NASA-TM-111666

FLIGHT DATA RESULTS OF ESTIMATE FUSION FOR SPACECRAFT RENDEZVOUS NAVIGATION FROM SHUTTLE MISSION STS-69

J. Russell Carpenter * and Robert H. Bishop †

A recently developed rendezvous navigation fusion filter that optimally exploits existing distributed filters for rendezvous and GPS navigation to achieve the relative and inertial state accuracies of both in a global solution is utilized here to process actual flight data. Space Shuttle Mission STS-69 was the first mission to date which gathered data from both the rendezvous and Global Positioning System filters, allowing, for the first time, a test of the fusion algorithm with real flight data. Furthermore, a precise best estimate of trajectory is available for portions of STS-69, making possible a check on the performance of the fusion filter. In order to successfully carry out this experiment with flight data, two extensions to the existing scheme were necessary: a fusion edit test based on differences between the filter state vectors, and an underweighting scheme to accommodate the suboptimal perfect target assumption made by the Shuttle rendezvous filter. With these innovations, the flight data was successfully fused from playbacks of downlinked and/or recorded measurement data through ground analysis versions of the Shuttle rendezvous filter and a GPS filter developed for another experiment. The fusion results agree with the best estimate of trajectory at approximately the levels of uncertainty expected from the fusion filter's covariance matrix.

INTRODUCTION

Most current rendezvous scenarios require many hours of ground tracking of both vehicles in order to generate inertial ephemerides for the two spacecraft accurate enough to compute rendezvous maneuvers by the chaser vehicle. The maneuver computations may occur on the ground, onboard the chaser, or both during most of the rendezvous. However, as the relative distance becomes ever smaller, less and less time is available for performing and correcting the maneuvers, so that safety and mission success concerns dictate that onboard targeting becomes the primary guidance. During this phase, the accuracy of the relative state estimates become much more important than the inertial, since the main effect of *inertial* navigation errors is inaccurate *long-term* propagation. Hence, the chaser may accomplish maneuvers in the proximity of the target using only a relative sensor producing relative state estimates. Fig. 1 shows a sketch of such a scenario involving accurate relative states but inaccurate inertial states.

More and more spacecraft now carry sensors capable of accurately determining the vehicles' inertial states, making available the opportunity for augmenting the rendezvous technique just described. For example, a Standard Positioning Service GPS receiver onboard the chaser vehicle produces inertial state estimates on the order of 100 meters accuracy. Constraining the inertial position of the chaser using the GPS anchors the accurate relative state derived from the relative sensor to its

*Aerospace Engineer, Aeroscience and Flight Mechanics Division, NASA Johnson Space Center, Houston, Texas 77058.

†Associate Professor, Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, Texas 78712.

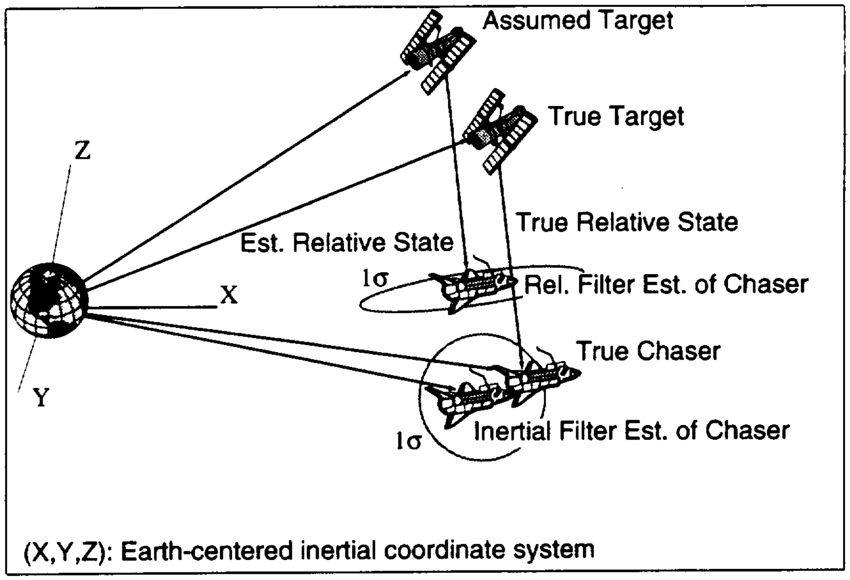

Figure 1: Sketch showing accurate relative state but poor inertial state estimates

proper position in inertial space, thereby yielding a target inertial state of comparable accuracy. The increased accuracy of the inertial states will more accurately locate the vehicles with respect to variations in the gravity potential, thereby improving predictions of their future trajectories. Such information provides more accurate maneuver targeting with fewer correction burns, thus saving consumables and other operational resources. Also, some relative state estimators contain suboptimalities resulting from implementation considerations that may degrade relative estimation accuracy more when inertial errors are large.

In particular, the Space Shuttle's rendezvous filter design makes the suboptimal assumption that it has a *perfect* target vehicle inertial state. Various unusual aspects of the filter's behavior have been attributed to this assumption, although officially NASA does not consider these characteristics to be performance or safety issues. However, NASA is retrofitting the Shuttle fleet with a Precise Positioning Service GPS receiver, and is considering how to use the GPS during Shuttle rendezvous. The *design philosophy* for integrating GPS into the Shuttle's avionics system is *to minimize impacts to existing systems* in hope of cost savings. The Shuttle/GPS integration design treats GPS states as if they are the inertial states periodically uplinked by ground operators during Shuttle missions. Once the onboard targeted phase of rendezvous begins, ground uplinks of inertial states typically cease. In fact, a procedure followed by the crew purposefully degrades inertial state accuracy to make updates to the relative state more visible, thereby confirming "correct" operation of the filter. When GPS becomes operational, the crew will be able to update the Orbiter inertial states with GPS at any time, but NASA has not yet decided whether or not they will actually perform such updates.

A parallel development in Shuttle rendezvous techniques is the increasing use of laptop computers with serial interfaces to the Orbiter avionics system and to various additional sensors, such as hand-held and payload-bay-mounted laser radars. The crew uses these computers and sensors as *situational awareness tools* during rendezvous maneuvers. Often, the crew uses these tools intensely during the manual phase of the rendezvous, which occurs during the last few hundred meters of the approach, because the avionics system's relative state becomes unusable due to accuracy limitations of its rendezvous radar. On the drawing board are plans for the laptop computer tool to use GPS data from the Orbiter, and also from the target, when available. When GPS data from both vehicles is available, the laptop computer tool should regularly exceed the inertial state accuracy of the Orbiter avionics system and perhaps the ground controller states, and yield comparable or better relative state accuracy than the existing systems. How the Shuttle will use the relative GPS data remains to be seen.

In many cases however, the target vehicle will have no GPS receiver, or no means of communicating the GPS data even if present to the Orbiter. In this scenario, fusion of the GPS and rendezvous radar data could provide accuracies approaching that of relative GPS. While this fusion could be accomplished with a standard, centralized Kalman filter processing both raw data types, such a filter would largely duplicate the effort expended by the existing Kalman filters resident in the avionics system and GPS receiver/processor. Replacing existing systems with a new filter would require a major costly re-verification of the entire relative navigation system. This is unreasonable, and would not likely occur. Fusing the state estimates of these filters instead, as advocated in this paper, may offer a better solution. Due to the excessive cost and risk associated with modifying the Orbiter flight software, either approach would almost certainly have to be implemented in the laptop computer tool. Although the laptop is in fact more capable than the avionics system computers, competition for its resources dictates that the most efficient strategy be used for any state estimation functions.

In a recent work, Carpenter and Bishop [1] present a solution to the problem of fusing two Kalman filters operating in parallel, in the context of spacecraft navigation. The basic fusion algorithm is identical to Bar-Shalom's [2], but was developed from a different point-of-view. The basic algorithm has been generalized to fusion of two filters with noncommon states and extended to allow feedback of the fused data to the filters while avoiding a singularity constraint (Ref. [3]). The generalized algorithm has been specialized to accommodate the suboptimal perfect target assumption in the Space Shuttle rendezvous filter (Ref. [4]). The proposed method requires only that a cross-covariance be maintained by the fusion filter, in contrast with other methods, such as those of Speyer (Ref. [5]), Kerr (Ref. [6]), Bierman (Ref. [7]), and Carlson (Ref. [8]), which require maintenance of a covariance for the fused state. This is significant, because in the absence of strong correlations between the filters, the cross-covariance may be eliminated or maintained using simpler algorithms than might be required for covariance maintenance. Further, the proposed rendezvous navigation fusion filter is well suited to the problem of retrofitting GPS onto the Space Shuttle because it avoids modifications to existing GPS and Space Shuttle navigation filters, unlike other approaches cited above, all of which require modifications of one sort or another to the local filters, such as computing an additional data vector or adjusting the local processors to eliminate cross-covariances.

Until the Fall of 1995, the new approach could not be tested with actual flight data, because of the few Shuttle missions which had flown GPS, none had successfully collected simultaneous GPS and rendezvous navigation data. The first time such an event occurred was on STS-69, which launched September 7, 1995. Two deploys and rendezvous were performed, one pair of which was with the Wake Shield Facility (WSF), the subject of two GPS experiments, a University of Texas at Austin Center for Space Research (UT/CSR) precise orbit determination experiment (Ref. [9]), and a joint NASA/European Space Agency (ESA) real-time relative GPS (RGPS) experiment (Refs. [10] and [11]). This paper presents results from fusing data from playbacks of the STS-69 mission data. These playbacks are generated from downlinked telemetry files processed in ground analysis versions of the Orbiter onboard navigation filter and the real-time relative GPS experiment filter operating in its single-vehicle mode. The results are evaluated by comparing them against precise orbit determinations resulting from the U.T. Austin Center for Space Research experiment (Ref. [12]).

