

# Navigation Sensor, Filter, and Failure Mode Simulation Results Using the Distributed Kalman Filter Simulator (DKFSIM)

Paul J. Lawrence, Jr., CAPT, USAF  
Michael P. Berarducci  
Wright Laboratory, Wright-Patterson AFB

*Abstract* -- The objective of this paper is to show simulation results obtained using the Distributed Kalman Filter Simulator (DKFSIM). Simulations were conducted at Wright Laboratory on a personal computer using the latest available version of this FORTRAN software package. The sensors modeled were a medium accuracy strapdown inertial navigation system (INS), a barometric pressure altimeter (BARO), the Global Positioning System (GPS), a Synthetic Aperture Radar (SAR), and finally, a terrain aided navigation (TAN) system. The SAR system includes an electro-optical (EO) type imaging model and a Precision Velocity Update (PVU) model. The TAN system combines radar altimeter measurements with digital terrain elevation data. The mission profile for each simulation included a low-level terrain following segment and a high dynamic combat maneuver segment typical for today's tactical fighter aircraft. The filter implementations used during this simulation sequence included (1) a single centralized Kalman filter (CKF) incorporating all of the sensor measurements, (2) a federated Kalman filter (FKF) with each sensor assigned to an independent local filter and the INS included as the system reference sensor, and (3) a cascaded Kalman filter (CASKF) where the GPS Kalman filter fed estimation information directly to a centralized Kalman filter. Simulations were performed where no failures were introduced in order to establish a baseline for the failed sensor cases. The first failure mode modeled was a GPS satellite clock failure consisting of a frequency bias shift plus ramp. The second failure mode modeled was an INS accelerometer failure consisting of a bias shift plus ramp. The failure simulations offer a coarse parametric range of failure magnitudes. Finally, observation of filter performance under varied conditions revealed advantages and other characteristics for each filter architecture modeled.

## INTRODUCTION

This paper presents the results from a series of simulations performed using the Distributed Kalman Filter Simulator (DKFSIM). DKFSIM was developed by Integrity Systems, Inc. [1], under U.S. Air Force funding to support the performance evaluations of several different filter architectures and "real world" sensor model conditions for advanced, multi-sensor navigation systems. Simulations were conducted at Wright Laboratory using a typical personal computer with

modest hardware resources and a mathematical package used for post-processing and plotting the output data files.

The sensor truth models used with this version of DKFSIM included a medium accuracy strapdown INS, barometric pressure altimeter, GPS, SAR, and TAN systems. The GPS model has additional features such as aircraft body and terrain masking, satellite selection based on best available Geometric Dilution of Precision (GDOP), and failure models to include gaps in the satellite coverage. The TAN system combines radar altimeter measurements with digital terrain elevation data.

Among the more attractive capabilities available within DKFSIM is the ability to accurately and efficiently model a suite of navigation sensors typical of advanced tactical aircraft. By setting switches and changing the values of the variables in the input data files, it is possible to simulate several Kalman filter architectures with no changes in the executable. Truth model variations are also available from within the specific sensor input data files, to include failure modes for the GPS and INS accelerometer and gyro components. Further, the structure of the source code provides a fairly direct approach to implementing a variety of different adaptive filtering concepts while demonstrating the fault tolerance capabilities of each filter design under test.

The estimation performances of the three filter architectures were evaluated by observing plots of the errors in the filter state estimates and the corresponding one-sigma values for the available filter states. The filter measurement residuals were evaluated to further characterize the system behavior and to aid in forming conclusions regarding possible failure detection techniques.

## TRUTH MODEL DESCRIPTIONS

The truth models are composed of the most predominant error sources and represent the highest level of fidelity modeled within the simulation. The truth models use an error state formulation of the INS in Earth-Centered-Earth-Fixed (ECEF) coordinates. However, all data is output in the appropriate coordinate frame, either local-level (wander azimuth) navigation (NAV) coordinates, east-north-up, or in the aircraft body frame, right wing-nose-up (XYZ). Because most sophisticated sensors have a built-in-test (BIT) capability for detecting gross errors and nonsense data, the hard-failure cases were not con-

sidered for this paper. However, for a more comprehensive perspective, these conditions should be investigated in order to define the upper boundaries between soft and hard failures. [2]

#### *Inertial Navigation System*

The INS model assumes the strapdown INS of medium accuracy are well modeled by less than 1 nm/hr circular error probable (CEP) in the horizontal position drift. Each INS has an orthogonal triad of ring laser gyros and accelerometers. [2]

For the medium-accuracy strapdown INS truth mode, there are a total of 42 time-correlated error sources, plus 3 accelerometer and 3 gyro wide-band noise sources. The baro altimeter truth model has 4 time-correlated error sources, plus 1 wide-band noise source. Second order baro-damping of the INS is implemented to limit the INS vertical channel errors. Herein, the term "INS" refers to the baro-damped INS. [3]

#### *Barometric Pressure Altimeter*

A barometric pressure altimeter is normally used to damp and bound the unstable dynamic behavior of the INS vertical channel. Among the natural sources of inaccuracies in a typical BARO are the current weather conditions, and dynamic conditions which may not be properly corrected. [2]

#### *Global Positioning System*

There are two subcomponents of interest when considering possible error sources: the GPS receiver and satellite models. The GPS receiver is assumed to be a two-frequency, five-channel, quartz clock P-code receiver. Selective availability is not considered here because the receiver is CRYPTO authorized. The simulator emulates code tracking loops which are carrier aided when possible. A single hemispheric antenna located on top of the aircraft is assumed. [2]

The GPS truth model assumes four-satellite operation. There are 29 time-correlated GPS error sources, plus 8 wide-band noise sources. The GPS truth model includes one channel random wide-band phase noise and one channel random frequency wide-band noise per satellite. [2]

The carrier measurements are assumed uncorrelated over time or between satellites. The code measurements are correlated over time because the code tracking loop is carrier-aided. With no jamming modeled, the code tracking loop time constants are short enough that its outputs are essentially uncorrelated. [2]

The GPS receiver clock is a critical component of the GPS receiver. It is modeled with white noise in the clock phase and frequency, with an additional acceleration-sensitive error for the clock frequency appropriate to the three physical axes of the oscillator crystal. The acceleration-sensitive error is crucial for accurately modelling GPS failures in the presence of high acceleration. [2]

#### *Synthetic Aperture Radar*

The SAR is considered an emissive imaging device. In the

EO mode, when the cursor is placed on top of a radar image, a fix is taken and the range and range-rate to the selected landmark can be obtained. Azimuth and elevation are also available, but are of lower quality and are often derived from the range/range-rate coordinates. [2]

The SAR focuses on a recognizable landmark for a few seconds as the aircraft moves relative to the landmark. The radar returns are processed by using very precise inertial data to tie the data together. This operation requires high-rate uninterrupted inertial data, preferably from an inertial measurement unit located near the SAR antenna. The raw SAR image is created in range/range-rate coordinates and then transformed to latitude and longitude, or azimuth and elevation coordinates, using the inertial solution. [2]

For the SAR EO model, each measurement is subject to a bias and random error. The landmark coordinates are also subject to a data base location error which is common to the EO sensors. A couple of typical EO sensors would be a forward-looking infrared (FLIR), and a laser ranger. [2]

The SAR can also be used in a Precision Velocity Update (PVU) mode. Effectively, this behaves much like a very precise Doppler radar, measuring the aircraft velocity relative to the ground in the aircraft body frame. In this mode, the SAR focuses on the ground to obtain a range-rate from the image. Resolution is not critical in this application, because the image is not being used for identification purposes. The output of the SAR PVU system is three velocity components in the body frame. Each velocity component is subject to bias misalignment and random noise sources. [2]

#### *Terrain Aided Navigation System*

The typical TAN system compares the measured ground height profile against a stored digital terrain data base kept on board the aircraft for real-time processing. Varying terrain gradients are required to gain observability. A perfectly linear terrain, such as a body of water or a gently sloping flat plain, is featureless from the point of view of the TAN and contains no observable latitude or longitude information. Conversely, the terrain must be locally linear for the Kalman filter to function properly. [2]

There are a total of 9 time-correlated error sources and 26 wide-band noise sources for the TAN sensor truth models, (radar altimeter, true terrain, and terrain map). The TAN truth model includes random measurement noise on the ground clearance measurement. [4]

## FILTER ARCHITECTURES

Three filter implementations were simulated for this paper. A centralized Kalman filter (CKF), a federated Kalman filter (FKF), and a cascaded Kalman filter (CASKF) configuration. The latter employs a GPS receiver integrated with the reference INS with a local Kalman filter, feeding a cascaded centralized filter. All three filters used the error state formulation

with the INS as the reference system, wherein the filters each estimate the error or drift of the inertial as well as other non-inertial sensor errors.

