AAS 08-103

ORBIT DETERMINATION OF SATELLITES IN LUNAR ORBIT USING AN OPTIMAL SEQUENTIAL FILTER

James Woodburn[*] and John H. Seago[†]

Determination of orbits about the Moon poses a number of extra challenges compared to those of Earth orbiting satellites. Most problematic for sequential estimation is the uncertainty in the dynamical environment. We describe deterministic and stochastic upgrades made to the Orbit Determination Tool Kit in pursuit of this goal. We present results from the processing of tracking data for Lunar Prospector over a one month period and STEREO during the lunar fly-by portion of the mission.

INTRODUCTION

Orbit determination is required for all operational satellite programs. The determination of orbits about bodies other than the Earth poses a number of extra challenges compared to those of Earth orbiting satellites. Measurement modeling must be performed at a higher fidelity, the accuracy and consistent use of planetary ephemerides becomes more important, and the relationships between multiple reference frames must be properly managed. However, these issues are deterministic and can be overcome by careful attention to detail.

More problematic is the uncertainty of the dynamical environment. For the case of lunar-orbiting satellites, the primary issue is the uncertainty in the lunar gravitational potential on the far side of the Moon. This type of dynamical uncertainty is typically dealt with in batch weighted least squares estimation by adjusting the fit span of the orbit determination run and solving for empirical accelerations with frequencies natural to the problem.

Our endeavor is to apply optimal sequential estimation methodologies to the problem of Moon-orbiting spacecraft. A successful application of this goal would result in a sequential estimator that could be continuously employed over the lifespan of the spacecraft. The dynamical uncertainty must be considered in a way that is properly connected to the physics of the problem. This is to satisfy the optimality criteria for the estimate and eliminate the need for operational tuning of the filter for each phase of a mission.

In this paper, we present a description of the deterministic and stochastic modeling upgrades that were recently made to the Orbit Determination Tool Kit (ODTK) software product developed by Analytical Graphics.[1] We also present sample results based on the processing of tracking data from the Lunar Prospector mission over a one month period, and results based on processing tracking data for the STEREO mission during the lunar fly-by portion of the mission. Finally, we

[*] Chief Orbital Scientist, Analytical Graphics, Inc., 220 Valley Creek Blvd., Exton, PA, 19341-2380.
[†] Astrodynamics Engineer, Analytical Graphics, Inc., 220 Valley Creek Blvd., Exton, PA, 19341-2380.

discuss the expected improvement from new lunar gravity fields developed from the far-side tracking of such missions as SELENE.

ODTK ENHANCEMENTS

A number of enhancements were required to process observations for lunar-orbiting spacecraft in ODTK. Many of the required changes were architectural or related to the user interface, and mainly designed to provide an abstraction of the primary gravitational body. Two major estimation enhancements were also required: providing a means to deal with the uncertainty in the lunar gravity field, and observation modeling for measurements from the NASA/JPL Deep Space Network (DSN).

The Lunar Gravity Field

The largest challenge to the operational use of a sequential filter for lunar-orbiting spacecraft is the uncertainty in the gravitational potential of the Moon. The Moon holds no appreciable atmosphere, and thus allows for spacecraft to orbit at very low altitudes. This is advantageous for gravity field determination; however, prior to the SELENE mission there had been no tracking data for satellites on the far side of the Moon. The gravity field models developed from tracking of modern lunar missions, such as Lunar Prospector, have the interesting characteristics of being of high degree and order - needed to capture the high frequency information available on the front side of the Moon - while also having large uncertainties on the far side (Figure 1).