ALGORITHM DESCRIPTION

This section briefly reviews the algorithm of Ref. [4], and describes some relevant details specific to implementing the algorithm for the flight data processing.

Fusion Algorithm

Because of the Shuttle rendezvous filter's perfect target assumption, the covariance of the target estimation errors, P_{t1}, is implicitly taken to be zero. Consequently, the covariance of the chaser state estimation errors, P_{c1}, and the relative state estimation errors, P_{rel1}, when expressed with respect to the same coordinate system, are identical. Ref. [4] exploits this insight to show that a general linear fusion of the augmented chaser-relative state from filter 1 and the chaser-only state from filter 2 can be performed, and showed that it is unbiased as long as the two filters being fused are unbiased. If the rendezvous filter has a target bias, as one expects it will since it does not update its target, the fusion removes this bias. Ref. [4] also showed that a feedback of the fused state and covariance to the rendezvous filter removes any chaser bias which originated from the biased target.

The main difference between the simplified system considered in Ref. [4] and this work is that the two filters being fused have additional states beyond merely chaser and target position and velocity. In addition to sharing the chaser states in common, both filters assume that the chaser dynamics include a time-correlated disturbance which accommodates unmodeled accelerations. Both filters also have several unique states which accommodate measurement biases. Ref. [3] shows how to accommodate such a case by extending the methods of Refs. [1] and [2]. Let $\boldsymbol{\xi}$ denote the common states, and $\boldsymbol{\omega}$ the unestimated parameters in filter 1 that one wishes to estimate via fusion (i.e. the target state). Then, using a straightforward extension of Refs. [1], [2], and [3], Ref. [4] shows that fusion of the augmented chaser-relative state from filter 1 and the chaser-only state from filter 2 can be performed as follows:

$$\begin{bmatrix} \hat{\boldsymbol{\xi}}_* \\ \hat{\boldsymbol{\omega}}_* \end{bmatrix} = \begin{bmatrix} (I - W_\xi) & 0 \\ 0 & -M_\omega \end{bmatrix} \begin{bmatrix} \hat{\boldsymbol{\xi}}_1 \\ \hat{\boldsymbol{\xi}}_1 - M_\omega^T \boldsymbol{\omega} \end{bmatrix} + \begin{bmatrix} W_\xi \\ M_\omega \end{bmatrix} \hat{\boldsymbol{\xi}}_2, \tag{1}$$

where $M_\omega = [I_6, \ O_{63}]$ is a matrix which selects the chaser elements from the common state vector. The states unique to the filters are not updated by the fusion directly, but their correlations with the common states may be updated from the fusion, and if this information is fed back to the filters, it will influence the filters' estimates of the unique states.

The recursion for the optimal fusion gain is based on the filters' covariance matrices for the common states, $P_{\xi\xi i}$, $i = 1, 2$, and a cross-covariance matrix, P_{12}, which accounts for any correlation between the filters. Fig. 2 shows a schematic summary of the fusion algorithm described in Ref. [4], which indicates the sequence of propagation, update, and reinitialization, along with the equations one must implement for each stage. For clarity, the figure shows both filters incorporating measurements at the same rate. Reinitializations occur every other update cycle, in keeping with the necessary condition described in Ref. [3]. In the figure, the subscripts c and t denote the chaser (Shuttle) and target vehicle states, respectively. Note that the covariance of the estimate computed by Eq. (1), P_*, the error covariance of the optimal fusion, does not form part of the fusion recursion, so does not have to be maintained by the fusion filter, unless it is going to be used for feedback to one or both of the filters.

Specific details of the implementation of the algorithm shown in Fig. 2 to the Space Shuttle rendezvous filter and the RGPS flight experiment filter may be found in Ref. [3].

Modifications for Flight Data Processing

One of the lessons commonly learned when processing real world data is to always expect the unexpected. Processing the STS-69 data reinforced this lesson. Although the simulation results of Ref. [3] indicate the fusion can accommodate up to 10σ inertial state dispersions, much larger dispersions can occur in a Shuttle rendezvous, and did on STS-69 during the WSF deploy.

Over the years, Kalman filter designers have invented a number of schemes for handling the vagaries of real world data. Two of the most successful of these schemes are innovations monitoring and data underweighting. In the former, the Kalman filter compares its innovations to their covariance matrix, and if the ratio of these exceeds a threshold, the filter rejects the measurement. This

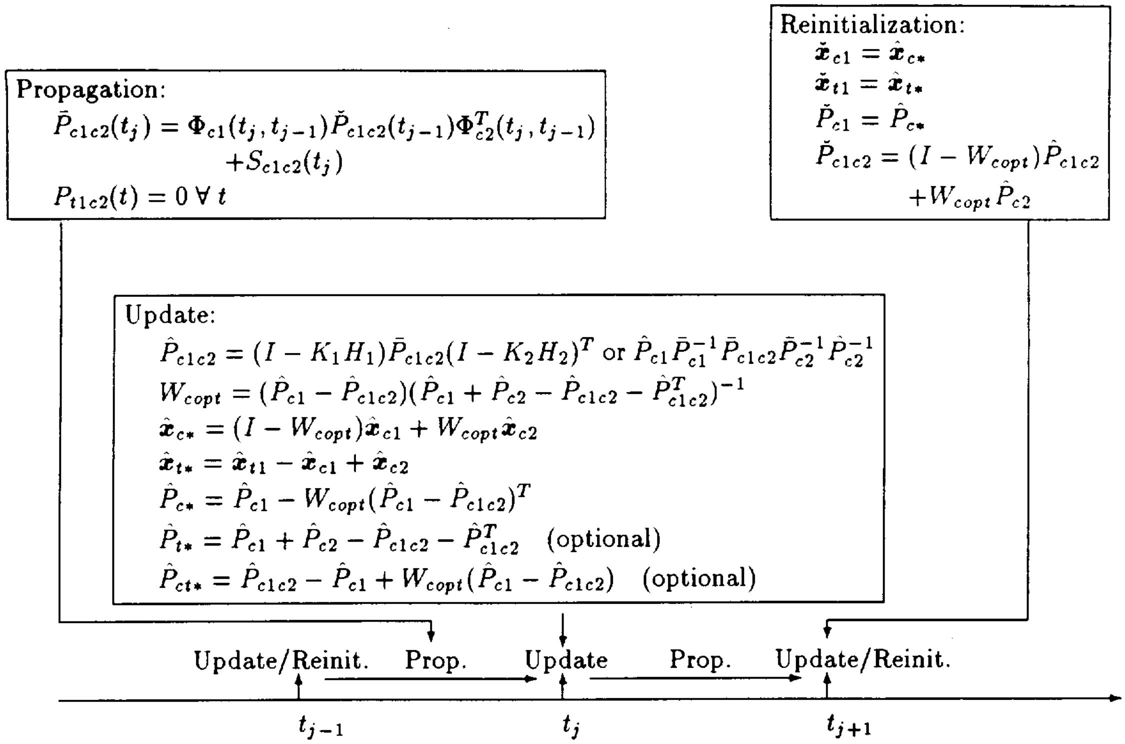

Figure 2: Summary of Algorithm for Fusion

scheme will often allow a Kalman filter to recover from transient conditions of poor geometry or degraded measurement accuracy, since as it rejects more and more measurements, the filter's covariance becomes larger. When the transient condition has passed, the filter's larger covariance allows it to reconverge even if its state has drifted from its correct value. Underweighting is a common scheme used to get standard Kalman filters to treat highly non-linear measurements with greater care. The extended Kalman filter's truncation of the measurement partials to first order often causes it to make larger updates than are warranted by the true geometry of a non-linear problem; the various underweighting schemes all aim to shrink the size of this update by reporting to the filter a larger innovations covariance.

In the context of the fusion of Eq. (1), the difference between the common states of filters 1 and 2, $\boldsymbol{d}_{12}$, serves the role of a residual, which one can see by rewriting the common state update as

$$\hat{\boldsymbol{\xi}}_* = \hat{\boldsymbol{\xi}}_1 - W_\xi[\hat{\boldsymbol{\xi}}_1 - \hat{\boldsymbol{\xi}}_2].$$

The inverse of the covariance of this difference,

$$P_{d12} = P_{\xi\xi1} + P_{\xi\xi2} - P_{\xi1\xi2} - P^T_{\xi1\xi2},$$

is computed as part of the fusion gain computation, as shown in Fig. 2. Hence, it is easy to compute the fusion's functional counterpart to the Kalman filter residual ratio, and check it against some threshold. Since successful fusion has been accomplished with 10σ inertial state dispersions, an appropriate fusion edit test is

$$\boldsymbol{d}^T_{12}P^{-1}_{d12}\boldsymbol{d}_{12} < 10.$$

If this condition is violated, do not perform fusion.

