Tuning and Robustness Analysis for the Orion Absolute Navigation System

Greg N. Holt[*]

NASA Johnson Space Center, 2101 NASA Parkway, Houston, TX 77058

Renato Zanetti[†]

Charles Stark Draper Laboratory, Houston, TX

Christopher D'Souza[‡]

NASA Johnson Space Center, 2101 NASA Parkway, Houston, TX 77058

The Orion Multi-Purpose Crew Vehicle (MPCV) is currently under development as NASA's next-generation spacecraft for exploration missions beyond Low Earth Orbit. The MPCV is set to perform an orbital test flight, termed Exploration Flight Test 1 (EFT-1), some time in late 2014. The navigation system for the Orion spacecraft is being designed in a Multi-Organizational Design Environment (MODE) team including contractor and NASA personnel. The system uses an Extended Kalman Filter to process measurements and determine the state. The design of the navigation system has undergone several iterations and modifications since its inception, and continues as a work-in-progress. This paper seeks to show the efforts made to-date in tuning the filter for the EFT-1 mission and instilling appropriate robustness into the system to meet the requirements of manned spaceflight. The results generally show Monte Carlo error performance bounded by the filter uncertainty for all phases of flight. Some future items of investigation are presented related to suspected anomalies in the trajectory truth reference file.

I.   Introduction

THE Orion spacecraft is to be NASA's next-generation Exploration vehicle for crewed missions beyond low-Earth orbit (LEO). As such, it is being designed to be capable of operating outside the comfortable confines of LEO. This introduces a variety of challenges, not least of which is the need for flexibility and robustness to a variety of environments including cislunar operations, low lunar orbit, and LEO. This paper seeks to describe the process of tuning the onboard navigation Kalman filter, including considerations for both performance and robustness. Whereas much of the focus and effort over the past year has been directed toward Exploration Flight Test 1 (EFT-1), the design is flexible enough to provide software extensibility for future exploration missions.

This paper is organized as follows: first, the overall architecture is presented along with some details of the Kalman Filter design. In section 2, the filter tuning methodology and high-level results are given. Next, the detailed results are given for each flight phase of the EFT-1 mission. In Section 5, the filter robustness analysis and design considerations are presented. Finally, a few remarks on future work are made.

The Orion Absolute Navigation System initial implementation was previously described by the authors.[?] Its underpinnings were derived by Hanak[?,?] and D'Souza.[?] Fundamentally, it is based off of work by such prolific designers as Bierman[?] and Schmidt.[?]

Filter performance is affected by many factors: data rates, sensor measurement errors, tuning, and others. This paper focuses mainly on the error characterization and tuning portion. Traditional efforts at tuning a navigation filter have centered around the observation/measurement noise and Gaussian process noise of the

---

[*]Navigation Engineer, Flight Dynamics Division/Mission Operations, AIAA Senior Member
[†]Senior Engineer, AIAA Senior Member
[‡]Navigation Engineer, Aerosciences and Flight Mechanics Division, AIAA Senior Member

American Institute of Aeronautics and Astronautics

Extended Kalman Filter. While the Orion Multi-Organizational Design Environment (MODE) team must certainly address those factors, the team is also looking at residual edit thresholds and measurement under-weighting as tuning tools. Tuning analysis is presented with open and closed loop Monte-Carlo simulation results showing statistical errors bounded by the 3σ filter uncertainty.

II.   System Architecture

The Orion Absolute Navigation System consists of inertial sensors and software that are used throughout the mission, from before launch to post-landing. As such, the design has to be robust to various environments and robust to various failures. The design consists of three Honeywell Orion IMUs (OIMUs), which are descendants of the Honeywell MIMU design. In addition, there are 2 Global Positioning System (GPS) receivers which provide both PVT and pseudorange and deltarange measurements. Finally, there are two star trackers to provide periodic attitude updates. This is shown in Figure 1.
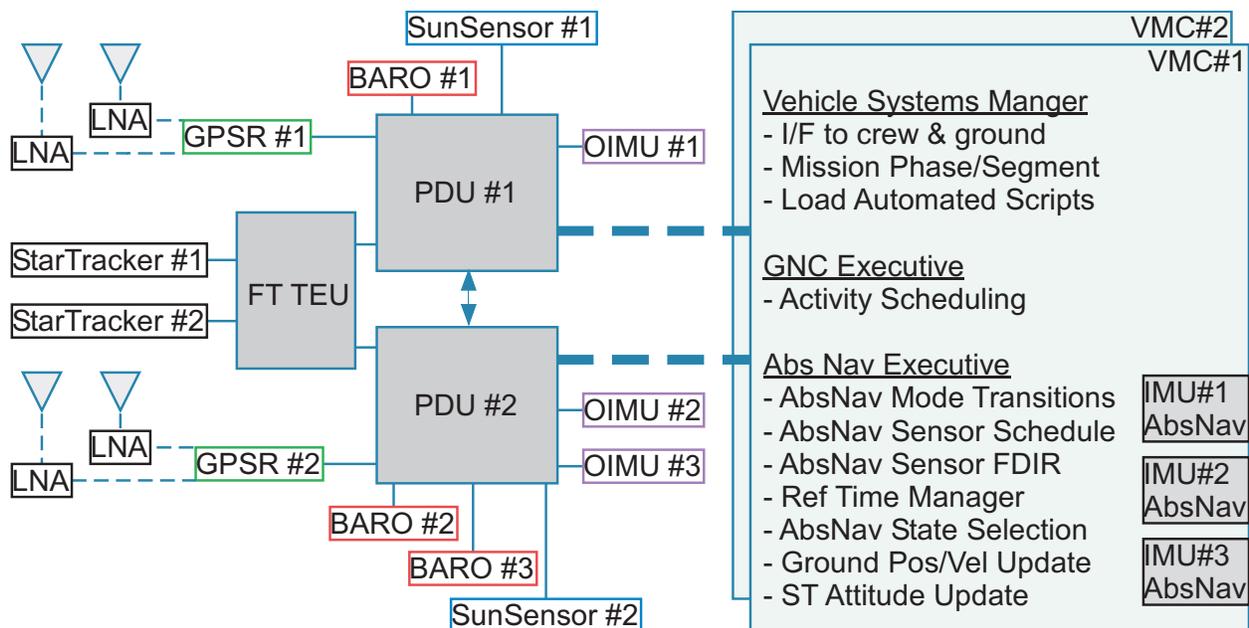