Each filter has the exact same model fidelity. The filter models provided for the different filter designs are reduced-order implementations of the truth models for each system. The filter models use an internal error state formulation of the INS in ECEF coordinates, with the state estimation output provided in local level and other frames.

The filter model for the INS includes 13 total filter states (the 3 position, 3 velocity, 3 attitude, baro alt, and 3 accelerometer bias states). The GPS filter model includes 2 states (1 clock phase bias and 1 clock frequency bias). The SAR EO filter model consists of 4 measurement bias states (range, range-rate, azimuth, and elevation). The SAR PVU filter model includes 3 velocity error states (essentially, Y component scale factor error, plus X and Z misalignment errors). The TAN filter model includes one measurement radar altitude bias error. [4]

#### Centralized Kalman Filter

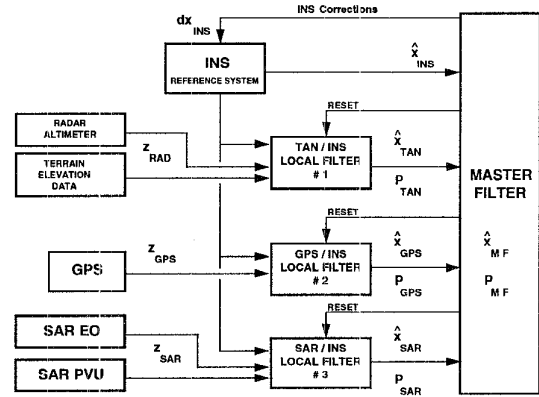
A single CKF using all of the external sensor measurements was simulated. The Kalman filter equations implemented for this filter design were the standard extended Kalman filter discrete-time measurement update and propagation equations. These are well established and well documented. [5]

The CKF filter error model states included the basic 9 INS states, (position, velocity, and attitude). Additionally, the BARO measurement bias, 3 INS accelerometer drift biases, and the GPS, SAR, and TAN filter states described above were included. All sensor measurements are provided to the filter with no pre-processing. DKFSIM requires the centralized filter to be tightly coupled to the GPS measurement process. [1] Namely, the raw GPS measurement data is fed directly to the filter. This is not a common configuration when existing aircraft integrated filter designs are considered. However, consideration has been given to employing a GPS receiver with an embedded INS as a replacement for the standard aircraft mission computer hardware. This would require measurement inputs from the other external sensors directly to the receiver, but the final result might be similar to the CKF design used herein.

#### Federated Kalman Filter

The federated filtering method described herein avoids the theoretical and practical difficulties of standard Kalman filtering. This is done by means of a simple, yet effective, information-sharing methodology. The advantages of information sharing implemented within the federated filter are increased data throughput by parallel operation of the local filters (Figure 1), enhanced system fault-tolerance by maintaining multiple component solutions, and improved accuracy and stability of cascaded filter operations. [4]

The basic concept of the information-sharing approach implemented by the federated filter is to divide the total system information among several component local filters. Next, per-



**Figure 1.** Federated Filter Application to a Multi-Sensor Navigation System

form the local time propagation and measurement processing, and then recombine the updated local information into a new total sum. [2]

Consider the partitioned filter architecture in Figure 1. Suppose the equivalent centralized filter solution can be represented by the covariance matrix  $\mathbf{P}_F$  and the state vector  $\hat{\mathbf{x}}_F$ ; the local filter solution by  $\mathbf{P}_i$  and  $\hat{\mathbf{x}}_i$ ; and the master filter solution by  $\mathbf{P}_M$  and  $\hat{\mathbf{x}}_M$ . Let the local filter index be designated by  $i=1 \dots N$ . If the local filter (LF) and master filter (MF) solutions are statistically independent, they can be optimally combined by the following additive information algorithm, where the inverse covariance matrix  $\mathbf{P}^{-1}$  is defined as the information matrix [4]:

$$\mathbf{P}_F^{-1} = \mathbf{P}_M^{-1} + \mathbf{P}_1^{-1} + \dots + \mathbf{P}_N^{-1} \quad (1)$$

$$\mathbf{P}_F^{-1} \hat{\mathbf{x}}_F = \mathbf{P}_M^{-1} \hat{\mathbf{x}}_M + \mathbf{P}_1^{-1} \hat{\mathbf{x}}_1 + \dots + \mathbf{P}_N^{-1} \hat{\mathbf{x}}_N \quad (2)$$

The key to the federated filtering method is to construct the individual LF and MF solutions so they can be recombined at any time by the above algorithm. This construction avoids the need to maintain LF/LF and LF/MF cross-covariances. Starting with the full centralized filter solution, we can then divide that solution so the LFs plus the MF each receive fractions  $\beta_k$  of the total information, where  $k$  is the total number of LFs plus the MF: [4]

$$\mathbf{P}_F^{-1} = \mathbf{P}_M^{-1} + \mathbf{P}_1^{-1} + \dots + \mathbf{P}_N^{-1} \quad (3a)$$

$$= \mathbf{P}_F^{-1} \beta_M + \mathbf{P}_F^{-1} \beta_1 + \dots + \mathbf{P}_F^{-1} \beta_N \quad (3b)$$

or more specifically,

$$\mathbf{P}_k^{-1} = \mathbf{P}_F^{-1} \beta_k \quad (4)$$

$$\hat{\mathbf{x}}_K = \hat{\mathbf{x}}_F \quad (5)$$

According to the conservation of information principle [6], in order to maintain constant total information across the sum in Eq. (3), the share-fraction values must sum to unity. Based on this development, the LF and MF solutions can be recombined to yield the correct total solution, without overcounting the common information. [4]

The discrete-time propagation process can also be performed via independent, parallel operations of the LFs and MF, as long as the common process noise information is divided in the same fashion as the fused solutions. The process noise covariance matrices are also governed by the information-sharing principle [4]:

$$\mathbf{Q}_F^{-1} = \mathbf{Q}_M^{-1} + \mathbf{Q}_1^{-1} + \dots + \mathbf{Q}_N^{-1} \quad (6a)$$

$$\mathbf{Q}_k^{-1} = \mathbf{Q}_F^{-1} \beta_k \quad (6b)$$

For the measurement update process, each local filter incorporates discrete measurements from its own unique sensor. Combining the results of the measurement update process yields the correct state estimate solution, essentially the same solution obtained by a centralized Kalman filter process. [4]

When certain implicit assumptions are satisfied, the FKF solution is the same as that of a single, CKF, and therefore is globally optimal. Simply, each filter employs a single share fraction value  $\beta_k$  for all of the full-system states and process noises, and the information fusion and reset operations are performed after every measurement cycle. Much less restrictive conditions can be accommodated, resulting in only a modest loss of optimality. Note that in such a case, the information-sharing fractions only apply to the common states among the LFs and MF, specifically the INS states. One is that each filter maintains only a subset of the full CKF state vector, as was the case for the FKF simulations described herein. [4]

For this paper, the federated filter simulations were conducted using the no-reset mode. Here, the local filters collectively maintain the system long-term memory, while the master filter provides short-term propagation of the fused solution after combining the local filter outputs. There is no information feedback from the MF to the LFs. Each LF retains its own, unique portion of the total system information. The no-reset mode permits the LFs to operate independently, as stand-alone filters, with estimation accuracies essentially at their normal levels. The LFs all send solutions to the MF for fusion at the same time. The MF propagates the fused solution to intermediate time points, but does not use it in the next fusion update. [4] Further, because the LFs run independently, the no-reset mode offers a more fault-tolerant system. None of the independent local filter solutions are receiving any information, good or bad, from any of the other filters. Thus, no corruption of the local filter solution occurs if any other sensor fails, with the exception of the INS.

