AAS 13-822

UNCERTAINTY CHARACTERIZATION FOR ANGLES-ONLY INITIAL ORBIT DETERMINATION

Christopher R. Binz* and Liam M. Healy*

When no information is known about a satellite's orbit, an initial orbit determination method must be used. Traditionally, this yields only a point solution, with no uncertainty information. This paper proposes that probabilistic information may be derived directly from the initial orbit solution, given measurement error characteristics. This information may be especially useful when observations of the object are very sparse, and any additional information is valuable. In this paper, we examine Gauss's angles-only method. Assuming the noise characteristics of the sensor are known, we estimate an empirical probability density function yielded by combining the required three measurements.

INTRODUCTION

When the orbit of a space object is unknown, a preliminary estimate may be made by combining measurements of the object in a process known as *initial orbit determination* (IOD). For measurements made via optical means (i.e. a telescope) from a known observing position, information is provided in two dimensions in the sensor's reference frame. As such, a minimum of three measurements is required to produce a rough estimate of the object's (six-dimensional) state. In a typical orbit determination processing scheme, this estimate is used as an initial seed to a batch least squares estimation process, which fits the orbit to a larger number of measurements, resulting in both a more refined orbit and the first expression of the uncertainty in the estimate.[1] Upon the collection of more measurements, the orbit can be refined with the usual methods (e.g. batch least squares, Kalman filtering). Typically, the uncertainty estimate from least squares is required to initialize a Kalman filter-like estimation scheme.

This paper examines the formation of the initial uncertainty estimate. The underlying motivation may be stated simply: to determine if there is useful information resulting directly from the IOD process beyond the rough point solution. Possible applications are more varied: obtaining an uncertainty estimate directly from a three-observation IOD and characterizing the obtainable quality of an IOD estimate are two examples that will be explored in this paper. One of the questions driving this work is, "can a *useful* IOD solution be obtained with poor quality sensors?"

We address this by exploring the distribution of IOD solutions given a set of measurements with quantifiable uncertainty distributions. This paper focuses on an optical sensor that is essentially binary—either the satellite is in the field of view or it is not. Specifically, the sensor is modeled with a uniform prior probability density function. This is not a common tool with which to make observations of satellites, but our motivation is twofold. First, the idea of *Ubiquitous Low-Cost*

*Mathematics and Orbit Dynamics Section, Naval Research Laboratory, Code 8233, 4555 Overlook Ave., SW, Washington, DC 20375-5355.

Sensors (ULCS) is that they are widely deployed yet significantly noisier than more expensive sensors. If the uncertainty is handled carefully, and a useful result is attainable, these cheap sensors may allow for significantly higher coverage than is possible with current costly sensors. Second, it can show what happens when one removes the assumption of normality that is so infused in the methods and techniques that are widely used in the orbit determination/estimation field. The IOD process is inherently nonlinear, and the assumption that the result will have a Gaussian uncertainty distribution will suffice is naïve.

There are several well-known methods for IOD with angles measurements only. Vallado offers an overview of three of the most popular techniques: Laplace's method, double r-iteration, and Gauss's method.[2] He asserts that Gauss's method produces reliable solutions, a claim largely based on previous work by Escobal[3] and Long, et al.[4] As a starting point, we chose to focus on Gauss's method for this study. We will not go into the details of this algorithm*, but it is worthwhile to list some of its salient features. First, unlike Laplace's method (which fits the orbit only to the middle observation), Gauss's method fits the orbit estimate to all three measurements. Second, the implementation of Gauss's method we use here is itself an iterative algorithm: after the initial estimate is made by finding the correct root of an eight-order polynomial, the solution is iterated on using the Herrick-Gibbs method to refine the range estimate.[2] Third, Long, et al. describes Gauss's method as most effective when the data are spread over an arc of less than 60° in mean anomaly, and we adhere to this guideline in our study.

By itself, an IOD solution resulting from angles measurements only is not especially accurate. Solutions to the angles-only IOD problem are generally used only as a starting point, e.g. for a differential correction routine. As stated previously, this process also produces an expression for the uncertainty in the estimate. However, the estimate it provides is only optimal (in the maximum likelihood sense) for measurements with normally-distributed error.[5] Additionally, the uncertainty is given as a covariance matrix, which does not contain higher order statistical information that may be significant. Various studies have proven that the nonlinearity in the various steps of the orbit estimation process invalidate the assumption of Gaussian uncertainty.[6,7] The inherent nonlinearity of the angles-only IOD process creates similar issues.

METHODOLOGY

Due to the algebraic complexity of Gauss's IOD method, we explore the solution space using Monte Carlo techniques. The problem is posed as follows: assuming three sets of uncertain angular measurements, what is the distribution of IOD solutions resulting from Gauss's method? To begin, we look at measurements of topocentric right ascension and declination over a single pass (ideal for Gauss's method) from a single ground station with a fixed, known location. We assume the issue of observation association has been resolved.

The issue of a uniform distribution in measurement space is not simply a matter of providing upper and lower bounds to the two angular measurements. Instead, we must produce a set of unit vectors that is uniform on the sphere, bounded by some maximum deviation (the field of view of the sensor). It is beneficial to begin with the unit vector pointed along the Z axis (towards the "North pole" of the sphere). In cylindrical coordinates (ρ, ϕ, z), the azimuthal angle ϕ is, in fact, uniformly distributed on $[0, 2\pi)$. Likewise, the distribution in z is uniform on $[\cos\Theta, 1]$, where Θ is the field of view of the sensor (the "half angle", to be precise). Now we may construct the set of uniform unit

*A detailed algorithm description may be found in any of the preceding references.

vectors in *Cartesian* space as $(\sqrt{1-z^2}\cos\phi, \sqrt{1-z^2}\sin\phi, z)$. All that remains is to rotate the set of vectors to the desired orientation (that is, back to the original orientation of the measurement).

Monte Carlo

By expanding our three angular measurements (i.e. topocentric unit vectors) into solid "cones" of N uniformly distributed unit vectors, we now have three sets of N measurements each, allowing for N^3 possible IOD solutions. In order to choose an N that is "large enough" that there is little statistical variation between runs, we borrow a technique from Healy.[8] Instead of generating the sample points with a pseudorandom number generator, we may use a low-discrepancy sequence to generate so-called quasi-random numbers.* Particularly, Sobol's algorithm[9] is used here to generate a uniform distribution in the measurement space. The key benefit of using a low-discrepancy sequence is that smaller sample sets are subsets of larger ones. This allows us to incrementally increase the sample set size and observe the statistical behavior of the results. In Figure 1, we look at the variance in each coordinate (using the set of modified equinoctial coordinates,[10] for reasons discussed later) with increasing sample size. From this, a value of $N = 30$ is chosen for these simulations, resulting in a solution space of $30^3 = 27000$.

