54-13
185864
p. 9

August 15, 1993

N94-14373

Failure Modes of Reduced-Order Orbit Determination Filters and Their Remedies

D. J. Scheeres
Navigation Systems Section

This article discusses ways in which failure can occur in reduced-order, orbit-determination filter, error covariance calculations. In the context of this article, reduced-order filters denote nonoptimal filters which include fixed levels of uncertainty in some parameters of the measurement models or in the spacecraft dynamical model which are not explicitly estimated in the filter equations. Failure is defined herein as an increase in the orbit determination covariance with the addition of data or as an unreasonable growth in the covariance with time, i.e., nonasymptotic behavior of the covariance. Some simple, known cases of failure are discussed along with their traditional remedies. In addition, more modern remedies are discussed which are currently under development at the Jet Propulsion Laboratory. The article first describes the known problems of reduced-order filters when they are employed for orbit determination, and their traditional remedies. Then, having defined these, the relevancy and desirability of the more modern remedies are made apparent.

I. Introduction

As the terminology used for reduced-order filters, which are sometimes called consider filters, is not uniform, it is important that the definition of the consider filter used herein be understood from the outset. What is called the consider filter is actually known as the consider option in JPL's orbit determination software system. In this approach, which is commonly used in spacecraft navigation at JPL, the relevant estimated state parameters (spacecraft position and velocity, etc.) are determined using only knowledge of other estimated state parameters; the filter is closed upon itself. It is acknowledged, however, that there are other parameters which may not be known exactly and which have an effect on measurements made of the state and on the dynamics of the state.

For a variety of reasons it is sometimes not desired to estimate these parameters. These reasons may be that there are no adequate models for the evolution of these parameters, that estimating them would require excessive computation time, or that, if estimated, the computed uncertainty in these parameters would be reduced far below the level warranted by model accuracy. Instead, in order to gauge how uncertainties in these parameters manifest themselves as uncertainties in the state estimates, the estimation error covariance of the state (the computed covariance) is modified to account for these errors in some way. It should be noted that this modified covariance, called the consider covariance, is not fed back into the filter. Thus the filter knows nothing about the consider contribution to the uncertainty in the state estimate. The mathematical definition of these filters is reviewed later in the article.

In this article a consider filter is defined to have failed when the addition of data yields an increase in the consider covariance, and hence less confidence in the estimate, or when the consider covariance begins to grow unreasonably large as time increases. These types of failure are usually alleviated by artificially increasing, or deweighting, the measurement noise variances of one or more data types so that additional data yield a smaller or equivalent consider covariance. Data weights are, ideally, the inverse of the noise variance for the measurements being employed. They also can serve, however, as free parameters with which one may emphasize or deemphasize certain data types. If this does not alleviate the failure, then the data arc may be restarted, essentially discarding old data. Two general modes of failure for consider filters are discussed. In analyzing these modes it becomes evident why and how a filter should be tuned. The ramifications of a consider filter that needs substantial tuning are also discussed. It is hoped that the reader can develop an intuitive feel for the modes of failure of consider filters from this article, and that it will shed light on some of the seemingly nonintuitive behavior that can be exhibited by this type of estimation scheme.

A special effort is made to discuss a more sophisticated sequential orbit-determination filter model being developed at JPL, the Enhanced Filter [1,2]. The use of this filter alleviates many problems associated with more traditional consider filters. At various points in this article mention is made of this approach.

II. Modes of Failure

The modes of failure discussed below correspond to two extreme conditions in which a reduced-order filter may operate. It is hypothesized that generic cases of failure are a combination of these modes. There may exist, of course, possible modes not discussed here. In the first case, process noise is not assumed to be present in the filter model. This is called the batch mode. This case corresponds to a simplified covariance analysis or, in a limiting case, to a situation where the frequency of the data measurements is large enough (or the characteristic time of the process noise is long enough) to render the process noise component essentially constant. In this case, the filter is modeled as a batch filter. The second mode is the converse of the first, namely, it is the case where process noise effects dominate the filter behavior. This is called the sequential mode. In this case, the filter is said to be saturated, and the uncertainty in the filter parameters has reached a quasi-steady state. The filter is modeled as a sequential Kalman–Bucy filter in this case.

III. Batch Mode

To begin, the classical results of batch covariance analysis with consider parameter effects are restated.