Figure 1. Orion Absolute Navigation Architecture - General View

For EFT-1 however, the Absolute Navigation System was scaled-down to comprise 2 OIMUs and 1 GPS receiver. Because of the short duration of the flight (two orbits), there will be no star tracker carried on-board. The EFT-1 configuration is illustrated in Figure 2.

III.   Filter Tuning

The measurements used in the Orion EKF are Integrated Velocity, GPS Pseudorange, and GPS Delta-range. On the pad, the Integrated Velocity measurement is used for IMU alignment/gyrocompassing, but not for the rest of the mission. The GPS antennas are obscured by the Launch Abort System shroud and therefore the GPS measurements are not available until several minutes into the ascent once the shroud is jettisoned. The GPS Pseudorange and Deltarange are used for the remainder of the mission.

A high-fidelity mission simulation was utilized for tuning the Orion EKF. The OSIRIS simulation uses a TRICK environment running the actual flight software in a variety of dispersed Monte Carlo scenarios.

The Orion filter design uses 24 Exponentially Correlated Random Variable (ECRV) parameters to esti-mate the accel/gyro misalignment and nonorthogonality. By design, the time constant and noise terms of these ECRV parameters were set to manufacturer specifications and not used as tuning parameters. They are included in the filter as a more analytically correct method of modeling uncertainties than ad-hoc tun-ing of the process noise. Tuning is explored for the powered-flight ascent phase, where measurements are
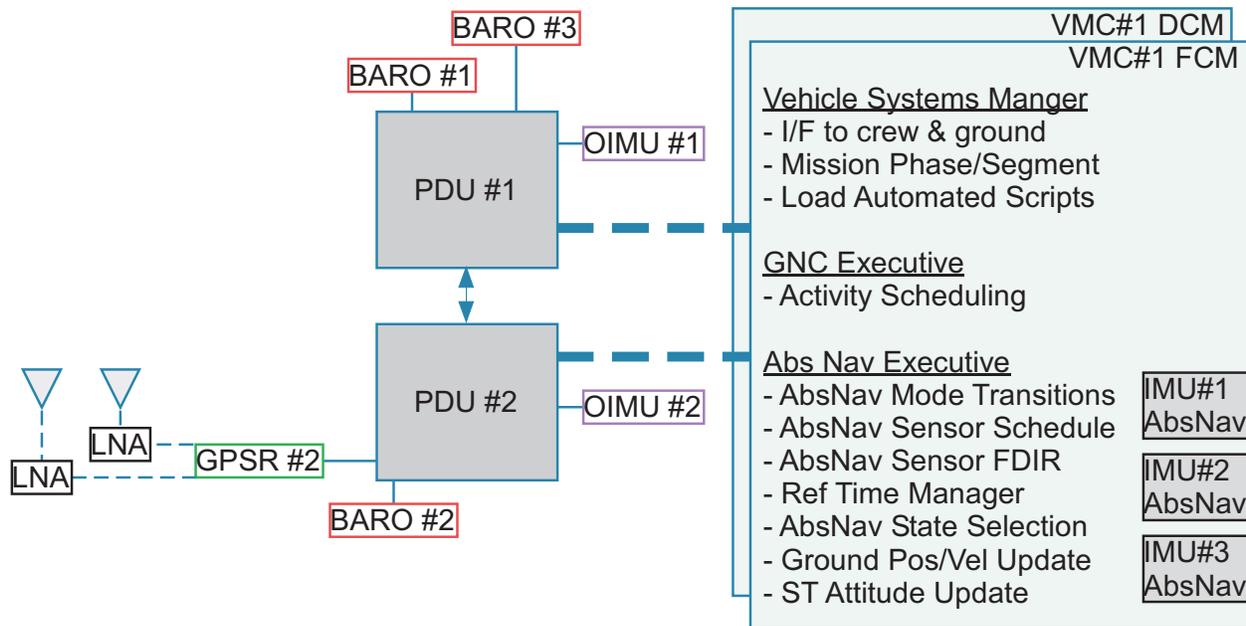
American Institute of Aeronautics and Astronautics

Figure 2. Orion Absolute Navigation Architecture for EFT-1

scarce and unmodelled vehicle accelerations dominate. On orbit, there are important trade-off cases between process and measurement noise. On entry, there are considerations about trading performance accuracy for robustness. Process Noise is divided into powered flight and coasting flight and can be adjusted for each phase and mode of the Orion EFT-1 mission. Measurement noise is used for the integrated velocity measurements during pad alignment. It is also used for Global Positioning System (GPS) pseudorange and delta-range measurements during the rest of the flight.

III.A. Methodology

As described previously, the IMU error states are modeled as ECRV states in the filter, so they are tuned using a process noise and time constant. The noise values were taken from the component specifications, and the time constants were chosen as 2 hours (for the IMU biases) and 4 hours (rest of the IMU error states) to be roughly constant over the duration of the EFT-1 mission. The other filter values to tune were the measurement underweighting, velocity process noise, angular process noise, and measurement weights. In general, the process noise was tuned by propagating the filter without measurements and ensuring the resulting error growth was bounded by the predicted uncertainty. The measurement noise was considerably more difficult to tune, since the measurements themselves were heavily time correlated, especially the pseudoranges. This is due to a poor ionosphere correction model and lack of low-elevation satellite masking. This behavior can be seen in the post-fit residuals in Figure 3.

