ANALYSIS OF RELATIVE MERITS OF UNSCENTED AND EXTENDED KALMAN FILTERS IN ORBIT DETERMINATION

James Woodburn,* Vincent Coppola†

The unscented and extended forms of the Kalman filter are compared in the context of orbit determination. Differences and similarities of the algorithms are identified with an emphasis on treatment of state-error uncertainty in the presence of uncertainty in the dynamics and measurement models. A hybrid filter which combines elements of both algorithms is proposed in search of an optimal combination of computational performance and accommodation of higher order effects during measurement processing. The filter variants are compared based on application to set of realistic orbit determination scenarios.

INTRODUCTION

The Extended Kalman Filter (EKF) has a long history in operational orbit determination. Satellite systems such as NASA's Tracking and Data Relay Satellites (TDRS) and the US Air Force's Global Positioning System (GPS) satellites use EKFs in daily operations. The Unscented Kalman Filter (UKF) is a more recently developed estimation algorithm that has been documented to have advantages over the EKF in a number of estimation applications including orbit determination[1234]. The advantages of using the UKF as compared to the EKF are attributed to the inclusion of higher order moments in the state error distribution. The UKF should generate a better description of the state-error and measurement-error distribution functions relative to the EKF when those error-distributions are non-Gaussian.

In this investigation, we take a systematic approach to investigate the differences in the EKF and UKF algorithms and evaluate the performance of both algorithms across a variety of realistic orbit determination scenarios. We emphasize that our study is specific to the estimation of orbits and impose restrictions on algorithm implementation consistent with the needs of operational orbit determination. Numerical comparisons will be performed using operational orbit determination software with realistic force and measurement models. Our goal is to produce a practical guide to indicate situations where one algorithm or the other may be preferred.

As a precursor to numerical comparisons, we start with a description of a generic filter recursion followed by the basic processing equations for the EKF and UKF algorithms. The equations are presented in a manner that helps identify similarities and differences between the two algorithms. We also provide a discussion of real-world challenges, such as the inclusion of appropriate dynamical process noise, that are critical to the successful use of sequential estimation processes in operations.

* Chief Orbital Scientist, Analytical Graphics Inc (AGI), woodburn@agi.com.
† Director, Technical Project Management, Analytical Graphics Inc (AGI), vcoppola@agi.com.

Our quantitative analysis begins with a common orbit determination scenario, where initial Gaussian error distributions remain Gaussian throughout the estimation of an updated orbit solution. This "easy" case is intended to demonstrate the equivalence of results obtained from the EKF and UKF algorithms. We then extend the analysis to several problems of practical concern where higher-order effects in the dynamics and/or measurement models become important. For each situation, we will evaluate where non-linear effects are most prevalent, and the effectiveness of each algorithm.

CONTEXT OF OPERATIONAL ORBIT DETERMINATION

In his description of optimal orbit determination, Wright lays out a number of tenets related to accommodating uncertainty in the orbital dynamics during orbit estimation[5]. While the application of these tenets in Wright's work was realized in the form of an EKF, most of the principles are equally relevant to the UKF. For the convenience of the reader, we repeat the subset of those tenets relevant to our current work here:

- Sequential processing is used to account for force modeling errors and measurement information in the time order in which they are realized.
- All state estimate models and state estimate model error representations are derived from appropriate force modeling physics and measurement sensor performance.
- All measurement models and measurement model error approximations are derived from appropriate sensor hardware definitions and associated physics, and measurement sensor performance.
- The state estimation structure is complete.

Both the EKF and UKF are sequential filters and therefore at least partially satisfy the first tenet. There is subtle meaning, however, that imposes a further restriction on the use of the algorithms. The requirement to account for force modeling errors in a temporally local manner means that it is not sufficient to wait until the end of a large measurement gap to apply the effects of dynamical process noise. Wright's tenets are based on the idea of appropriately connecting process noise to the physical processes which are the source of the uncertainty. In practice, this philosophy leads to filter implementations requiring minimal operator intervention after initial calibration. Long time update periods are divided into a series of smaller time updates with process noise added to the state-error covariance matrix after each small update. This allows the incremental additions of uncertainty to accumulate according to the physics of the problem over the remainder of the time update interval. The use of physically-connected process noise models is central to the problem of producing realistic state-error covariance as a product of orbit determination[6].

The second and third tenets above motivate the type of model used to represent uncertainty in physical models for orbit evolution and measurement representation. For example, most random parameters have time varying values but the evolution process is either unknown or lacks sufficient input information to be accurately modeled. The known physics, however, dictate how variations in the unknown parameter affect the model, the expected range of values for the parameter and how quickly the value might change. In such cases, the application of appropriately configured stochastic processes is a reasonable approach.

The final tenet above reflects the desire to estimate all unknown elements of the dynamics and measurement models based on the importance of a complete state space to achieving realistic covariance. The argument for the use of a complete state space is simple, to have an accurate model of the uncertainty in the output of a process, we need to properly characterize the uncertainty of the inputs and of the process itself. Yet, this tenet is often not followed operationally. Reasons for

disregard include concerns about observability in least squares algorithms and concerns about computational performance and the size of outputs in filters. We include completeness of state space in our discussion to since it is relevant to the computational performance of the filters, most notably the UKF.

Stochastic sequences

The use of stochastic sequences to model uncertain parameters in dynamical and measurement models is a common practice in sequential orbit determination. With selection of an appropriate stochastic model and proper calibration, the estimation process can be configured to add state-error uncertainty in a manner consistent with the associated physics in terms of timing, size of effect and rate of change[7]. For example, we might select a random-walk process to emulate the drift of a free-running clock since, while the rate of drift may be bounded, the accumulated drift is not. Conversely, we choose sequences with bounded uncertainty to model stochastic force modeling parameters such as a solar pressure or drag coefficient. For the purpose of this study, we selected to use an exponentially-correlated Gauss-Markov sequence with exponential decay to model unknown force and measurement model parameters[8].

The stochastic sequence is generated using the following formula during simulation of the random variable, V,

$$V_{k+1} = \Phi(t_{k+1}, t_k)V_k + \sqrt{\left(1 - \Phi^2(t_{k+1}, t_k)\right)}\, w(t_k).$$

The definitive time transition is defined as an exponential decay given as,

$$\Phi(t_{k+1}, t_k) = e^{\alpha\lfloor t_{k+1} - t_k \rfloor},$$

where α is the diffusion coefficient, which is set by the analyst to specify the volatility of the sequence. Note that α is typically set via the specification of a half-life representing the time required for the estimate to decay to half of its current value. The Gaussian white-noise random variable, w, is configured to have an analyst selected variance of σ_w^2,

$$E[w(t_k)w(t_k)] = \sigma_w^2.$$

Requirements for implementation in a filter process are completed with the identification of the process noise function,

$$\mathrm{Q}(t_{k+1}, t_k) = \left(1 - \Phi^2(t_{k+1}, t_k)\right)\sigma_w^2.$$

We note that this particular stochastic sequence is appropriate for modeling physical parameters where the mean, range of variation and time scale of variation are relatively well known. These criteria are easily met inside our simulation framework.

ALGORITHM REVIEW AND COMPARISON

In order to facilitate an understanding of the differences in the results obtained from the EKF and the UKF, it is useful to review a description of the two filter algorithms. In general, filtering algorithms can be described as a recursive machine of time and measurement updates that serve to generate estimates of an unknown state with an associated description of the error associated with that state estimate. Time updates serve to propagate the state and associated state-error estimates through time. Measurement updates fold in observational information to improve the state estimate. The generic flow and structure of a filtering algorithm is depicted in Figure 1. Note that the overall

time update is depicted as the composite of multiple sub-steps and the overall measurement update is depicted as processing multiple measurements at the same epoch.

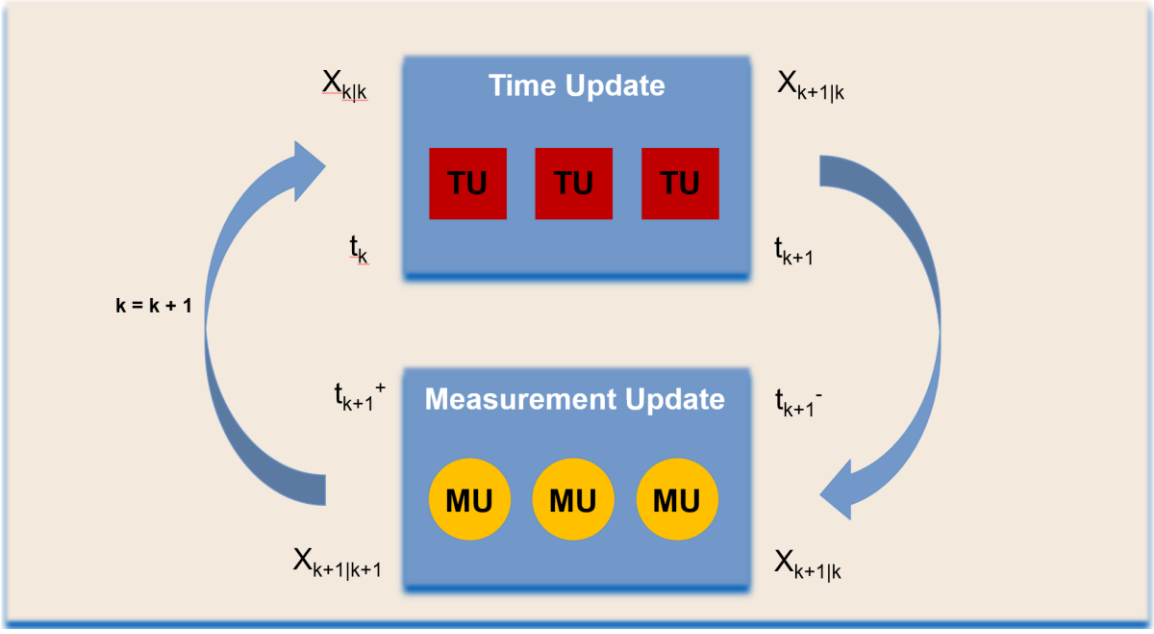

Figure 1. Generic Filter Recursion Diagram.

The time transition operation of the time update moves the state estimate forward to the time of the next measurement, to the time of the next node on a processing time grid, to the time of an event, such as a maneuver, or to any other time as required by the filter process. A time update may be followed by a measurement update, another time update or some type of specialized processing associated with an event. Because the time transition model is imperfect or incomplete due to force modeling uncertainty and approximations, the time update must also include appropriate dynamical process noise to accommodate the increase in the uncertainty of the state across a time transition.

The inclusion of measurement information produces a discontinuous change in the state estimate and the associated description of uncertainty. The uncertainty associated with elements of the state observable from a particular measurement decreases with the inclusion of the measurement information. State elements may be directly observable, where the measurement model formulation includes the state element explicitly, or indirectly observable through correlation with one or more elements of the state that are directly observable. It is common for more than one measurement to be processed at a given point in time. In these cases, the analyst or the algorithm designer elects to process the multiple measurements at a common epoch sequentially, simultaneously or via a custom processing scheme. Sequential processing uses the *a posteriori* state from each processed measurement at the common epoch as the *a priori* state for processing the next measurement. In simultaneous processing, also referred to as vector measurement processing, the same *a priori* state is used for processing all measurements at the common epoch and a single, combined update is made to the state and state-error uncertainty.

An important part of the measurement update procedure is the determination of measurement acceptance or rejection. The rejection of "bad" observation data is critical to avoid corruption of the estimation state. For example, the acceptance of bad measurements can move the state estimate

away from the true state while reducing the state-error uncertainty in a manner that results in the rejection of all future observations, a condition known as filter divergence.

In the case of orbit determination, a typical estimation state contains a number of sub-states including six elements describing the orbit, parameters associated with the modeling of forces on the satellite, and parameters associated with taking measurements of the satellite. A simple estimation state might include the position and velocity of the satellite, the ballistic coefficient, the solar pressure coefficient, a transponder delay, and biases for range, azimuth and elevation measurements taken from three ground stations. In this example, as shown in Table 1, the state space contains 18 elements. It should be noted that the size of the state space can grow significantly when the orbits of multiple satellites are estimated simultaneously, a larger network of tracking stations is involved, or more detailed force and measurement models with additional unknown parameters are used.

Table 1. Example Orbit Determination State Content.

State Element	
Position X	Station 1 Azimuth Bias
Position Y	Station 1 Elevation Bias
Position Z	Station 1 Range Bias
Velocity X	Station 2 Azimuth Bias
Velocity Y	Station 2 Elevation Bias
Velocity Z	Station 2 Range Bias
Solar Pressure Coefficient	Station 3 Azimuth Bias
Drag Coefficient	Station 3 Elevation Bias
Transponder Delay	Station 3 Range Bias

Common definitions

The orbit determination process produces an estimate, $\hat{x}$, of an unknown state, x, which includes the orbit of the satellite plus other parameters. The state estimate evolves according to a set of non-linear, ordinary differential equations whose solution is given as

$$\hat{x}_{k+1|k} = X\left(t_k, x_{k|k}, t_{k+1}\right). \tag{1}$$

The subscript on the state, $k+1|k$, denotes two discrete time references. The index prior to the vertical separator denotes the time of the current state and the index after the vertical separator denotes the time of the last processed measurement. Equation (1), therefore, represents the computation of the state at time t_{k+1} as a function of the state at time t_k, where the state at time t_k includes all measurements up to and including those at time t_k. The orbit determination process requires measurements, y, to improve the estimate of the orbit. Measurement values are a function of the estimated state. The measurement model representation to compute an expected value of the measurements, $\hat{y}$, from the current state estimate is given as

$$\hat{y}_{k+1} = Y\left(t_{k+1}, \hat{x}_{k+1|k}\right). \tag{2}$$

Note that in the case of measurement notation, the subscript serves to indicate the discrete time index of the measurement. Measurement residuals, also referred to as innovations, are computed as the difference between the actual and the expected values of the measurement,

$$\Delta y_{k+1} = y - \hat{y}_{k+1}. \tag{3}$$

Extended Kalman Filter

The EKF is a well-known adaptation of the linear Kalman Filter algorithm for use on problems that include non-linear state dynamics, non-linear measurement modeling or both. The EKF algorithm requires an *a priori* state estimate and associated state-error covariance as initial conditions. The state-error uncertainty in the EKF algorithm is described by the estimate of the mean and the state-error covariance, the first two moments of the uncertainty distribution. In the case of a Gaussian error distribution, knowledge of the first two moments is sufficient to provide a complete description of the distribution.