A. Batch Consider Covariance Definitions

Assume the state model to be

$$x_k = \Phi(t_k, t_{k-1})x_{k-1} \tag{1}$$

and the measurement model to be

$$z_k = H_{x_k}x_k + \nu_k \tag{2}$$

with $E(\nu_i\nu_j^T) = R_i\delta_{ij}$. The term x_k is the state at time t_k; $\Phi(t_k, t_{k-1})$ is the state transition matrix; z_k is the measurement at time t_k; ν_k is the measurement noise; R_i is the variance of the measurement noise, and δ_{ij} is the usual Kronecker delta function. Then the classical approach to covariance computation, in the absence of process noise, can be stated as [3, p. 27]:

$$P_k = \left[\sum_{j=0}^{N}\Lambda(t_j, t_k)\right]^{-1} \tag{3}$$

where P_k is the state-estimation error covariance matrix at time t_k, the $\Lambda(t_j, t_k)$ are information matrices propagating measurements taken at time t_j to time t_k, and N is the total number of measurements. Essentially, the filter reduces all measurements to the same epoch and then applies a weighted least-squares estimation process. Here the time at $j = 0$ may either be an initial measurement or an a priori information matrix.

Explicitly, the information matrices are defined as

$$\Lambda(t_j, t_k) = \Phi^T(t_j, t_k) H_{x_j}^T R^{-1} H_{x_j} \Phi(t_j, t_k) \qquad (4)$$

where R is the covariance of the measurements at time t_j (often assumed to be diagonal and constant), H_{x_j} is the partial derivative of the observation vector with respect to the state at time t_j, and $\Phi(t_j, t_k)$ is the state transition matrix which describes the linear evolution of state perturbations about the nominal trajectory from time t_k to t_j.

Note that the information matrices $\Lambda(t_j, t_k)$ are at least positive semidefinite, and if the proper types of measurements (including possible a priori information) are taken, the sum of the information matrices will be positive definite for some $j \geq 0$; thus the inverse of the sum will exist if enough measurements are taken. As the information matrices are positive semidefinite, it is obvious that additional measurements will always cause the size of P_k to decrease or remain constant. Thus, to minimize the trace of the P_k matrix, or $tr(P_k)$, a common measure of the size of the covariance matrix, it is useful to increase the number of measurements N over any given data arc. Due to the nature of the dynamics of the state, the size of the information matrix will often vary with $| t_j - t_k |$ depending on the specific dynamics of the state (governed by the state transition matrices).

A consider parameter is a parameter whose exact value is not known, which has some influence on the measurements or dynamical environments of the state, and is not estimated. If such parameters are present, then the covariance of the state must be augmented to account for the uncertainty of the consider parameters. The resultant covariance matrix is called the consider covariance[1]:

$$P_{c_k} = P_k + S_k \Pi S_k^T \qquad (5)$$

where P_k is the computed state covariance matrix as defined above, Π is the covariance matrix of the consider parameters, and S_k is the sensitivity matrix. The sensitivity matrix can be defined as the partial derivative of the estimated state with respect to the consider parameters. For this discussion, it is useful to view it in a more elemental form:

[1] S. C. Wu, W. I. Bertiger, J. S. Border, S. M. Lichten, R. F. Sunseri, B. G. Williams, P. J. Wolff, and J. T. Wu, *OASIS Mathematical Description, Version 1.0*, JPL D-3139, Jet Propulsion Laboratory (internal document), Pasadena, California, pp. 6–11, April 1, 1986.

$$S_k = -P_k \sum_{j=0}^{N} \Phi^T(t_j, t_k) H_{x_j}^T R^{-1} \left[H_{x_j} \theta(t_j, t_k) + H_{c_j} \right] \qquad (6)$$