Although the algorithm for Gauss's IOD method is not particularly resource-intensive, the combinatorial explosion in the number of solutions resulting from the Monte Carlo approach begs a more efficient means of estimating the distribution. One approach is to use weighted, deterministic *sigma points* (as in the unscented transform[11]) to characterize the distribution in measurement space. Then we may follow the same process of solving for every possible IOD solution (now a much smaller number). The resulting distribution of sigma points characterizes the first two moments of the actual distribution.[11] The process of using the unscented transform to quantify uncertainty has been explored by others.[12] We explore this in more detail below.

Least Squares

For the sake of comparison, we use the traditional approach of fitting the measurements using the least squares method. For this, we process the same three measurements further with a batch least squares process (using the IOD solution obtained using the center of each observation distribution as the initial guess).† This produces a modified state estimate as well as an uncertainty estimate expressed by a covariance matrix. It is important to note that this process assumes that the measurements distribution may be described completely by the mean (center point of the measurement) and covariance. The least squares fit is done using only the three measurements because we are assuming these are the only measurements currently available, and it allows for more direct comparison to the IOD solution distribution. A discussion of the measurement covariance matrix may be found near the end of this section.

Sigma Points

Implementing our sigma point method in measurement space necessitates another approximation. Although it is defined to be able to handle higher-order moments,[11] we choose to simplify the process in this paper and use the common practice of utilizing only the first two moments of the original distribution. Thus we again return to the restriction of using only the first two moments to

*`http://en.wikipedia.org/wiki/Low-discrepancy_sequence`

†Lacking any other information, we must use the mean value of the measurement.

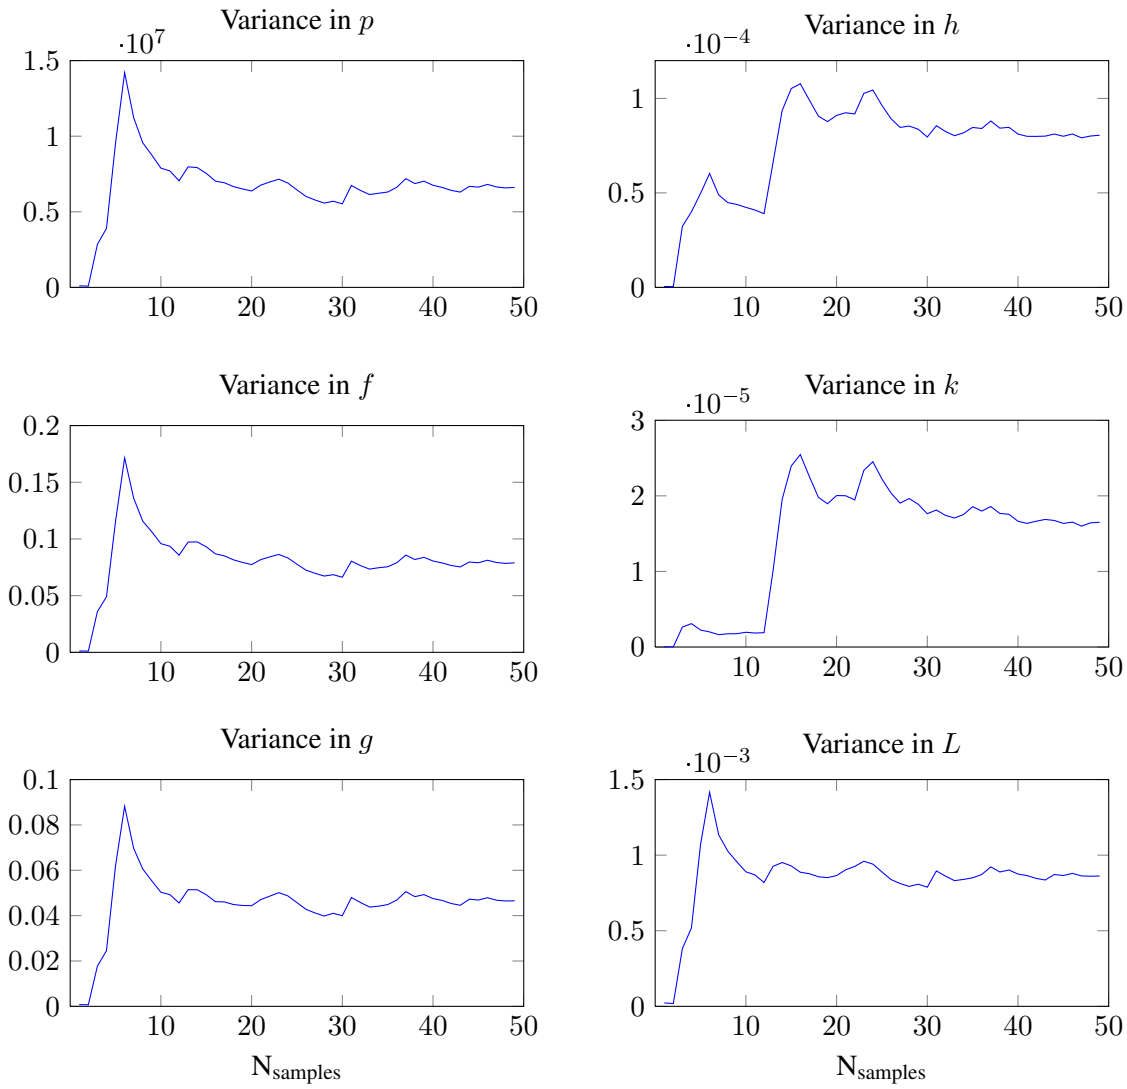

Figure 1. Behavior of variance in each coordinate with increasing sample size.

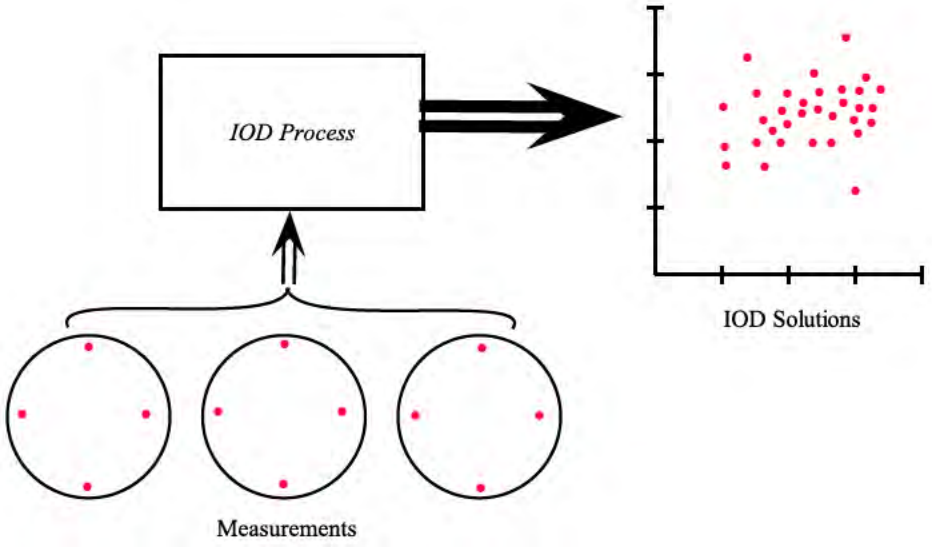

Figure 2. Overall process for describing IOD solution-space distribution with sigma points.

describe the initial distribution, losing some of the information about the initial distribution in the process. The benefit of doing this is that it allows for more direct comparison with the output from the least squares solution, which also only produces a covariance as an uncertainty estimate.