In the EKF time update, the state estimate is propagated forward using the non-linear dynamical model, Equation 1, and the corresponding state-error covariance is propagated using the linear state-error transition matrix as shown in Equation 4.

$$P_{k+1|k} = \Phi(t_{k+1}, t_k) P_{k|k} \Phi^T(t_{k+1}, t_k) + \mathrm{Q}(t_{k+1}, t_k). \tag{4}$$

The state-error transition matrix, $\Phi(t_{k+1}, t_k)$, is developed using a linearized approximation of motion in the vicinity of the current state estimate and requires the computation of the Jacobian of the state derivative with respect to the state. The definitive propagation of the covariance matrix, P, is augmented by the addition of the process noise matrix, $\mathrm{Q}(t_{k+1}, t_k)$, which accommodates deficiencies in the dynamical model and is essential to the effective operation of the EKF. The computation and application of $\mathrm{Q}(t_{k+1}, t_k)$ is discussed in more detail in the sequel.

During the EKF measurement update, the expected value of the measurement is computed using the non-linear measurement representation, Equation 2, based on the current state estimate. The measurement-error covariance, $\tilde{R}_{k+1}$, is computed as the sum of the linear mapping of the state-error covariance to measurement coordinates and the measurement white noise variance, R_{k+1},

$$H_{k+1} = \frac{\partial Y(t_{k+1}, \hat{x}_{k+1|k})}{\partial \hat{x}_{k+1|k}}, \tag{5}$$

$$\tilde{R}_{k+1} = H_{k+1} P_{k+1|k} {H_{k+1}}^T + R_{k+1}. \tag{6}$$

Mapping of state-error covariance into measurement coordinates requires the computation of the Jacobian, H, of the measurement representation with respect to the state. The Kalman gain, K_{k+1}, is computed and used in the update equations for the state and state-error covariance as

$$K_{k+1} = P_{k+1|k} H_{k+1}^T \tilde{R}_{k+1}^{-1}, \tag{7}$$

$$\hat{x}_{k+1|k+1} = \hat{x}_{k+1|k} + K_{k+1} \Delta y_{k+1}. \tag{8}$$

$$P_{k+1|k+1} = (I - K_{k+1} H_{k+1}) P_{k+1|k}. \tag{9}$$

Augmentation of the EKF for the purpose of improved performance in the presence of higher order effects is a standard practice for applications where such effects are expected. Common mitigation strategies include the use corrective terms associated with the second-order Gaussian filter as well as the application of measurement underweighting which applies an analyst specified limit to reduction of the state-error covariance by a single measurement[9]. The second-order Gaussian filter requires the computation of Hessian matrices of the dynamics and measurement models[10].

Unscented Kalman Filter

The UKF algorithm uses a judiciously selected set of sample points, often called sigma points, to represent the uncertainty distribution of the state. The transition of the state estimate across time and the evaluation of measurement representations are treated as coordinate transformations accomplished using the unscented transform procedure. The unscented transform uses only non-linear evaluations of the state propagation and measurement models and therefore does not require the evaluation of Jacobians. State-error uncertainty in the UKF algorithm is expressed as the estimate of the mean and the state-error covariance. At some steps in the processing algorithm, the state-error and measurement uncertainty are described by a collection of samples, which can contain information on higher order moments, but this sample-based representation is not maintained across the entirety of the algorithm. For example, the UKF algorithm requires an *a priori* state estimate and associated state-error covariance as initial conditions but replaces the state-error covariance with a set of sample points for propagation over the time update interval. Samples are used again to model the measurement, but the result of the measurement update is an updated state estimate and state-error covariance.

A familiarity with the unscented transform is needed to understand the UKF. We will use the scaled variant of the unscented transform that was developed to overcome numerical and localization issues associated with the sigma point selection method in the standard unscented transform[4]. The scaled, unscented transform begins with a mean, $\bar{x}$, and covariance, P_{xx}, in a particular set of coordinates. The mean and covariance are replaced by a collection of samples including the mean and sigma points that have the same statistics as the original mean and covariance. These sigma points are computed as scaled column vectors of the lower-triangular Cholesky decomposition of the covariance matrix. The scale factor used in the construction of the sigma points is a function of state size, n, and two user-selectable parameters: κ, which is a free parameter from the standard unscented transform, and α, which is a scaling control introduced to retain locality in the sigma points. The scaling control, α, is set to a small number to reduce the distance between the sigma points and the mean of the distribution. An additional free parameter, β, is included in the weight of the mean sample when used in the computation of the transformed covariance. The value of β may be chosen to affect the matching of higher order moments where $(\beta = 2)$ is optimal for Gaussian distributions. A convenience parameter, λ, may be defined to allow for a presentation of the scaled algorithm that mirrors the standard unscented transform[11]. The procedure for using the scaled unscented transform to convert a mean and covariance from the set of coordinates, x, to a new mean and covariance in another set of coordinates, y, where x and y are related by the non-linear transformation defined by $y = F(x)$, is as follows:

$$\lambda = \alpha^2(n + \kappa) - n \tag{10}$$

$$\chi_0 = \bar{x} \tag{11}$$

$$\chi_i = \bar{x} + \left(\sqrt{(n + \lambda)P_{xx}}\right)_i \quad i = 1, \dots, n \tag{12}$$

$$\chi_i = \bar{x} - \left(\sqrt{(n+\lambda)P_{xx}}\right)_i \quad i = n+1, \ldots, 2n \tag{13}$$

$$\Upsilon_i = F(\chi_i) \tag{14}$$

$$W_0^m = \lambda/(n+\lambda) \tag{15}$$

$$W_i^m = W_i^c = 1/\{2(n+\lambda)\} \quad i = 1, \ldots, 2n \tag{16}$$

$$W_0^c = \lambda/(n+\lambda) + (1-\alpha^2+\beta) \tag{17}$$

$$\bar{y} = \sum_{i=0}^{2n} W_i^m \Upsilon_i \tag{18}$$

$$P_{yy} = \sum_{i=0}^{2n} W_i^c (\Upsilon_i - \bar{y})(\Upsilon_i - \bar{y})^T \tag{19}$$

As noted above, the UKF algorithm leverages the unscented transform in both the time and measurement steps of filter recursion. In the time update, the UKF propagation of the mean estimate follows the scaled unscented transform to create a transformed mean state at the end time of the propagation. Each sample of the state distribution at time t_k is propagated using the non-linear dynamical model given in Equation 1 to time t_{k+1}. The mean estimate is then formed as the weighted sum of the transformed samples as shown in Equation 23. The time transition of the state-error covariance is performed as the sum of the outer products of the differences between the samples and the mean as shown in Equation 24, but with the addition of process noise, Q, to account for uncertainty in the non-linear propagation model.

$$\chi_k^0 = \hat{x}_{k|k} \tag{20}$$

$$\chi_k^i = \hat{x}_{k|k} \pm \left(\sqrt{(n+\lambda)P_{k|k}}\right)_i \qquad \begin{matrix} i = 1, \ldots, 2n \\ i = 1, \ldots, n \text{ uses} + \\ i = n+1, \ldots, 2n \text{ uses} - \end{matrix} \tag{21}$$

$$\chi_{k+1}^i = X\left(t_k, \chi_k^i, t_{k+1}\right) \tag{22}$$

$$\hat{x}_{k+1|k} = \sum_{i=0}^{2n} W_i^m \chi_{k+1}^i \tag{23}$$

$$P_{k+1|k} = \sum_i W_i^c \left(\left(\chi_{k+1}^i - \hat{x}_{k+1|k}\right)\left(\chi_{k+1}^i - \hat{x}_{k+1|k}\right)^T \right) + \mathrm{Q}(t_{k+1}, t_k) \tag{24}$$

In the general UKF formulation, the transformed state samples, χ_{k+1}^i,may then be used directly in measurement processing or formed into a sample mean and covariance, where the new covariance is resampled prior to measurement processing. Inclusion of process noise as described in Equation 24 dictates that our UKF implementation use the latter strategy of re-sampling an updated state-error covariance. The reasons for our selection of the formulation shown above for the inclusion of process noise in the UKF are discussed in the sequel.

During measurement processing, the scaled unscented transform is applied to compute the mean and error-covariance of the expected measurement. Each state sample is used to generate a measurement sample. The sample mean and error-covariance are then computed from the set of measurement samples. The cross-covariance between the state and measurement coordinates is also computed for use in the Kalman gain computation. Finally, the Kalman measurement update is performed to generate the updated state and state-error covariance. The UKF measurement update sequence is given as

$$\chi_{k+1}^{0} = \hat{x}_{k+1|k} \tag{25}$$

$$\chi_{k+1}^{i} = \hat{x}_{k+1|k} \pm \left(\sqrt{(n+\lambda)P_{k+1|k}}\right)_i \qquad \begin{matrix} i = 1, \dots, 2n \\ i = 1, \dots, n \ uses + \\ i = n+1, \dots, 2n \ uses - \end{matrix} \tag{26}$$

$$\Upsilon_{k+1}^{i} = Y\left(t_{k+1},\, \chi_{k+1}^{i}\right) \quad i = 0, \dots, 2n \tag{27}$$

$$\hat{y}_{k+1} = \sum_{i=0}^{2n} W_i^m \Upsilon_{k+1}^{i} \tag{28}$$

$$\tilde{R}_{k+1} = \sum_{i=0}^{2n} W_i^c \left(\left(\Upsilon_{k+1}^{i} - \hat{y}_{k+1} \right) \left(\Upsilon_{k+1}^{i} - \hat{y}_{k+1} \right)^T \right) + \mathrm{R}_{k+1} \tag{29}$$

$$P_{k+1}^{xy} = \sum_{i=0}^{2n} W_i^c \left(\left(\chi_{k+1}^{i} - \hat{x}_{k+1|k} \right) \left(\Upsilon_{k+1}^{i} - \hat{y}_{k+1} \right)^T \right) \tag{30}$$

$$K_{k+1} = P_{k+1}^{xy} \tilde{R}_{k+1}^{-1} \tag{31}$$

$$\hat{x}_{k+1|k+1} = \hat{x}_{k+1|k} + K_{k+1} \Delta y_{k+1} \tag{32}$$

$$P_{k+1|k+1} = P_{k+1|k} - K_{k+1} \tilde{R}_{k+1} K_{k+1}^T \tag{33}$$

Process noise inclusion in the UKF

Literature on the UKF offers two options for the inclusion of process noise: a purely additive approach where a process noise matrix is added to the covariance determined from definitive propagation, or augmentation of the state vector to include process noise states. The latter option, while preferred by some authors since it avoids a re-sampling operation to create new sigma points after the covariance is updated, appears to only be applicable to cases where process noise can be applied in a purely additive sense at the end of a time update without concern for the method of uncertainty evolution over the time update interval. The state vector augmentation approach does not alter the trajectories of the sigma points, it merely utilizes the mechanics of the unscented transform to affect the state error covariance to avoid an additive operation that necessitates re-sampling of the sigma points[12].

The use of stochastic sequences within the dynamical model establishes a physical-connection in the models of dynamical uncertainty, linking the inclusion of process noise to its effect on the orbit and allowing for the selection of time scales appropriate to individual physical phenomena.

Sigma points, constructed at the singular epoch at the start of a time update and propagated using a definitive dynamical model, cannot adequately emulate the increase in uncertainty over time associated with such random parameters. Proper evolution of the state error covariance requires that process noise effects be included in a temporally and spatially local sense to the source of the dynamical uncertainty.

Finally, we note that, in addition to process noise associated directly with state entries, it is often desirable to include process noise to accommodate inaccuracies in dynamical models that cannot be estimated in a practical manner in state space. One such example is the accommodation of errors in the central body gravity field[13]. Under this circumstance, reforming the covariance and adding Q as shown in Equation 24 is required to retain proper physical and temporal connection of the dynamic modeling errors. We therefore find the augmented state vector approach to process noise in the UKF to be incompatible with the tenets of optimal orbit determination. It can be argued, however, that in cases of large state uncertainty, the proper handling of process noise is a minor consideration in the overall evolution of the orbit-error uncertainty. For this reason, we will evaluate both options in the UKF time update: breaking long time updates into a series of shorter time updates accommodating temporally correct process noise inclusion and mapping all process noise effects to the end of the time update interval for inclusion at the final time.

Algorithm comparison

The EKF and UKF both use the first (mean) and second (covariance) order moments of the state-error distribution to represent the state uncertainty. The UKF will also generate information describing higher moments during portions of the processing algorithm. The measurement update operation, however, produces an updated mean and state-error covariance without higher order moment information. The state-error covariance is resampled to start the next iteration of the recursion. A brief summary of the most apparent algorithmic characteristics of the EKF and UFK is given in Table 2.

Table 2. Algorithm Comparison.

	EKF	UFK
Initial conditions	Mean/Cov	Mean/Cov
Non-linear propagations in time update	1	2n+1
Measurement evaluations per observation	1	2n+1
State transition Jacobian	Yes	No
Measurement Jacobian	Yes	No
Computational load	Moderate	High
Implementation complexity	Moderate	Low
Higher order moments	No	Partial
Preferred coordinates	Yes	Difficult

The EKF can be implemented using a variety of state coordinate options. It has been shown, for example, that the use of orbit elements in the state space of the EKF has advantages in terms of estimation accuracy and convergence when compared to an implementation in Cartesian coordinates[14]. The main benefit of preferred coordinate selection in state space is retention of a state-error

distribution that is Gaussian or nearly Gaussian as uncertainty increases. In theory, the UKF could also be implemented utilizing orbit elements in state space, but the addition and subtraction operations required in the construction of sigma points and in the reconstruction of a transformed mean and covariance can be problematic for state elements with restricted domains or specialized addition rules. For example, a deviation in the eccentricity of the orbit cannot be allowed to result in a negative value and the subtraction of mean anomaly values must be cognizant of the fact that the difference between 180 and -180 degrees is actually zero.

Comparisons of the equations used in the time and measurement updates of the EKF and UKF are given in Figures 2-3. It is clear from an examination of the equations that the UKF will require considerable more computational resource than the EKF when the non-linear processes of state propagation and measurement modeling are complex and as state size becomes large. We note that the need for (2n+1) non-linear state propagations during each time update is a strong motivator to ignore the tenet of complete state space, but doing so would significantly reduce the likelihood of achieving a realistic error-covariance. The additional complexity in the implementation of the EKF, due to the need for partial derivatives, is less obvious, but is still a relevant comparison point.