III.B. End-To-End Performance

The filter tuning performance was analyzed for a closed loop End-to-End mission scenario. The position components are most impacted by GPS atmospheric errors, while the velocity components are most affected by unmodelled accelerations and IMU model errors. There is still ongoing investigation into the velocity error signature as it is suspected this may be an artifact of unrealistic accelerations supplied by the simulation from the trajectory dynamics file.

This section contains the results of 90 Monte Carlo runs of a complete mission, from pre-launch to landing. These runs confirm the validity of the chosen parameters. Figure 4 shows the performance of the usual eleven states (position, velocity, attitude, GPS receiver clock bias and drift) as well as the performance of the IMU error states. In these plots, the green section is the coarse align, the orange is the fine align, the first blue segment is ascent, the pink section is orbital flight, and the final blue section is entry and landing.
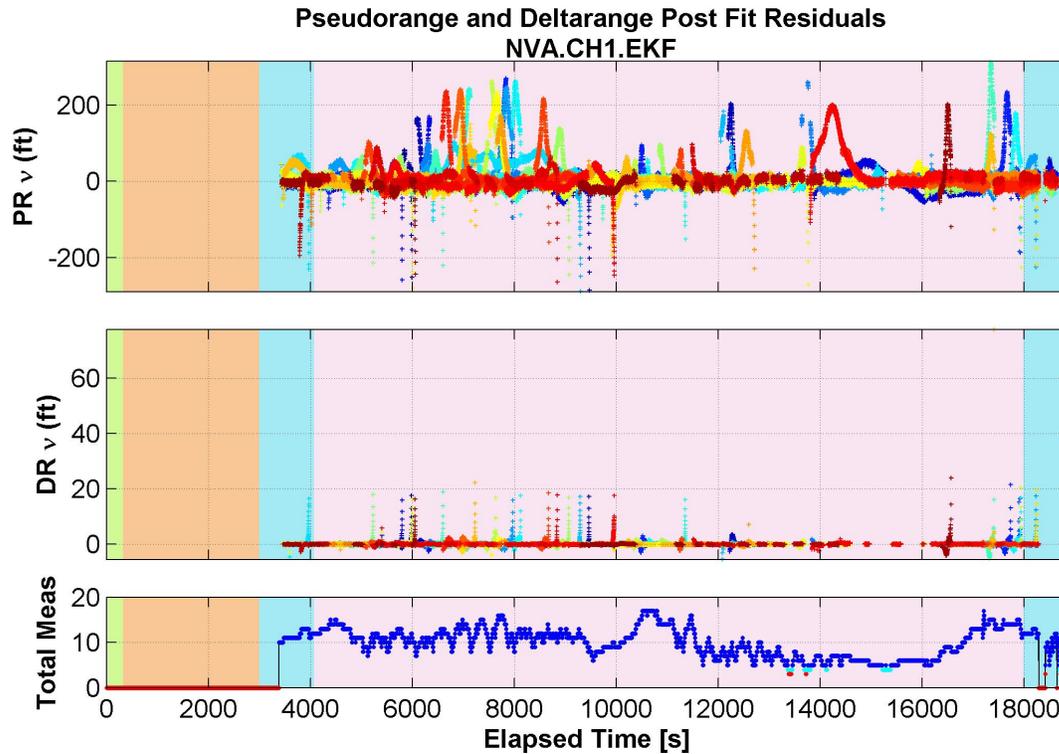
American Institute of Aeronautics and Astronautics

Figure 3. Postfit Measurement Residuals

III.C. Flight Phases

III.C.1. Prelaunch

This is the phase prior to launch when the vehicle is on the pad. The only measurement used during fine-align mode is integrated velocity (IV). The pseudorange (PR) and delta range (DR) measurement types are not used during this phase. The main purpose of fine align is to better estimate the attitude and the IMU states, therefore none of these states are considered and they are all estimated.

During fine align we sense the opposite of gravity which is a large acceleration. The accelerometer threshold is the same for all flight phases and its value is determined during tuning of the orbit phase. The velocity random walk of the accelerometer is very small, therefore the translational process noise is small and dominated by the unmodelled gravity. The angular process noise is chosen from the gyro's angular random walk spec value. The GPS clock bias and drift noises are chosen from the expected clock physical properties. The coast values for the process noise are not used in this phase, they are only used during the orbit phase. They are kept the same across flight phases and tuned for orbit.

Figure 5 shows the performance of 450 monte carlo runs in a Launch Hold scenario which transitions back to fine align for a longer hold prior to launch. It demonstrates the choice of parameters provides good performance even in the worst-case alignment conditions. In this scenario the initial position error on the pad is dispersed. The attitude, position, and velocity are all shown to be well bounded for the entire duration of the alignment and launch hold.