The process of using sigma points for an IOD problem is unique in that it does not follow the typical paradigm of N m-dimensional sigma points propagated through a function $g(\cdot)$ that yields N m-dimensional propagated points. Instead, we have three sets of N two-dimensional sigma points (one set for each measurement), and put it through our function which yields N^3 "propagated points". So, some modification to the usual unscented transform method is required. The overall process is shown in Figure 2.

First, we choose a weighted set $\boldsymbol{X}$ of $2n+1 = 5$ sigma points for each measurement distribution, following the method described by Hartikainen, et al.[13] For given mean $\boldsymbol{m}_j$ and covariance $\boldsymbol{P}_j$ of the j^{th} measurement ($j = 1, 2, 3$), the points are

$$\boldsymbol{X}_j = \begin{bmatrix} \boldsymbol{m}_j & \cdots & \boldsymbol{m}_j \end{bmatrix} + \sqrt{c} \begin{bmatrix} \boldsymbol{0} & \sqrt{\boldsymbol{P}_j} & -\sqrt{\boldsymbol{P}_j} \end{bmatrix}, \tag{1}$$

where $c = \alpha^2\sqrt{n+\kappa}$ is a function of the dimension of the measurement space ($n = 2$) and two positive constants α and κ, which are tunable parameters (in this work, we use $\alpha = 0.5$ and $\kappa = 1$).[13] The set of weights for the means W_m and for the covariance W_c are calculated as

$$\begin{aligned} W_m^{(i)} &= W_c^{(i)} = \lambda/(n+\lambda), & i = 0 \\ W_m^{(i)} &= W_c^{(i)} = \lambda/(2(n+\lambda)), & i = 1, \ldots, 2n, \end{aligned} \tag{2}$$

where $\lambda = \alpha^2(n+\kappa) - n$.[13] As all three sets of measurements are of the same dimension $n = 2$, the weights will be identical between sets, and thus must only be calculated once. Once every possible set of sigma points goes through the IOD process, the mean and covariance of the "propagated" set

$\boldsymbol{m}_p$ and $\boldsymbol{P}_p$ may be calculated by

$$\begin{aligned}\boldsymbol{m}_p &\approx \sum_{i=0}^{2n} W_m'^{(i)} \boldsymbol{y}^{(i)} \\ \boldsymbol{P}_p &\approx \sum_{i=0}^{2n} W_c'^{(i)} (\boldsymbol{y}^{(i)} - \boldsymbol{m}_p)(\boldsymbol{y}^{(i)} - \boldsymbol{m}_p)^T,\end{aligned} \tag{3}$$

where $W_m'^{(i)}$ and $W_c'^{(i)}$ are $(2n+1)^3 = 125$ element weight vectors, consisting of the product of the weights of each specific measurement that went into the IOD process (e.g. for the first IOD solution consisting of sigma points with index zero from all three sets, the weights are $W_m'^{(0)} = W_m^{(0)} W_m^{(0)} W_m^{(0)}$). The vectors $\boldsymbol{y}$ are the solutions to the IOD problem—e.g., if the IOD process is represented by the function $g(\cdot)$, then $\boldsymbol{y} = g(\boldsymbol{X}_1, \boldsymbol{X}_2, \boldsymbol{X}_3)$.

As both the least squares and the sigma point methods require an expression for the measurement covariance, discussion of how we calculate this is warranted.* The complexity of generating a uniform distribution in this measurement space (as described in the Methodology section) makes this a complicated problem. Although we expect an analytical solution method exists (and intend to pursue this further in the future), the approach taken here is to simply generate a large number of samples and compute the sample covariance. The sample generation is a vectorized set of calculations that finishes in a trivial amount of time in Matlab, and this does not significantly affect the total run time of any of the methods shown.

Due to the relatively delicate convergence behavior of Gauss's IOD method, there will often be a group of samples that are far from the actual orbit solution, and occasionally outside the confines of elliptical Earth orbits. If we make the prior assumption that the object is an Earth-orbiting satellite,† this is easy to counteract—we can simply filter out points that fall outside certain bounds, such as $0 \leqslant e < 1$ and $R_\oplus < a < 10R_\oplus$, where e is eccentricity, a is semimajor axis, and $R_\oplus$ is the radius of Earth. Filtering these points, however, has statistical implications. In the full Monte Carlo case, in which there is a large number of samples, analyzing only the filtered result is typically enough to provide an accurate description of the distribution. However, in the case that uses only a few sigma points, filtering has a significant impact on the distribution. In the test case presented, we do not encounter the issue of convergence, although a diverged case is shown and discussed briefly.

Due to the relatively delicate convergence behavior of Gauss's IOD method, there will often be a group of samples that are far from the actual orbit solution, and often outside the confines of elliptical Earth orbits. This is easy to counteract—we can simply filter out points that fall outside certain bounds, such as $0 \leqslant e < 1$ and $R_\oplus < a < 10R_\oplus$, where e is eccentricity, a is semimajor axis, and $R_\oplus$ is the radius of Earth. Filtering these points, however, has statistical implications. In the full Monte Carlo case, in which there are many thousands of samples, looking at the filtered result is typically enough to provide an accurate description of the distribution. However, in the case that uses only a few sigma points, filtering has a significant impact on the distribution. In the test case presented, we do not encounter the issue of divergence, although an example case is shown and discussed briefly.

*The mean value is trivial, as it is just the center of the distribution.

†Which is probabilistic information in itself.

Table 1. Reference orbital elements.

Orbit Element	Value (km and deg)
Epoch	1 Jan. 2013 17:00:00
a	6878.00
e	1e-5
i	45.0
ω	0.0
Ω	0.0
ν	0.0

Test setup

Simulated observations of a satellite moving under two-body dynamics are taken at ten second intervals from a single ground station with a known location. Table 1 shows the reference orbit used. The selection of which three measurements to choose for the IOD is a deep topic in itself; after some trial and error we choose three equally-spaced points at 150 second intervals.

The sensor model is a telescope with a half angle field of view of 0.5°. However, instead of modeling the measurement with, for instance, a bivariate Gaussian distribution centered on the boresight, we assume a uniform distribution over the entire field of view. This simulates a sensor that can detect when the satellite is in its field of view, but provides no information beyond that. The simulated measurements are thus corrupted with uniformly distributed noise. The telescope is located at 30° North latitude, -80° East longitude, and 0 km altitude.

RESULTS

We may view the distribution in solution space directly by looking at plots of the sample population from the Monte Carlo runs. For this, we choose to plot pairs of the *modified equinoctial element* set (p, f, g, h, k, L), as defined by Walker et al.[10] It is useful to use this element set for two reasons: it is nonsingular for any eccentricity or inclination, and it is a more natural system in which to describe orbits than, for instance, Cartesian coordinates.