The repeated quantities are defined as before. The new quantities are defined as follows: H_{c_j} is the partial derivative of the observation vector with respect to the consider parameters at time t_j, and $\theta(t_j, t_k)$ represents the influence that the consider parameters have on the evolution of the state. The term $\theta(t_j, t_k)$ is defined as the particular solution to the nonhomogeneous equation

$$\dot{\theta}(t, t_k) = A(t)\theta(t, t_k) + B(t) \qquad (7)$$

$$\theta(t_k, t_k) = 0 \qquad (8)$$

where θ is a matrix of size $n \times p$ where n is the number of elements in the state and p is the number of consider parameters. $A(t)$ is the partial derivative of the equations of motion with respect to the state and $B(t)$ is the partial derivative of the equations of motion with respect to the consider parameters. The solution of this equation expresses the influence the consider parameters have on the state, and can be stated in terms of the state transition matrix:

$$\theta(t, t_k) = \int_{t_k}^{t} \Phi(t, \tau) B(\tau) d\tau \qquad (9)$$

$$\theta(t_k, t) = -\Phi(t_k, t)\theta(t, t_k) \qquad (10)$$

Thus the sensitivity matrix may be expressed as

$$S_k = M_k + \bar{\theta}_k \qquad (11)$$

where two new quantities are defined. These are

$$M_k = -P_k \sum_{j=0}^{N} \Phi^T(t_j, t_k) H_{x_j}^T R^{-1} H_{c_j} \qquad (12)$$

$$\bar{\theta}_k = \left[\sum_{j=0}^{N} \Lambda(t_j, t_k) \right]^{-1} \sum_{j=0}^{N} \Lambda(t_j, t_k) \theta(t_k, t_j) \qquad (13)$$

The term M_k denotes the effect of the consider parameter uncertainty as the uncertainty acts through the measurement alone; this term is called the measurement consider contribution. The effect of this term is usually to add a nearly constant uncertainty to the consider covariance. Its effect may dominate over $\bar{\theta}_k$ in many situations where the consider parameters do not enter strongly into the state dynamics, yet this term will in general not lead to unbounded growth in the consider covariance.

The matrix $\bar{\theta}_k$ represents the effects of the consider parameters on the state dynamics. This term may lead to covariance divergence and some of the counterintuitive orbit determination results which are associated with the use of consider filters. As the information matrices Λ are positive semidefinite, they may be viewed as analogues to masses or to weights, hence $\bar{\theta}_k$ may be viewed as analogous to the center of mass of the accumulated information matrix:

$$\left[\sum_{j=0}^{N} \Lambda(t_j, t_k)\right] \bar{\theta}_k = \sum_{j=0}^{N} \Lambda(t_j, t_k)\theta(t_k, t_j) \qquad (14)$$

Alternatively it may be viewed as an average or mean of the function $\theta(t_k, t_j)$ weighted by the information matrices. This is especially clear if the number of measurements becomes large, leading to

$$\bar{\theta}_k \approx \left[\int_{t_0}^{t_N} \Lambda(\tau, t_k) d\tau\right]^{-1} \int_{t_0}^{t_N} \Lambda(\tau, t_k)\theta(t_k, \tau) d\tau \qquad (15)$$

which is a classical definition of averaging. The term $\bar{\theta}_k$ is called the dynamic consider contribution.

B. General Properties of the Sensitivity Matrix

The sensitivity matrix with its measurement and dynamic contributions has some properties that may be stated without recourse to specific example. Some of these are listed below.

1. Performance Characterization. The performance of the consider covariance may be characterized by $tr(P_{c_k}) = tr(P_k) + tr(S_k \Pi S_k^T)$. If the consider parameter covariance Π is diagonal, as is often the case, then $tr(S_k \Pi S_k^T) = \sum_{i=1}^{p} \Pi_i \|S_{k_i}\|^2$, where $\| - \|$ is the usual Euclidian norm, Π_i is the ith consider parameter variance, and S_{k_i} is the ith column of the sensitivity matrix. As noted previously, $tr(P_k)$ always decreases with additional data. As will be shown, $\|S_k\|$ may increase with additional data, implying that the consider filter has failed. Note that in this formulation (diagonal Π matrix) the consider parameters may be discussed independent of each other.

2. Data Weighting Properties. A classical result associated with the sensitivity matrix is its invariance with respect to a common scaling of the weights (noise variances) of all data types. Note that in the definition of both the measurement and dynamic consider contributions, the inverse of the data noise matrix, R^{-1}, appears in both the numerator and denominator (or their matrix generalizations). Thus, should the data noise matrix R be scaled by some factor, the sensitivity matrix will be invariant with respect to this scaling. This implies that if there is only one data type being used, the consider contribution is independent of the data noise (assuming no a priori information). Thus for single-data-type consider filters the tuning strategy of artificially deweighting the data cannot be used to alter the sensitivity matrix, and data editing must be used instead. This will not be the case when process noise is added to the filter.