EKF

$$\hat{x}_{k+1|k} = X(t_k, \hat{x}_{k|k}, t_{k+1})$$

$$P_{k+1|k} = \Phi(t_{k+1}, t_k) P_{k|k} \Phi^T(t_{k+1}, t_k) + \mathrm{Q}(t_{k+1}, t_k)$$

UKF

$$\chi_k^0 = \hat{x}_{k|k}$$

$$\chi_k^i = \hat{x}_{k|k} \pm \left(\sqrt{(n+\lambda)P_{k|k}}\right)_{i,i+n}$$

$$\chi_{k+1}^i = X(t_k, \chi_k^i, t_{k+1})$$

$$\hat{x}_{k+1|k} = \sum_i W_i^m \chi_{k+1}^i$$

$$P_{k+1|k} = \sum_i W_i^c \left((\chi_{k+1}^i - \hat{x}_{k+1|k})(\chi_{k+1}^i - \hat{x}_{k+1|k})^T \right) + \mathrm{Q}(t_{k+1}, t_k)$$

Figure 2. Time Update Operations Comparison.

EKF

$$\hat{y}_{k+1} = Y(t_{k+1}, \hat{x}_{k+1|k})$$

$$\Delta y_{k+1} = y - \hat{y}_{k+1}$$

$$H_{k+1} = \frac{\partial Y(t_{k+1}, \hat{x}_{k+1|k})}{\partial \hat{x}_{k+1|k}}$$

$$\tilde{R}_{k+1} = H_{k+1} P_{k+1|k} {H_{k+1}}^T + R_{k+1}$$

$$K_{k+1} = P_{k+1|k} H_{k+1}^T \tilde{R}_{k+1}^{-1}$$

$$\hat{x}_{k+1|k+1} = \hat{x}_{k+1|k} + K_{k+1} \Delta y_{k+1}$$

$$P_{k+1|k+1} = (I - K_{k+1} H_{k+1}) P_{k+1|k}$$

UKF

$$\Upsilon_{k+1}^i = Y(t_{k+1}, \chi_{k+1}^i)$$

$$\hat{y}_{k+1} = \sum_i W_i^m \Upsilon_{k+1}^i$$

$$\tilde{R}_{k+1} = \sum_i W_i^c \left((\Upsilon_{k+1}^i - \hat{y}_{k+1})(\Upsilon_{k+1}^i - \hat{y}_{k+1})^T \right) + \mathrm{R}_{k+1}$$

$$P_{k+1}^{xy} = \sum_i W_i^c \left((\chi_{k+1}^i - \hat{x}_{k+1|k})(\Upsilon_{k+1}^i - \hat{y}_{k+1})^T \right)$$

$$K_{k+1} = P_{k+1}^{xy} \tilde{R}_{k+1}^{-1}$$

$$\hat{x}_{k+1|k+1} = \hat{x}_{k+1|k} + K_{k+1} \Delta y_{k+1}$$

$$P_{k+1|k+1} = P_{k+1|k} - K_{k+1} \tilde{R}_{k+1} K_{k+1}^T$$

Figure 3. Measurement Update Operations Comparison.

Hybrid formulation

It has been seen in a prior study that, when represented in advantageous coordinates, the state error uncertainty distribution can remain Gaussian or near Gaussian even in the case of large uncertainties while mapping of uncertainty to measurement space is problematic[15]. Figures 4-5 serve to illustrate the difference in covariance realism between state and measurement coordinates in a scenario where the state uncertainty is very large. The magenta dots represent samples of an originally Gaussian distribution which have been propagated for sufficient time to illustrate the curvature of the sample points along the nominal orbit. The yellow ellipsoid, which looks like a line in Figure 5, depicts the linear mapping of the error covariance into measurement coordinates. Unfortunately, the yellow ellipsoid depicting the linear propagation of the covariance in Cartesian coordinates is hidden by the green ellipsoid in Figure 4. The green ellipsoids depict the covariance resulting from the unscented transform in Cartesian and measurement coordinates in Figures 4 and 5 respectively. Finally, the blue ellipsoid, which also looks like a line, depicts the linear propagation of the covariance represented in curvilinear coordinates in both figures. Though hard to see from the figures, the state-error representation in curvilinear coordinates provides an excellent representation of the distribution of sample points with the blue ellipsoid containing all of magenta points. The linear mapping to measurement coordinates, however, provides a poor representation of the uncertainty in possible measurement values. The unscented transform does a more satisfactory job of capturing the range of values in measurement space.

The observations above combined with the fact that the EKF and UKF employ the same generic filter recursion of alternating time and measurement updates, suggests an interesting possibility. It would seem desirable to create a hybrid filter that uses the time update from the EKF and the measurement update from the UKF since the concern over non-linear effects in the filter is mainly justified in the measurement update. The hope for the hybrid filter would be that run-time performance would be close to the EKF, since the time update is typically the most computationally expensive operation during orbit determination, while providing a better criterion for measurement acceptance and an improved update when processing measurements under the condition of large state uncertainty.

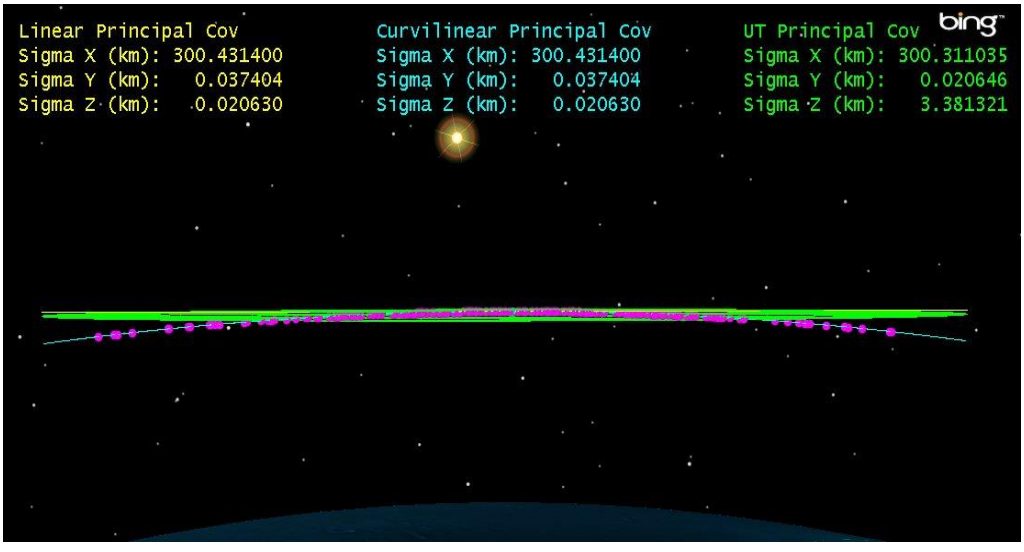

Figure 4: Example of non-linear effects on orbit-error distribution

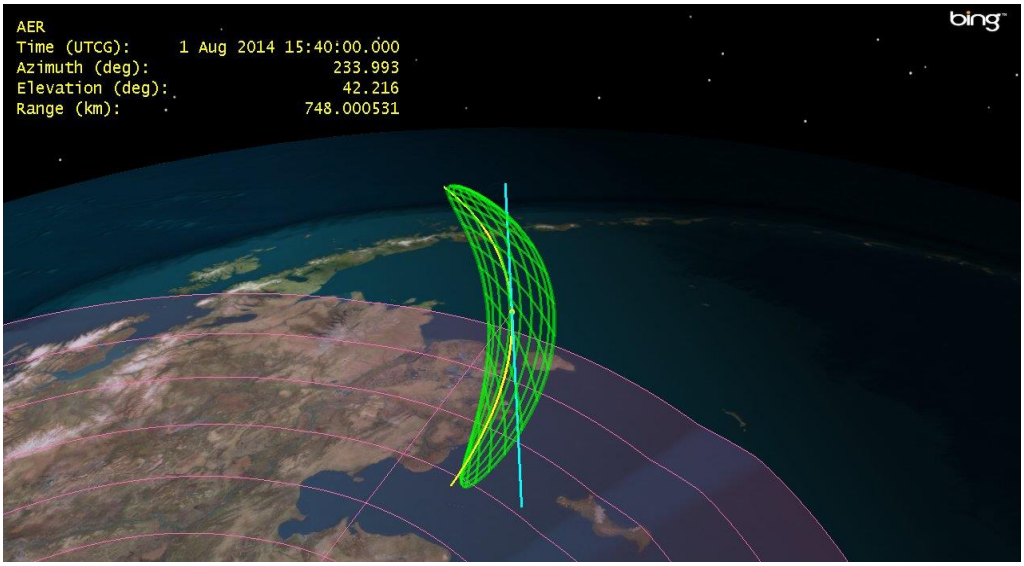

Figure 5: Example of non-linear effects on measurement-error distribution

METHOD OF COMPARISON

Evaluation of the relative merits of the EKF, UKF and Hybrid formulations will be based on the set of test cases defined in Table 3. Ground tracking was provided from a network of 4 ground stations generating range, azimuth and elevation measurements. Test cases were designed in a manner where drivers of uncertainty, mainly gaps in tracking data, were increased until the EKF lost track custody in at least one observable type. Specific estimation challenges include large initial tracking gaps which allow the along-track uncertainty to grow prior to initial tracking to sizes on the order of the tracking baseline, tracking of the HEO primarily around perigee, and a long gap in tracking immediately following a significant maneuver. The orbit uncertainty resulting from these challenges was increased to ensure that the EKF results are poor in all but the LEO Baseline test case. We note that while the LEO Baseline case provides a scenario where all filter configurations performed well, the levels of uncertainty still exceed those typically seen for cooperative tracking.

Table 3. Test Case Descriptions.

Test Case	Description
LEO Baseline	300 Km altitude, 70x70 gravity, drag (20% uncertainty, 180 min half-life) Moderately sparse ground-based tracking frequency (range, azimuth, elevation)
LEO Maneuver	500 Km altitude, 70x70 gravity, with drag (20% uncertainty, 180 min half-life) Large apogee raising maneuver (100 m/s intrack, 2% magnitude, 2 deg uncertainty) 10 hour gap in ground-based tracking after the maneuver (range, azimuth, elevation)
Sparse HEO	550 x 40000 Km altitude, 40x40 gravity, SRP (20% uncertainty, 48 hour half-life) Large initial uncertainty, intermittent ground tracking (range, azimuth, elevation) primarily near perigee
LEO Formation	Two satellites in a 500 Km altitude orbit with drag, separated by 4 – 8 Km. Large initial uncertainty, inter-satellite tracking (range, right ascension, declination, 10 min every 3 hours)

Test cases were exercised using up to 5 filter variants as listed in Table 4. The augmented EKF variants were not evaluated on the LEO Baseline case where the standard EKF performed well. The criteria on which the filter variants were evaluated are given in Table 5. These criteria are mainly qualitative in nature as we are more interested in behaviors than numerical results from a particular set of inputs.

Table 4. Filter Variant Descriptions.

Variant	Description
EKF	EKF formulation using equinoctial element formulation in measurement update
UKF	UKF formulation using $= 1/\sqrt{n}$,$\kappa = 0, \beta = 2$. Process noise is mapped to the end of the time update window to avoid resampling.
UKF PN	UKF formulation using $= 1/\sqrt{n}$,$\kappa = 0, \beta = 2$. Process noise is added on a regular time grid during the time update window requiring resampling.
Hybrid	EKF time update, UKF measurement update with $\alpha = 1/\sqrt{n}$,$\kappa = 0, \beta = 2$
EKF 2nd	EKF augmented with measurement update corrections from Gaussian 2nd order filter
EKF UW	EKF augmented with measurement underweighting to allow a reduction in covariance of no more than a factor of 10 per processed measurement

Table 5. Comparison Criteria.

Criterion	Description
Computation Time Ratio	The ratio of the computation time of the current estimation process divided by the computation time required for the EKF
Orbit Custody	The acceptance of **all** types of measurements over time during filtering
Position Uncertainty	The normal, tangential and cross-track components of the formal uncertainty from the filter solution. The maximum value of any component is reported in the summary tables
Covariance Consistency	The inclusion of normal, tangential and cross-track position component errors within the 3-sigma formal uncertainty from the filter solution

RESULTS

Relevant results are provided for each test case. Graphs of residual ratios, defined as residuals divided by the square root of the measurement-error variance, are used to demonstrate orbit custody. Nominally, 99 percent of residual ratio values fall within a range of +/- 3. Measurement editing criteria were set to reject measurements falling outside this range. In cases with ground based tracking, range residual ratios are shown in red, azimuth in blue and elevation in green. The +/- 3 bounds used for residual editing are also displayed on the plots. An excess of residual ratios outside of the editing bounds indicates that the filter has lost custody of the orbit solution, at least for a particular measurement type. Position uncertainty graphs are used to identify the orbit-state uncertainty estimated by each filter. Position uncertainty is represented in terms of 1-sigma values in the normal (to the tangent direction), tangential and cross-track directions. Graphs of position error

ratios, defined as position error components divided by the square root of the associated position error variance, help identify when the solution accuracy is consistent with the estimated uncertainty.

LEO Baseline

The LEO Baseline case was designed to have higher than typical uncertainty for a cooperatively tracked satellite, but without any severe spikes. All filters performed equally well with qualitatively identical results as summarized in Table 6 and depicted by the plots in Figure 6. We additionally note that the UKF results showed no qualitative difference between cases where the process noise was injected on a regular time grid, thus requiring resampling of the covariance, and when the process noise was applied at the end of the overall time update intervals.

Table 6. LEO Baseline Results.