Unfortunately, in the data from STS-69, *all* potential fusions violate this threshold, by quite a large margin. Disregarding the edit test quickly leads to divergence. The cause of this difficulty is the rendezvous filter's perfect target assumption, which during the WSF deploy caused large inertial errors. Reasoning that both filters are adequately screening their own measurements with their own internal residual monitoring, one concludes that causes other than the perfect target assumption for exceeding a fusion edit threshold as large as ten are unlikely. So, in the context of this problem, exceeding an edit test whose threshold is set very large is an indication that the rendezvous filter's chaser state errors are much larger than its covariance matrix.

Table 1 displays an underweighting scheme which can accommodate this situation. This algo-

1. **WHILE** $d_{12}^T P_{d12}^{-1} d_{12} > 10$
2. $P_{\xi 1} \leftarrow 2P_{\xi 1}$
3. $P_{\xi 1 \xi 2} \leftarrow \sqrt{2} P_{\xi 1 \xi 2}$
4. $P_{d12} = P_{\xi 1} + P_{\xi 2} - P_{\xi 1 \xi 2} - P_{\xi 1 \xi 2}^T$
5. **INCREMENT** *counter*
6. IF *counter* > 20, *fault* = .ON., **BREAK**
7. **END WHILE**
8. **IF** *fault* = .OFF., **DO FUSION**

Table 1: Rendezvous fusion underweighting algorithm

rithm provides satisfactory relative and inertial performance in cases for which fusion would not otherwise be possible. Its robustness to so-called "soft" failures of the Kalman filters, in which measurements are not rejected by the filters but they nevertheless diverge is of course open to question. However, a fundamental assumption in the derivation of the fusion algorithm is that both filters are stable and operating optimally. Although the operational features introduced in this section extend this assumption to many realistic situations of non-optimality in the filters, fusion ultimately cannot fix a broken filter.

RESULTS

As pointed out by Schutz, et al. (Ref. [9]), a hardware failure caused the WSF GPS receiver to stop tracking approximately 19 hours after deploy[1]. Hence, although fusion can be performed for the rendezvous, a precise best estimate of trajectory (BET) is not available during this time period for comparison. During the WSF deploy, on September 11 (day of year 254), the Orbiter only performed a few relatively short duration arcs of rendezvous radar tracking of the WSF. One of the longer arcs for which there is a BET is a 20 minute data set running from about 12:35 to 12:55 Universal Time Coordinated (UTC). During this radar track, only radar range and range-rate measurements were flagged good by the Orbiter flight computers. Although a somewhat longer data set with all radar data types available would be superior, this data appears to be adequate for the purpose of validating the simulation results of Ref. [4] with flight data.

This section depicts two of several variations on the data set that have been examined. In both, playbacks of both filters start simultaneously near the beginning of the data set, and both

[1] This failure has since been isolated to a defectively documented power converter in the external data recorder which provided the interface between the receiver and the WSF.

cases process the same downlisted and/or recorded measurement data. The variation is among two different pairs of initialization vectors for the filters. The first case, "as flown," initializes the filters as they were during the flight, to get some sense of how a fusion processor operating in real time during this mission would have performed. The second case is motivated by the planned future capability for Orbiter crews to initialize the rendezvous filter with a chaser vector based on the GPS state. If a valid target state exists, the GPS state may update it too, by subtracting the existing relative state from the new GPS chaser state (Ref. [13]). The second case initializes the rendezvous filter in this manner to see what additional benefit fusion could provide throughout the subsequent tracking period.

Overview of Results

The reader may glean an overview of the results of the two variations from Figs. 3-5. These plots show the trajectories generated by each of the filters and the fusion filter, along with the BET trajectories from Ref. [12], relative to the Orbiter rendezvous filter's assumed target position. This relative motion is shown in the target orbit plane, with motion downtrack from the target on the horizontal axis and radially from the target on the vertical axis, in units of meters. Fig. 5 depicts a zoomed in view detailing the behavior of the chaser estimates for the second case.

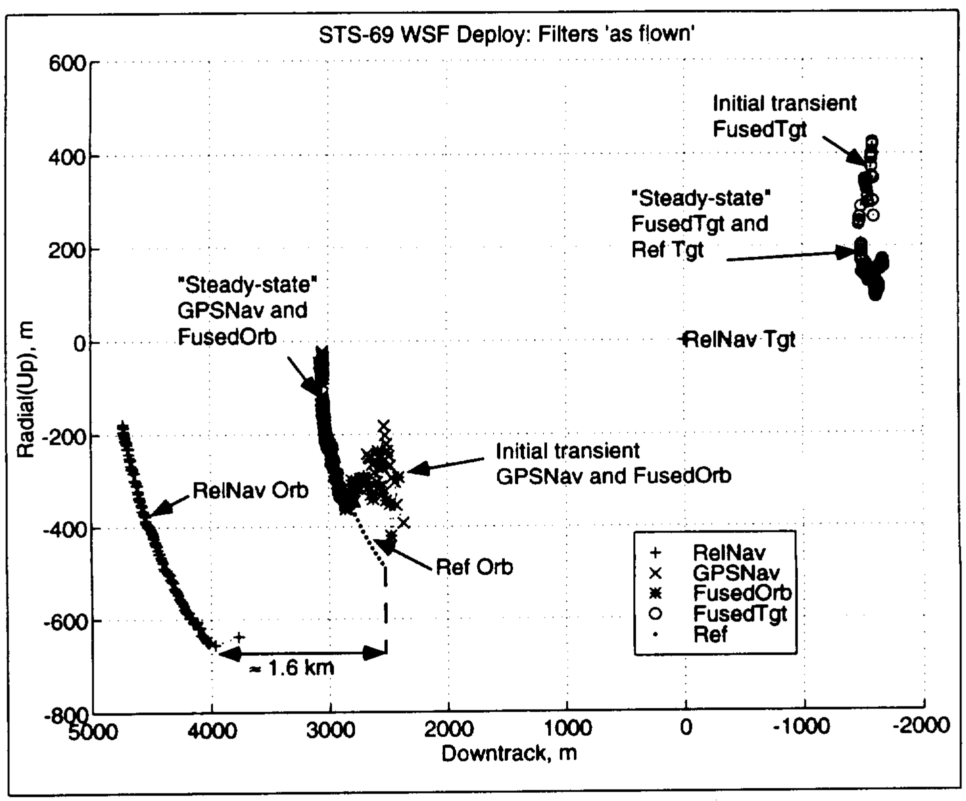

Figure 3: Comparison of estimates relative to target assumed by RelNav, filters initialized "as flown"

Noted on the plot of the as flown case, Fig. 3, is a downtrack bias of nearly 1.6 kilometer in the Orbiter rendezvous filter. The STS-69 mission had a particularly clear reason for the existence of such a bias. The WSF was a rare example, for Shuttle missions, of a maneuvering target vehicle. Its maneuver consisted of a continuous, low magnitude thrusting period to ensure separation from the Orbiter without contamination by Orbiter jet plumes. The large downtrack error of the Orbiter rendezvous filter is probably due to this maneuver, since not only does the filter assume perfect knowledge of the target, it assumes the target is passive. After WSF deploy, the filter propagated the WSF as a passive vehicle, when in fact it was maneuvering. By the time the radar data used in this study became available to the filter, about one hour, ten minutes after deploy, its erroneous propagation of the target produced a large residual in the first measurements. Since the filter could not adjust its target state to eliminate this residual, it moved the Orbiter downtrack to accommodate it. Because of the large downtrack errors, the fusion's edit ratios, based on the difference between the rendezvous and GPS filter's estimates, cause the fusion filter to invoke its underweighting scheme. The fusion inflates the rendezvous filter's covariance to such a degree that the fusion gain strongly weights the GPS state. During the initial transient of the GPS filter, the fusion commits errors of a few hundred meters for the chaser and for the target. Once the GPS filter converges however, it and the fusion appear to accurately track the BET reference.

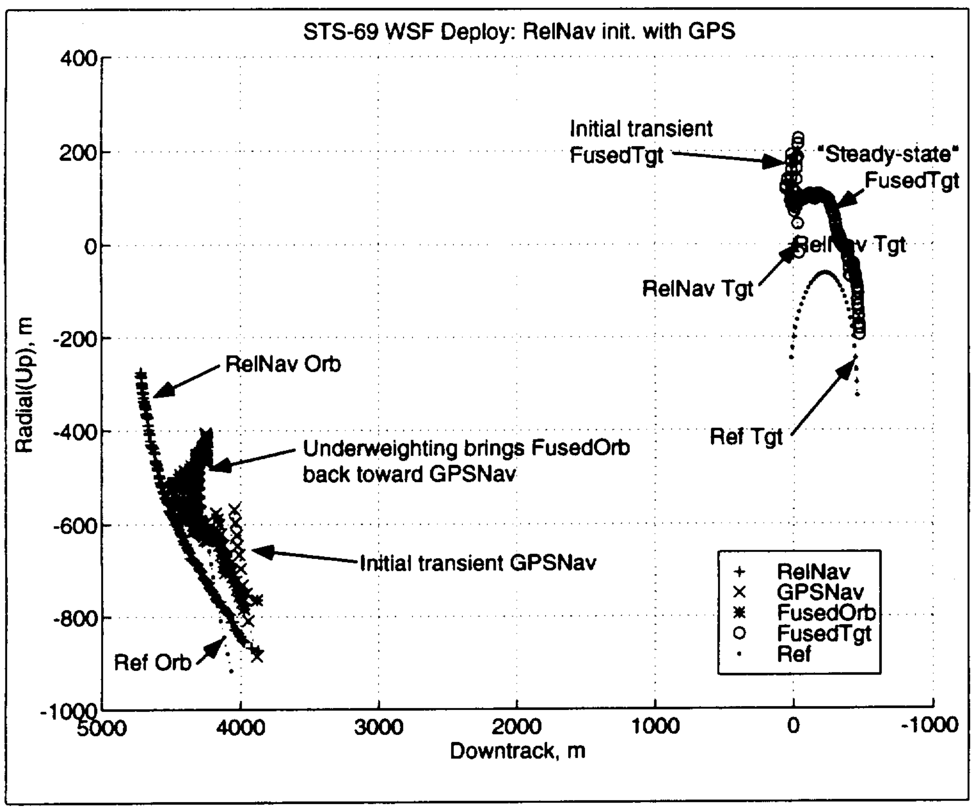

Figure 4: Comparison of estimates relative to target assumed by RelNav, RelNav filter initialized with GPS states