If there are two or more data types present, then the consider contribution may be altered by scaling the data weights relative to each other. This may be a useful procedure if one data type induces a large consider contribution with respect to the other data types. Then by deweighting that one data type, the consider contribution of this data type is reduced relative to the other data types. Unfortunately, in performing the deweighting, the computed covariance of the state will grow. If the data type in question is essential to the performance of the filter, then deweighting that data type may lead to poorer performance.

3. Measurement Consider Contribution. The measurement contribution M_k [Eq. (12)] is related to the uncertainty of the measurements with respect to the consider parameters. It has no terms that generically increase with time. While its magnitude may be large for effects such as the tropospheric-calibration-error or station-location-error contribution for the radio Doppler data type, or a range bias uncertainty for the radio range data type, it tends to remain at a fixed level.

4. Dynamic Consider Contribution. The dynamic contribution $\bar{\theta}_k$ [Eq. (13)] is related to the uncertainty induced in the state dynamical model by the consider parameters. As discussed earlier, it may be viewed as an average of the quantities $\theta(t_k, t_j)$ weighted by their respective information matrices $\Lambda(t_j, t_k)$. Recalling the definition of the θ [Eq. (9)], note that it is an integral over the time interval $t_k - t_j$. Hence, if this time interval grows large (and the information content of the matrix Λ does not decrease as the time from the measurement increases) these individual terms may grow large. Thus, if conditions are right, the mean value of these terms will begin to grow, leading to an increase in the size of the consider covariance if the computed covariance does not decrease swiftly enough. Some specific instances where such a failure may occur are discussed in the next subsection.

C. Some Batch Failure Modes and Their Remedies

The usual remedy for a large consider covariance, as discussed above, is to deweight the data type with the largest consider contribution. This approach will usually allow one to reduce the consider covariance somewhat, yet may not always be the most insightful remedy. It also requires tuning by the navigator, often a difficult and imprecise procedure.

1. High Information Content Divergence. One particular failure mode often encountered in practice is due to measurement data with both a dynamic consider contribution and a large information content. The immediate result of such data is to improve the computed covariance and to have only a small effect on the dynamic consider contribution, as the integral time spans will be short. If the subsequent data have a smaller information content, then the averaging process in computing the dynamic consider contribution will emphasize the terms related to the large-information-content data. As time progresses, these terms will grow at least linearly in time due to the integration effects and may cause the consider covariance to grow in an unbounded manner.

A simple example of this is a planetary (or planetary satellite) flyby navigation scenario, in which the ephemeris errors of the planet are considered and not estimated. While the spacecraft is flying by the planet the information content of data tends to be very large due to the increased dynamics of the spacecraft. However, the considered position errors of the planet will translate directly into an uncertainty in the model of forces acting on the spacecraft. At closest approach the dynamic consider contributions may be negligible, but soon after they will grow with time, as the post-flyby data will contain relatively less information.

The usual remedy to this failure mode is to restart the data arc using the minimized consider covariance at (or soon after) flyby as the a priori covariance matrix. Another remedy is to use the flyby information to estimate the planet position and reduce the otherwise fixed uncertainty in these parameters.

2. Flat-Information-Content Divergence. Another failure mode may be identified with information matrices having a uniformly constant information content. In this instance the dynamic consider contribution weights all its terms equally. Thus the older data will begin to dominate the mean as they grow at least linearly with time.

An example of this effect would be a spacecraft in an outer solar system cruise period, where the nongravitational accelerations acting on the spacecraft are treated as consider parameters. In such a system, without new data with a larger information content to increase the relative size of the information matrices, the consider covariance will grow.

There are several remedies to this mode of failure, the simplest being to discard old data after some time span. Unfortunately, this remedy also discards potentially useful information. Remedies that are more modern are either to model nongravitational accelerations as process noise terms in the dynamics, which forces the filter to give new information relatively greater weight, or to estimate the deterministic model parameters which represent the nongravitational accelerations.

D. Implications of Failure Modes

The fact that, in some cases, the sensitivity matrix may become large due to either the measurement or dynamic contribution indicates something more: that there is a significant information content in the measurements concerning the state estimate which could be exploited if it were not for the effects of the consider parameters. In fact, this is often why these parameters were considered in the first place, as they were estimated too quickly and too well with conventional filters and simplified models.