Filter	Computation Time	Orbit Custody	Position Uncertainty	Covariance Consistency
EKF	1	Y	~10 Km	Y
UKF	28	Y	~10 Km	Y
Hybrid	1.2	Y	~10 Km	Y

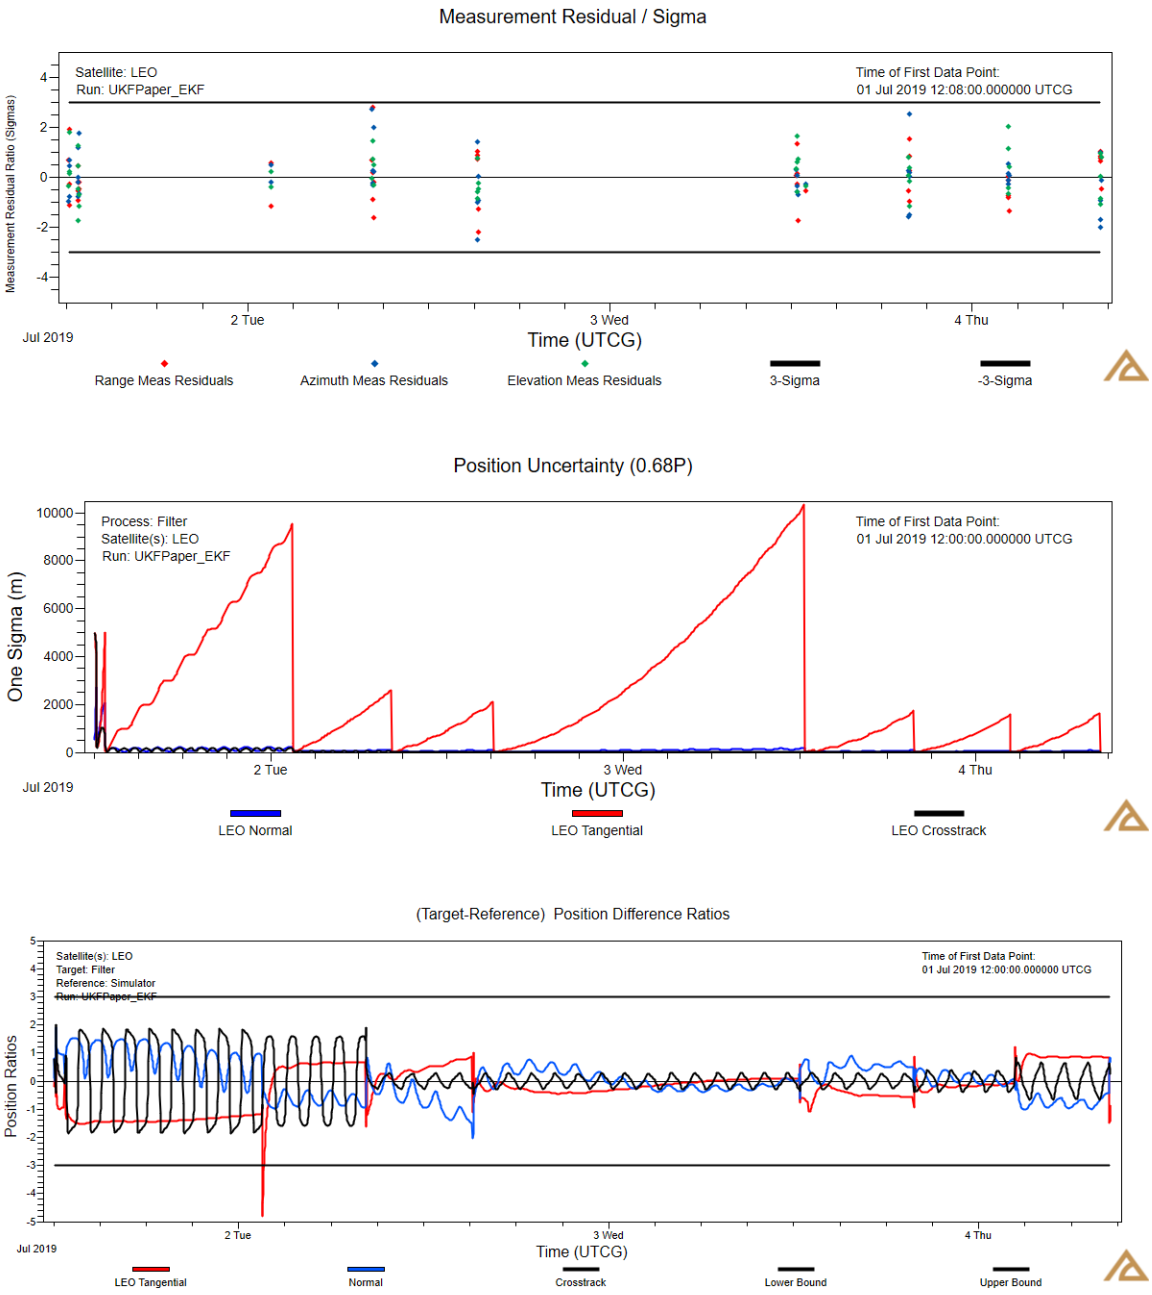

Figure 6. LEO Baseline Results.

Sparse HEO

The Sparse HEO case was designed to be extremely challenging for the filters to gain custody of the orbit. A propagation over three revolutions is included prior to the first observations which occur near perigee where the along-track uncertainty reaches a maximum. A graph of the measurement locations as function of altitude is presented in Figure 7. The EKF and Hybrid formulations performed poorly in this case and failed to retain orbit custody. The UKF provided a better solution, but did require adjustment of the test case due to instances where one or more samples reentered

during the time update. A strategy to deal with such errors would be needed in an operational environment. A tabulated summary of the results is given in Table 7. The UKF PN run produced a notable result, the inclusion of PN at sub-intervals within the time update resulted in a large expansion in the size of the orbit-error covariance. An explanation of the mechanism which generated the rapid growth is given below. Estimation results obtained from the various filter configurations are presented in Figures 8-11.

Table 7. Sparse HEO Results.

Filter	Computation Time	Orbit Custody	Position Uncertainty	Covariance Consistency
EKF	1	N	~150 Km	N
UKF	23	Y	~150 Km	Y
UKF PN	23	Y	~950 Km	Y
Hybrid	1	N	~150 Km	N
EKF 2nd	1	N	~150 Km	N
EKF UW	1	N	~150 Km	N

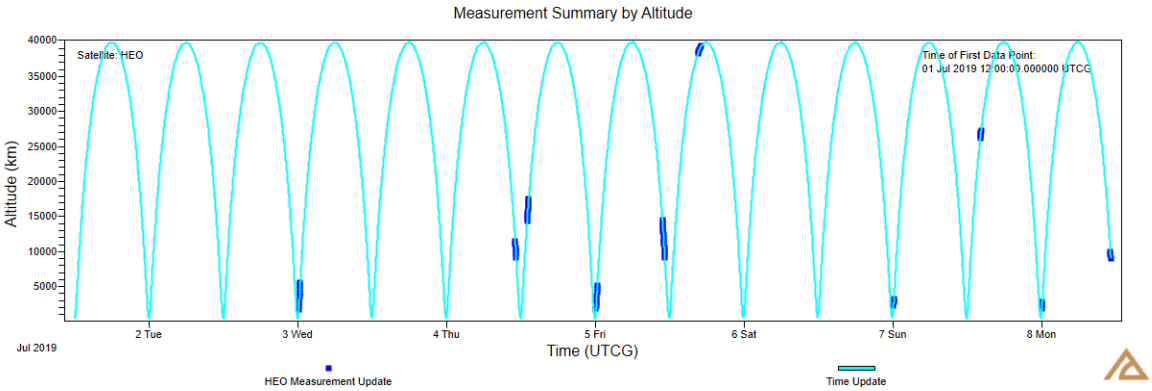

Figure 7. Sparse HEO Measurement Altitudes.

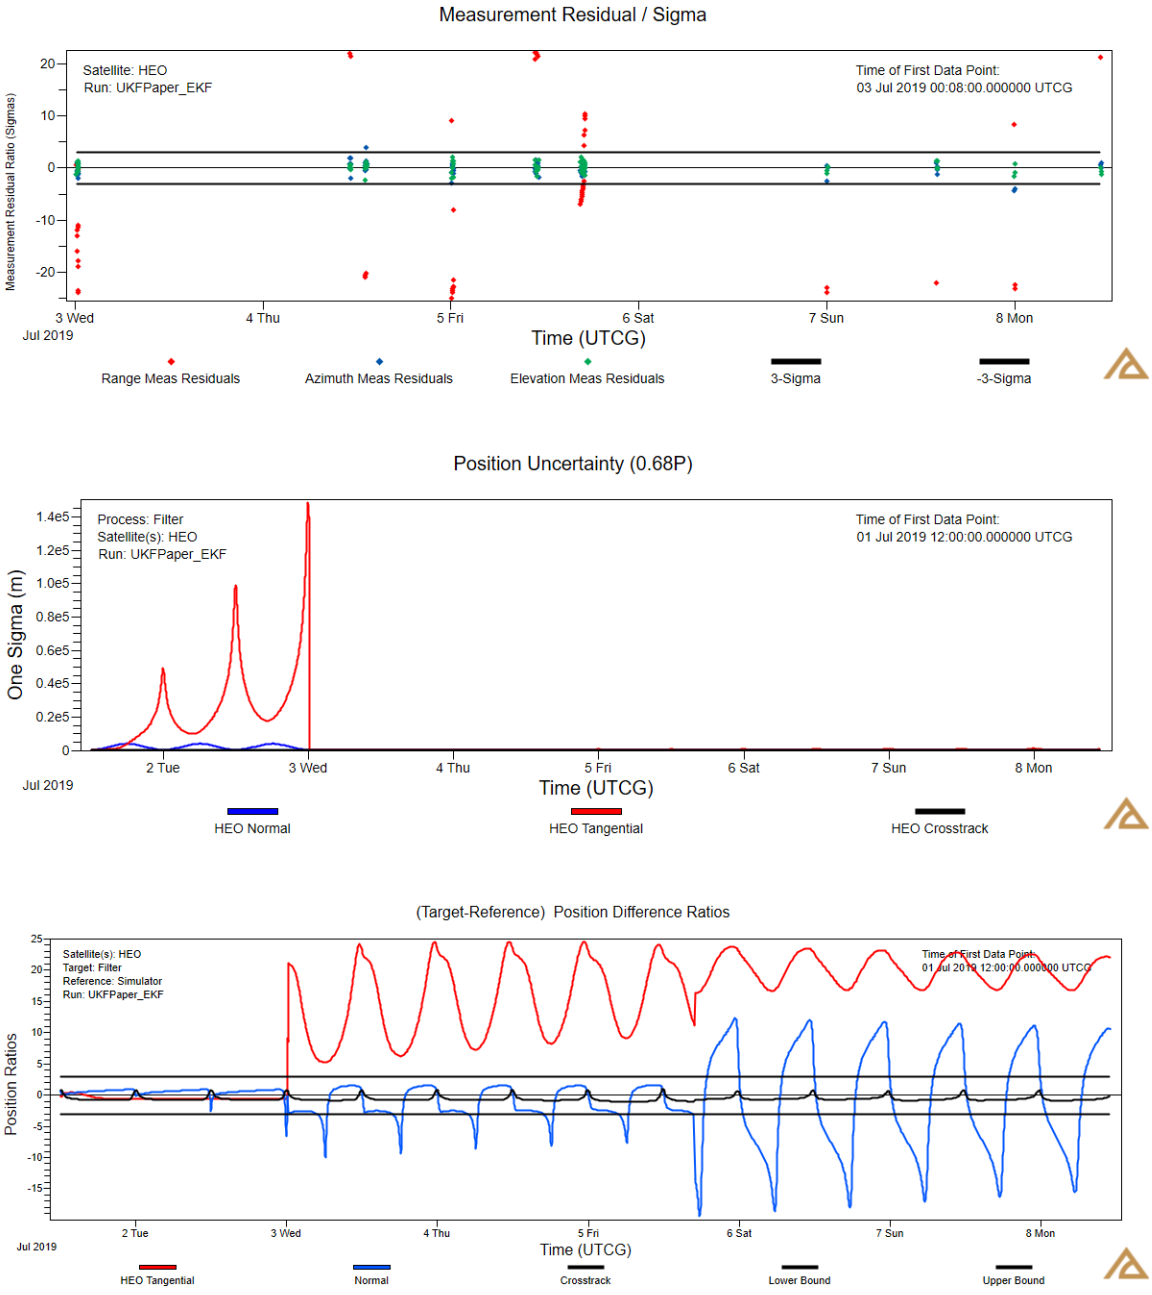

Figure 8. Sparse HEO EKF Results.