III.C.2. Ascent

The ascent phase is divided in two parts, the first when GPS measurements are not enabled, and the second when they are. The only difference between Fine Align and Ascent Without GPS is that IV measurement processing flag is set to zero in the latter. Integrated velocity (IV) is only used on the pad, while GPS measurements (PR and DR) are utilized during this phase. To date there is no simulation or any other evidence that attitude or IMU state estimation would corrupt the EKF solution. Therefore none of these
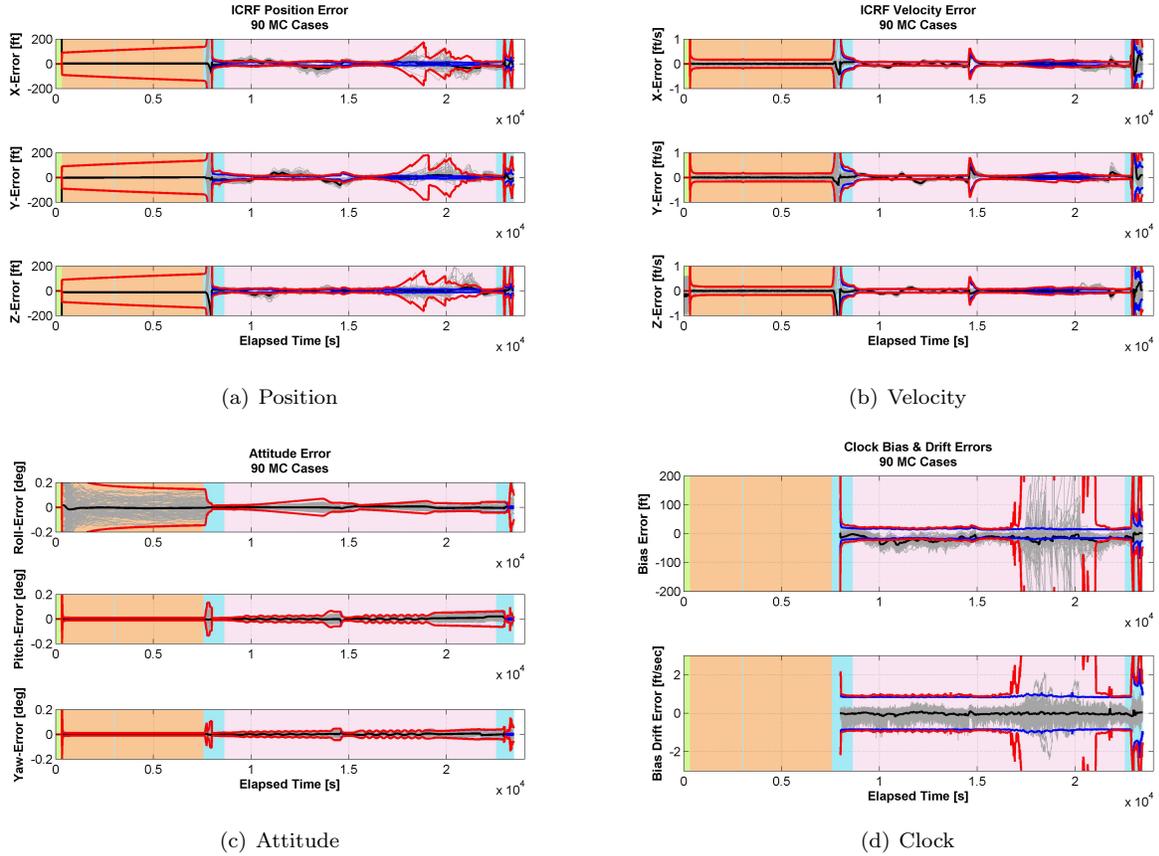
American Institute of Aeronautics and Astronautics

(a) Position

(b) Velocity

(c) Attitude

(d) Clock

Figure 4. End-to-End Performance

states are considered and they are all estimated.

The maximum number of processable measurements is set to 12 which is a large enough number to obtain very good performance while keeping the throughput reasonably low. To avoid possible transient issues the first PR measurement is not processed. After a long blackout the covariance becomes very large and the nonlinearity of the DR measurement creates convergence issues. Through numerical simulation it was determined that allowing for multiple PR ($\sim$30s) to be processed before incorporating a DR mitigates this issue because the PRs shrink the uncertainty before DRs are introduced. If a satellite is not present for a single cycle, the counters are not reset and if the satellite comes back is immediately used as a measurement. If the satellite is absent for more than a cycle the counters are reset.

PR and DR noise standard deviations are obtained from GPSR performance. The numbers are kept high to include residual ionosphere error. A larger residual edit threshold is used to compensate for the higher noise values. While we allow for the usage of the PR variance output from the GPS receiver, the EKF has a PR variance floor that will almost always prevent us from using the GPS receiver output value except in the case of very large values. Underweighting is applied when the estimated measurement has an uncertainty greater than $100ft$.

For the EFT-1 mission we want to test GPS clock stability, clock filter state restarts, and high altitude GPS processing, so we do not want to exercise a GPS clock reset using the clock bias reinitialization timer. Therefore, this value is set to a number larger than the expected duration of the mission.

Initial uncertainty of the clock bias and drift are obtained from GPSR specifications. When accelerating fast under the chutes the attitude dynamics is not accurately represented by the 4Hz IMU buffer. Therefore PR and DR are inhibited above a certain angular velocity, the values are determined from simulation analysis.