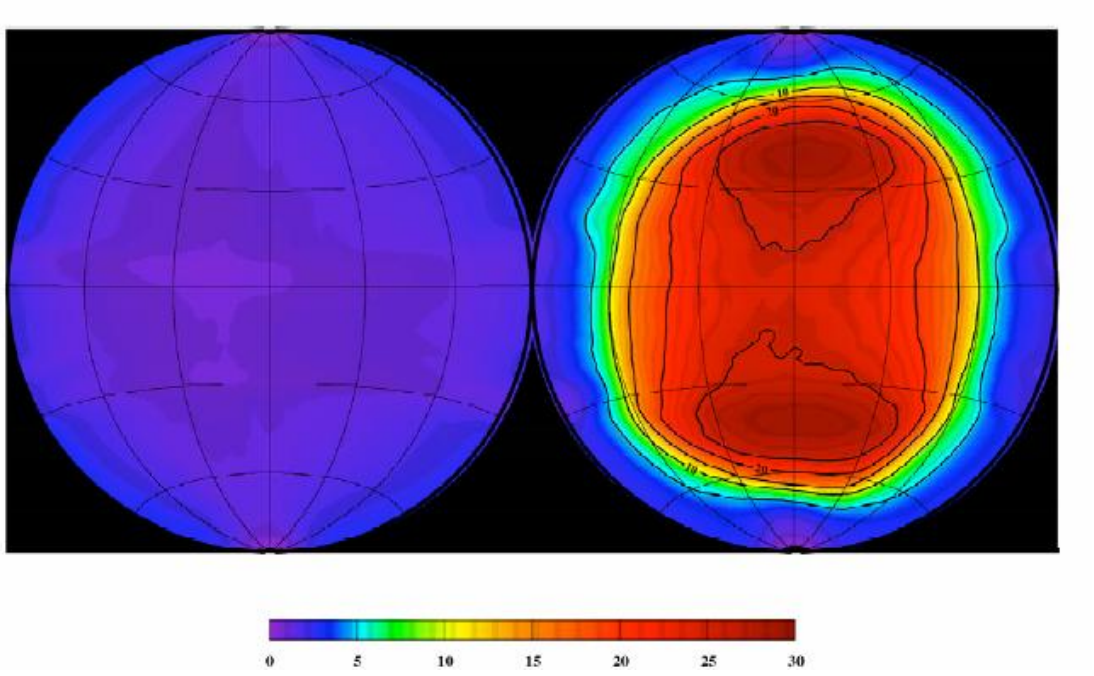

Figure 1 Geoid uncertainty for the Moon (m), LP150Q[2]

When using traditional batch weighted-least-squares estimators in the presence of significant force modeling uncertainties, fit spans are adjusted to find an optimal trade between incorporating more observations, via longer fit spans, and matching the trajectory to the dynamical model which is more readily accomplished with shorter fit spans. The effectiveness of the adjustment is usually gauged by overlap comparisons and residual analysis.

A different strategy is used with a sequential filter over the life of a mission. The desire is to always restart the filter with the final state estimate from the previous filter solution. If this goal is achieved, then all observations are effectively processed as one long, recurrent orbit estimate. This lofty goal becomes complicated, however, by the need to adequately model the uncertainties in the dynamical model. Otherwise, if the filter becomes overly optimistic, all tracking data will be rejected and the filter will diverge.

Dynamical uncertainties are considered by using a dynamical process-noise model during the filter time update. In simple terms, the state error covariance (which approximates the uncertainty in the estimated state) becomes inflated between observations times to compensate for imperfect knowledge of the forces on the satellite. The dynamical process model for uncertainty in the gravity field used during this study is described in the biblography.[3, 4, 5]

The gravity process-noise model used in ODTK averages gravity errors from all latitudes and longitudes over a shell of constant altitude. For the case of the Moon, this averaging causes the process-noise model to be infrequently optimistic (when the satellite is over the back side of the Moon) and pessimistic much more of the time (when the satellite is not over the back side of the Moon). The use of such a process-noise model for lunar-orbiting satellites therefore results in a generally pessimistic covariance. Luckily, a pessimistic covariance keeps the filter stable - we demonstrate this in the sequel by processing one month of low lunar-orbit tracking data in a single arc.

DSN measurement models

Measurements from the DSN are processed using observation models for two-way sequential range, two-way and three-way Doppler, and two-way and three-way total-count-phase measurements. In the context of DSN measurement modeling, a two-way measurement involves a signal which is transmitted and received by the same ground station. A three-way measurement involves a signal which is transmitted by a different ground station than the one receiving it. The transmitted signals may be ramped or constant in frequency. The mathematical modeling of these measurement types follows Moyer, except as noted in the sequel.[6] The measurement modeling is consistent with DSN operations using Block V receivers following the completion of the Network Simplification Program.

The DSN sequential range measurement is an ambiguous range. The length of the unambiguous part of the observable is supplied in the tracking-data file in conjunction with the observation value. Normal DSN procedures involve an internal calibration of antennae prior to each tracking pass. The pre-pass calibration serves to remove most of the bias from the ranging data. The remaining bias is small, on the order of 14 range units or less, and independent from one pass to the next.[7] While ODTK does not estimate a state-parameter called "pass-dependent measurement bias," the time between passes for a particular DSN antenna is long compared to the time between observations; this allows the exponentially correlated Gauss-Markov bias state estimate available in ODTK to absorb the remaining measurement bias.