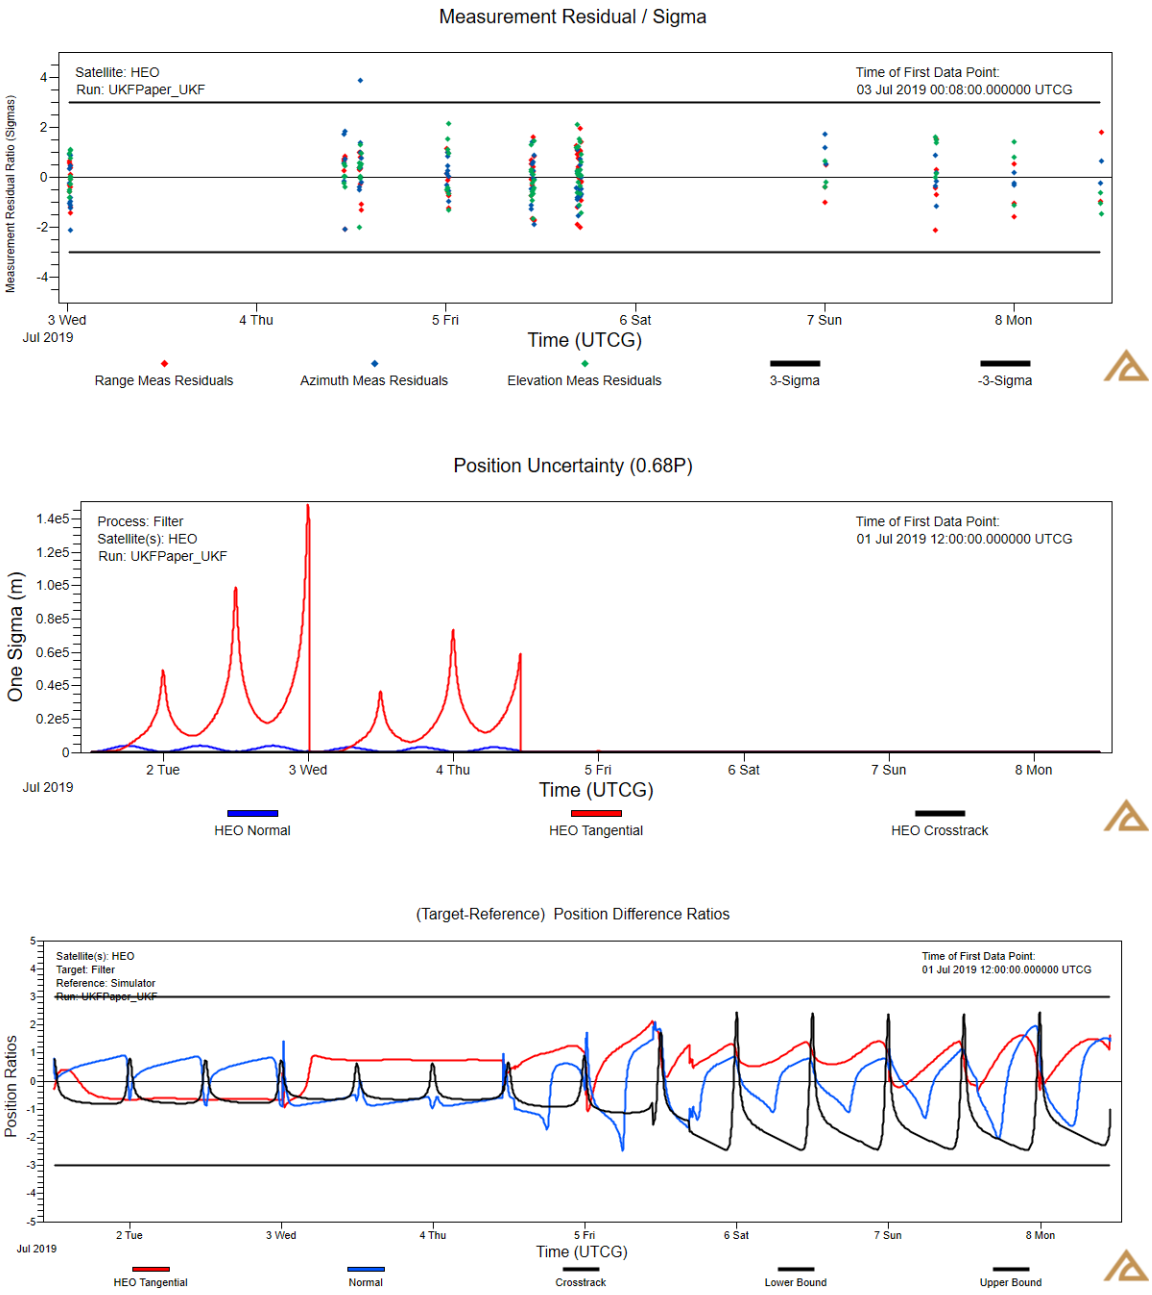

Figure 9. Sparse HEO UKF Results

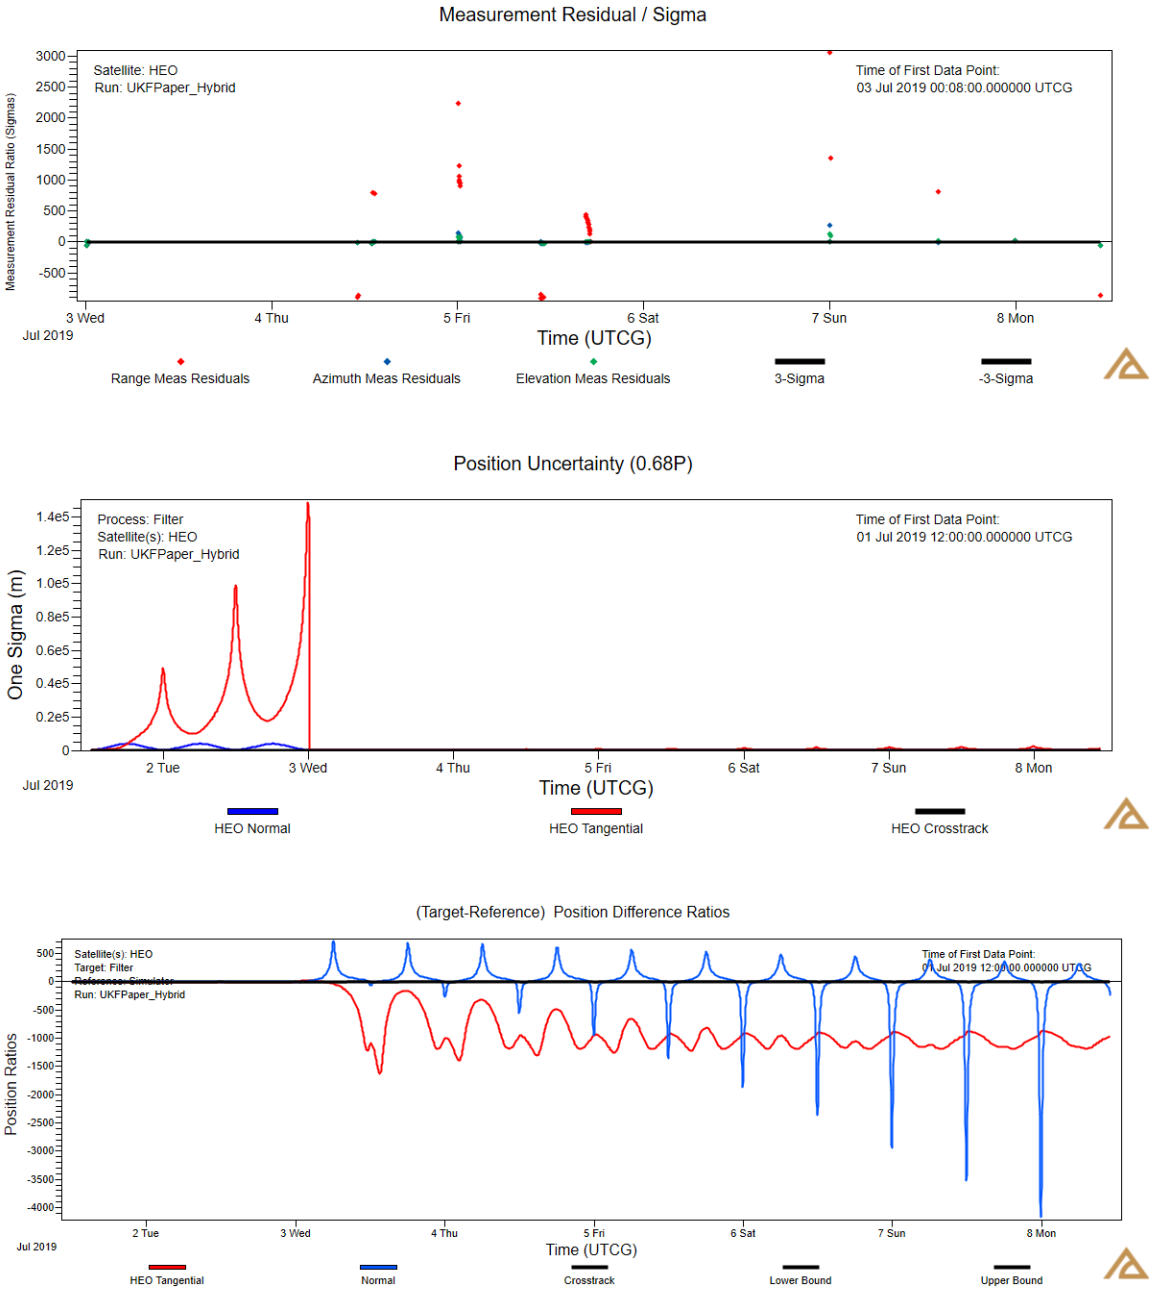

Figure 10. Sparse HEO Hybrid Results

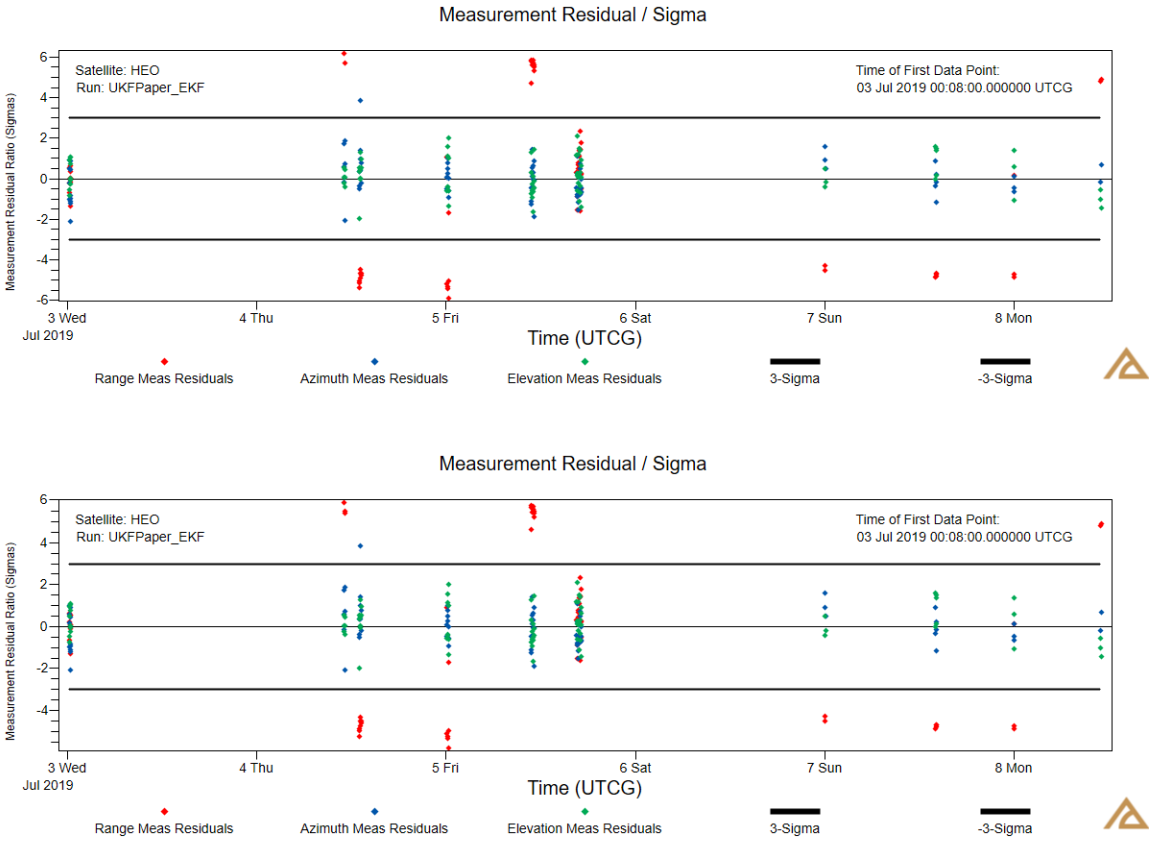

Figure 11. Sparse HEO EKF 2nd and EKF LW Residual Ratios.

A large expansion in the orbit-error covariance in the Sparse HEO case occurred when process noise was added at intermediate times during time update intervals. The driving mechanism behind the unexpected growth in orbit-error covariance was the generation of new sigma points after the process noise inclusion at times when the Cartesian representation of the covariance was a poor fit to the actual orbit error distribution. In such cases, the dimension of the position covariance in the orbit plane but normal to the velocity direction is increased, during construction of a sample covariance, to accommodate the curvature in the uncertainty distribution. When the resulting covariance is resampled, the new sigma points span the larger extent of the constructed sample covariance. As a result, the new sigma points have a greater spread in the in-plane normal direction than the sigma points prior to process noise inclusion, even if the process noise inclusion itself provides no expansion in the in-plane normal direction. The new larger spread of sigma points evolves to an even larger distribution through non-linear propagation until the next process noise update is made and the process continues. This phenomenon is most prevalent in the perigee region where the curvature in the orbit uncertainty distribution is most curved. This effect is illustrated in Figure 12.

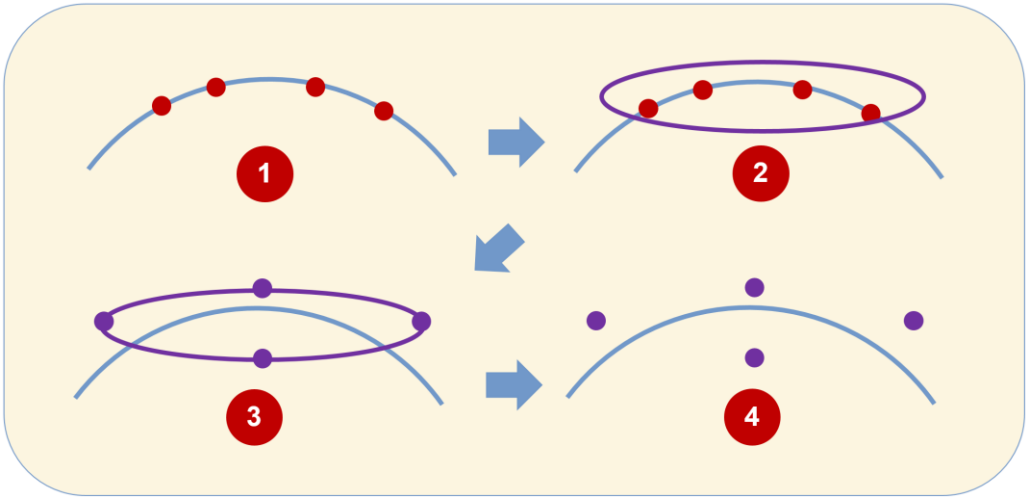

Figure 12. Resampling Effect on Uncertainty Propagation in UKF.

LEO with maneuver

The LEO Maneuver case was designed to test the ability of the filters to regain custody of the orbit after a large maneuver followed by a significant gap in tracking. The 10 hour measurement gap after the maneuver allows the orbit uncertainty to grow rapidly due to the velocity uncertainty injected with the maneuver. As with the Sparse HEO case, the EKF and Hybrid formulations performed poorly in this case and failed to retain orbit custody. Estimation results obtained from the various filter configurations are presented in Figures 13-16. A tabulated summary of the results is given in Table 8.

Table 8. LEO With Maneuver Results.