The FKF was implemented with the MF having only the 13 INS common states. Local filter 1 (LF1) had the TAN sensor assigned, and the filter model included the 13 common INS

states, plus one TAN measurement radar altitude bias error state. Local filter 2 (LF2) had the GPS sensor assigned, and the filter model included the 13 common INS states, plus the two GPS sensor specific states, clock phase bias and clock frequency bias. Local filter 3 (LF3) had the SAR sensor assigned, and the filter model included the 13 common INS states, plus four SAR EO measurement bias states, and three SAR PVU velocity error states.

#### *Cascaded Kalman Filter*

The differentiation between the cascaded Kalman filter formulation and the federated formulation revolves around the amount of information passing between the distributed processing nodes. Cascaded forms flow from one filter to another, adding sensor information or combining subfilter information at each step, achieving a final global estimate. In a cascaded design, each set of sensor outputs is completely reduced by its associated filter into a new state estimate with covariance. This constitutes the sum total of the information allowed to pass to any of the other filters. A rough criterion for good cascading is that the premeasurement residual variance is negligibly affected by the states omitted in the partitioning. This corresponds to a lack of correlation over time and measurement. Because residuals are useful for fault detection and isolation (FDI), a large portion of the fault detection and isolation capabilities of the cascaded system are limited to the particular filter in which the faulty sensor data is processed. Nonetheless, cascaded filters are natural in retrofit operations where a preexisting filter must accommodate a new sensor. [2]

Stability for cascaded filter systems presents a significant problem. The processing of locally filtered sensor data by a subsequent centralized, or master, filter violates a fundamental theoretical requirement of Kalman filtering. The measurement errors must be "white". To properly incorporate such prefiltered measurements in the master filter, a model of the prefiltering process is required. [2]

The cascaded configuration depicts a loosely coupled pair of filters, one GPS sensor specific local filter providing the 13 common INS states and corresponding covariance values to the centralized master filter. The loosely coupled architecture is characterized by relatively simple incorporation of the GPS local filter solution. The CASKF incorporates the LF state estimates as a measurement, and uses the LF's covariance as the measurement error covariance. Minimal message-handling and control logic is required. Double-counting old GPS data forces the master filter to become too optimistic regarding its own accuracy. The CASKF incorporates the GPS output as a normal vector measurement, including the GPS output covariance as the associated measurement noise covariance. [4]

The advantage of the loosely coupled cascaded filter is its simplicity. The cascaded Kalman filter uses essentially two standard Kalman filters. Some additional measurement processing software is required unless the measurement set is cross-correlated, and then the measurements should be decor-

related. [4] By setting the update cycle time step within the CASKF to a value exceeding that of the GPS estimation time-correlation constant, the CASKF will essentially receive "white" measurement information from the GPS filter.

The CASKF was implemented with the four SAR EO measurement bias states, 3 SAR PVU velocity error states, and 1 TAN measurement radar altitude bias error state, plus the 13 common INS states. The GPS receiver model was implemented as it was in the FKF design, having 13 common INS, plus the two GPS sensor specific states, clock phase bias and clock frequency bias.

### SENSOR MEASUREMENT RESIDUALS

The potential Failure Detection, Isolation, and Reconfiguration (FDIR) capabilities of the filters under test represent the area where the greatest opportunities exist regarding adaptability and fault tolerance. One possible fault detection scheme would be to monitor the measurement residuals of a specified filter architecture. The residual sequence of interest, commonly referred to as innovations, is theoretically a white Gaussian sequence of zero-mean with covariance being a function of the observation matrix  $\mathbf{H}(t_i)$ , the error state covariance  $\mathbf{P}(t_i^-)$ , and the measurement noise covariance  $\mathbf{R}(t_i)$ :

$$E\{\mathbf{r}(t)\mathbf{r}^T(t)\} = \mathbf{H}(t)\mathbf{P}(t)\mathbf{H}^T(t) + \mathbf{R}(t) \quad (7)$$

During operation of the filter, the actual residual sequence can be monitored and compared to this description. If the description appears to be violated consistently, then one can deduce that something has occurred to invalidate the model within the filter. Otherwise, if the violation occurs in only one component of a vector residual process, then it can be assumed that the measuring device generating that particular residual component is the source of the difficulty. [4] Thus, monitoring the residuals has the potential to provide accurate fault detection information. The federated filter provides the capability to monitor the master filter fusion residuals for failure indications, as well as the local filter measurement residuals. Although the above equation describes a vector measurement update process, the actual implementation in the DKFSIM source code was accomplished as a scalar update, one at a time at a given time tag.

In order to verify the capability of capturing the measurement residuals (Table 1) for future use in fault detection tests, the local filter measurement residual output data files were investigated. Within each local filter, the residuals are subjected to a "reasonableness" test prior to being processed by the filter [5]. If the filter rejection threshold is surpassed, the measurement should be rejected, or filter corruption could occur if the measurement is used to update the estimates. Also, by checking the residuals and comparing them to the flight profile, some additional insight may be gained when viewing the behavior of the errors in the state estimates.

**Table 1 Measurement Residual Listing by Sensor**

---



---

<b>GPS - 4 Satellite Operation</b>
Satellites 1, 2, 3, 4
Pseudorange Measurement
Pseudorange-rate Measurement
<b>SAR EO</b>
Range Measurement
Range-Rate Measurement
Azimuth Measurement
Elevation Measurement
<b>SAR PVU</b>
X (Pitch Axis) Velocity Component Measurement
Y (Roll Axis) Velocity Component Measurement
Z (Yaw Axis) Velocity Component Measurement
<b>TAN</b>
Radar Altimeter Ground clearance
<b>16 Total Sensor Measurement Residuals</b>

---



---

The tolerance values for sensor measurement residual rejection was set to 7 times the one-sigma values for the GPS, and 100 sigma for the TAN, SAR EO and SAR PVU. These values were set artificially high for test purposes. This forces the filters to accept the measurements until the residual magnitudes become very large.

### SIMULATION STUDIES

Our efforts focused primarily on conducting performance analyses for each filter, averaging over five random Monte Carlo runs for a variety of different flight conditions and operating conditions. In the "real world" environment, any one of these filters would be driven by sampled data measurements from actual sensors. It is thus necessary to have an accurate statistical portrayal of estimation errors committed by each filter in the "real world" environment. A performance analysis fulfills this objective by replacing the "real world" with a reasonably complete and accurate mathematical model. Having access to the true values from the truth model allows something which is typically denied us. We can generate the "true" error committed by the Kalman filter in attempting to estimate the quantities of interest. [5]

The objective of a performance analysis is to statistically characterize the error process. Because of using stochastic processes as the basic modeling entity, we are more interested in the statistical, or ensemble average, behavior of the error process than a single sample output. One means of gathering this statistical information is by conducting a Monte Carlo study, where many samples of the error process are generated and the sample statistics are computed directly. If enough samples are generated, this should approximate the process statistics very well. [4]

For this paper, the error state space Kalman filter was used. Thus, the means of all processes are often assumed to be zero

for all time. A statistical description of this error process can be achieved by computing sample statistics, and averaging over the number of runs conducted. Usually, as a minimum, twenty is a reasonable number of simulations. A five run analysis is much smaller than desired, however, an ensemble average of just five runs still reduces the randomness of a single run by about 15-20 percent. A more useful statistic might be the root-mean-square (rms), which presents the non-zero mean and the standard deviation. Conducting twenty, or more, runs and performing an ensemble average should verify the means of the random processes will converge to zero. An exception to this supposition would be that nonrandom bias errors exist which have not been modeled and have proven to be significant. If searching for bias errors, or verifying models, with enough simulations conducted, averaging the runs will reduce the randomness of the processes to a level which could aid in detecting a nonzero bias condition. [5]