A deltarange timing threshold of 1.5 seconds prevents the flight software from trying to incorporate a DR measurement using two non-consecutive PR measurements (GPS measurements are available at 1Hz).
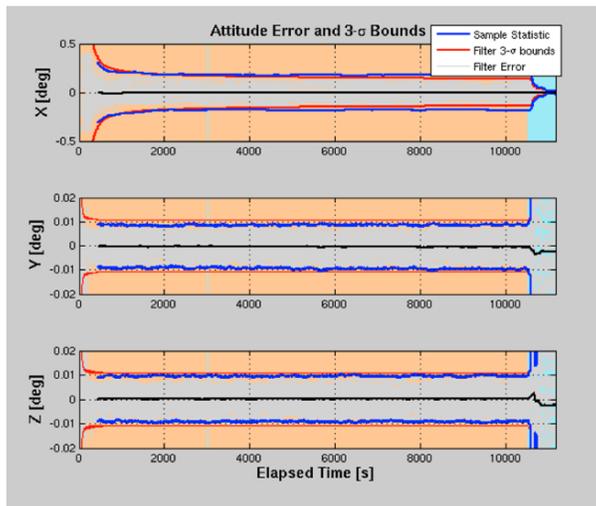
Figure 6 shows the performance of 450 Monte Carlo runs, this is a nominal scenario to demonstrate the choice of parameters provides good performance. These runs are performed without dispersing the GPS
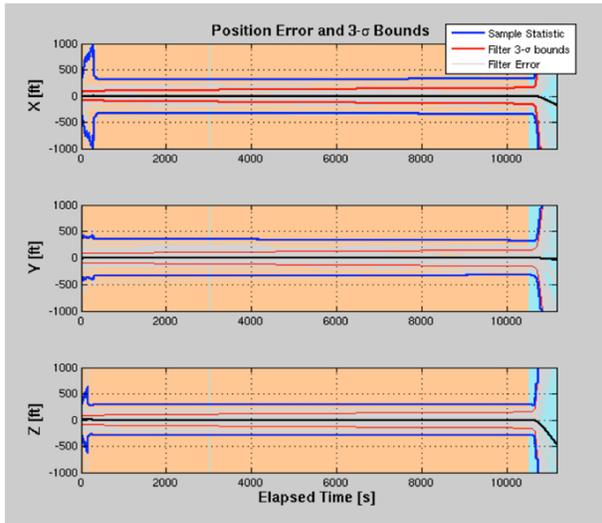
(a) Position
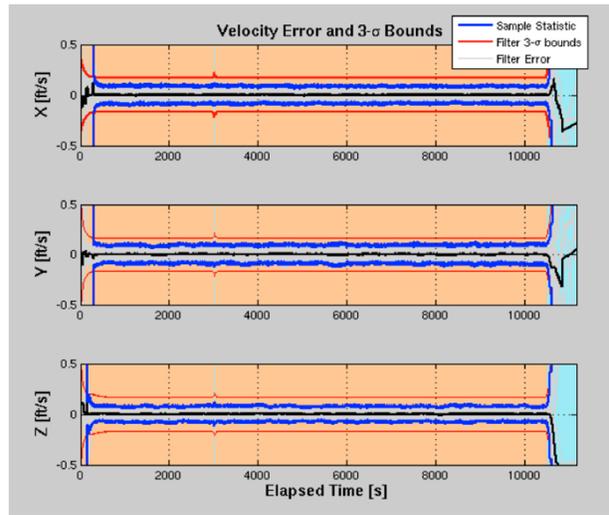
(b) Velocity

(c) Attitude

Figure 5. Prelaunch Performance

constellation in order to obtain the same set of measurement for each Monte Carlo run. As before, the errors seem well bounded and convergence times are reasonable.
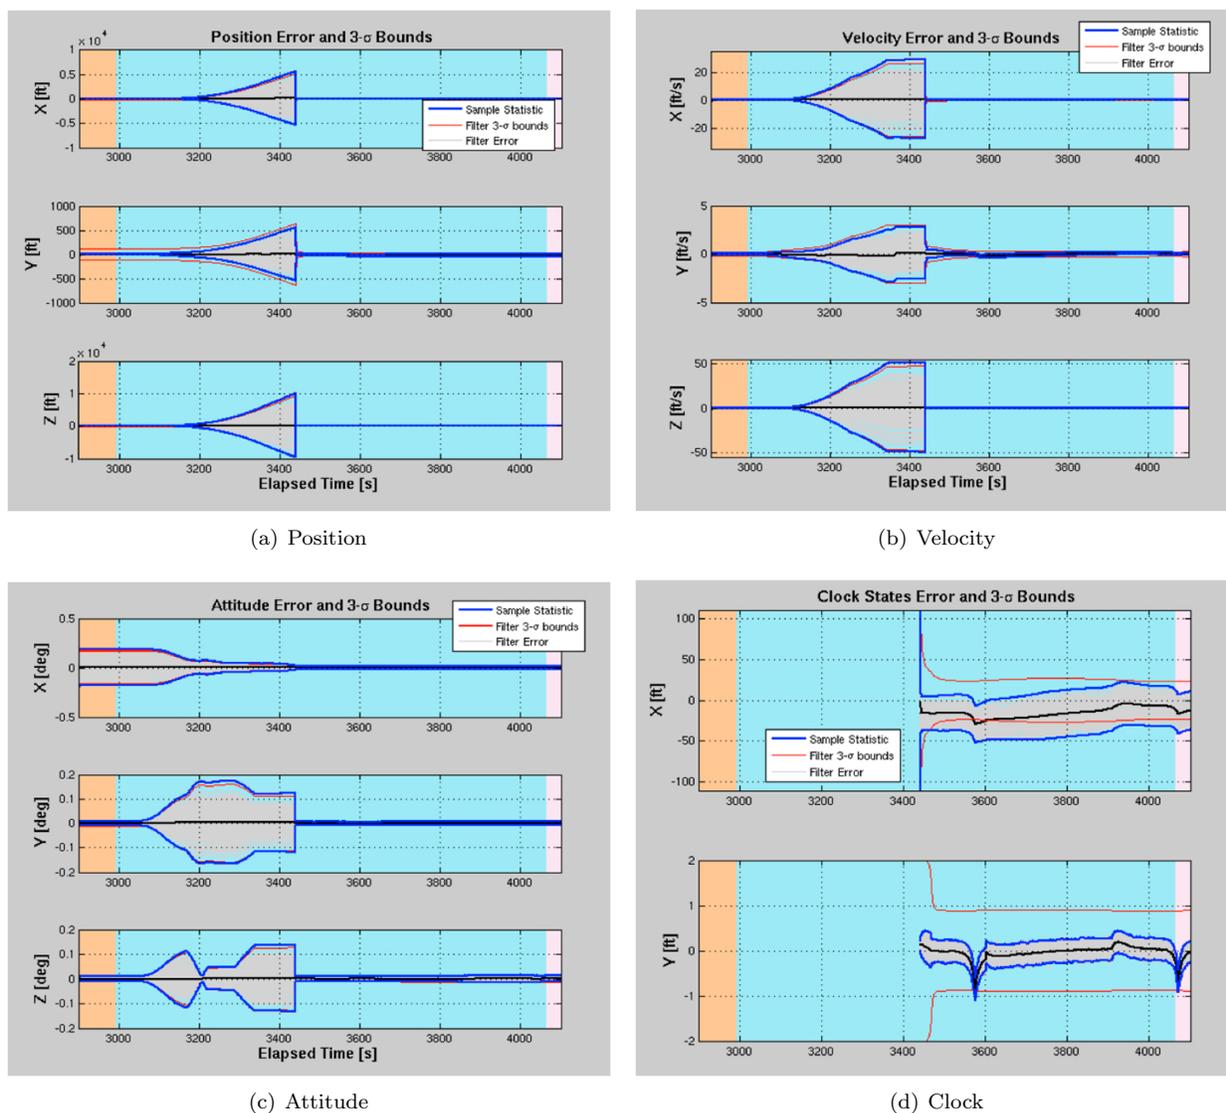