Filter	Computation Time	Orbit Custody	Position Uncertainty	Covariance Consistency
EKF	1	N	~220 Km	N
UKF	26	Y	~220 Km	~Y
Hybrid	1.1	N	~220 Km	N
EKF 2nd	1	N	~220 Km	N
EKF UW	1	N	~220 Km	N

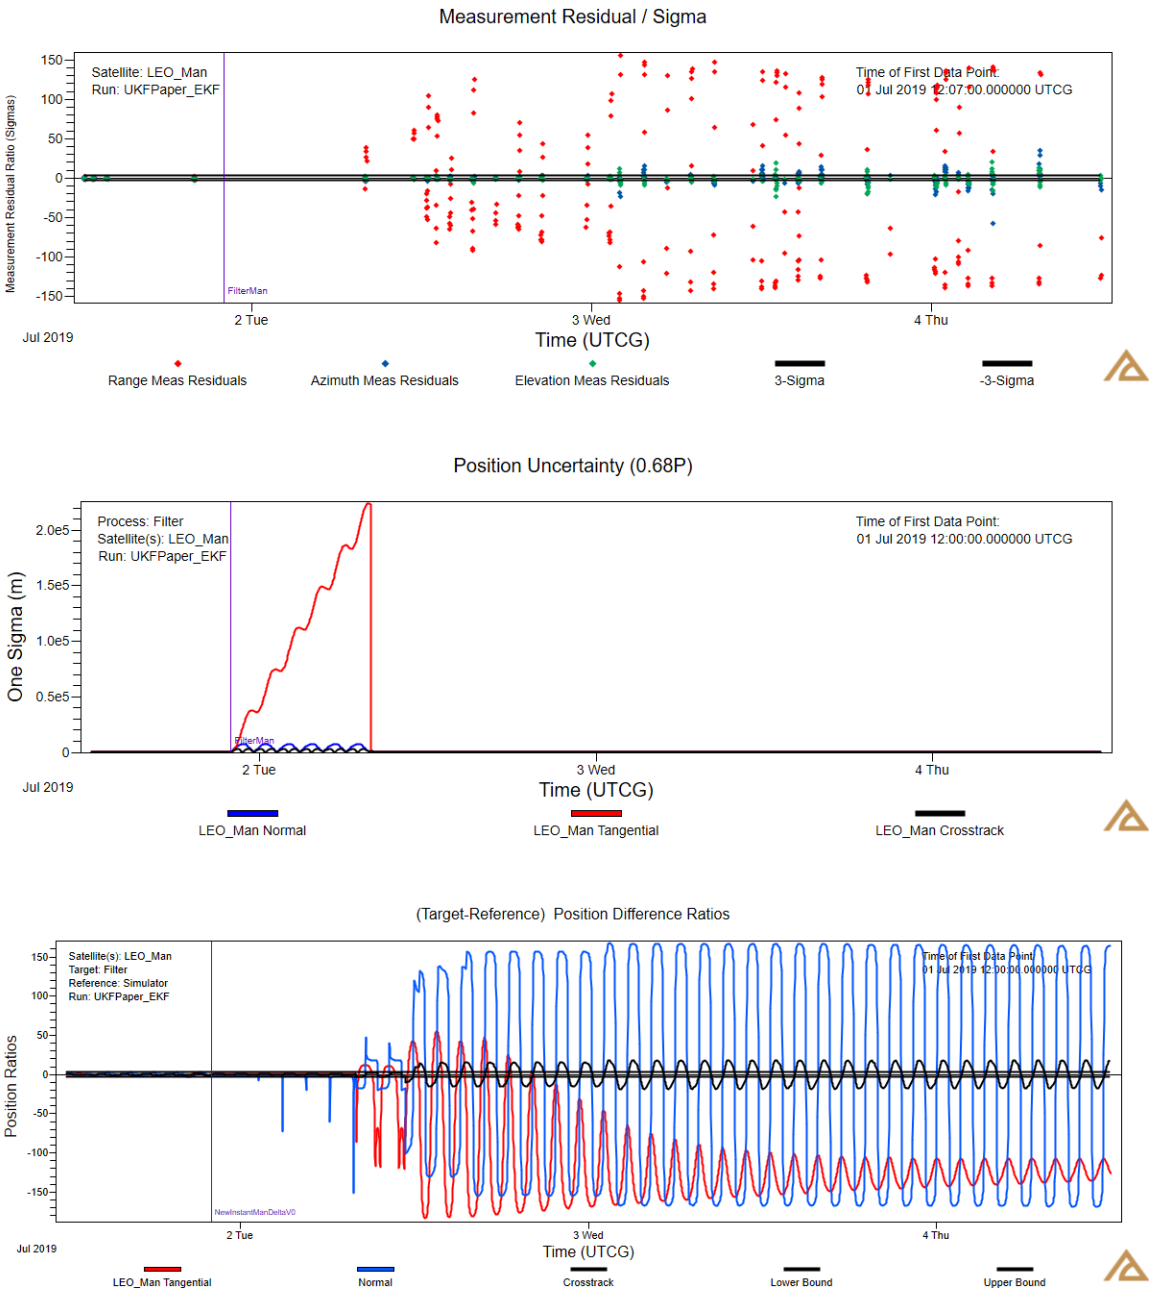

Figure 13. LEO Maneuver EKF Results.

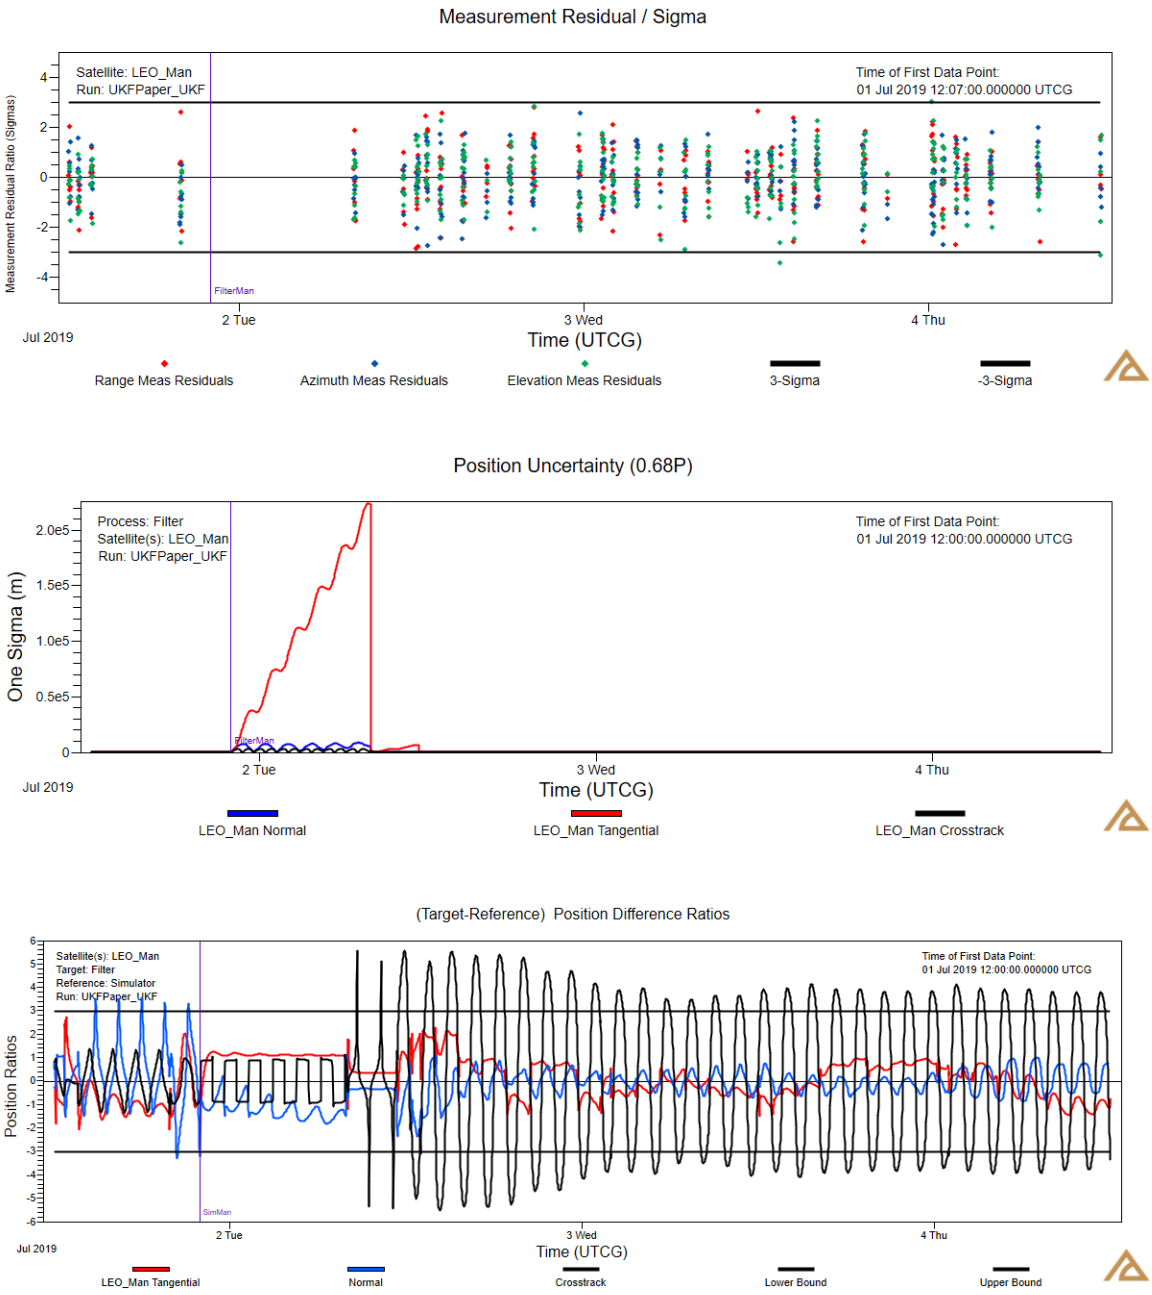

Figure 14. LEO Maneuver UKF Results.

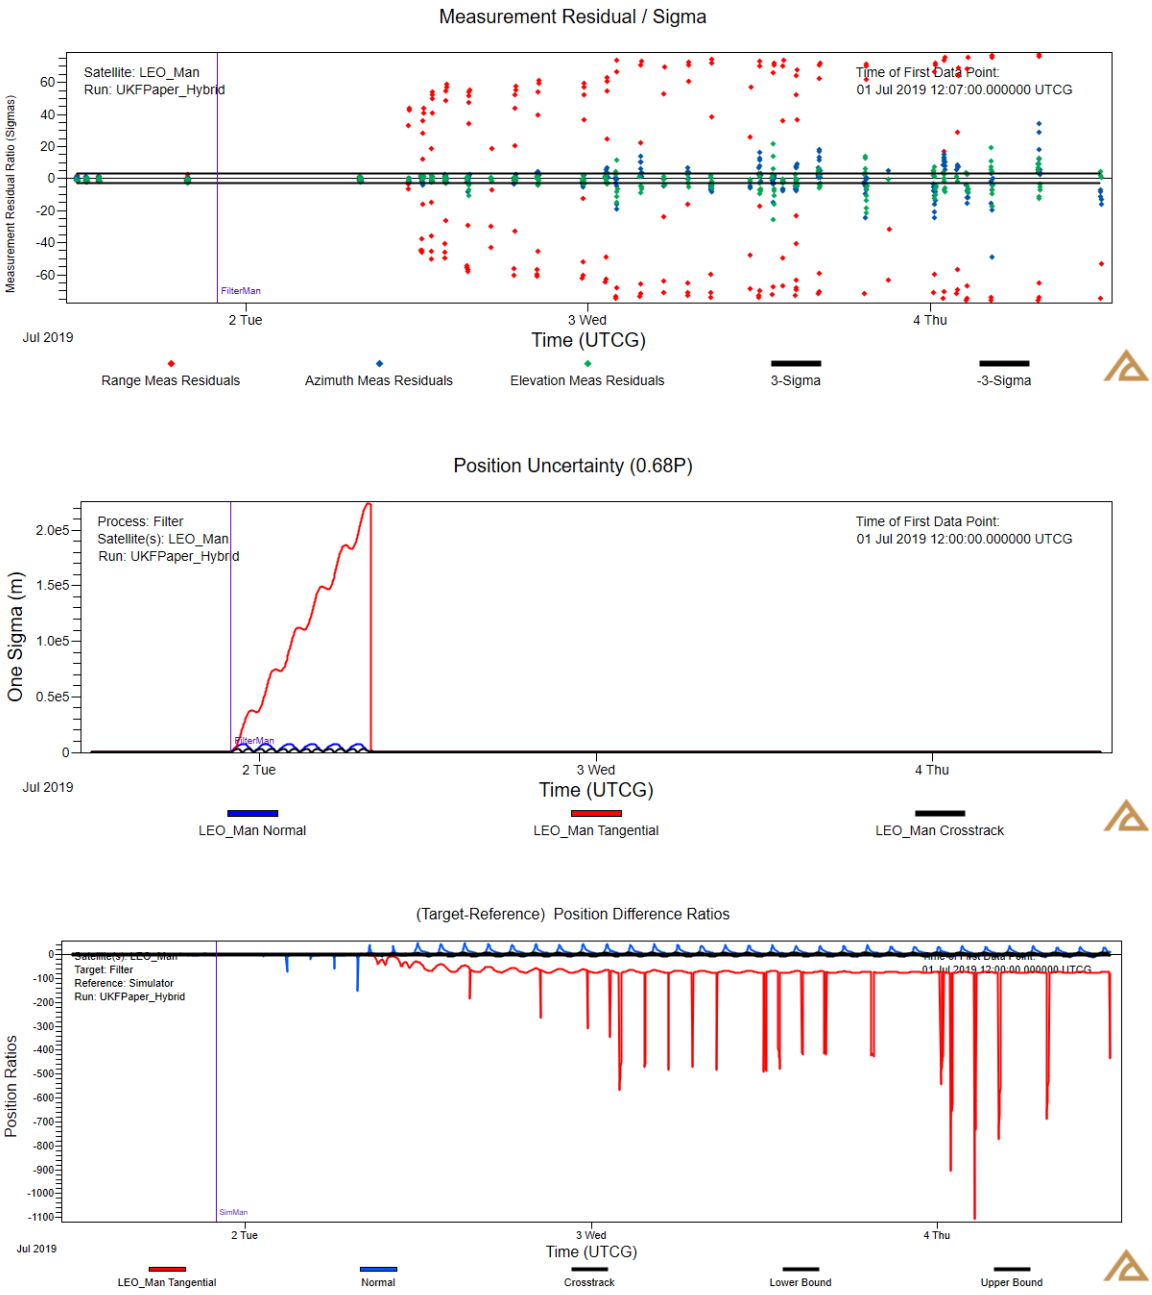

Figure 15. LEO Maneuver Hybrid Results.