When performing the filter tuning, it is useful to compare the actual rms errors committed by the filter to the filter's own representation of its errors. For zero-mean processes the one-sigma values of the processes should converge to the rms values of the processes. [5]

Although a covariance performance analysis is computationally more efficient for generating accurate statistics than a Monte Carlo study (depending on the number of runs), there is an advantage to employing the Monte Carlo approach. Sign errors in the filter algorithm that may not be readily apparent in a covariance analysis due to squaring effects become evident in the Monte Carlo study. Also, effects of nonlinearities, such as device saturation or neglected terms when obtaining linear perturbation equations cannot be evaluated by a covariance approach. This must be done by performing a Monte Carlo study. [5]

#### *Mission Profile*

A typical tactical fighter flight profile, derived from the standard Strapdown Performance Study (SPS) mission profile, was used for this analysis. A Monte Carlo analysis was conducted for each filter, for the unfailed sensor case, for each failure scenario, and was accomplished over the duration of the mission from 0 to 7200 seconds. This mission profile includes an ingress to and including a low level terrain following and terrain avoidance segment (descent and run-in to target at Mach 0.90 and 500 ft, including a 92 second period of 4g jinks with heading changes of not more than +/-60 degrees), a target bombing run (execute a 4g pull-up to 4,426 ft maximum), climb-out and level-off for exiting the target range, and finally, an air combat segment (lasting 90 seconds, 3 climbing turns at 5g and 10 degrees/second).

The entire SPS mission profile was used for all simulations to provide more of a challenge to the integrated navigation systems. Filter stabilities are often stressed most dramatically, and are either proven or failed during the high-dynamic flight conditions.

#### *Description of Performance Simulations*

Eight simulation scenarios were accomplished for each of the three filter designs. These were conducted over the full length of the SPS trajectory using the INS, GPS, SAR, and TAN as the sensor suite. First of all, the filters were simulated with no failure conditions imposed on any of the sensors. Due to the large number of test cases for this paper, the sample mean of the filter state estimation errors and corresponding filter sigmas, and measurement residuals and corresponding residual sigmas, was conducted over 5 Monte Carlo runs for each test case. All plots included in this paper represent the 5 run ensemble average of each simulation scenario performed. The truth models were initialized before the zero time-tag to allow settling of the truth models prior to data storage. The five random number sequences were exactly the same for each simulation. All filter designs were selectable by setting switches and variables found within the input data files.

The federated filter was implemented using LF1, LF2, LF3, and the MF, all working cooperatively. The GPS measurements, assigned to LF2 had an execution time step for the GPS measurements of 2 seconds. The TAN measurements, assigned to LF1, began at the 200 second point in the mission and then continued for the rest of the run at a rate of one measurement every 10 seconds. A more typical rate would be about 2-4 Hz, however, one measurement every 10 seconds essentially reduced the simulation time. The radar minimum height for a low-altitude warning indication was set to 50 ft. The maximum operating height for the radar was set to 50,000 ft. The SAR model was developed to effectively provide EO and PVU measurement data to LF3 during the mission.

The execution time step for the INS was set to 2 seconds. Some of the selectable parameters within the INS input data file are the Baro damping gain for position and velocity, correlation times for the gyro and accelerometer biases, and the sigmas for various error sources, all of which govern the accuracy of the INS. Because the INS is baro-damped, the Baro time step was also set to 2 seconds.

An error state accuracy demonstration of the FKF in the no-reset and fusion-reset modes should characterize the system as functional, efficient, and operating correctly under normal and highly dynamic conditions. The filter's estimation accuracy is determined by differencing the master filter state estimate and the truth state, obtained from state estimate output data files and the real-world models output data files. The error state filter formulation used in DKFSIM estimates the errors in the navigation and attitude information using the difference between the INS and external sensor information. [5]

## FAILURE MODES

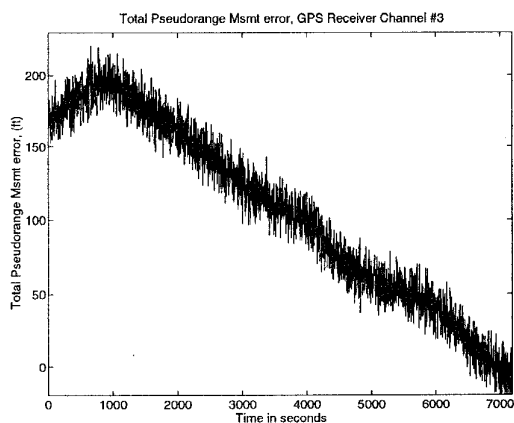
#### *Unfailed Sensor Truth Models*

Generation of the unfailed sensor truth model data sets was necessary to define the foundation for the failed conditions imposed on the GPS and INS sensors. By viewing the unfailed

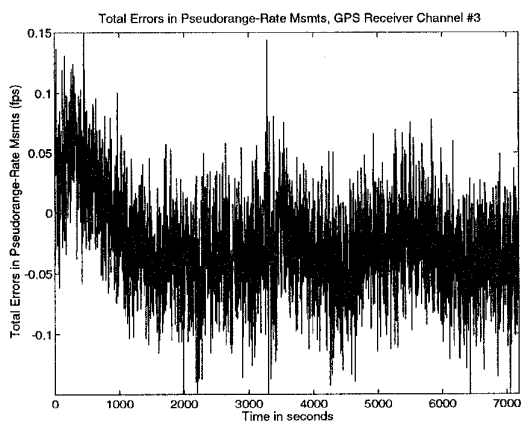
case plots, reasonable failure magnitudes were established based on the spread of the data across the entire mission.

Because the filters are implemented in error state form, it is necessary to have the truth model data in error form also. The true trajectory data file is available for comparison of whole-valued variables, such that the true sensor errors can be obtained and subsequently differenced with the filter estimates. This allows for a straight-forward differencing between the filter estimates and the truth model error values. The end result is an indication of the filter's ability to estimate the desired states, namely the filter's estimation error of the sensor errors.

The GPS pseudorange (PR) measurement error, for all four satellites, did not exceed 225 ft, and had a relative spread of the error of about 30 ft, Figure 2. The pseudorange-rate (PRR) measurement error plots exhibited a spread of about 0.25 fps, Figure 3. The total user clock frequency offset was usually



**Figure 2.** GPS Receiver Channel #3, Pseudorange Measurement Error, (ft), No Failures.

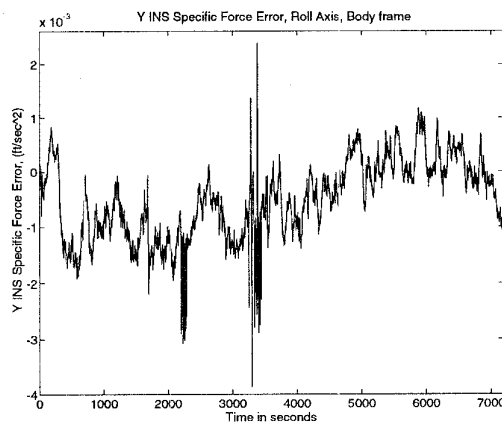


**Figure 3.** GPS Receiver Channel #3, Pseudorange-Rate Measurement Error, (fps), No Failures.

less than 0.05 fps (in equivalent range-rate units) in magnitude for the entire mission, with the exceptions occurring during the dynamic maneuvers.

The GPS satellite identification plots allowed an opportunity to introduce a failure based on satellite selection. The satellite measurement validity indexes indicated loss of visibility in several places, but only for short periods of time and usually around the high dynamic portions of the mission.

The INS error plots and whole-value plots matched the flight profile of the SPS trajectory very well. The position drift components were well behaved, oscillatory, and never exceeded 5000 ft in magnitude. The vertical channel is baro damped using a second order model. The vertical position and velocity drifts were much less than the other two Nav frame components, North and East. The INS specific force error components were all about  $1e-3 \text{ ft/sec}^2$  ( $\sim 31\mu\text{g}'\text{s}$ ), Figure 4.