Figure 3 shows the population of solutions from the Monte Carlo process. The "truth" orbit, which is known here because we simulated it, is shown in these plots with a black cross. By inspection of the sample population we see that the solution distribution is distinctly non-Gaussian, which is particularly evident in the p, L and h, k planes. The histograms along the axes illustrate this further. We also see the degree of uncertainty inherent in the IOD process (even with a relatively modest 0.5° field of view): solutions in the semilatus rectum range from thousands of kilometers below the true solution to tens of thousands of kilometers above.

The histograms in Figure 3 also illustrate the error inherent in using the maximum *a posteriori* estimate in solution space: while there is a relatively clear maximum in $p, L, f,$ and g, it is biased from the truth solution, and in g and k, a clearly-defined maximum simply does not exist.

The degree to which the sample population follows a multivariate normal distribution may be

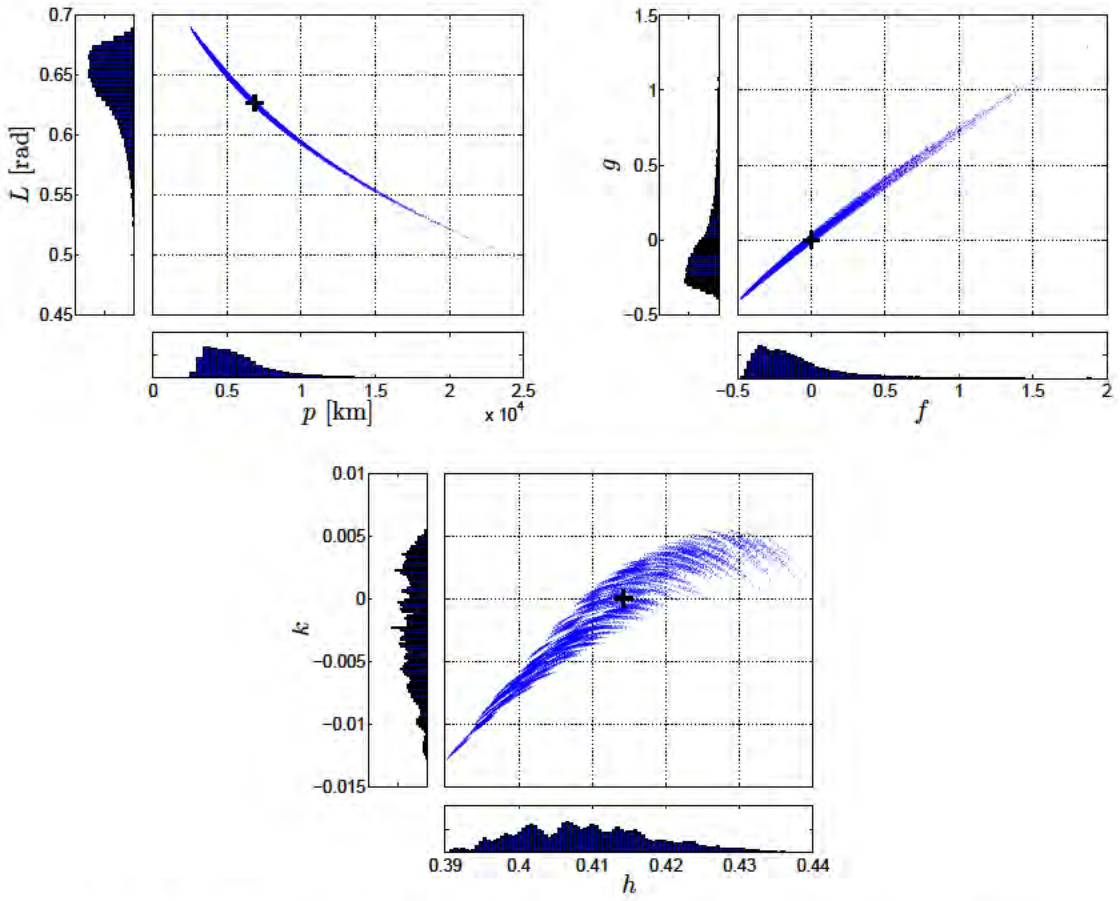

Figure 3. Distribution of IOD solutions in modified equinoctial element space plotted with the true solution (black cross).

further illustrated using the χ distribution, a technique also used by Sabol, et al.[6] For a normally distributed random vector $\boldsymbol{X} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, the probability distribution of

$$k = \sqrt{(\boldsymbol{X} - \boldsymbol{\mu})\boldsymbol{\Sigma}^{-1}(\boldsymbol{X} - \boldsymbol{\mu})^{-1}} \tag{4}$$

follows the χ distribution.[14] k is also known as the Mahalanobis distance, a normalized metric that is essentially a measure of how many standard deviations a random sample is from the underlying normal distribution. For a six-dimensional random vector, the cumulative distribution function (CDF) F of the Mahalanobis distance k has the form $F(k) = 1 - 1/8(k^4 + 4k^2 + 8)e^{-.5k^2}$.[6]

Using the sample mean and covariance of the Monte Carlo population ($\boldsymbol{\mu}_{MC}, \boldsymbol{\Sigma}_{MC}$), we may plot the Mahalanobis distance of the samples from the multivariate normal distribution $\mathcal{N}(\boldsymbol{\mu}_{MC}, \boldsymbol{\Sigma}_{MC})$ and use this as a measure of the normality of the solution distribution. Figure 4 shows this, with the ideal Gaussian CDF curve plotted in red.

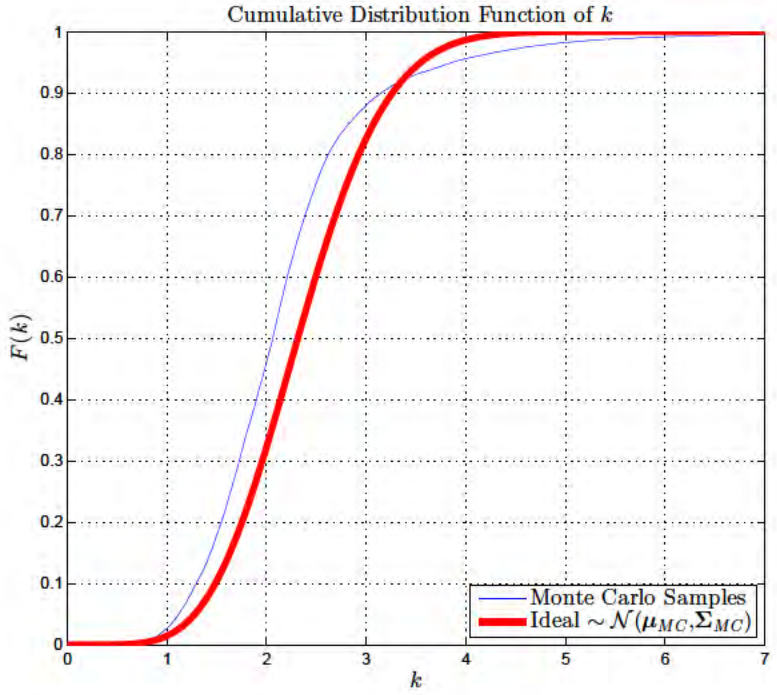

Figure 4. Cumulative distribution function (CDF) of the Mahalanobis distance k of samples from the Monte Carlo distribution to the normal distribution described by the sample mean and covariance of the population.