The DSN total-count-phase measurement is a measure of total received cycles of the carrier from a specified epoch. The epoch for the start of the count is supplied in the tracking-data file in conjunction with the observed value. Sequential total-count-phase measurements are differenced during the reading of the tracking-data file so that the resulting measurement processed in ODTK is the received cycles over the current count interval. This differencing operation, which simplifies the modeling of the measurement in a sequential filter, makes the processing of total-count-phase observations equivalent to processing Doppler observations. No bias estimate is

typically used when processing total-count-phase measurements, but a bias may appear if the spacecraft is spin stabilized.

The DSN Doppler measurement is a measure of total received cycles of the carrier over the specified Doppler count interval. The count interval is supplied in the tracking-data file in conjunction with the observed value. The Doppler measurement differs slightly from the total-count-phase measurement in that the time tag on the Doppler measurement is placed in the middle of the count interval while the time tag on the total-count-phase measurement is placed at the end of the count interval. Although the model for this measurement type has been implemented in ODTK, it does not appear that Doppler measurements are routinely reported any longer because the total-count-phase measurements provide the same information content.

To process DSN measurements, ODTK was enhanced to read observations from the TRK 2-34 tracking data format.[8] The TRK 2-34 format contains a wealth of information about the configuration of the antenna while observations were taken. Of primary importance to the processing of the total-count-phase measurements is the frequency ramp information contained in the TRK 2-34 file. Ramps are provided for each tracking station and consist of a starting time and frequency and a frequency rate. Modeling of the total-count-phase observations requires a very precise computation of the transmission interval and proper indexing into the frequency ramp information to correctly determine the number of cycles transmitted during that interval.

LUNAR PROSPECTOR

Approximately one month of Lunar Prospector (LP) tracking data were processed for a time interval starting on 1 March 1999 and ending on 30 March 1999. The approximate orbit parameters for LP during this time frame are given in Table 1. During the tracking-data interval, the Moon completed slightly more than one complete revolution about the Earth. This is significant because all possible tracking geometries were experienced over the period of the tracking data. Under normal operations, the tracking data would have been processed in chunks at a pre-determined frequency (*i.e.*, once per day) to provide orbit updates for mission planning purposes. For our analysis, however, we processed the entire 30 days of data in a single filter run. Since an operational filter would be periodically restarted from data records containing the final filter state from the prior filter run, the results are the same as what would have been achieved by processing the data over many shorter intervals.

Table 1 Orbit Elements for Lunar Prospector

Epoch: 28 Feb 1999 00:00:00.00 UTC			
Semi-major axis	1767.6 km	Periapsis Height	20.1 km
Eccentricity	0.0053	Apoapsis Height	39.0 km
Inclination (J2000)	66.0 deg	Inclination (TOD)	89.9
Ascending Node (J2000)	193.4 deg	Ascending Node (TOD)	194.3
Argument of Periapsis (J2000)	93.9 deg	Argument of Periapsis (TOD)	101.0

The force modeling options selected for LP are shown in Table 2, and the LP100K gravitational field was augmented with matched gravity process-noise inputs generated using the algorithm described in Wright *et al.* (2008).[3]

Table 2 Force modeling for Lunar Prospector

Lunar Gravity			
	Field	LP100K	100 x 100
	Solid Tides	On	$k_2 = 0.03$
3rd Body Gravity			
	Sun		
	Earth		
Solar Pressure			
	Satellite model	Spherical	Cp = 1.2
	Shadow model	Dual Cone	

The LP tracking data contained observations from DSN stations 16, 27, 46, and 66. The locations of these tracking stations are illustrated in Figure 2. Stations 16 and 27 are located at Goldstone. Station 66 is in Madrid and station 46 is located at Canberra and station 66 is located at Madrid. All stations provided Doppler, with DSS27 also providing a small amount of ranging data.

GTDS outputs were also made available for the time period covered by the tracking data. These output files sometimes show the processing of observations from additional stations that were not included in the provided data sets, and sometimes do not show the processing of data from some stations that were included.

