate. However, using a covariance modification method that strategically inflates, but doesn't over-inflate the covariance, a modified mean equinoctial covariance can serve as an accurate measure of orbital uncertainty over extended periods of time. The accuracy of this new uncertainty measure persists for 12 hours or more, providing a powerful tool relevant to space-situational awareness needs.

Starting with a space environment composed completely of UCTs, the simulated and real data cases performed for this study show that, implementing this new uncertainty measure, about 78% of those objects can be acquired by the fence, an IOD can be executed, the fence can reacquire the satellite and can correctly identify the satellite, and can subsequently refine the orbital state. All this happens in about 12 to 36 hours, in the case of a 13600 object scenario, for the specified fence performance and with a single fence site. For a more powerful single fence, the same should hold true for the full NASA 2030 catalog.

ACKNOWLEDGMENTS

The authors would like to acknowledge the efforts of Kevin Stout to propagate covariance in the GTDS ephemeris generator and to write tools for assessing the covariance containment using real observations and truth orbits. The authors would also like to thank those who provided insight and guidance in this work: Sid Sridharan, Mark Dickson and Paul Cefola. This work was sponsored by the Department of the Air Force under contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by the United States Government.

REFERENCES

[1] M. R. Akella J. L. Junkins and K. T. Alfriend. Non-gaussian error propagation in orbital mechanics. *Journal of the Astronautical Sciences*, 44, 1996.

[2] Alfriend K. T. Sabol C. Hill, K. Covariance-based uncorrelated track association. *AIAA/AAS Astrodynamics Specialist Conference and Exhibit. Honolulu, Hawaii. Aug 18-21*, 2008.

[3] Sukut T. Hill K. Alfriend K. T. Wright B. Li Y. Schumacher P. Sabol, C. Linearized orbit covariance generation and propagation analysis via simple monte carlo simulations. *20th AAS/AIAA Space Flight Mechanics Meeting. San Diego, CA. Feb 14-17*, 2010.

[4] Cappellari J. O. Jr. Velez C. E. Fuchs A. J. Long, A. C. Goddard trajectory determination system (gtds) mathematical theory revision 1. *NASA GSFC Document FDD/522-89/001 GSFC Code 550, Computer Sciences Corp. Document CSC/TR-89/6001*, 1989.

[5] R. Cefola P., Proulx. Application of the semianalytic satellite theory to shallow resonance orbits. *AAS/AIAA Space Flight Mechanics Meeting. Houston, TX. Feb 11-13*, 1991.

[6] William E. Weisel. *Modern Orbit Determination.* Beavercreek, OH: Aphelion Press, 2003.