While the shape of the curves do not match (an expected result, given the histograms from Figure 3), it is worth noting that the sample points are almost totally contained by a 6σ ellipsoid. This suggests that while there is significant information in the higher order moments of the underlying distribution, the first two moments describe it reasonably well, and a slightly inflated covariance gives a reasonable approximation.

Next, we compare the full distribution to the results obtained from the least squares method. Figure 5 shows the Monte Carlo sample population plotted in blue, the truth solution in black, and

the least squares mean and 1σ covariance ellipse in magenta. While the covariance is aligned along roughly the correct directions, the estimated mean is significantly biased, a direct consequence of the error in the measurements. Since even a small perturbation in measurement space can significantly affect the resulting IOD solution (as can be seen in the Monte Carlo population), the assumption

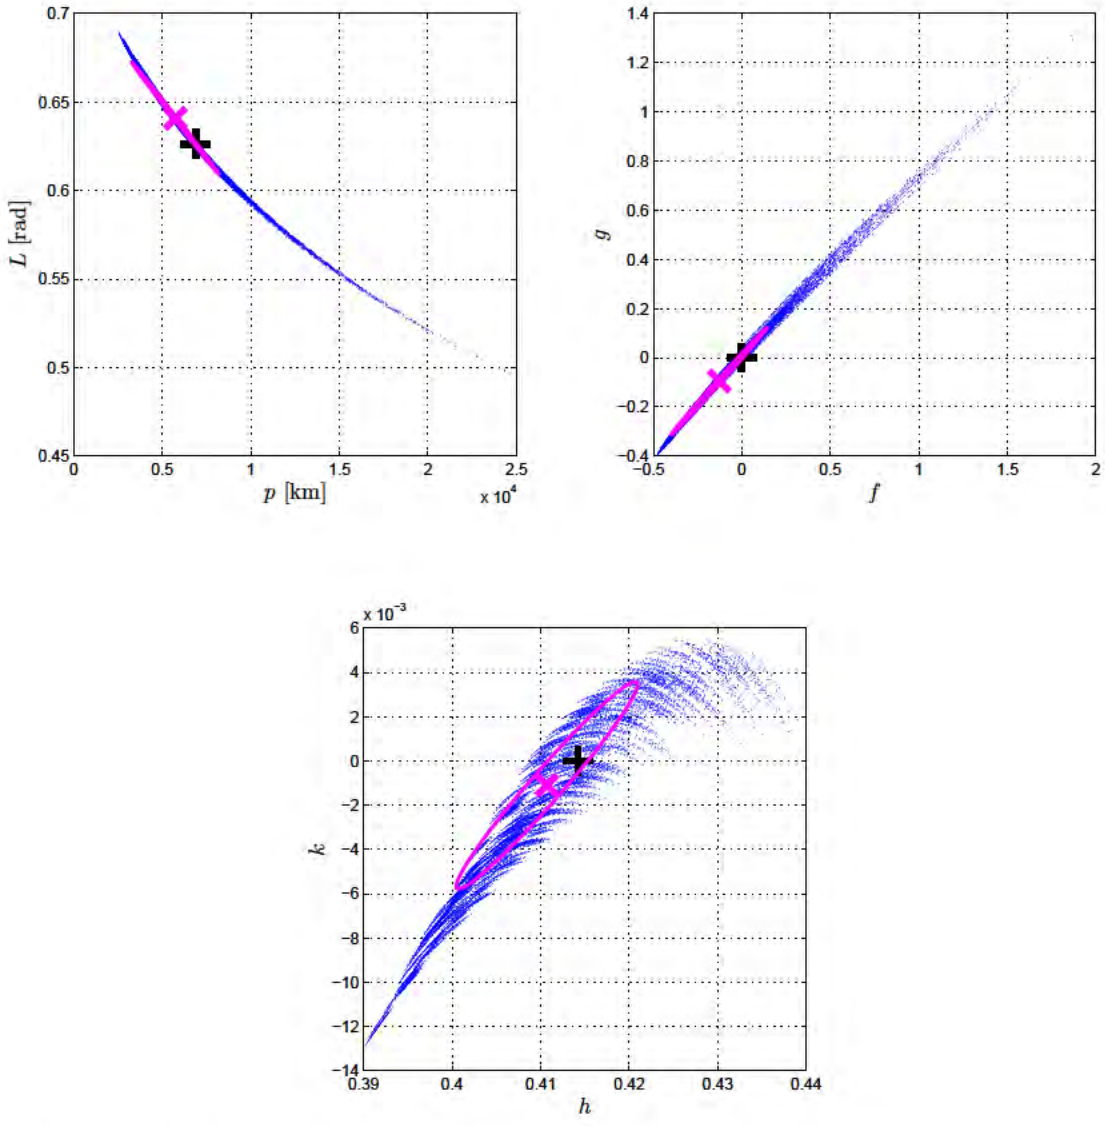

Figure 5. Plots showing IOD solution distributions in modified equinoctial element space (blue points), along with the estimate from least squares (magenta) and truth orbit (black cross).

Plotting the Mahalanobis distance using the estimated mean and covariance from least squares illustrates just how poorly the estimates describe the distribution (Figure 6). Even with a 10σ ellipsoid, less than 60% of the population is contained. This means that not only is the uncertainty estimate a poor descriptor of the underlying distribution, it is also overly optimistic.

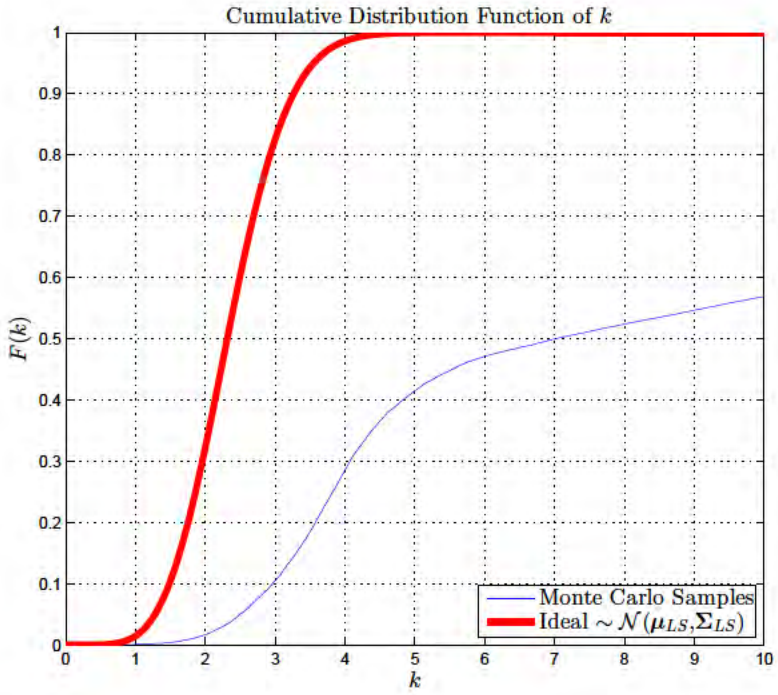

Figure 6. CDF of the Mahalanobis distance of Monte Carlo samples from the normally distributed estimate obtained with the least squares method.