(a) Position

(b) Velocity

(c) Attitude

(d) Clock

Figure 6. Ascent Performance

III.C.3. Orbit

The choice of acceleration threshold is documented by Zanetti, Clark, and Holt.[7] The accelerometer bias is constant for a particular trajectory, but random across an ensemble of trajectories. To obtain a worst-case estimate we take three times the standard deviation of the bias, and treat the result as a constant. Each error source applies to each component of the velocity. We assume that the velocity white noise and the velocity random walk are independent. The same GPS parameters are used during orbit as are used during ascent. The only change is that a slightly lower value is used for the underweighting coefficient.

Figure ?? shows the performance of 450 Monte Carlo runs, this is a nominal scenario to demonstrate the choice of parameters provides good performance. These runs are performed without dispersing the GPS constellation in order to obtain the same set of measurement for each Monte Carlo run. Although some of the errors lie outside the three-standard deviation bounds estimated from the filter's covariance matrix, this is a nominal condition. The errors quickly converge to values within the requirement for accuracy of estimation of on-orbit position and velocity. The acceleration models in the trajectory simulation file are suspected to not match the velocity, so the oscillating error behavior is likely at least partially due to this inconsistency.

American Institute of Aeronautics and Astronautics

III.C.4. Entry

Figure 7 shows the performance of 450 Monte Carlo runs, this is a nominal scenario to demonstrate the choice of parameters provides good performance. These runs are performed without dispersing the GPS constellation in order to obtain the same set of measurement for each Monte Carlo run.