**Figure 4.** INS, Y Component Specific Force Error, No Failures, ( $\text{ft/sec}^2$ ), Body Frame.

The SAR EO exhibited nominal error magnitudes of about 100 ft in range, and about 0.15 fps in range-rate. The SAR PVU velocity error components maintained relative orders of magnitude of about  $1e-3$  fps or less.

The TAN system radar altitude bias error quickly converged from about 50 ft in error to a value around 10 to 15 ft. The filter's measurement validity index showed few radar measurements were rejected, as expected. Ultimately, the total radar altimeter bias error plus measurement noise never exceeded 100 ft.

Some significant mission characteristics were evident in the plots. The areas of high dynamic maneuvers depicted occurred at about 2300, 3400, and 6400 seconds. The most severe maneuvers occurred at about the 3400 second time.

#### *GPS Satellite Frequency Range-Rate Failure*

The GPS satellite frequency (range-rate) failure was imposed upon satellite identification number 2, GPS receiver channel 3. An initial offset bias was added, then an additional constant-slope ramping bias was added to simulate a potential soft fail-

ure condition (Figures 5 and 6). Satellite ID #2 was selected because the receiver uses this satellite for its solution from 2200 to 5200 seconds. This allowed for stabilization of the navigation filter for 2000 seconds before the failed satellite comes into view, 50 minutes of failed conditions, and an opportunity for the filter to restabilize upon changing to an unfailed satellite. The filter response to the change in satellites, going from using a bad to a good one, is a critical aspect in the accuracy of the navigation solution. When the failed satellite departs at about 5200 seconds, the newly acquired satellite has no failures. Also, the high dynamics of the mission occurred at times where the filters might be the most vulnerable to divergence.

The primary goal of implementing the failed conditions was to establish a coarse parametric description of the failed conditions. The most significant parameters being the lowest detect-

able failure magnitudes, and the largest failure magnitude just short of a hard failure indication. For this paper, the magnitudes of the GPS satellite frequency (range-rate) failures varied from about 0.05 fps to 0.10 fps for the initial PRR offset pseudorange-rate bias, and then from  $10e-6$  ft/sec<sup>2</sup> to  $1.0e-3$  ft/sec<sup>2</sup> for the frequency drift rate. A total of 4 simulations per filter were conducted within this failure range.

The most revealing plots allow the characteristics of both failed and unfailed magnitudes to be visible for the PR and PRR measurement errors. Here the previously unfailed magnitudes are not overcome by the magnitude of the failure, somewhat the nature of an insidious soft failure.

The GPS receiver clock failure was another failure mode available for simulation within DKFSIM, but limitations prohibited covering all of the failure modes in this paper.

#### INS Roll Axis Accelerometer Failure

The INS roll axis accelerometer also had an initial offset bias imposed plus a constant-slope ramping bias intended to represent an accelerometer soft failure. The failure was turned on at 2000 seconds and continued throughout the remainder of the mission. Again, the objective was to show a parametric range of failed condition magnitudes from between an undetectable failure to the hard-over failed condition. The aircraft body frame designation is defined by X axis - right wing, Y axis - nose, and Z axis - up.

The magnitudes of the failures were  $2.0e-4$  ft/sec<sup>2</sup> for the initial offset bias used across all INS failures, and from  $5.0e-7$  ft/sec<sup>3</sup> to  $2.5e-5$  ft/sec<sup>3</sup> for the constant-slope ramping bias. This corresponds to about a  $6.2$   $\mu$ g offset, and a maximum bias drift rate of  $0.78$   $\mu$ g/sec. This minimum bias drift rate equates to about  $0.016$   $\mu$ g/sec second. The most visible contrast between the failure magnitudes and the unfailed magnitudes was obtained using  $2.0e-4$  ft/sec<sup>2</sup> initial offset bias, plus an  $11.5e-6$  ft/sec<sup>3</sup> ramping bias, Figure 7.

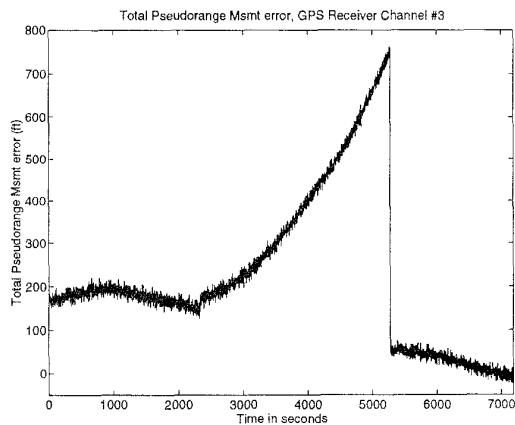


Figure 5. GPS Receiver Channel #3, Pseudorange Measurement Error, (ft), Satellite ID #2 Failure.

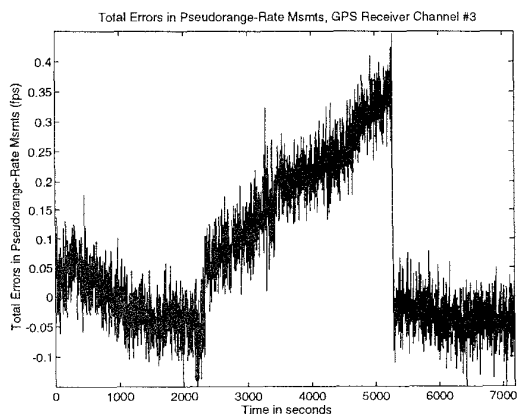


Figure 6. GPS Receiver Channel #3, Pseudorange-Rate Measurement Error, (fps), Satellite ID #2 Failure.

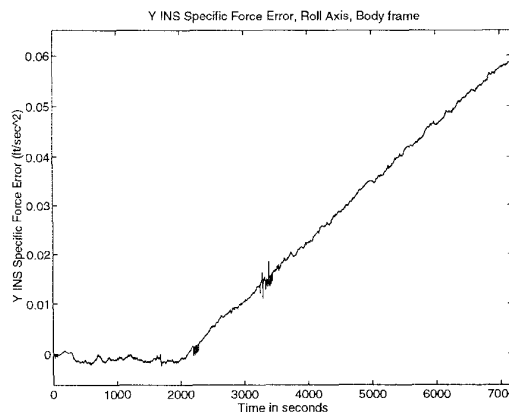


Figure 7. INS, Y Component Specific Force Error, Roll Axis Accelerometer Failure, (ft/sec<sup>2</sup>), Body Frame.



## RESULTS AND OBSERVATIONS

The following discussions address the most significant points observed from the simulations conducted. Selection of the plots for inclusion in this paper is based on the recognizable contribution from each. The focus is set upon the failure modes because the filters' responses to the failure modes are the most informative. The different filter configurations were chosen to show the capabilities of DKFSIM, and from there, one can observe the potential advantages of one filter over another.

The behavior of the filter states and filter measurement residuals for all of the filter designs describes each filter's individual performance. With no failures, the GPS measurement residuals were all well behaved and remained almost entirely within the one-sigma bounds for the entire mission.

As was the case for all filter designs, the one-sigma values for the pseudorange residuals had a fairly constant value of about 25 ft, for all 4 satellites. The magnitudes of the residuals typically remained below 20 ft. The one-sigma values for the pseudorange-rate residuals had a nominal value of about 0.1 fps, and magnitudes were well within that range for most of the mission.

The one-sigma values for the SAR EO azimuth and elevation residuals were on the order of  $1e-3$  radians. The range measurement stayed mostly inside the 100 ft one-sigma bounds. The range-rate residuals had nominal values of about 0.1 fps. The magnitudes of the SAR PVU velocity measurement residuals typically remained at or below 0.5 fps. The TAN radar altimeter measurement residuals were well within the one-sigma bounds of about 65 ft.

For all of the unfailed cases, the magnitudes and one-sigma values of the available sensor measurements residuals were fairly typical of an efficient estimator. There were no unusual trends noted for the unfailed case.