In the case which initializes the rendezvous filter with the GPS state, which Figs. 4 and 5 show, one sees that the error in GPS state used for initialization has an error from the BET of a couple hundred meters, large, though not out of the question for a Standard Positioning Service orbital position fix. As in the as flown case, the error in initial conditions causes inertial divergence of the rendezvous filter, although to a smaller degree than the as flown case. The fusion chaser estimate lies mostly somewhere between the rendezvous filter and the GPS filter estimates, indicating that it is weighting the two estimates approximately equally. During about the last half of the playback,

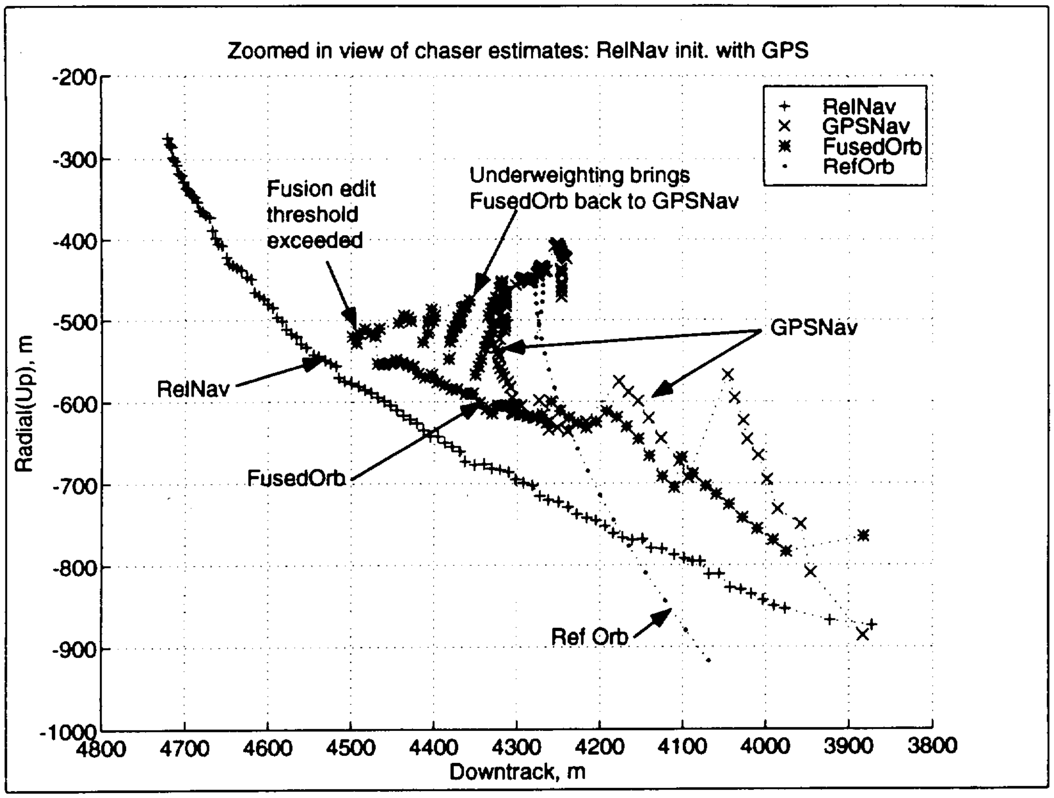

Figure 5: Detail of Fig. 4, chaser estimates relative to target assumed by RelNav, RelNav filter initialized with GPS states

the fusion begins to more closely track the rendezvous filter in the downtrack direction (Fig. 5). However, when the difference between the rendezvous filter and the GPS filter becomes too large for the fusion's edit test, it begins to underweight the rendezvous filter, and the fusion moves back towards the GPS, narrowing the inertial estimation error. The fusion also begins to more accurately track the target at about this point (Fig. 4). The fusion's target estimate initially has a radial bias which is somewhat worse than the rendezvous filter, but appears to track the downtrack motion of the target more accurately than the rendezvous filter. Eventually, the fusion's target estimate approaches the BET reference target to within about 100 meters.

Estimation Error Comparisons

Figs. 6-13 are plots comparing the estimation errors of the rendezvous filter and the fusion. For these plots, Ref. [12] provides the truth trajectory. The figures depict estimation errors of two types: Orbiter inertial state errors, and relative state errors. The plotting program propagates the UT/CSR BET data over up to one second using a spherical Earth force model to synchronize the data time tags. For all the plots, heavy dots signify the sampled differences between filter and BET, and lighter dotted lines connect these marks to clarify trends. The figure also plots the square roots of the diagonals of the filter covariances with dashed lines to indicate the envelope of 1σ errors the filters assume. Time is indicated in minutes elapsed since the day of year and UTC listed below the plots. The inertial comparison plots show the differences in components instantaneously aligned with the UT/CSR state vector's radial, downtrack, and crosstrack directions. For relative comparisons, the three coordinates are aligned with the line of sight vector from target to chaser, its normal in the plane defined by the two vehicle's position vectors, and an out of plane vector, normal to the other two.

The first four comparison plots, Figs. 6 through 9, depict the as flown case. One sees in Fig. 6 that the rendezvous filter commits a downtrack position error of up to 1600 meters, while its assumed standard deviation is on the order of 100 meters. These errors are are completely corrected by the fusion, which commits estimation errors generally in agreement with its covariance matrix, as Fig. 7 shows. During this entire case, the fusion is almost completely relying on the GPS filter due to underweighting of the rendezvous filter, so most of the transient is merely the GPS filter's convergence transient. Figs. 8 through 9 show that the relative state estimation errors are comparably small for both the rendezvous filter and the fusion, although the relative position error covariance for the fusion is much larger, reaching a 1σ level of over 20 kilometers in the out-of-plane position component as Fig. 9 shows. The comparably large size of the fusion's relative covariance is actually indicative of the proper size of the target covariance, which is estimated by the fusion filter, but assumed to be zero by the rendezvous filter. The large target covariance of the fusion filter indicates the size of the rendezvous filter covariance after the fusion underweights it enough to accommodate 1.6 kilometers in downtrack error.

Figs. 10 through 13 show the estimation errors for the second case. Note in the inertial error comparisons, Figs. 10 and 11, that the fusion's chaser 1σ envelope is better than half the rendezvous filter's; this is due to the use of independent information from the GPS filter. However, during the first half of the playback, the fusion's errors consistently exceed their 1σ envelope in downtrack position, and the rendezvous filter's downtrack errors exceed their envelope throughout the playback, reaching nearly 600 meters in downtrack position by the end. Only when the fusion begins editing and underweighting the rendezvous filter do its errors come down to expected levels. One could choose a smaller state vector difference threshold for the edit test to bring the fusion's errors down sooner. Doing so would be undesirable though, since a smaller edit threshold could cause the algorithm to erroneously underweight the rendezvous filter whenever a transient from the GPS filter caused the state vector difference to be large. The basis of the underweighting scheme is the assumption that really large differences can be ascribed only to the suboptimal perfect target assumption of the rendezvous filter. For smaller differences, such a contention loses its authority. In terms of relative state performance, which Figs. 12 and 13 show, the fusion is somewhat noisier than the rendezvous filter in line of sight position once it starts to downweight the rendezvous filter's states. However, the fusion more than doubles its 1σ error envelopes as the errors increase, honestly reporting its progress.

Tuning the Fusion Filter

The data from the filters exhibit very little cross-correlation. The cases shown above use only a small initial cross-covariance, with zero process noise correlation. The initial cross-covariance is an empirical estimate, determined from averaging sample cross-covariance estimates for several overlapping segments of common state estimation errors. Note that the rendezvous filter's unmodeled acceleration covariance is typically an order of magnitude smaller than the GPS filter's, and although the GPS filter unmodeled acceleration states often approach 100 milligees, the rendezvous filter unmodeled acceleration states rarely exceed 0.1 milligee. A number of attempts to manually tune the process noise cross-covariance showed either unnoticeable performance differences, or outright divergence when too much correlation was applied, causing the difference covariance to become negative definite. Ref. [3] shows that, for this problem, even when the filters' process noises are maximally correlated, ignoring the cross-covariance causes very little performance degradation. One may conclude that when the filters subject to fusion are not obviously correlated by design, neglecting the cross-covariance may be a harmless omission.