(a) Position

(b) Velocity

(c) Attitude

(d) Clock
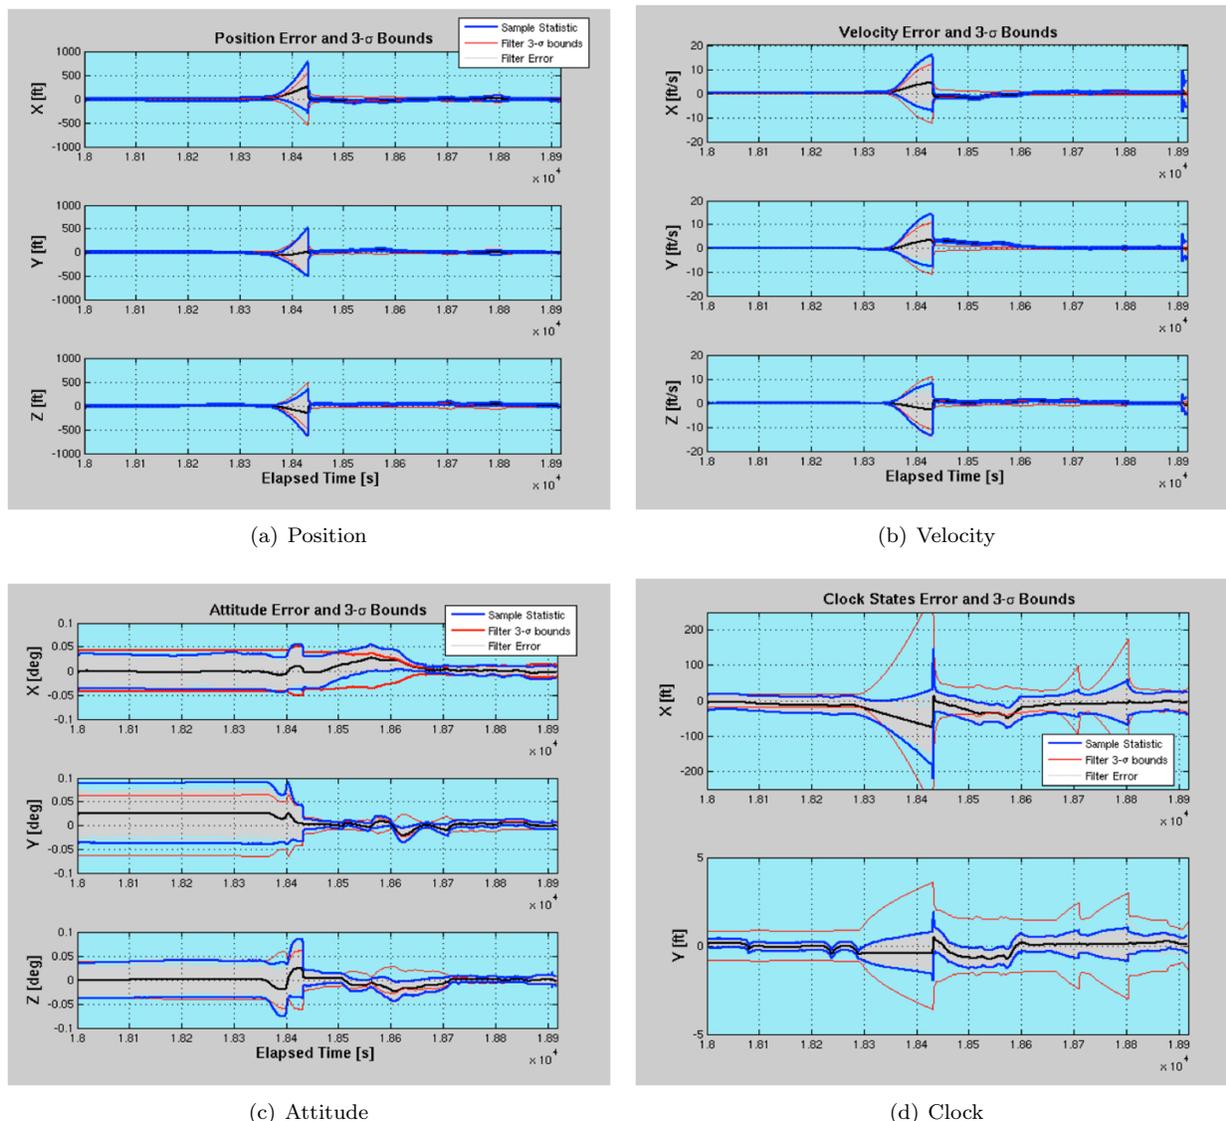
Figure 7. Entry Performance

IV. Robustness

The robustness effort has been focused on maintaining filter convergence and performance in the presence of unmodelled error sources. These include unmodelled forces on the vehicle and uncorrected errors on the sensor measurements. Orion uses a single-frequency, non-keyed GPS receiver, so the effects due to signal distortion in Earth's ionosphere and troposphere are present in the raw measurements. Results are presented showing the efforts to compensate for these errors as well as characterize the residual effect for measurement noise tuning. Another robustness tool in use is tuning the residual edit thresholds. The trade-off between noise tuning and edit thresholds is explored in the context of robustness to errors in dynamics models and sensor measurements. Measurement underweighting is also presented as a method of additional robustness when processing highly accurate measurements in the presence of large filter uncertainties.

American Institute of Aeronautics and Astronautics

IV.A. Entry

The entry robustness performance concentrated on two areas: the loss of GPS signals during entry and the reacquisition of GPS signals during entry after an extended outage. Special consideration of filter robustness is given to the Entry phase since the Orion MPCV will not be actively controlled during ascent and most of the orbital coast and thus the navigation results are only for performance evaluation.

IV.A.1. Loss of GPS During Entry

The first robustness test considered was the loss of GPS at Entry Interface that is not regained for the rest of the descent and landing. The expected behavior for appropriate tuning was to see that the errors were bounded by the uncertainty in the filter states. Again, 90 Monte Carlo simulations were run dispersing the IMU errors, GPS constellation, atmospheric conditions, etc. The simulation starts near the second orbit apogee at the separation from the launch vehicle upper stage. The entry phase is shown in light blue, starting 5 minutes before entry interface. The plasma blackout begins soon after plunging into the atmosphere, approximately 3000 seconds into the run. The position error is shown in Figure 8(a). The Monte Carlo samples are well bounded by the filter 3σ uncertainty values. The velocity errors are shown in Figure 8(b). Again, the error from the Monte Carlo runs are well bounded by the filter 3σ uncertainty values. The attitude errors are shown in Figure 8(c). As expected with very low process noise, the covariance growth is minimal even without measurement updates and the errors remain bounded. Finally, the clock errors are shown in Figure 8(d). The process noise on the clock terms is tuned conservatively large to allow for unforeseen variations in clock performance on the day of flight. The errors are then, as expected, bounded by the 3σ uncertainty after the GPS measurements are lost.