### Centralized Kalman Filter

The first soft failure mode implemented was a satellite frequency (range-rate) initial offset bias of 0.05 fps, plus an additional ramp of  $1e-5$  ft/sec<sup>2</sup> ( $\sim 0.31\mu g/s$ ), affecting the PRR measurements. As expected, the PR measurement error plots exhibited the integral relationship between the PR and the PRR measurements.

When this GPS satellite failure was imposed at 2200 seconds, the horizontal position errors quickly increased above the one-sigma values, but never exceeded 100 ft for east or north axes. However, the vertical position peaked at about 125 ft. The NAV frame velocity components all shifted but never approached the 3-sigma values for the errors. Because the magnitude of the induced failure offset biases never reached significant proportions, this failure can easily be estimated out by the filter.

Most notably, when the failed satellite passed and a healthy satellite was acquired, the GPS measurement residuals became relatively large. The cumulative effect from the soft failure affected the filter performance by as much as 50 ft as indicated by the residual magnitudes. This failure indicates that a magnitude of this size would likely go undetected by the filter.

The hardest satellite frequency (range-rate) failure implemented, involved the same initial offset bias, but instead used a much higher rate of change for the ramping bias. A slope of  $1.0e-3$  fps per second displayed the most drastic reactions. The filter horizontal position estimate diverged by as much as 5800 ft. This phenomenon occurred upon acquisition of the healthy satellite. This level of accumulated error easily exceeded the 7 sigma tolerance threshold set for filter measurement rejections. Therefore, the healthier set of GPS satellites were incorrectly considered unusable by the filter. The SAR and TAN measurements also falsely indicated failed sensor performance and appeared to show a trend of oscillatory divergence. The filter position states all exhibited huge magnitudes several hundred sigma above the filter's expectations. The filter's one-sigma values were not adjusted to attempt to track the divergence. However, it does appear that the filter was capable of recovering after several thousand seconds.

An important occurrence under these conditions was the revealing plot of the north INS position estimation error in the CKF, Figure 8. The satellite came within range at 2200 seconds and departed at about 5200 seconds. The huge position error shows the filter's inability to detect or estimate this level of bias out of the navigation solution.

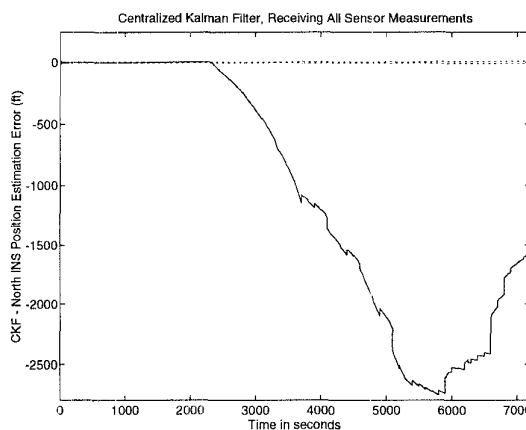
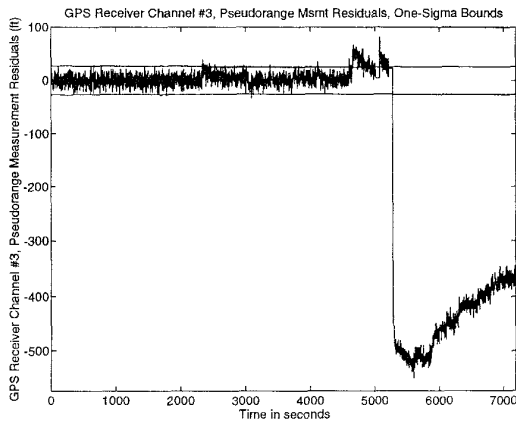


Figure 8. Centralized Kalman Filter, North INS Position Estimation Error, (ft), Largest GPS Failure.

The second largest GPS failure implemented involved a ramping bias slope of  $1.0e-4$  fps per second. In this case, the unfailed conditions are still distinguishable among the larger more dominant error magnitudes. Under these conditions, the

PRR measurements from the satellite on receiver channel #3 did not exhibit any indication of a possible failure. It was the satellite #3 PR residual that had the most severe reaction, Figure 9. It occurred upon acquisition of the new satellite at 5200 seconds. The only other measurement residual to suffer any difficulties was the SAR EO range-rate. From Figure 9, one can easily see the filter tracking the error.



**Figure 9.** Centralized Kalman Filter, GPS Receiver Channel #3, Pseudorange Measurement Residual, (ft), Second Largest GPS Failure.

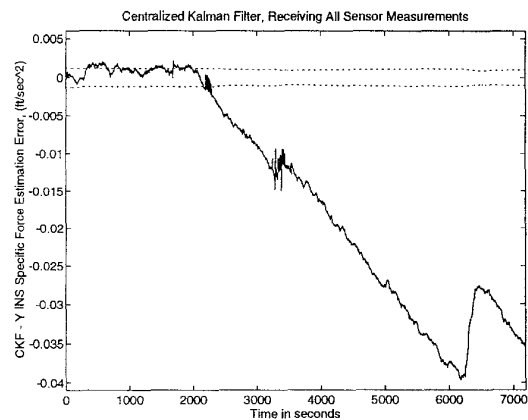
The INS accelerometer failure imposed was the same kind of bias set, an initial offset plus a constant-slope ramping bias added to that. The smallest bias offset values were  $2.0 \times 10^{-4}$  ft/sec<sup>2</sup> ( $\sim 6.2 \mu\text{g/s}$ ) for the initial offset, and a slope of  $5.0 \times 10^{-7}$  ft/sec<sup>2</sup> ( $\sim 0.015 \mu\text{g/s}$ ) per second. This was simply indicative of an accelerometer which might drift faster than usual. The CKF detected no failures at this magnitude.

The largest INS accelerometer failure magnitude was set at the same initial offset bias, plus an additional  $2.5 \times 10^{-5}$  ft/sec<sup>3</sup> ( $\sim 0.78 \mu\text{g/sec}$ ). When the aircraft performs the final high dynamic maneuver at about 6400 seconds, essentially all of the filter estimates react violently. Because the aircraft is flying straight and level, the filter is misinterpreting the specific force error as platform attitude drift, and is adjusting the platform tilts to compensate. The two error components are indistinguishable until the observability is improved by performing a maneuver. This is true only for the strapdown INS. A gimbaled INS cannot distinguish the two components at all, because it maintains a fixed inertial orientation.

The filter's peak horizontal position errors range to about 500 ft, where the actual position drift exceeded 5000 ft. The filter's velocity errors appear to be minimal until the last dynamic maneuver. The two areas, other than position, that cannot be estimated well by the filter under these conditions consist of the attitude drift in the north and east directions, and the X and

Y components of the specific force drift. These errors grow significantly, and because the filter is unaware of the failure, the filter's one-sigma values do not grow with them. While it may not have diverged as a result of this failure, a similar failure having a larger magnitude could cause the filter to diverge.

The last INS accelerometer failure implemented was intended to nearly divide the results obtained from the smallest and the largest magnitudes. The slope of the ramping bias was  $11.5 \times 10^{-6}$  ft/sec<sup>3</sup> ( $\sim 0.36 \mu\text{g/sec}$ ). The position error did not exceed 125 ft. However, even though the CKF has the benefit of the SAR and TAN, the same results prevailed indicating an undetected failure. One can easily see the filter's inability to track the most insidious failure, that being an INS failure, Figure 10. Comparing this plot to the INS specific force error plot from the truth model, it is obvious the errors are totally unseen by the filter.



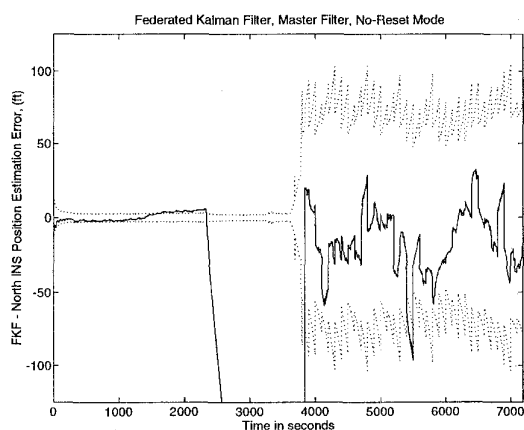
**Figure 10.** Centralized Kalman Filter Estimation Error, Y Component Specific Force Error, Roll Axis Accelerometer Failure, (ft/sec<sup>2</sup>), Body Frame.