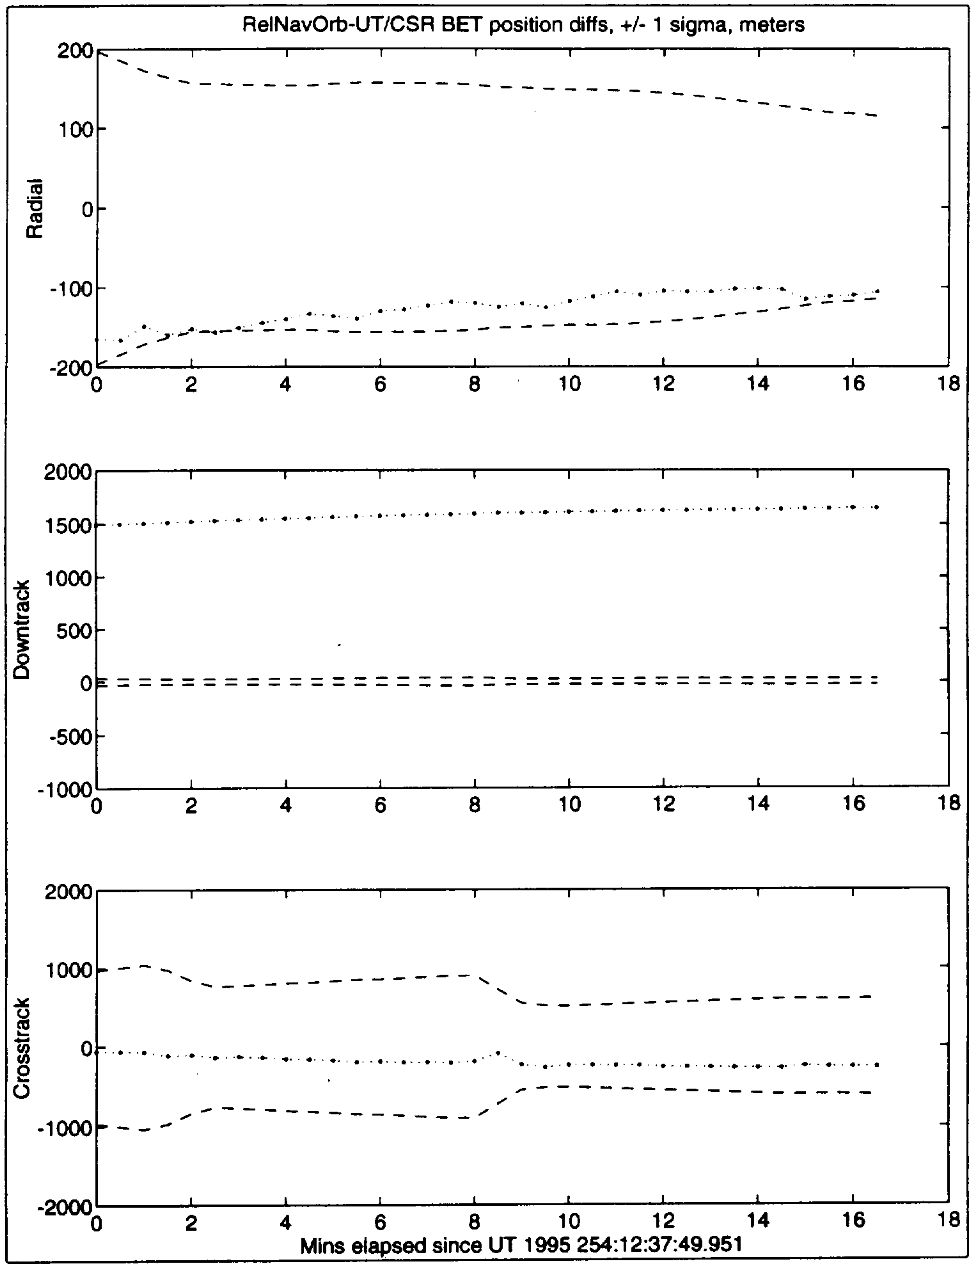

Figure 6: Rendezvous filter inertial position errors, filters initialized "as flown"

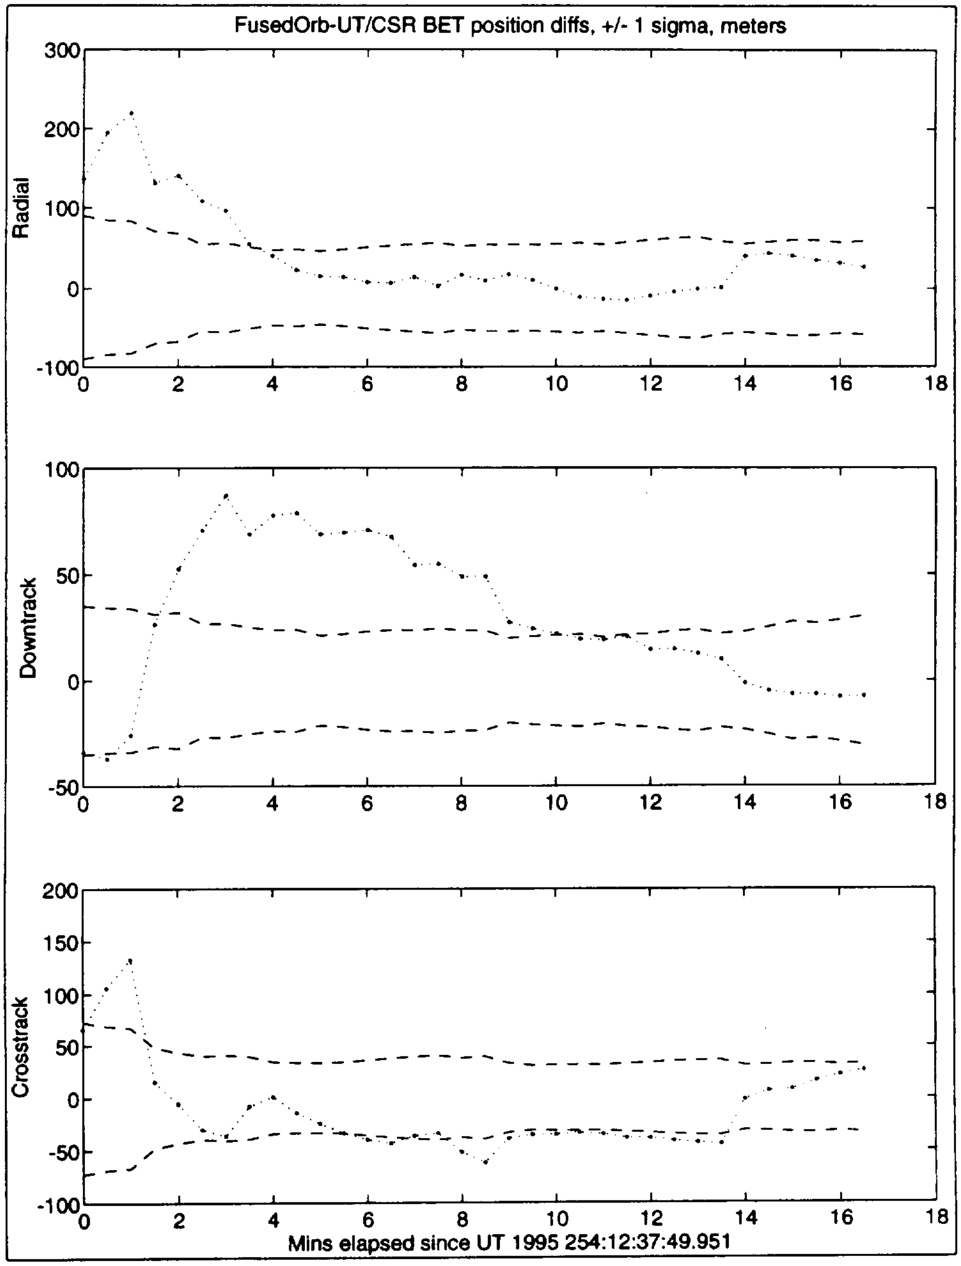

Figure 7: Fusion filter inertial position errors, filters initialized "as flown"