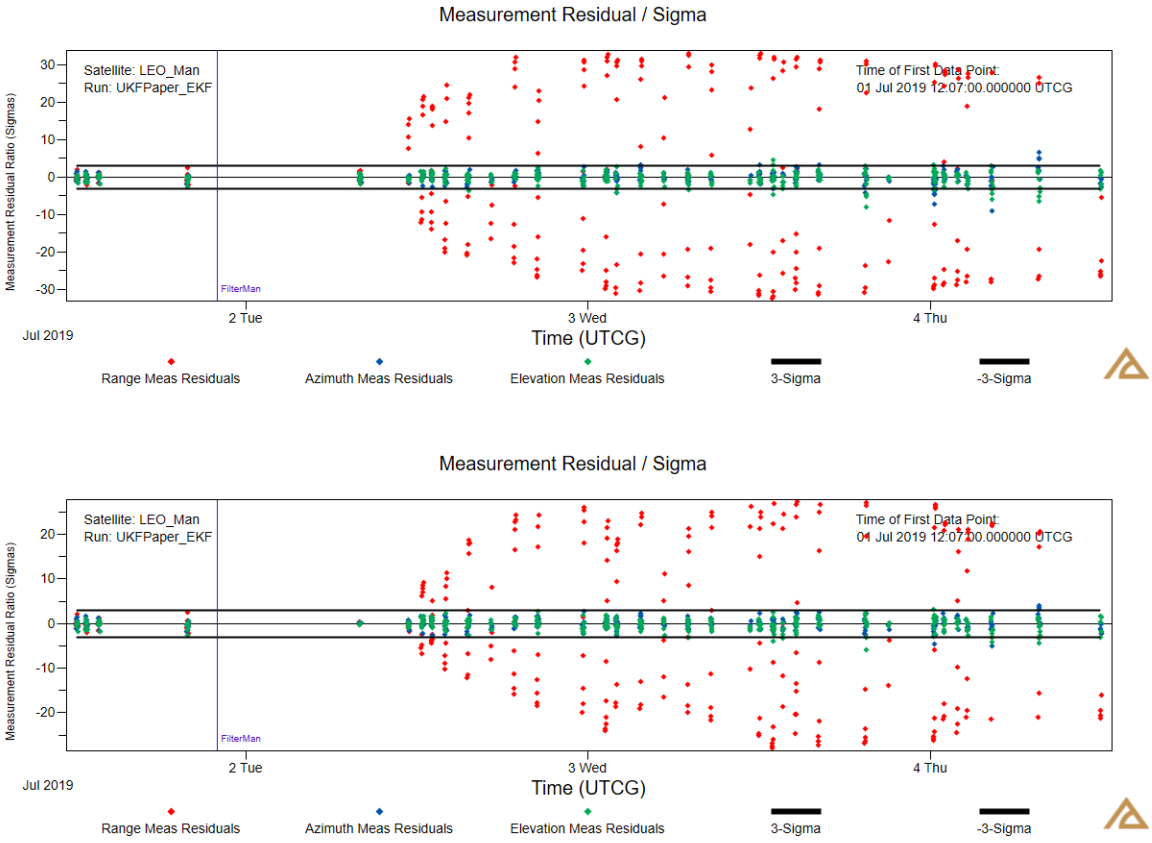

Figure 16. LEO Maneuver EKF 2nd and EKF LW Residual Ratios.

LEO Formation

The LEO Formation case was designed to test the ability of the filters to gain custody of the orbits of both orbits from conditions where the initial position uncertainty was of similar size to the baseline, 4-8 Km, of the inter-satellite measurements. Measurements over short baselines, such as those typically associated with formation flight, experience non-linear effects under conditions of much smaller state uncertainty than measurements on longer baselines. In this case, the position uncertainty, while still large compared to most cooperative tracking scenarios, is slightly smaller than the LEO Baseline case described previously. While the standard EKF struggled to lock onto the orbits, both EKF augmentation strategies were effective as were the UKF and Hybrid formulations. A tabulated summary of the results is given in Table 9. Estimation results obtained from the various filter configurations are presented in Figures 17-20.

Table 9. LEO Formation Results.

Filter	Computation Time	Orbit Custody	Position Uncertainty	Covariance Consistency
EKF	1	N	~7 Km	N
UKF	31	Y	~7 Km	Y
Hybrid	1.2	Y	~7 Km	Y

EKF 2nd	1	Y	~7 Km	Y
EKF UW	1	Y	~7 Km	Y

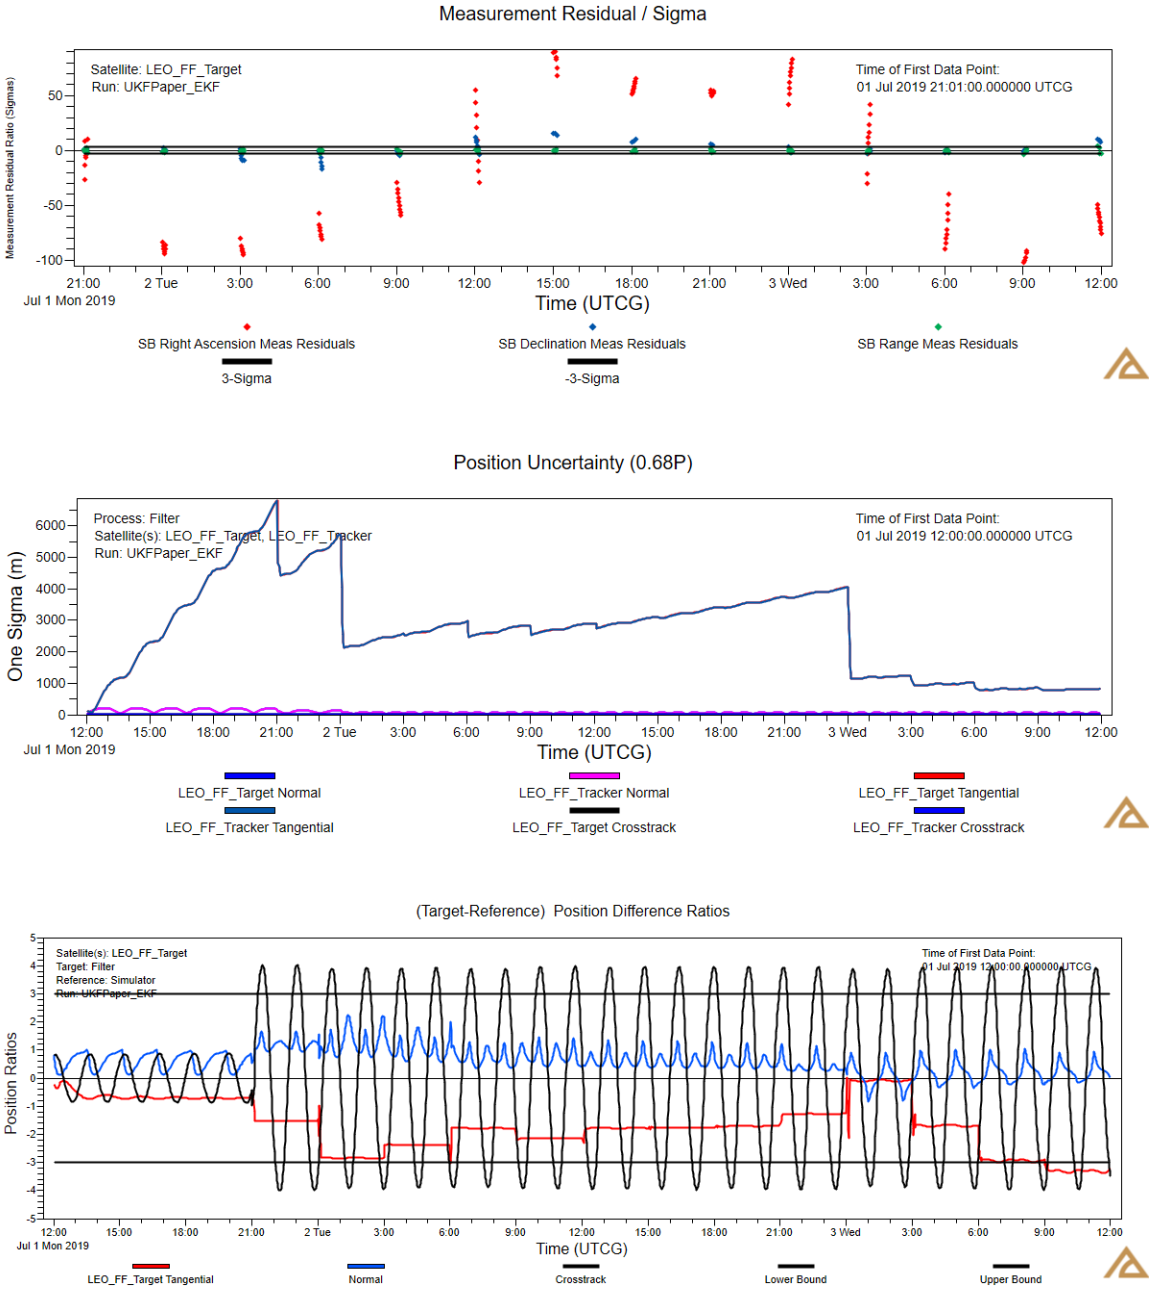

Figure 17. LEO Formation EKF Results.

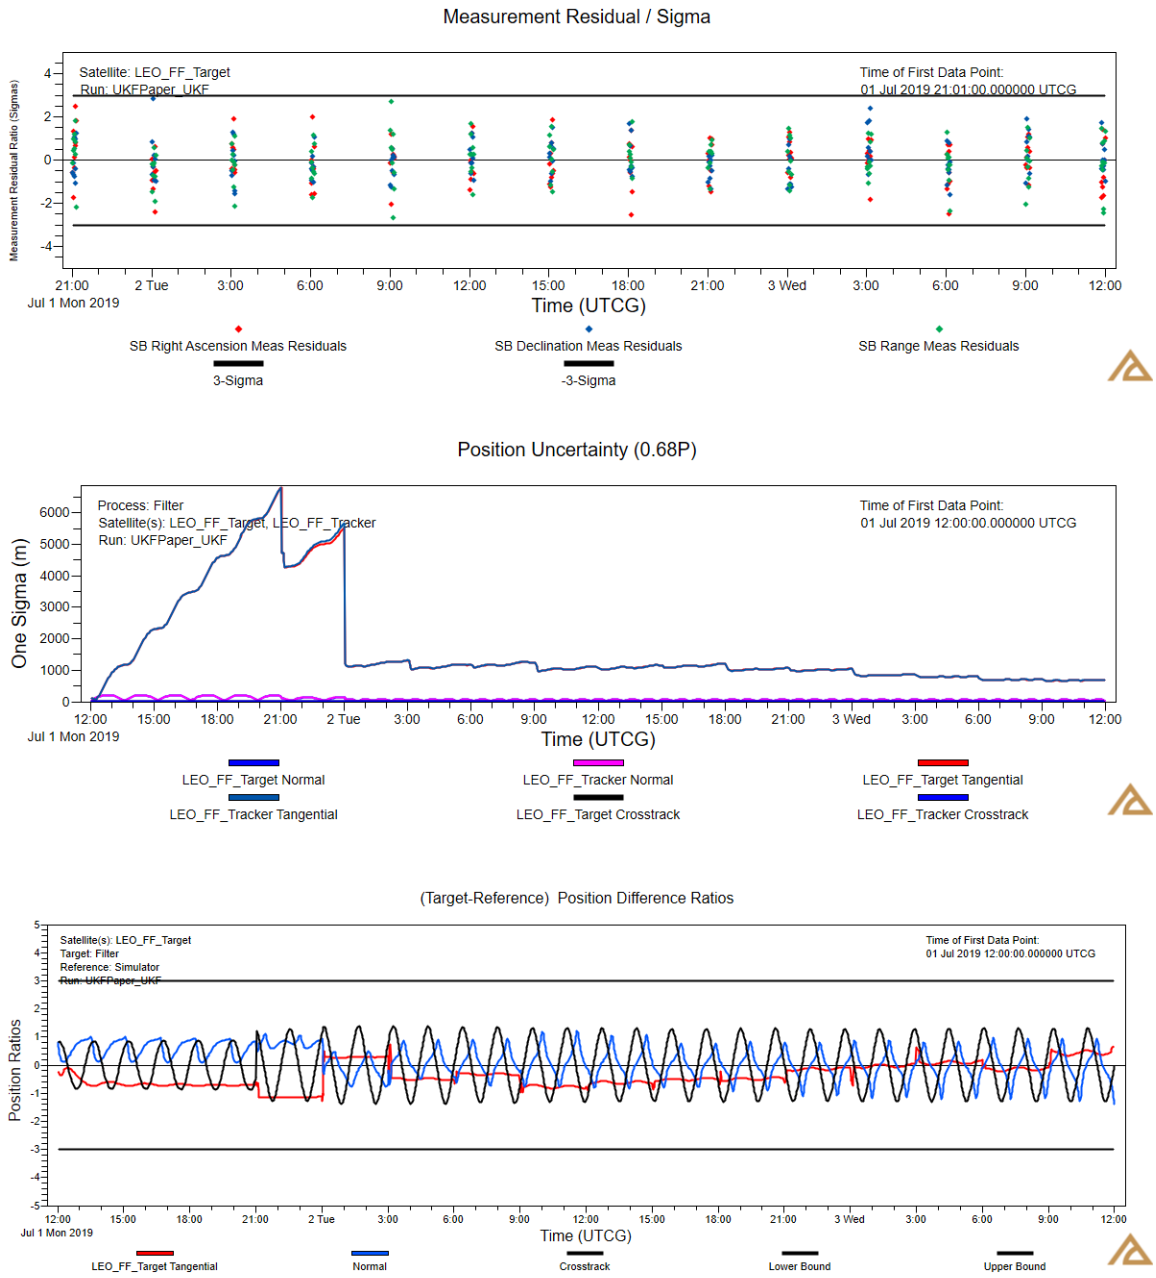

Figure 18. LEO Formation UKF Results.