Some recent investigations into improved filter modeling have turned this fact into a useful strategy [1,2]. This improvement entails building a more sophisticated model for the consider parameter effects, especially the random components, and estimating this expanded parameter set, most of the members of which are treated as stochastic parameters. If appropriate noise variances and correlation times are chosen, the estimation process may not yield significant improvement in the knowledge of these parameters, yet may extract much of the information contained in the measurement partials of these parameters.

That the stochastic parameter uncertainties do not improve significantly in this approach is not an inherent problem, and can be understood in the following context. It is usually the case that these consider parameters are better determined by various intensive, off-line calibration techniques. With regard to station locations, for example, these techniques would take the form of many careful very long baseline interferometry measurements of quasars. Thus, if the current filter's information matrix with respect to these parameters is weighted against the accumulated data (which defined the a priori uncertainties in these parameters), the current information will have a negligible impact on the uncertainty. The introduction of stochastic parameters to account for modeling imperfections serves as a way to introduce such effects while exploiting the current measurement information to improve the state covariance.

IV. Sequential Mode

Now the case where process noise has been added into the filter is discussed. This alters the covariance computation scheme to a Kalman–Bucy approach. The basic definitions for this type of computation are briefly reviewed.

A. Sequential Consider Covariance Definitions

Assume that the state model is given by

$$x_k = \Phi_{k,k-1} x_{k-1} + \Gamma_{k,k-1} u_{k-1} \tag{16}$$

and that the measurement model is the same as in the batch mode [Eq. (2)]. Then computation of the state covariance when process noise is present is defined as follows:

$$P_k = [I - K_k H_{x_k}] \bar{P}_k \tag{17}$$

$$K_k = \bar{P}_k H_{x_k}^T \left[H_{x_k} \bar{P}_k H_{x_k}^T + R_k \right]^{-1} \tag{18}$$

The quantity P_k is the state (or computed) estimate error covariance computed at time t_k as before. Note that Eq. (17) is only valid when the gain K_k is the optimal gain matrix. The quantity H_{x_k} is the partial derivative of the observation vector at time t_k with respect to the state. The quantity K_k is the Kalman gain matrix computed at time t_k, and R_k is the measurement noise covariance. Finally, the matrix $\bar{P}_k$ is the computed covariance from time t_{k-1} mapped to the present time. Involved in the mapping are both the dynamics of the state and the uncertainty due to process noise in the intervening period:

$$\bar{P}_k = \Phi_{k,k-1} P_{k-1} \Phi_{k,k-1}^T + \Gamma_{k,k-1} Q_{k-1} \Gamma_{k,k-1}^T \tag{19}$$

Equations (17), (18), and (19) define the matrix Riccati equation for the discrete filter [4]. In Eq. (19) the matrix $\Phi_{k,k-1}$ is the state transition matrix from time t_{k-1} to time t_k, and P_{k-1} is the computed covariance at time t_{k-1}. The matrix $\Gamma_{k,k-1}$ maps the process noise state uncertainty from time t_{k-1} to time t_k. The matrix Q_{k-1} is generally the covariance of a Gaussian white noise sequence u_i where

$$E\left[u_i u_j^T\right] = Q_i \delta_{ij} \tag{20}$$

The term δ_{ij} is the usual Kronecker delta function. The process noise model may also be correlated in time.

The definition of the $\Gamma_{k,k-1}$ matrix is very similar to the $\theta(t_k, t_{k-1})$ function defined previously. Explicitly

$$\Gamma_{k,k-1} = \int_{t_{k-1}}^{t_k} \Phi(t_k, \tau) B_n(\tau) d\tau \tag{21}$$

In this equation $B_n(\tau)$ is the partial derivative of the state dynamics with respect to the process noise parameters.

As before, the consider covariance at time t_k is simply

$$P_{c_k} = P_k + S_k \Pi S_k^T \tag{22}$$

The formula for the covariance P_k is given by Eqs. (17), (18), and (19). The definition of the sensitivity matrix is now given as [3, p. 177]