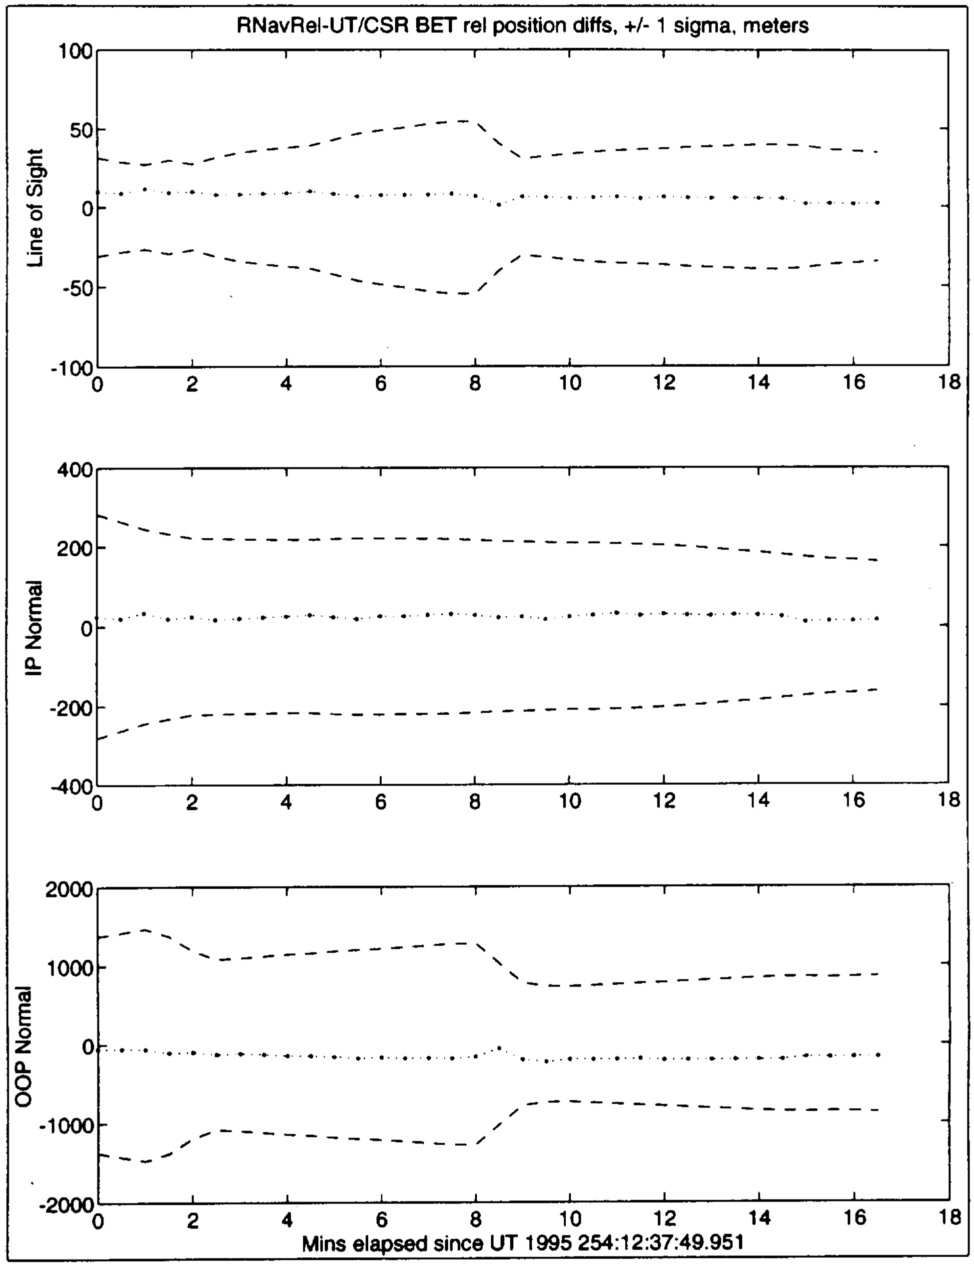

Figure 8: Rendezvous filter relative position errors, filters initialized "as flown"

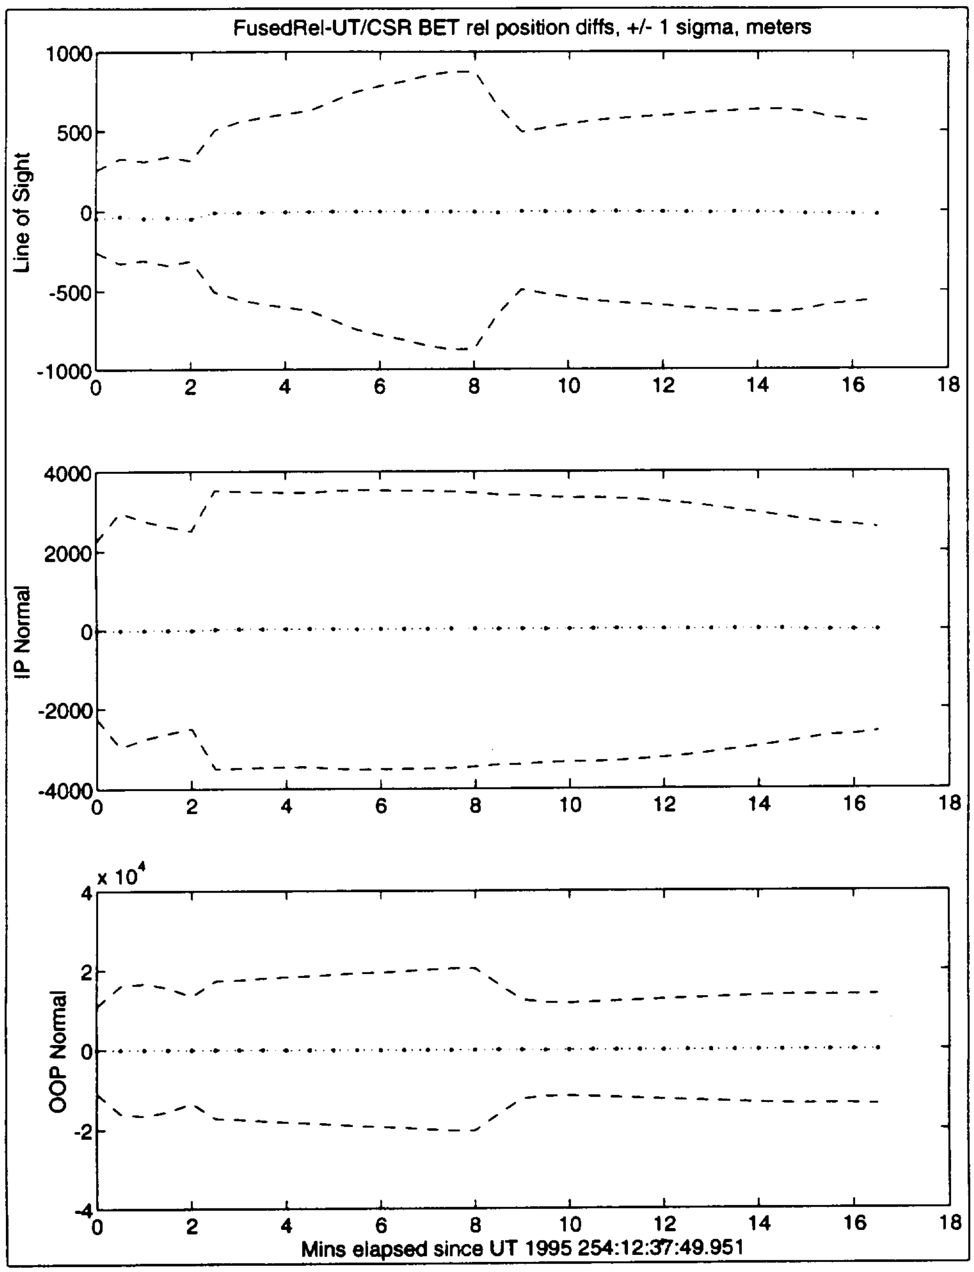

Figure 9: Fusion filter relative position errors, filters initialized "as flown"