IV.A.2. Reacquisition of GPS After Extended Outage

The next robustness scenario tested was acquiring GPS signals after an extended outage. As a stressing case, the simulation was run with no GPS measurements for the entire flight until after the plasma blackout. By this time, of course, the filter uncertainty will have grown quite large, and the ability to incorporate very accurate measurements will be a test of the chosen measurement underweighting values. The position errors are shown in Figure 8(e). It is noteworthy that the errors are still more pronounced after reacquisition than in the nominal entry case. Velocity errors are shown in Figure 8(f). Convergence times are reasonable but, as in the position errors, the values are higher than seen in a nominal entry case. High dynamics during guided entry and during the chute deploy sequence are suspected as contributing factors. Attitude error is shown in Figure 8(g). Due to low process noise on the attitude terms, the uncertainty is not too large when GPS is acquired. The convergence is very reasonable and the performance is soon comparable to that seen on a nominal entry case. Finally, clock errors are shown in Figure 8(h). The clock states are not estimated when no GPS measurements are present, so the plot begins at the point where GPS is acquired. The Orion EKF bootstraps its clock bias and drift estimate from the receiver kinematic solution, so the initial performance is good and stays well bounded for the duration of the entry.
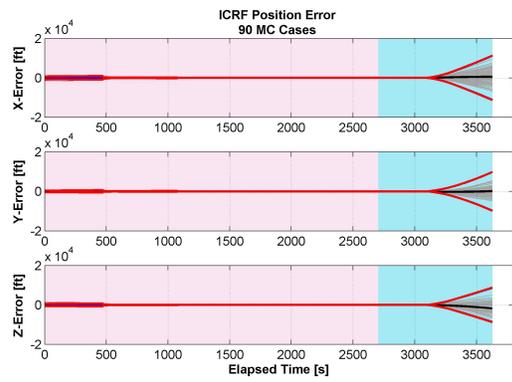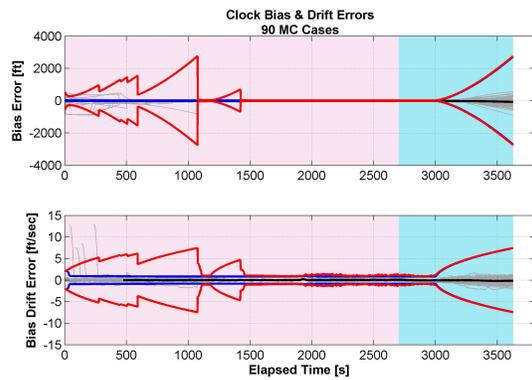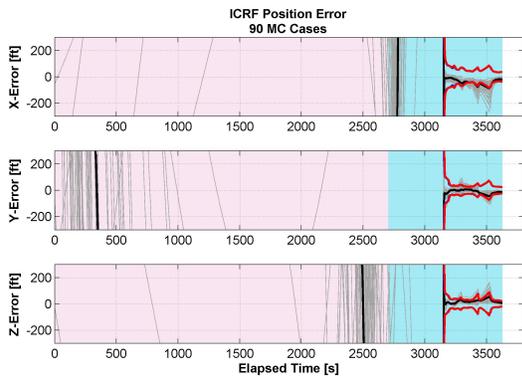
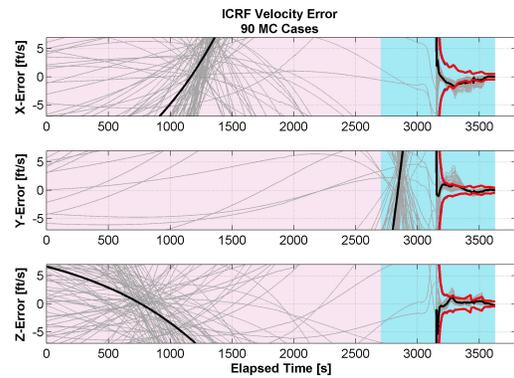(a) Position with GPS failure
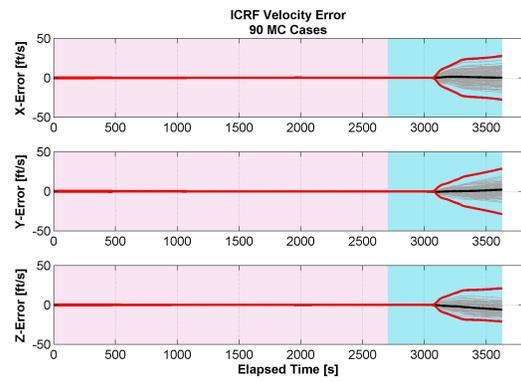
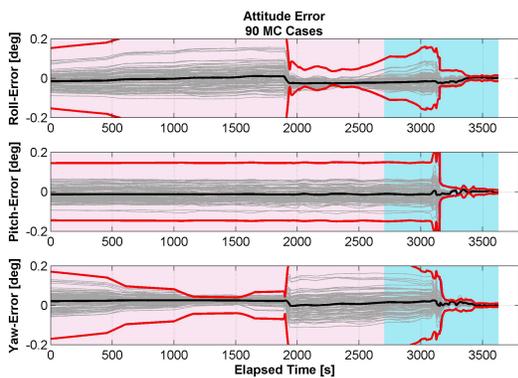(b) Velocity with GPS failure
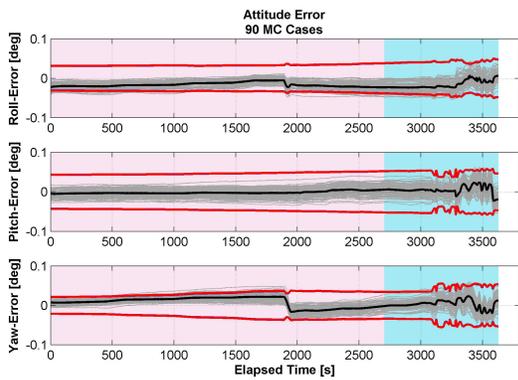
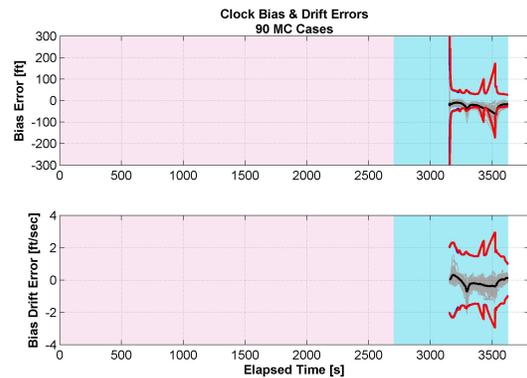(c) Attitude with GPS failure

(d) Clock with GPS failure

(e) Position with GPS reacquisition

(f) Velocity with GPS reacquisition

(g) Attitude with GPS reacquisition

(h) Clock with GPS reacquisition

Figure 8. Entry Performance Robustness Results

V.   Future Work

There are still many parts of the filter tuning and simulation environment that are under investigation. The oscillating velocity error signature during Low Earth Orbit is suspected to be at least partially due to numerical issues with the input trajectory file. Future implementations of the filter design will include masking of GPS signals with high ionospheric interference, which should mitigate the position excursions outside the 3σ filter uncertainty.

Acknowledgments