$$S_k = [I - K_k H_k] \bar{S}_k - K_k H_{c_k} \tag{23}$$

$$\bar{S}_k = \Phi_{k,k-1} S_{k-1} + \theta_{k,k-1} \tag{24}$$

In the above formulae, H_{c_k} is the partial derivative of the observation vector with respect to the consider parameters at time t_k, and $\theta_{k,k-1}$ represents the sensitivity of the spacecraft dynamics due to possible errors in the consider parameters. Both of these terms were defined in the previous section. Note that there is not uniform agreement on the calculation of $\bar{S}_k$. Slightly different forms for this mapping may be derived depending on whether the consider covariance is first computed and then propagated forward in time, or the state estimates are first mapped in time and then the consider covariance is computed. The reason for this difference, which does not exist for the computed covariance, is due to the peculiarities of the consider covariance when mapped. This article follows Ref. [3] and Footnote 1 and uses the first definition given, that the consider covariance is first computed and then mapped. This definition is more in line with the use of the Kalman–Bucy filter. Note that if no process noise is present in the dynamics ($Q_i \equiv 0$), then the above definitions are equivalent to the definitions given in Section III.A.

B. General Properties of the Filter and Sensitivity Matrix

The Kalman–Bucy sequential filter with process noise exhibits some significant differences in behavior versus the batch filter without process noise. The major difference is that in the sequential filter old data are forgotten as time goes on, as the accumulated effect of process noise is to render measurements made some time ago not valuable. This occurs because the current covariance is continually made relatively larger by the process noise terms in the filter model, and hence the gain matrix K is increased to rely more heavily on the current data.

Some implications of this forgetting process for the sensitivity matrix follow. First of all, this process implies that the sensitivity matrix is no longer invariant under uniform scaling of measurement noise. Thus, data may be deweighted to influence the sensitivity matrix, even if only one data type is used. This was not the case with the batch mode. Another implication is that the dynamic consider contribution, which could lead to unbounded growth of the consider covariance in the batch mode, has a reduced effect. This is related to the forgetting property of the filter. As the time from the measurement increases, and as the potential dynamic consider contribution increases, the relative importance of this measurement is effectively deweighted automatically by the filter. Thus the addition of process noise to the dynamics largely eliminates the dynamic consider contribution, as the filter deweights old data autonomously. The dynamic contribution may still influence the covariance, but its effect should not be as pervasive or long lasting as in the batch mode.

In the sequential mode the major consider contribution comes from the measurement errors or the measurement consider contribution, as described earlier in Section III.B.3. This is seen from the update equation for the sensitivity matrix, Eq. (23), where the measurement partial derivative with respect to the consider parameters is directly scaled by the gain matrix, with no other attenuations acting on this term until it is mapped to the next time step.

C. Saturation of the Sequential Filter

One of the unique features of a sequential filter which contains process noise terms is that after enough measurements have been processed, a quasi-steady state condition will be achieved in which the variances of the estimated parameters reach a lower bound and do not decrease further. This is not the case for a batch filter in which no process noise is present. Assuming that the process noise effects are fairly large, then the elements of the mapped computed covariance, $\bar{P}_k$, may be significantly larger than the corresponding elements of the computed covariance P_{k-1}. If the filter has reached a quasi-steady state, which is assumed, then the update of the mapped computed covariance $\bar{P}_k$ to P_k must reduce the size of P_k to the same order as the previous computed covariance, P_{k-1}. This implies that the gain matrix K_k is relatively large, or that the update matrix $I - K_k H_{x_k}$ is relatively small. In this instance, increasing the data weight (decreasing the assumed measurement noise variance) will reduce the size of the current computed covariance, P_{k-1}, yet will have little effect on the computed covariance as mapped forward in time, $\bar{P}_k$. Thus, while increasing the data weight may yield reduced state uncertainties at the measurement time, it will not reduce the uncertainties after they are mapped forward in time, i.e., the filter is insensitive to the data noise. Of course, if the data weight were decreased by a sufficient amount, a significant degradation would eventually be observed.

D. A Possible Failure Mode

Now consider a case where the process noise is large. From the above discussion, the gain K_k is relatively large and the update matrix $I - K_k H_{x_k}$ is relatively small. Thus, an approximate result for the form of the sensitivity matrix at time t_k is

$$S_k \approx -K_k H_{c_k} + \cdots \tag{25}$$

leading to a consider covariance of

$$P_{c_k} \approx P_k + K_k H_{c_k} \Pi H_{c_k}^T K_k^T + \cdots \tag{26}$$

The neglected terms are of negligible size for a first-order analysis. It is clear now that if the filter places too much emphasis on the current measurement, i.e., increases the size of K_k, then the consider contribution may become unreasonably large. This behavior is due entirely to the measurement consider contribution. Some specific physical effects that would lead to this type of uncertainty when treated as consider parameters would be transmission-media calibration errors or unknown biases in the measurements. The traditional remedy for this failure mode is to deweight the current measurement data. As noted above, in this case it is possible to decrease the consider contribution by deweighting even if only one data type is present, an option not available in the batch mode. Again, if more than one data type is present, a better choice would be to deweight each measurement datum relative to the others in an attempt to retain precision.