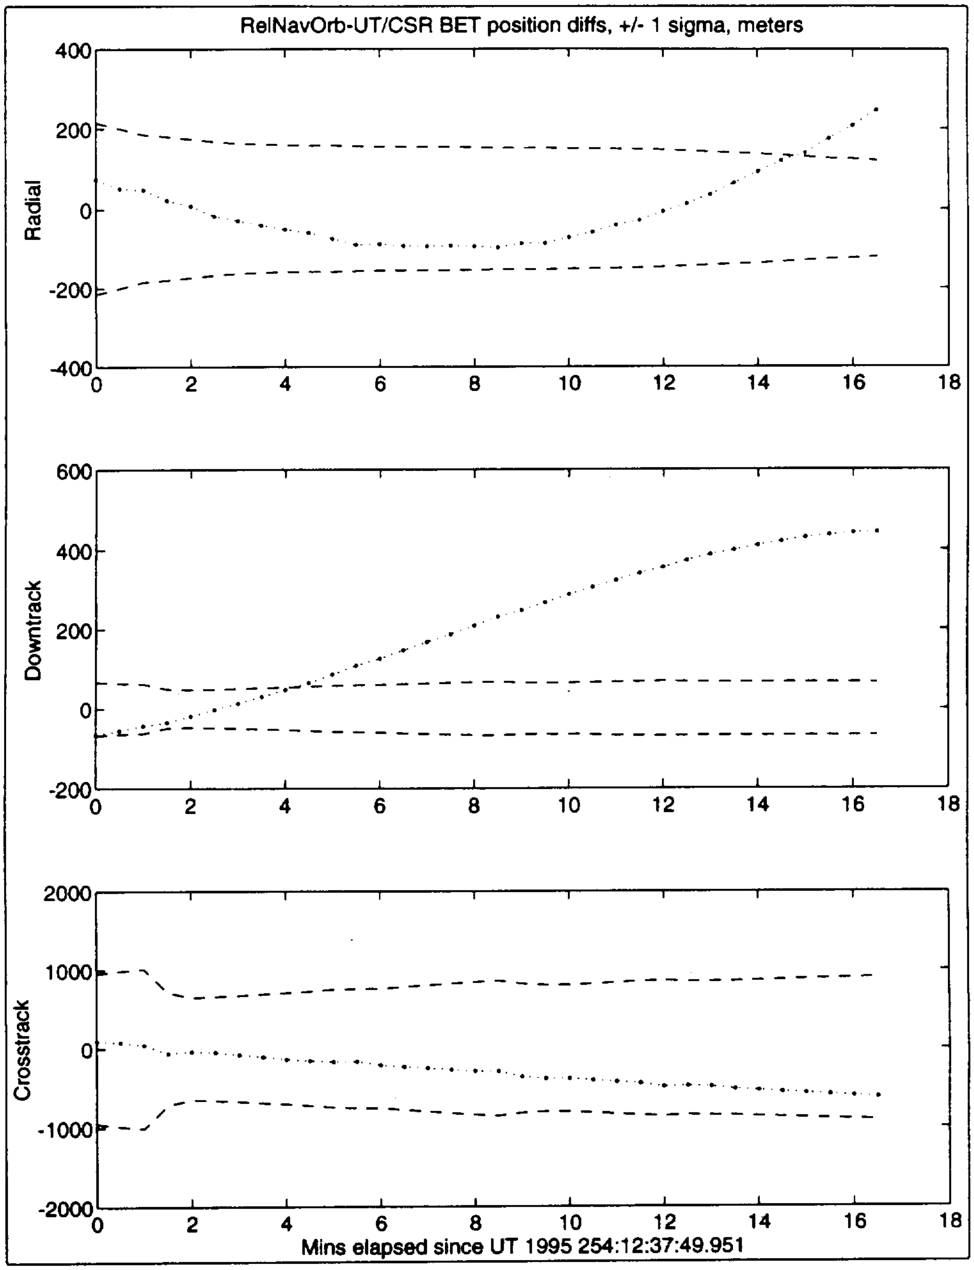

Figure 10: Rendezvous filter inertial position errors, filters initialized with onboard GPS solutions

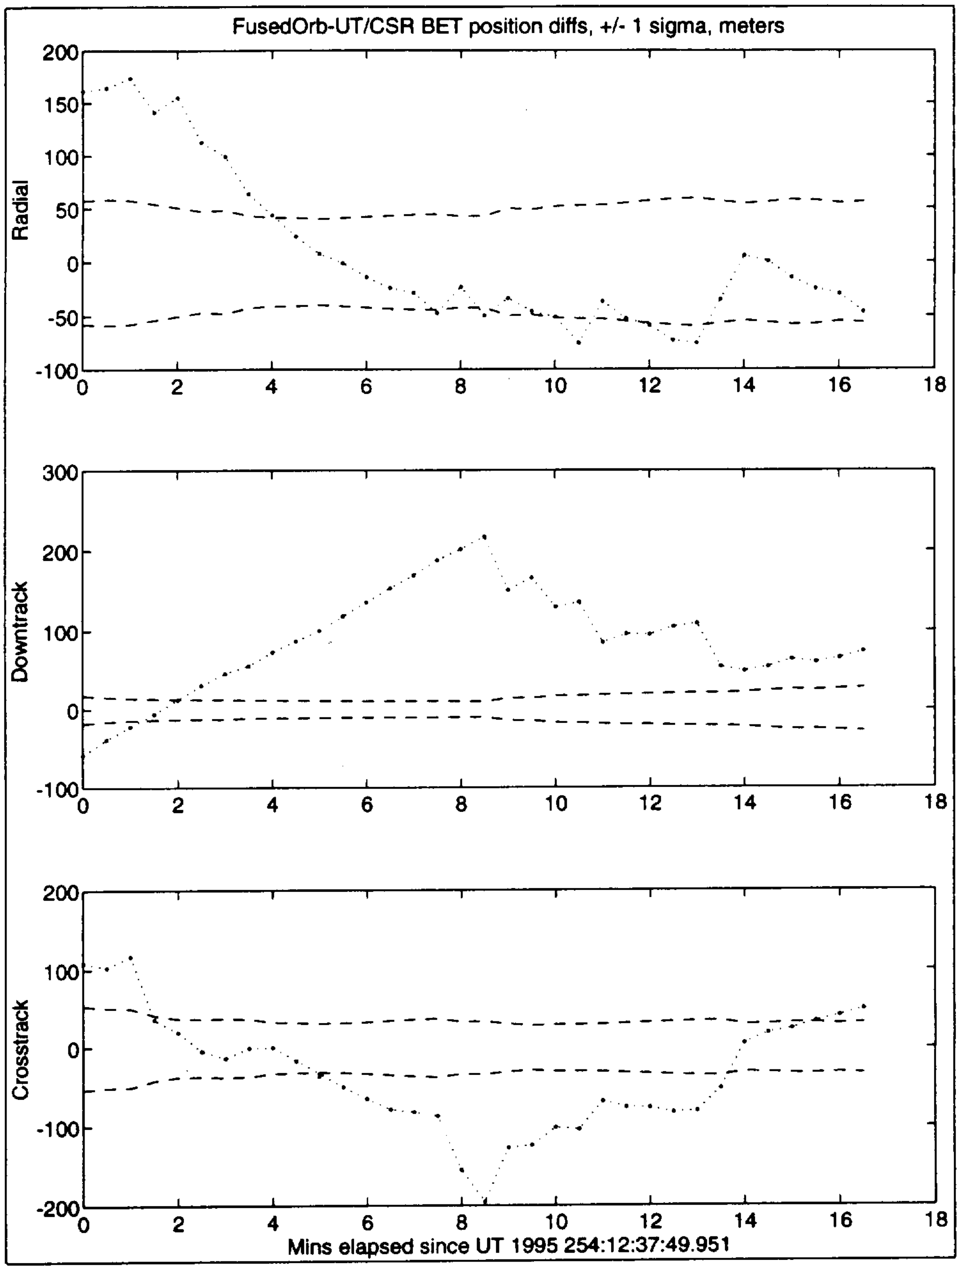

Figure 11: Fusion filter inertial position errors, filters initialized with onboard GPS solutions

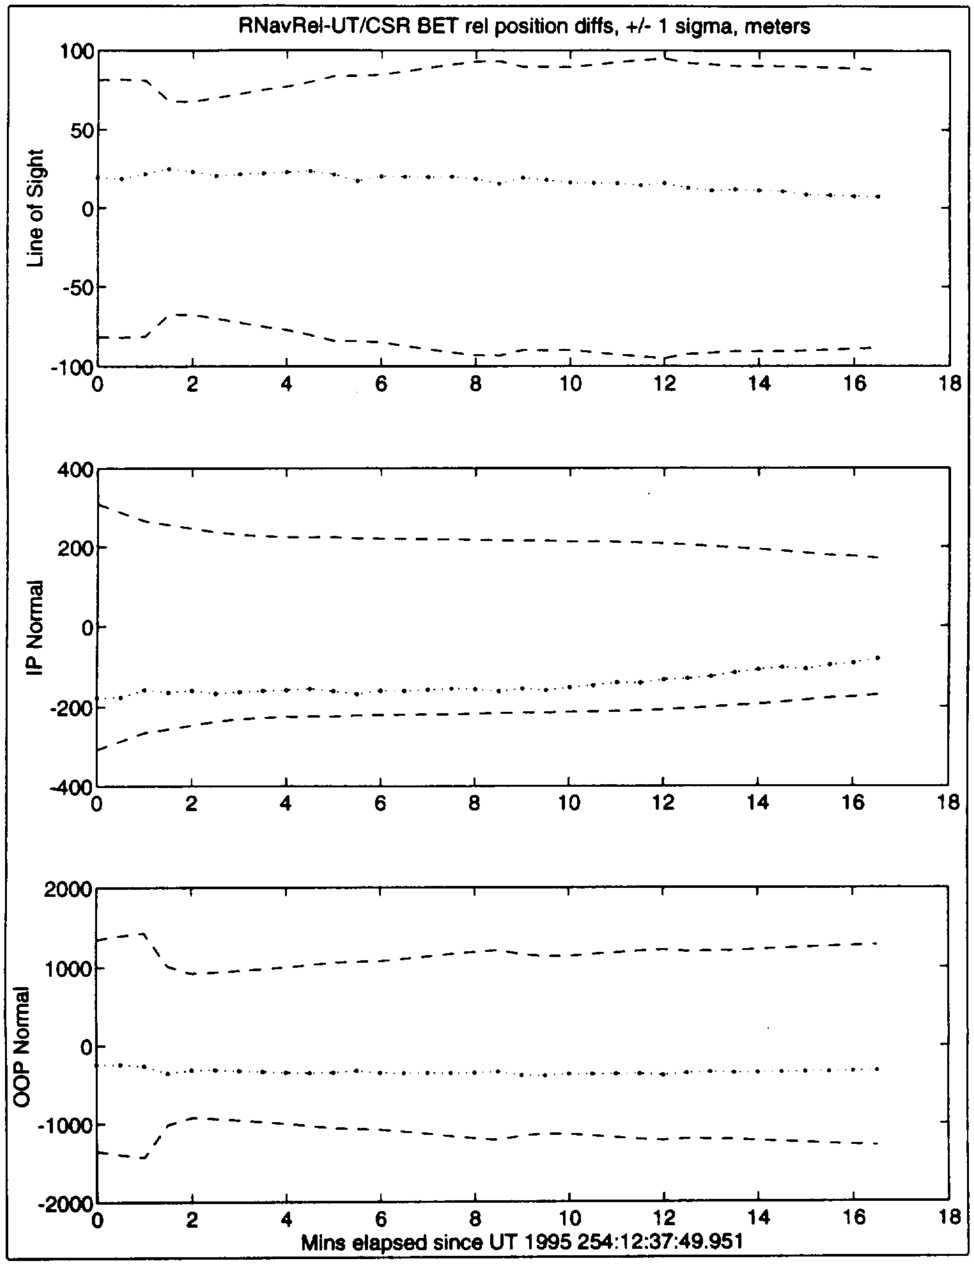

Figure 12: Rendezvous filter relative position errors, filters initialized with onboard GPS solutions