#### Federated Kalman Filter

The FKF received the exact same failure mode and unfailed scenarios as the CKF. The random number seeds were identical for all of the random processes and magnitudes were exactly the same.

The minimum magnitude GPS satellite failure produced modest results, as was the case before with the CKF. Although there was a distinguishable difference in the accuracies of the filters, the east and north position errors were almost identical, while the vertical position estimation error indicates a difference of only about an 8 to 10 ft maximum accuracy advantage of the CKF over the FKF. Also, it is evident that the navigation solution contributed by GPS LF2 has the greatest accuracy of any of the local filters. The MF must automatically weight this solution more because of its significantly smaller covariance matrix. This is clearly reflected by the likeness of the master filter's estimation error plots to those of LF2.

The largest magnitude GPS satellite failure offers a much clearer perspective into filter behavior. The FKF was initially tracking the error, and reached a maximum horizontal position estimation error of about 1200 ft, where the actual INS horizontal position drift grew to about 5000 ft. The MF's fusion residual test detected the errors and rejected the GPS LF2 solution. Figure 11 shows the recovery of the FKF's estimate of the north position error after detecting the failure in the GPS. The estimation error and one-sigma bounds are consequently a combination of the SAR and TAN local filter values at about 3800 seconds. This is clearly indicative of the advantage of the FKF over the CKF with respect to fault tolerance. The CKF diverged to almost 3000 ft in north position estimation error for the same failure.



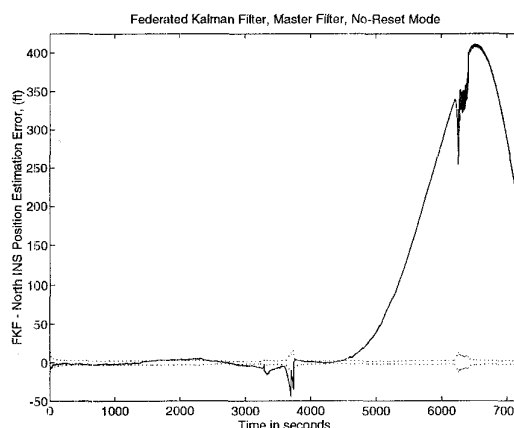
**Figure 11.** Federated Kalman Filter, Master Filter North INS Position Estimation Error, (ft), Largest GPS Failure.

Both the CKF and FKF filters exhibited improved observability in the north and east directions during the maneuver at 3200 seconds, except, as a consequence of the maneuver, the vertical velocity components had significantly larger errors. The attitude estimation error was much better for the FKF. And, the INS specific force drift components were estimated better. Because the MF does not process the sensor specific error states for output, in order to compare the CKF and FKF user clock frequency bias estimation error, it be obtained from LF2. LF2 apparently performed much more efficiently when estimating the bias associated with the satellite failure. The highest point on the plot is about 1.4 fps in total estimation error, which is lower than the CKF's highest value of estimation error. However, after the faulty satellite is removed from the receiver's solution, note that the filter recovery is much longer for LF2 than for the CKF. This can be expected because the CKF had access to the SAR and TAN measurements.

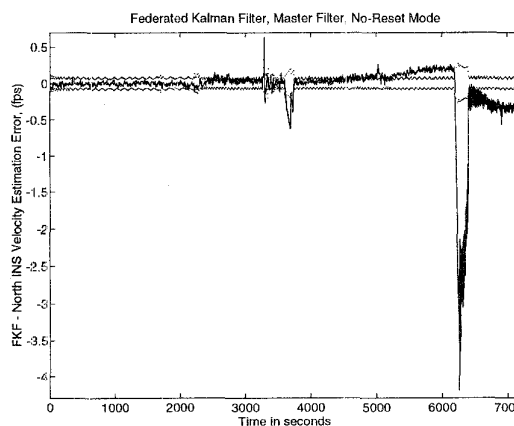
For the FKF, because the failure occurred in the GPS alone, the SAR and TAN systems were unaffected and performed

normally. The CKF, incorporating all measurements together, clearly experienced several false indications of failed sensors. Because the local filters are independent, and the failure is non-inertial, the failure in the GPS is not corrupting the local filter solutions for LF1 and LF3. However, from the master filter fusion residual data file, the fusion residual rejection rate is sufficiently high for the GPS and SAR. This should aid in the isolation of the failed sensor. Some additional logic is required to coincide with the FKF operation if the decision process for fault detection, isolation, and recovery (FDIR) is to be implemented.

Characteristic of the failed conditions, the north INS position (Figure 12) and velocity (Figure 13) estimation errors suffer



**Figure 12.** Federated Kalman Filter, Master Filter, North INS Position Estimation Error, (ft), Accelerometer Failure.



**Figure 13.** Federated Kalman Filter, Master Filter, North INS Velocity Estimation Error, (fps), Roll Axis Accelerometer Failure.

from the growing magnitudes of the imposed roll axis accelerometer bias, and subsequently, the high dynamic maneuvers performed by the aircraft. The FKF's north position estimation error appears to remain acceptable, about 400 ft, and eventually converges.

Comparing the results from the accelerometer failures, the FKF exhibited very similar results to the performance achieved by the CKF. There were significant position errors with little or no indication from the measurement residuals that a failure had occurred. The TAN local filter exhibited very large estimation errors after the high dynamic maneuver was performed near the end of the mission. As was the case for the CKF, the accelerometer drift bias had accumulated to such an amount, that the filter could not keep pace. The estimation errors were quite reasonable until this maneuver. Similar to the CKF, in straight and level flight, the FKF cannot tell the difference between the attitude and specific force biases, so it adjusts its estimates for the tilts to compensate. When the maneuver is accomplished, the observability of the accelerometer bias failure goes up. Attitude and specific force estimation errors are affected most adversely from beyond this point because the filter could not distinguish between them. The position and velocity estimates lose their track either just before or just as the final maneuver is accomplished at 6400 seconds.

#### Cascaded Kalman Filter

The cascaded centralized Kalman filter (CASKF) does not have access to the raw GPS measurement data, nor does it attempt to model the GPS error states. Certainly, one should not expect similar results for this filter as those results obtained for the CKF and FKF. For the smallest magnitude GPS satellite failure, the position and velocity estimation errors appear very close to the other filters', but the CKF continues to exhibit a slight advantage in estimation accuracy.

For north and east position estimation error plots, the CASKF appears to have very high confidence in its estimates, even though the errors peak at about 90 ft. The GPS receiver falsely influences the centralized filter to a higher confidence level than what is appropriate. The CASKF velocity estimation error plots exhibit approximately the same minimum boundary as the CKF and FKF, but inconsistencies in the measurement processes force a wider variation in the estimates than should be necessary. The SAR EO range-rate, along with all three position states, all three specific force states, experience this condition. Only the TAN measurement residuals show any hint of an affect from the failed condition.

Although the plots are not presented here, the magnitudes of the Y INS specific force estimation errors are on the same order as the previous filters, but the corresponding one-sigma bounds are one third that of the CKF. Again, the filter was inappropriately influenced by the GPS receiver solution.

The largest magnitude GPS satellite failure is quite revealing in this case. The filter is able to detect the divergence because the artificially small one-sigma values allow a quicker rejection

of the GPS receiver solution, Figures 14 and 15. This is not a result of good filter design. It is a result of the overconfidence generated by the filter-driven-filter, where the driving filter is the GPS receiver with a very small covariance matrix. This reaction did not occur for the smaller magnitude failures. The plots were reduced in scale to capture the relative characteristics prior to the rejection. And, even though the SAR measurement residuals were well behaved until the maneuver at 3200 seconds, the filter's error became evident because of the improved observability during that maneuver. It was at this point that the filter one-sigmas were opened up to allow reconvergence. And, while this may not be optimum performance, the CKF drifted off in position error, in all three directions, to over 2500 ft. The CASKF recovered at a maximum of 600 ft, and the error was less than 100 ft for the re-




Figure 14. Cascaded Centralized Kalman Filter, North INS Position Estimation Error, (ft), Largest GPS Failure.