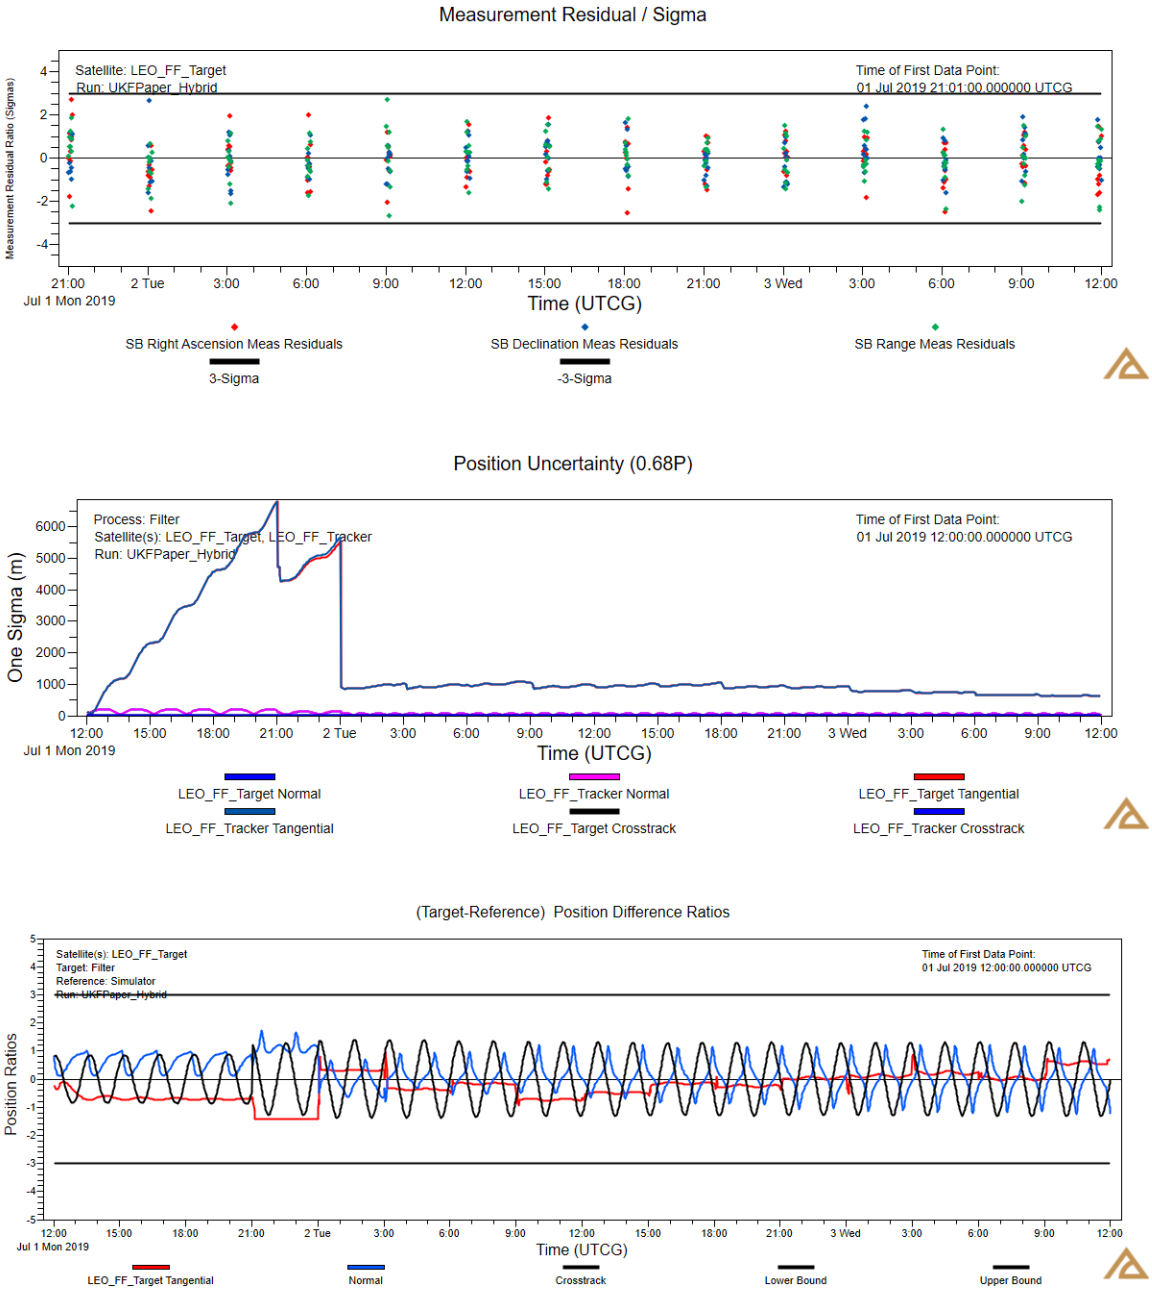

Figure 19. LEO Formation Hybrid Results.

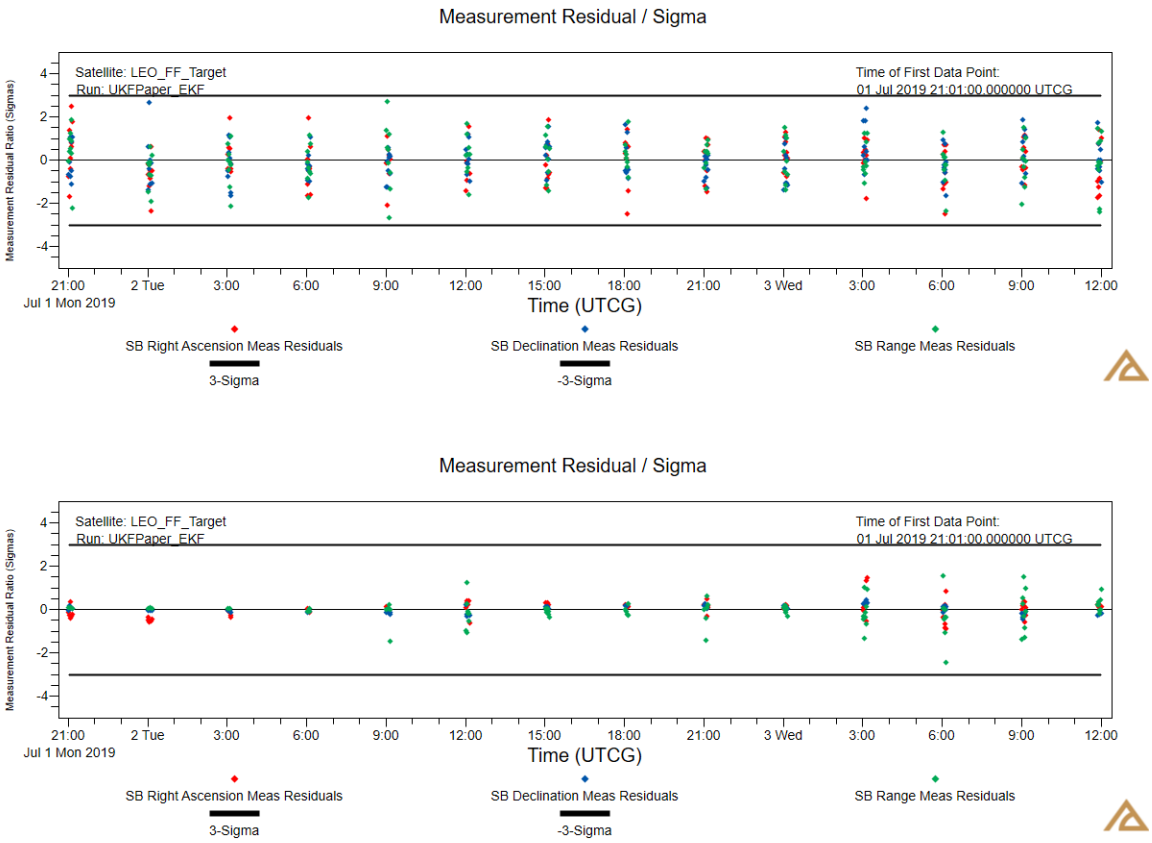

Figure 20. LEO Formation EKF 2nd and EKF LW Residual Ratios.

CONCLUSION

By the design of this study, the EKF was pushed beyond its limits in the examined cases. In an operational setting, we would recommend using an initial orbit determination method followed by a least squares estimator to initialize filter solutions. In this study, however, we created a more challenging initialization environment for the filters to test the limits of the algorithms. Algorithm comparisons were performed using operationally vetted dynamical and measurement models. Our EKF/UKF/Hybrid implementations all include necessary noise modeling and estimation states for operational determination. The resulting state size is larger than typical for examinations of the UKF and this fact is reflected in the comparative run times.

The overall summary of our recommendations for filter selection based on the relevance of non-linear effects in the measurement and dynamical model is provided in Table 10. For this summary, we have grouped the EKF and its augmentations together under the EKF label and the two variations in the application of the process noise in the UKF under the UKF label. We define the *Typical* category as the vast majority of orbit determination scenarios where non-linear effects are small. The *Meas NL* category includes the set of problems where non-linearity effects are mostly restricted to the measurement model. Finally, the *Meas NL Dyn NL* category includes problems where non-linear effects are prevalent in both the measurement model and the dynamics.

The UKF is seen to provide an improvement over the standard EKF in cases where non-linear effects in the measurement model become significant. We note two specific sub-cases: when the

orbit-state uncertainty is well modeled by the Cartesian representation of the state-error covariance and the other where the orbit-state uncertainty is large enough that the Cartesian representation of the covariance is not a good representation of the error distribution. The first situation can be realized when a short measurement baseline is involved, such as the LEO Formation test case. In this situation, EKF augmentation strategies such as the 2nd order Gaussian filter corrections to the measurement update or measurement underweighting are effective and we find them to be preferred to the UKF due to runtime considerations. The EKF is also preferred in cases where non-linearity in the measurement model is not significant due to large differences in run times. The Hybrid formulation was also effective under this circumstance since the samples generated from the EKF state-error covariance in Cartesian coordinates still provided a good representation of the error distribution. In the second scenario, where the orbit-state uncertainty is also large, the EKF augmentation methods are not as effective and the Hybrid filter performs poorly due to sampling a poor representation of the orbit-state uncertainty. In these cases, we prefer the UKF as it can retain custody of the orbit under more stressing conditions. We note, however, that the process noise strategy used in the UKF time update can have drastic effects on the covariance evolution. Finally, we note a preference for the application of process noise at the end of the UKF time update for cases of large orbit uncertainty.

Table 10. Preference Summary.

Filter	Typical	Meas NL	Meas NL, Dyn NL
EKF	**Preferred**	**Preferred**	**Less Capable**
UKF	**Capable**	**Capable**	**Preferred**
Hybrid	**Capable**	**Capable**	**Less Capable**

We note that UKF, while requiring significantly more computations than the EKF, could be designed to perform many of those computations in parallel. This strategy is difficult to evaluate in terms of computational performance, however, since effectiveness will be influenced by the number of processing cores available and the desire to perform orbit determination on multiple objects at the same time, as is typical in a SSA application. We also specifically excluded the processing of Doppler measurements in this analysis due to the large runtime penalty experienced using the UKF. In filter implementations, integrated Doppler measurements require the backwards propagation of the orbit state over the accumulation interval of the measurement, typically 10 – 60 seconds. The sampling strategy used in the UKF also opens the door to issues such as the reentry of sample points in low perigee scenarios and non-uniqueness of solutions based on selection of different values for the free parameters in the unscented transform.

The Hybrid formulation developed in this study performs sampling of the Cartesian representation of state-error covariance from the EKF time update in support of the UKF measurement update. In future work, we will investigate a change of the sampling and associated measurement update to work in coordinates which are better for state-error representation. We expect this modification to improve the range of application of the Hybrid algorithm while retaining its computation time advantage over the UKF. We also plan to evaluate estimation results in a more quantitative manner.

.

REFERENCES

[1] Julier, S.J., Uhlmann, J. K., and Durrant-Whyte, "A new approach for filtering nonlinear systems," 1995 American Control Conference, Seattle, WA, pp. 1628–1632.

[2] Julier, S.J., Uhlmann, J. K., and Durrant-Whyte, "A New Method for the Nonlinear Transformation of Means and Covariances in Filters and Estimators," Technical Note in IEEE Transactions on Automatic Control, Vol. 45, No. 3, March 2000, pp. 477–482.

[3] Julier, S.J., and Uhlmann, J. K., "A new extension of the Kalman Filter to Non-linear Systems," AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls, Orlando, FL, 1997.

[4] Julier, S.J., and Uhlmann, J. K., "The scaled unscented transformation," 2002 American Control Conference, Anchorage, AK, pp. 4555–4559.

[5] Wright, James R., Optimal Orbit Determination, Paper AAS 02-192, AAS/AIAA Space Flight Mechanics Meeting, San Antonio, Texas, 27-30 Jan., 2002.

[6] Vallado, D. A. and Seago, J.H., "Covariance Realism," AAS 09-304, 2009 AAS/AIAA Astrodynamics Specialist Conference, Pittsburgh, PA, August 2009.

[7] Johnson, T.M., "Orbit Prediction Accuracy Using Vasicek, Gauss-Markov, and Random Walk Stochastic Models," Paper AAS 13-280, 2013 AAS/AIAA Space Flight Mechanics Meeting, Kauai, Hawaii, January 2013.

[8] Wright, J. R. et al, "Orbit Determination Tool Kit: Theory & Algorithms", Technical report, Analytical Graphics Inc., 2013.

[9] Lear, W. M., "Kalman Filtering Techniques," NASA Technical Report JSC-20688, Johnson Space Center, Houston, TX, 1985.

[10] Jazwinski, A. H., *Stochastic Processes and Filtering Theory*, Dover, 2007.

[11] Lee,D., Alfriend,K.T., "Precise Real-Time Satellite Orbit Estimation Using the Unscented Kalman Filter", AAS 03-230, 13th AAS/AIAA Space Flight Mechanics Meting, Ponce, Puerto Rico, February 2003.

[12] Wu, Y., Hu, D., Wu, M., Hu, X., "Unscented Kalman Filtering for Additive Noise Case: Augmented vs. Non-augmented", 2005 American Control Conference, Portland, OR, June 2005.

[13] Wright, James R., Woodburn, James, Truong, Son, Chuba, William, "Orbit Gravity Error Covariance," AAS 08-157, AAS/AIAA Space Flight Mechanics Meeting, Galveston, TX, January 2008.

[14] Woodburn, J. and Coppola, V., "Effect of Coordinate Selection on Orbit Determination," AAS 13-825, 2013 AAS/AIAA Astrodynamics Specialist Conference, Hilton Head, SC, August 2013.

[15] Woodburn, J., Tanygin, S., "Coordinate Effects on the Use of Orbit Error Uncertainty", 2014 International Symposium on Space Flight Dynamics, Laurel, MD, May 2014.