E. Implications of the Failure Mode

One remedy of the failure mode that was discussed in Section IV.D is to deweight the current data, a process which must usually be accomplished through trial and error, especially when multiple data types are present. It would be highly desirable, though, to weight the measurements at their inherent accuracy and employ a filter model which accounts for error sources which are otherwise treated as consider parameters; such a filter will then automatically place the proper amount of emphasis on each measurement, which is determined by the assumed behavior of the modeled error sources. This is the approach behind the enhanced filter model, mentioned earlier, which employs stochastic and deterministic parameters to approximately represent the error sources affecting the measurements.

As an example of the potential benefit of enhanced filter modeling, an error covariance analysis was performed using an enhanced filter to reduce radio Doppler and ranging data in interplanetary orbit determination scenarios derived from the Mars Observer and Mars Environmental Survey Pathfinder missions [2]. In this study, the enhanced filter used to reduce the measurements contained stochastic process models to approximately represent the principal error sources affecting the measurements. The results predicted that orbit-determination accuracy improvement of factors of 2 to 4 could be realized over a conventional approach using a reduced-order filter, in which the measurement error sources were treated as consider parameters. Again, appropriate process noise terms must be added into the estimated-measurement error parameters to represent the actual random and uncertain nature of these error sources, and to prevent the filter from developing an unrealistic knowledge of them.

V. Summary and Conclusion

In this article some of the classical failure modes of reduced-order (consider) orbit-determination filters were discussed. By developing an informal description of these failure modes in terms of the information content of measurements, their error sources, and associated uncertainties, some relatively fundamental aspects of the well-known problems associated with the use of consider filters were established. A detailed description of specific failures of these types of filters is normally given in terms of mission-specific parameters and events; thus, it is difficult to go into greater detail about failure modes without recourse to specific examples.

The traditional motivation for the use of consider parameters was to simplify and speed up the filtering process in an age when computer runs often took hours, and computational efficiency and model simplifications were significant driving conditions on orbit determination procedures. The drawbacks of consider filters, i.e., the non-intuitive behavior and failure modes that were sometimes encountered, were accepted as the cost of using simplified models to increase efficiency. The recent advent of low-cost, high-speed computer workstations makes it possible to eliminate these undesirable consider contributions through the use of more sophisticated filter modeling (i.e., the enhanced filter), an important development. Enhanced filters may yield substantial improvements in orbit determination accuracy, with both existing data types and new data types proposed for future use. In particular, this makes it possible to achieve greater accuracies with simpler data types (such as those generated by the DSN's Doppler system) in future small, low-cost interplanetary missions, for which minimizing the resources and effort needed to support navigation is highly desirable. In the future, enhanced filters may enable very high accuracies for more ambitious missions in which relatively complex navigation techniques and data types may be employed.

Acknowledgment

The author appreciates the discussions with and suggestions by Sam Thurman concerning this article. His input gave a clear focus and intent to this work.

References

[1] S. W. Thurman and J. A. Estefan, "Radio Doppler Navigation of Interplanetary Spacecraft Using Different Data Processing Modes," paper AAS 93–163, presented at the AAS/AIAA Spaceflight Mechanics Meeting, Pasadena, California, February 22–24, 1993.

[2] J. A. Estefan, V. M. Pollmeier, and S. W. Thurman, "Precision X-Band Doppler and Ranging Navigation for Current and Future Mars Exploration Missions," paper AAS 93-250, presented at the AAS/GSFC International Symposium on Space Flight Dynamics, NASA Goddard Space Flight Center, Greenbelt, Maryland, April 26–30, 1993.

[3] G. J. Bierman, *Factorization Methods for Discrete Sequential Estimation*, San Diego, California: Academic Press, 1977.

[4] R. S. Bucy and P. D. Joseph, *Filtering for Stochastic Processes with Applications to Guidance*, New York: Chelsea Publishers, p. 57, 1987.