Finally, Figure 7 shows the estimate from the sigma point method against the Monte Carlo population. Upon inspection, we see that the 1σ covariance ellipses have larger minor axes and slightly different orientations when compared with the least squares estimates. Additionally, the mean estimate is significantly closer compared to the estimated least squares mean.

Figure 8 serves to illustrate how much better the sigma point method fits the solution space than the least squares method. Comparing this with Figure 4, we see that the sigma point method describes the distribution nearly as well as the Monte Carlo method, which uses orders of magnitude more points. Both of these, however, are limited to describing the distribution using only the first two statistical moments, and thus information is lost and the fit is not ideal. This is an unsurprising result: the unscented transform has been widely shown to be effective even in cases involving highly nonlinear propagating function, which is essentially what the IOD process serves as here. However, this again illustrates that a slightly inflated covariance is a reasonable approximation to the underlying distribution.

The error (with respect to truth) in mean values from the various methods are displayed in Table 2. This again serves to illustrate the utility of the sigma point method, as the error is significantly lower than even the Monte Carlo method, and orders of magnitude better than the least squares approach in some coordinates.

This table also shows how inaccurate the IOD process can be. Even for the best estimates provided by the sigma point method, the difference from truth is significant. However, this is not an entirely uncommon problem in angles-only initial orbit determination, and a better initial estimate

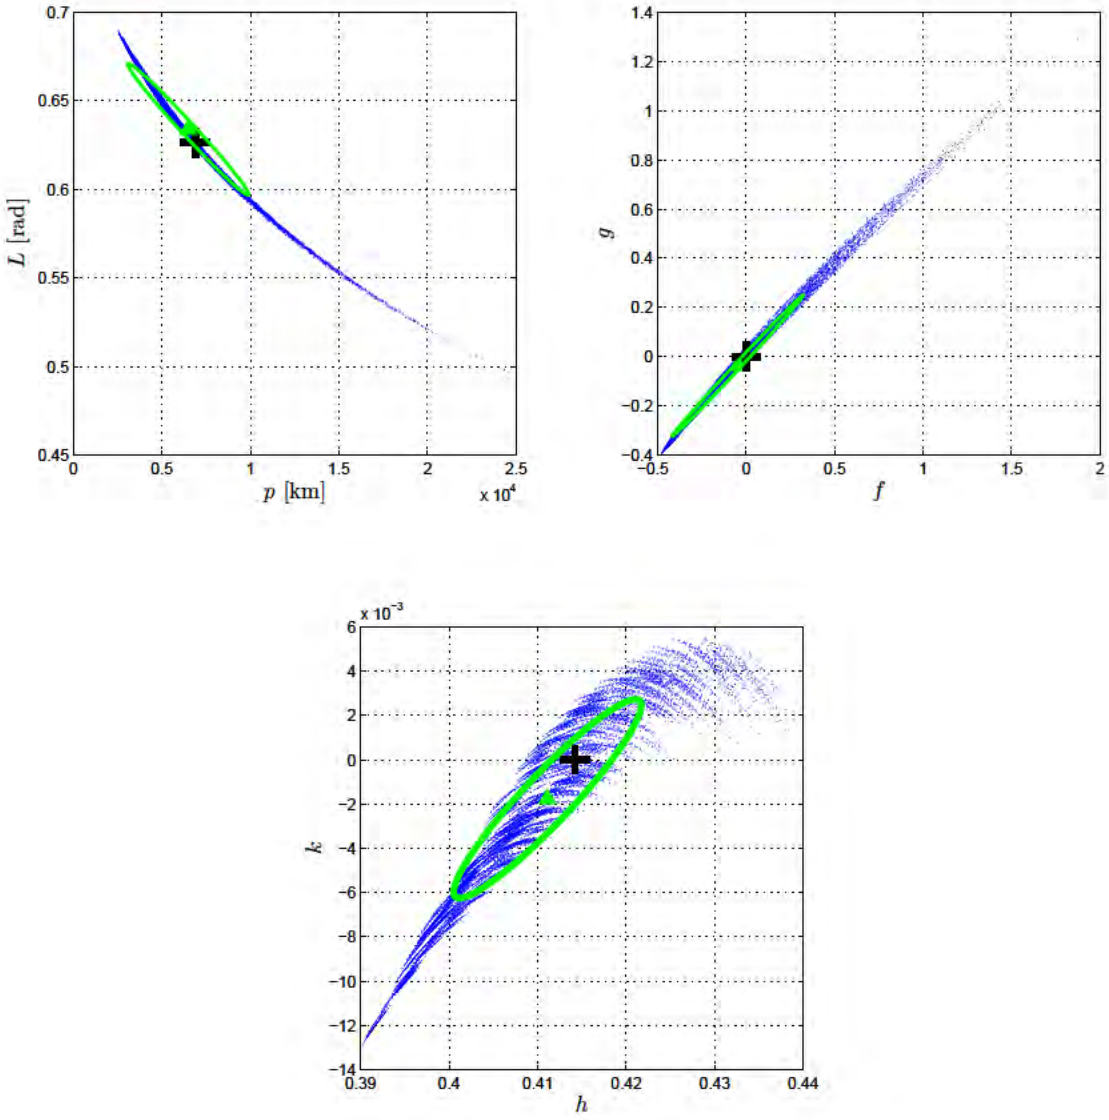

Figure 7. Plots showing IOD solution distributions in modified equinoctial element space (blue points), along with the estimate from the sigma points (green) and truth orbit (black cross).

Table 2. Actual error in estimated mean values from the Monte Carlo, least squares, and sigma point methods.