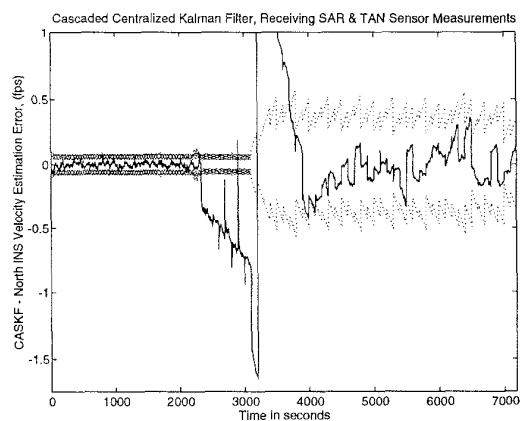


Figure 15. Cascaded Centralized Kalman Filter, North INS Velocity Estimation Error, (fps), Largest GPS Failure.

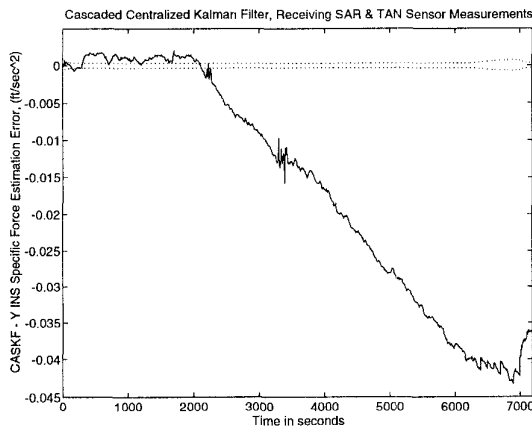
remainder of the mission. This corresponds to rejection of the GPS receiver solution, and the resulting dependence on the SAR and TAN solutions alone. The filter tolerance levels were not altered to favor the CASKF, although by its design, they are artificially reduced.

For clarification, the maximum error in Figure 14 falls to about the -600 ft point. The maximum error in Figure 15 goes up to about +2.5 fps.

For the second largest magnitude GPS satellite failure, the SAR and TAN simply did not have the accuracy to maintain track, and the CASKF's estimates were overcome by the failure. As one might expect, the CASKF's position error grew to over 350 ft, but the filter never attempted to adjust the one-sigma values to compensate, probably because the error went undetected.

When comparing the worst-case INS roll axis accelerometer failure estimation error plots, the CASKF tracks the additional bias drift fairly well, until the aircraft performs the maneuver at 3400 seconds. At this point, the drift rate overcomes the filter's ability to compensate. And when the aircraft performs the final maneuver at 6400 seconds, the estimation error explodes to over 2000 ft. Again, this occurs because the filter cannot distinguish the specific force errors and the platform tilt errors during straight and level flight. It therefore adjusts its estimates of the platform attitude to compensate.

For the medium-range magnitude accelerometer failure, the Y INS specific force drift estimation error clearly depicts the filter's inability to track the added bias drift, Figure 16.



**Figure 16.** Cascaded Centralized Kalman Filter Estimation Error, Y INS Specific Force Error, (ft/sec<sup>2</sup>), Medium-Range Accelerometer Failure.

## CONCLUSIONS

The filter performances were much as expected. The CKF had nominally better estimation accuracy than the other two

designs. The FKF in the no-reset mode, benefits from the isolated local filters by having the ability to detect and isolate non-inertial sensor failures. The CASKF is over-confident in its position estimates, even though the execution time step for the master filter was large enough to offset the time-correlated GPS filter estimation outputs. The high dynamics within this trajectory enhance the test scenario by stressing the filters, even under normal unfailed conditions. On the one hand, there is better observability, but on the other, the unmodeled g-sensitive errors are revealed. When failures are implemented, the strengths and weaknesses of each filter under test can be observed.

The improved observability during the high dynamics brought forth the filters' inability to estimate the specific force errors due to the INS roll axis accelerometer failures simulated. This occurred for all three filters. None could distinguish that failure from the platform tilt errors during straight and level flight. The FKF estimation errors remained reasonable throughout all of the simulations conducted.

The truth models used with DKFSIM have proven to be medium to high fidelity representations of the sensor systems. The vast amount of information available from the simulation output data files allows for problem definition when designing and testing potential filter candidates. The GPS truth model is impressive, offering an efficient and realistic truth model simulation. Realistic conditions modeled permitted GPS outages, signal reacquisition delays, and GPS satellite frequency or receiver clock failures, for the simulations.

The filter models are representative of current filter implementations, using the appropriate dynamic equations and the most recent developments in filtering algorithms and data manipulation. The availability of the failure modes further enhances the capabilities of the simulator. Two failures were modeled and the expected results were achieved. As a demonstration of the simulator, this might suffice, but for design purposes, the intent is to discover the results from new filter designs. Therefore, based on the simulations, truth models output data, and the filter output data, DKFSIM should be considered an excellent design tool.

Upon observing the filters' behaviors, a broader study of the failure types and magnitudes might prove beneficial. However, a span offered by the simulations performed define a relatively coarse parametric range and hopefully will bracket any failure simulations with magnitudes smaller than those presented herein. Anything more than the largest failure magnitudes simulated would likely show results indicating sensor "hard failures", for both the GPS and INS failures.

## Summary

DKFSIM is an excellent simulator for designing and testing candidate filter architectures. Having several different variable conditions, the filter architecture may be stressed to meet various performance requirements. Ultimately, the results from these simulations offer evidence of an effective simulator hav-

ing realistic sensor truth models and the capability to simulate pertinent filter designs and failure modes.

Finally, the FKF demonstrated its expected fault tolerance when the GPS satellite frequency failure was introduced. The CKF exhibited difficulty in detecting any failure. Without adding some form of FDIR, the CKF is quite vulnerable. The CASKF demonstrated its recovery capabilities upon introducing the largest GPS satellite frequency failure. However, this response was enabled by the intrinsic design deficiency of this architecture, having an artificially small covariance matrix. Consequently, it is evident that the federated Kalman filter is the filter of choice. There is a very small loss of estimation accuracy, compared to the advantages gained by the inherent fault tolerance of the FKF architecture.

#### *Acknowledgements*

The authors would like to recognize Dr Neal Carlson for his contribution to our efforts in generating this paper, and for his unending patience. Also, we would like to offer our appreciation to Mr Jason Baldwin for his contribution to the completion of this paper.

#### REFERENCES

1. Carlson, Neal A., *User's Manual for the Distributed Kalman Filter Simulator (DKFSIM)*, Integrity Systems, Inc., Belmont, MA, 15 February, 1993.
2. Denaro, Loomis, and Geier. *Common Kalman Filter Development*. Document No. WRDC-TR-89-1152. Wright-Patterson AFB, OH, Feb 1990.
3. Lawrence, Paul J., Jr., and Berarducci, Michael P., "Comparison of Federated and Centralized Kalman Filters With Fault Detection Considerations." *IEEE PLANS, Position Location and Navigation Symposium, 1994*: 703-710 (Apr 1994).
4. Carlson, Neal A., *Distributed Kalman Filter Architectures, Phase II, Parts A and B*. Document No. WL-TR-95-1096. Integrity Systems, Inc., Belmont, MA, April 1995.
5. Maybeck, Peter S., *Stochastic Models, Estimation, and Control, Volume 1*. San Diego, CA: Academic Press, 1979.
6. Carlson, Neal A., "Federated Square Root Filter for Decentralized Parallel Processes". *IEEE Transactions on Aerospace and Electronic Systems*, 26: No. 3, 517-525, (Apr 1989).