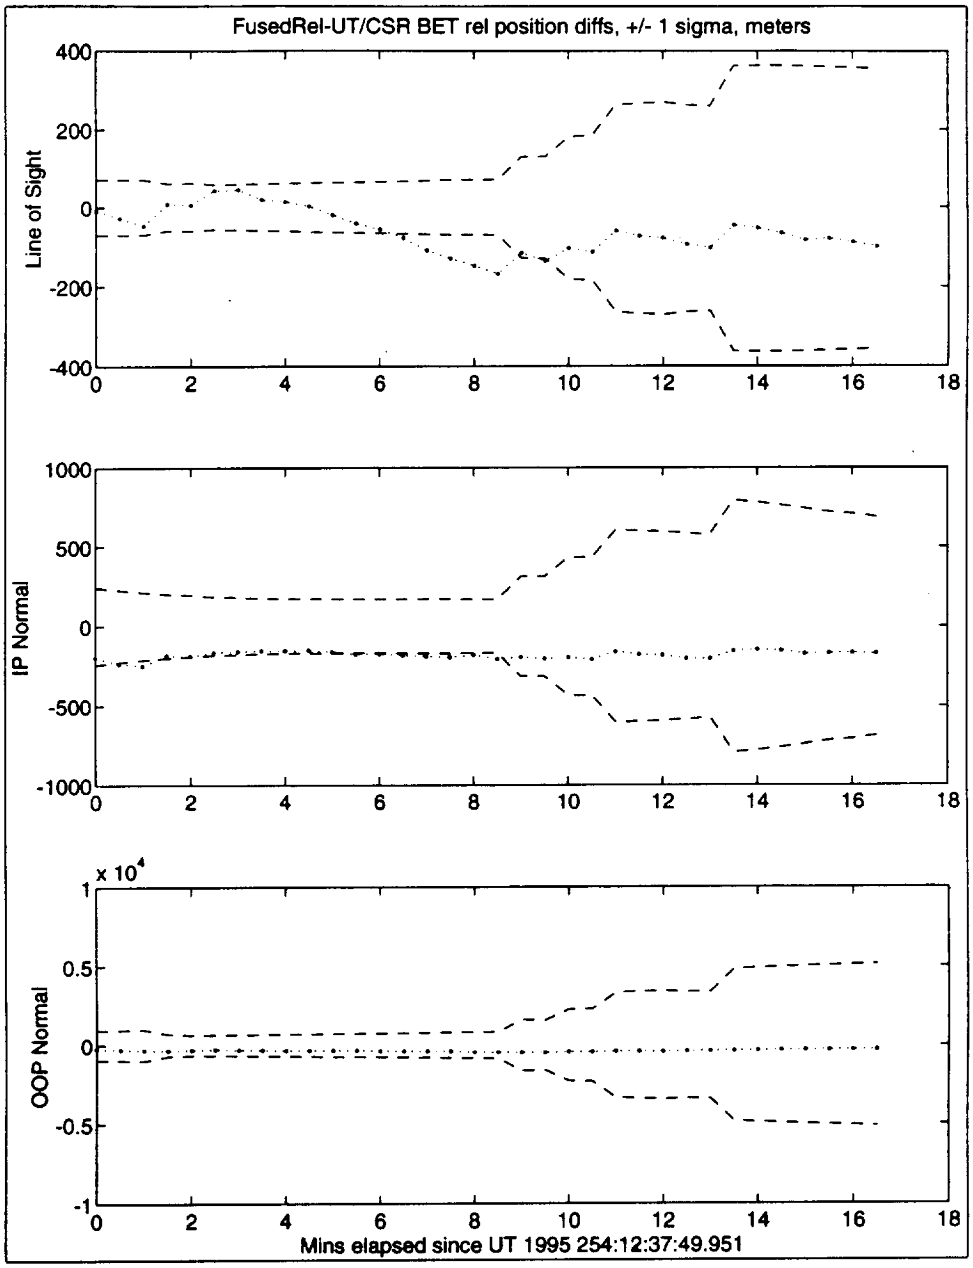

Figure 13: Fusion filter relative position errors, filters initialized with onboard GPS solutions

CONCLUSIONS

A data fusion filter for rendezvous navigation with GPS aiding can achieve the aim of optimally combining accurate inertial state estimates derived from a GPS filter with the suboptimal rendezvous filter used by the Space Shuttle, as flight data results from STS-69 show. The fusion produces substantially improved inertial state estimation accuracy compared with the suboptimal rendezvous filter, while producing relative states that are only marginally less accurate than the rendezvous filter. Even though the constraints imposed by testing with data playbacks makes the feedback scheme impractical, the fusion is able to successfully accommodate downtrack errors of over one mile in the rendezvous filter, thanks to a filter underweighting scheme. The underweighting scheme may be automatically scheduled by an innovations monitor based on local filter state vector differences. Although the data fusion filter can accommodate correlations between the filters which arise due to common process noise in their unmodeled acceleration states, the STS-69 data show that such correlations are very weak, and may be ignored. However, the flight data set examined in this test involves only coasting flight by the Space Shuttle. As new data sets become available for which there are accurate reference trajectories, it should be possible to investigate whether or not powered flight segments produce correlations between the filters. This is a topic for further study.

Although no operational implementation of the rendezvous navigation fusion scheme is planned at this time, the technique provides an alternative to integrated inertial/relative rendezvous filtering approaches which could be developed to support automated and piloted space rendezvous missions, such as those of ESA's Automated Transfer Vehicle, various X-vehicles being designed by the US aerospace industry, or an upgraded Space Shuttle. Since relative state accuracy is more important to rendezvous than a globally optimal solution, the existing Space Shuttle rendezvous filter is a verified and proven asset which could be successfully exploited, if it could be easily integrated with other sensors for near and far range tracking such as laser rangers and GPS. The rendezvous fusion technique provides such a bridge, which could allow maximal use of and minimal impact to existing filters in an integrated rendezvous navigation solution.

ACKNOWLDGEMENTS

Without the efforts of the following people, the data sets used for testing and evaluation of the fusion algorithm could not have been collected: P.A.M. Abusali, Kevin Lee, Jim Newman, Ray Nuss, Penny Saunders-Roberts, Christie Schroeder, and Bob Schutz. Also, this work benefited from many fruitful discussions with Peter Kachmar, Moises Montez, and Nassib Nabaa.

REFERENCES

[1] Carpenter, J. R., and R. H. Bishop, "Estimate Fusion for Lunar Rendezvous," *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Monterrey, CA, August 1993, pp. 1-10.

[2] Bar-Shalom, Y., and T. E. Fortmann, "Multisensor Track-to-Track Fusion," *Tracking and Data Association,* Academic Press, Inc., Boston, 1988, pp. 266-272.

[3] Carpenter, J. R., "A New Data Fusion Filter for Rendezvous Navigation with Global Positioning System Aiding," Ph.D. Dissertation, University of Texas at Austin, August 1996.

[4] Carpenter, J. R., and R. H. Bishop, "Target State Estimation for Suboptimal Rendezvous Navigation via Fusion of Inertial and Rendezvous Filter Estimates," Paper No. AAS 95-214, *AAS/AIAA Spaceflight Mechanics Meeting*, Albuquerque, NM, February 1995.

[5] Speyer, J. L. "Computation and Transmission Requirements for a Decentralized Linear-Quadratic-Gaussian Control Problem," *IEEE Transactions on Automatic Control,* Vol. AC-24, No. 2, Apr. 1979, pp. 266-269.

[6] Kerr, T. H. "Decentralized Filtering and Redundancy Management for Multisensor Navigation," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-23, No. 1, Jan. 1987, pp. 83-119.

[7] Bierman, G. J., and Belzer, M. R., "A Decentralized Square Root Information Filter/Smoother," *Proceedings of the 24th Conference on Decision and Control*, 1985.

[8] Carlson, H. A. "Federated Square Root Filter for Decentralized Parallel Processes," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-26, No. 3, May 1990, pp. 517-525.

[9] Schutz, B., et al., "GPS Tracking Experiment of a Free-Flyer Deployed from Space Shuttle," *Proceedings of ION-GPS '95*, The Institute of Navigation, September 1995, pp. 229-235.

[10] Hinkel, H. D., Y.-W. Park, and W. Fehse, "Real-Time GPS Relative Navigation Flight Experiment," *Proceedings of the National Technical Meeting*, The Institute of Navigation, January 1995, pp. 593-601.

[11] Park, Y., et al., "Flight Test Results from Real-Time Relative GPS Flight Experiment on STS-69," Paper No. AAS 96-182, *AAS/AIAA Spaceflight Mechanics Meeting*, Austin, TX, February 1996.

[12] Schroeder, C., B. Schutz, and P. Abusali, "STS-69 Relative Positioning GPS Experiment," Paper No. AAS 96-180, *AAS/AIAA Spaceflight Mechanics Meeting*, Austin, TX, February 1996.

[13] Kachmar, P., W. Chu, and P. Neirinickx, "Streamlined Shuttle GPS Capability," Shuttle Software Change Request No. 91051K, NASA Johnson Space Center, Sept. 15, 1995.