Method	δp [km]	δf	δg	δh	δk	δL [rad]
Monte Carlo	-813.99	-0.09369	-0.07366	-0.004473	-0.002288	0.01284
Least Squares	-1134.6	-0.1270	-0.09517	-0.003496	-0.001098	0.01444
Sigma Points	-386.72	-0.04688	-0.03690	-0.003104	-0.001751	0.007613

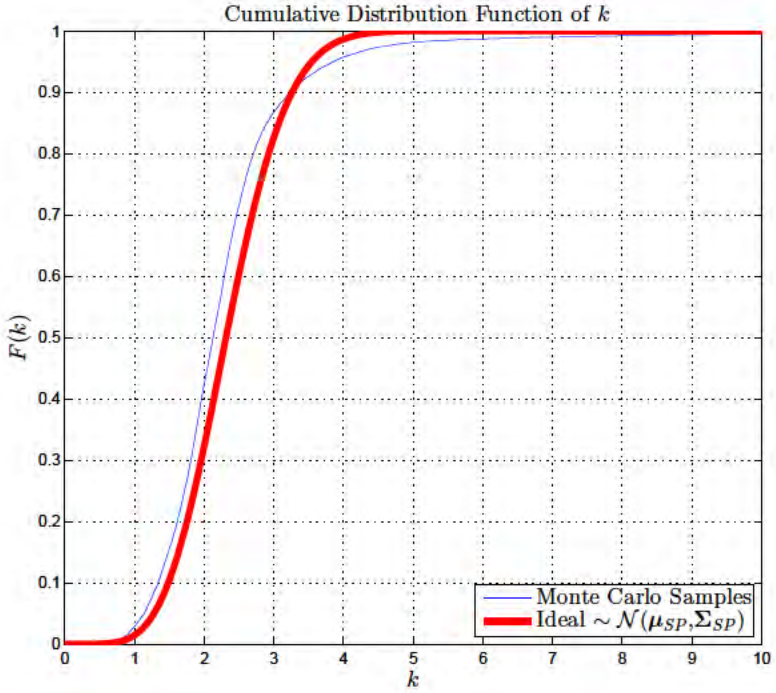

Figure 8. CDF of the Mahalanobis distance of Monte Carlo samples from the normally distributed estimate obtained with the sigma point method.

could lead to more accurate results once more measurements are processed.

Convergence

In this context, "convergence" means that the distribution in solution space is relatively compact, with well-defined structure, as in the previous figures. When the process diverges, the solution space tends to spread out much more. However, when zooming in, we see essentially the same structure as in the converged cases (see Figure 9). This suggests that if one makes the assumption that the observed object is an Earth-orbiting satellite, and filters the solution space with the bounds that follow this assumption, the methods presented here could yield useful results. However, this type of filtering introduces its own issues, particularly with our sigma point technique, as discussed in the Methodology section.

Through multiple simulations, we found that only cases with a sensor half angle field of view of 0.6° or less consistently converge. Cases with a field of view of $0.6 < \Theta \leqslant 1.0$ exhibit occasional convergence, and fields of view higher than 1° consistently diverge, albeit with the same sub-structure as shown in Figure 9.

CONCLUSIONS

This paper presents preliminary investigations into the solution space of an angles-only initial orbit determination problem. For a sensor with a uniform error distribution over its field of view (a binary sensor), Monte Carlo techniques were used to populate the solution space. In addition,

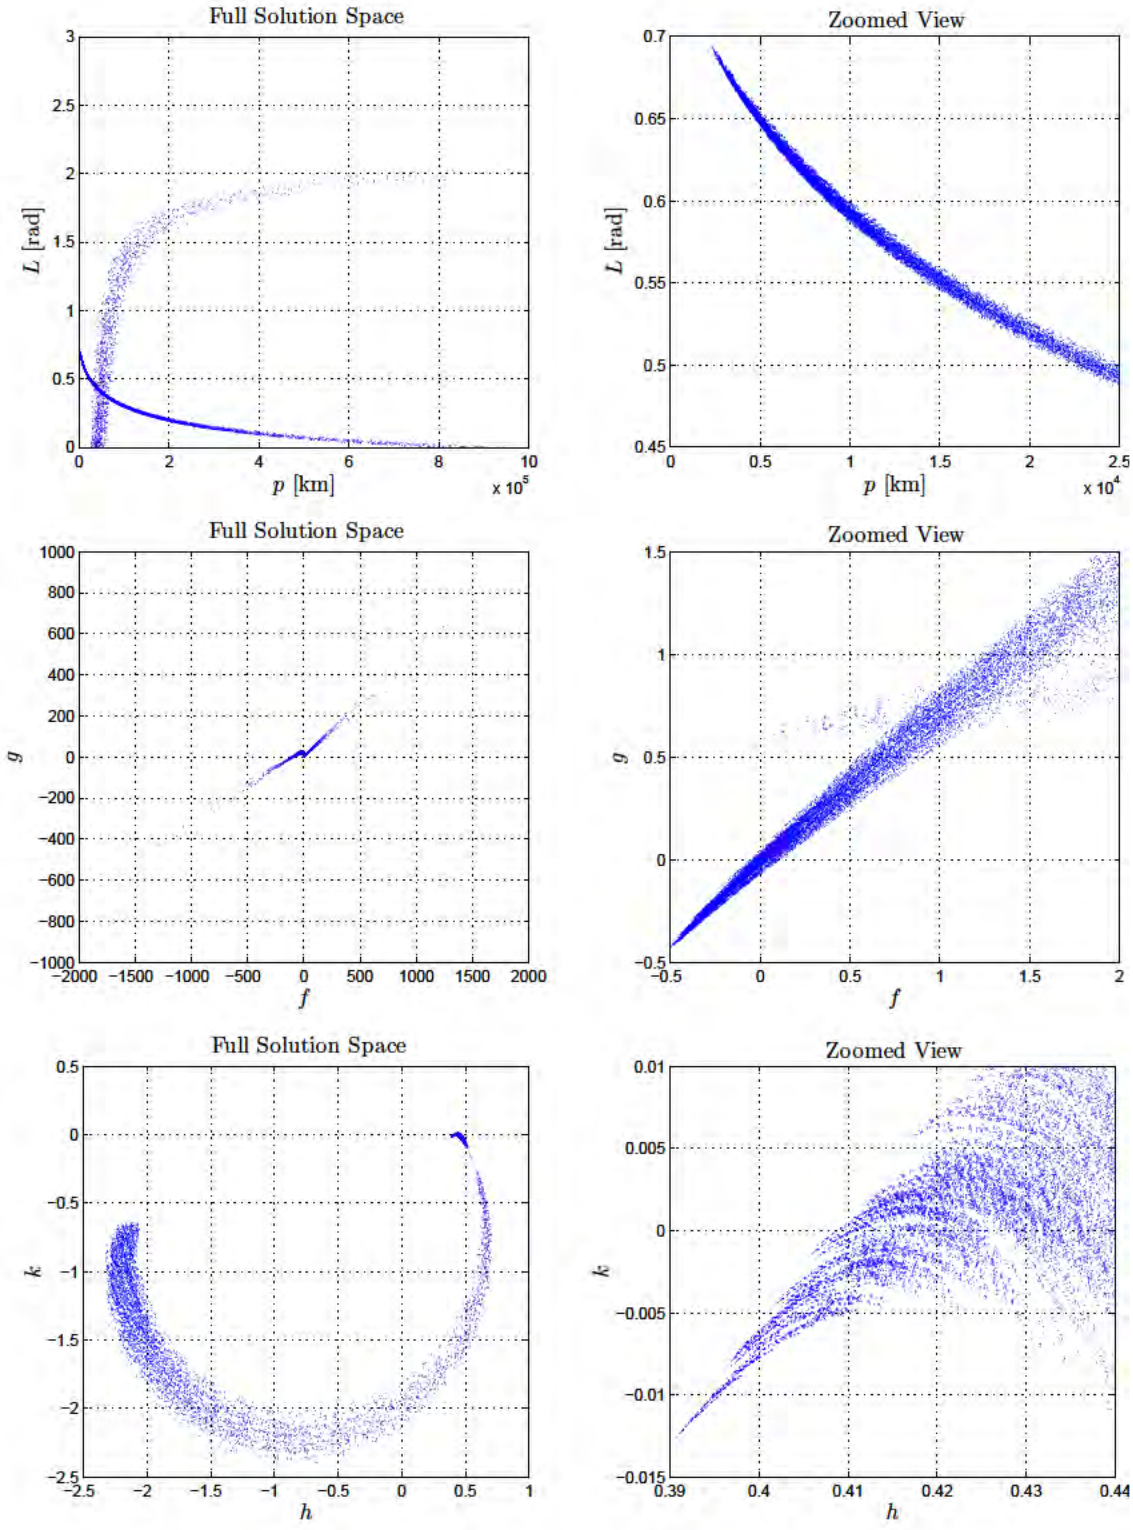

Figure 9. Solution space exhibiting divergence. This simulation was run with a sensor field of view of 1°.

a more traditional least squares estimation approach was used to calculate an orbit and associated uncertainty estimate. Finally, a sigma point technique based on the unscented transform was used to describe the uncertainty transformation from measurement space to IOD solution space.

The least squares technique performed relatively poorly, due in no small part to conflicting assumptions with respect to the underlying uncertainty distributions. There was significant error in the orbit estimate, and the resulting covariance did a poor job of matching or even containing the sample population distribution. The sigma point method, however, did a much better job describing both the orbit estimate and its associated uncertainty. This shows the utility of using the ideas of the unscented transform in a different way: instead of using sigma points to describe the nonlinear dynamics of a system, we can use them just as effectively to describe nonlinear measurement transformations.

Although the actual error in the estimates was still high (indicative of the sensitivity of Gauss's IOD method), the idea is that a well-defined initial estimate (paired with an equally informative uncertainty description) provides a better basis for further measurement processing. This would be particularly important for situations in which measurements are very sparse. Additionally, these results are promising for the concept of *Ubiquitous Low-Cost Sensors* in that they show that we can move toward useful solutions with poor-quality measurements.

FUTURE WORK

This paper performs only preliminary investigations into this problem, and there is a great amount of future work. The initial impetus for this paper was to see how useful of an IOD solution one could obtain given only binary optical measurements, and while the uniform measurement distribution produces interesting results, applying these techniques to a set of measurements with normally-distributed errors will perhaps be more realistic. Of course, Gauss's method for initial orbit determination is only one of several, and assuming the technique is amenable to acting as a "black box" transformation function, the sigma point methods presented here should adapt easily.

While the sigma point technique produced good results here, the relationship of the transformed points to the original set should be investigated in a more statistically rigorous fashion. Although it is similar to the methodology of the unscented transform/filter, it is a fundamentally different application, and the same relationships may not hold in all cases. The re-use of samples from each measurement could also introduce some undesirable correlation in solution space, and this should be investigated further.

The introduction of sigma points here could also be beneficial to a filtering scheme for processing more measurements. That is, instead of simply using the estimated mean and covariance, propagating the individual solution-space sigma points in a process not unlike the unscented Kalman filter (UKF) may lead to interesting results for IOD with sparse measurements. We expect to pursue this idea further in a future paper.

This paper presented only a single simulation case in the interest of brevity, but it is important to see how these issues change with different orbits and observation geometries. At the very least, a deeper investigation into the convergence behavior is warranted. On the topic of convergence, it is worth exploring the technique of filtering "inadmissible" solutions, particularly how this applies to the sigma point method. A valid filtering technique also opens the door for IOD with very wide field of view sensors, which dovetails with the ULCS concept.

Setting aside the obvious issue of observation association, an interesting application of this work

could be to look at overlapping distributions, say, from two different ground stations or simply two different sets of measurements from the same ground station. Even slight changes in the principal axes of the solution space could narrow down the number of possible solutions dramatically.

REFERENCES

[1] J. Sharma, G. H. Stokes, C. v. Braun, G. Zollinger, and A. J. Wiseman, "Toward operational space-based space surveillance," *Lincoln Laboratory Journal*, Vol. 13, No. 2, 2002, pp. 309–334.

[2] D. A. Vallado, *Fundamentals of Astrodynamics and Applications*. Hawthorne, CA: Microcosm Press, 3rd ed., 2007.

[3] P. R. Escobal, *Methods of Orbit Determination*. Huntington, NY: Krieger Publishing Co., Inc., 2nd ed., 1976.

[4] A. C. Long, J. O. Capellari, Jr., C. E. Velez, and A. J. Fuchs, "Goddard trajectory determination system (GTDS) mathematical theory, Revision 1," tech. rep., NASA Goddard Space Flight Center, 1989.

[5] A. Gelb, ed., *Applied Optimal Estimation*. MIT press, 1974.

[6] C. Sabol, T. Sukut, K. Hill, K. T. Alfriend, B. Wright, Y. Li, and P. Schumacher, "Linearized orbit covariance generation and propagation analysis via simple Monte Carlo simulations," *Spaceflight Mechanics 2010*, Vol. 136 of *Advances in the Astronautical Sciences*, San Diego, CA, Univelt, Inc., 2010. AAS 10-134.

[7] J. L. Junkins, M. R. Akella, and K. T. Alfriend, "Non-Gaussian error propagation in orbital mechanics," *Journal of the Astronautical Sciences*, Vol. 44, No. 4, 1996, pp. 541–563.

[8] L. M. Healy, "Lambert targeting for on-orbit delivery of debris remediation dust," *Spaceflight Mechanics 2012*, Vol. 143 of *Advances in the Astronautical Sciences*, San Diego, CA, Univelt, Inc., 2012. AAS 12-249.

[9] I. Sobol, "Uniformly distributed sequences with an additional uniform property," *USSR Computational Mathematics and Mathematical Physics*, Vol. 16, No. 5, 1976, pp. 236–242.

[10] M. Walker, B. Ireland, and J. Owens, "A set of modified equinoctial orbit elements," *Celestial Mechanics*, Vol. 36, No. 4, 1985, pp. 409–419. doi:10.1007/BF01227493.

[11] S. Julier and J. Uhlmann, "New extension of the Kalman filter to nonlinear systems," *Signal Processing, Sensor Fusion, and Target Recognition VI*, Vol. 3068 of *Proc. SPIE*, July 1997. doi:10.1117/12.280797.

[12] L. Angrisani, P. D'Apuzzo, and R. Schiano Lo Moriello, "Unscented transform: a powerful tool for measurement uncertainty evaluation," *Instrumentation and Measurement, IEEE Transactions on*, Vol. 55, No. 3, 2006, pp. 737–743, 10.1109/TIM.2006.873811.

[13] J. Hartikainen, A. Solin, and S. Särkkä, *Optimal filtering with Kalman filters and smoothers*, 2011. `http://becs.aalto.fi/en/research/bayes/ekfukf/documentation.pdf`.

[14] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA: Artech House, Inc., 1st ed., 1999.