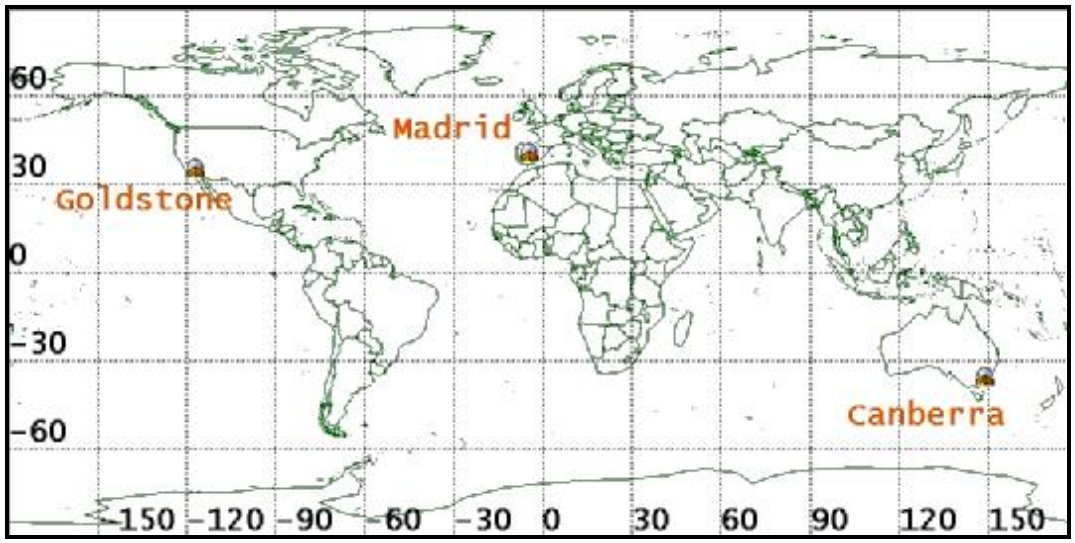

Figure 1 Location of DSN tracking stations used for Lunar Prospector

The LP tracking data were supplied in the Universal Tracking Data Format (UTDF). The Doppler data in UTDF are not represented as a cycle count on the carrier frequency as reported in TRK-2-34 format. Since the tracking data came in UTDF format, ODTK processed the data using its nominal two-way range and two-way Doppler observation models. The main differences in the

nominal two-way range and the DSN sequential range models are: the absence of a model for the gravitational bending of the signal path, and the solution of the light-time equation in the geocentric frame, instead of the solar system barycentric frame. For a lunar orbiter, the differences in the observation models are small.

A summary of the tracking data processed for LP is given in Table 3. The indicated numbers of measurements include only measurements that were accepted by the filter. All tracking was performed in S band.

Table 3 Tracking data summary for Lunar Prospector

Station	Measurement Type	# Observations
DSS16	Doppler	18430
DSS27	Range	11570
	Doppler	18961
DSS46	Doppler	32562
DSS66	Doppler	37212

Precise tracking-station locations were provided at the start of this study, but the measurement white noise and bias uncertainties for the tracking stations and the transponder delay were not known. These values were iteratively determined by repeatedly processing the tracking data and tweaking input values until the resulting residual patterns seemed reasonable.

The initial value of the transponder delay was determined by plotting range residuals from orbit-determination runs fit only to Doppler data. The results of this calibration process are shown in Table 4. An observed bias was added to the two-way Doppler measurements to account for the fact that Lunar Prospector was spin stabilized.[9] The same calibration values were used for all stations. Additional calibration could be done on a per-station basis, but little accuracy gain would probably have been achieved due to the significant process noise associated with lunar orbits.

Table 4 Measurement statistics

Measurement Type	Constant Bias	Bias Est.	Bias Sigma	White noise
Range	0.0 m	Yes	10 m	1.5 m
Doppler	-0.112 cm/s	No	N/A	0.1 cm/s
Transponder delay	405 m			

A number of quality indicators is used to scrutinize the orbit determination results obtained using ODTK. Two of the most important are the apparent randomness of the residuals, and satisfaction of the McReynolds filter-smoother consistency test. Histograms of normalized residuals - residuals divided by the square root of the measurement error variance - are presented in Figures 4a-4b. The theoretical normal curve is superposed on the histograms. Residual scatter

plots for typical range and Doppler passes are shown in Figures 5a-5b. There is no apparent indication from these figures of significant measurement modeling errors. It is, however, interesting to note the expansion of the 3-sigma measurement uncertainty envelope for the Doppler passes. This expansion is largely due to the pessimistic gravity error process-noise model.[3] This expansion is also the main reason for the severely peaked Doppler histograms in Figure 4b.

The McReynolds filter-smoother consistency statistic is computed by dividing the difference between the filter and smoother estimates by the difference between the filter and smoother root variances for specific state elements.[10] This test statistic should be normally distributed, so that under ideal conditions, the test statistic should remain within ± 3 for 99% of the time. If it is highly abnormal, then this is interpreted as a defect in the dynamical and/or measurement modeling.

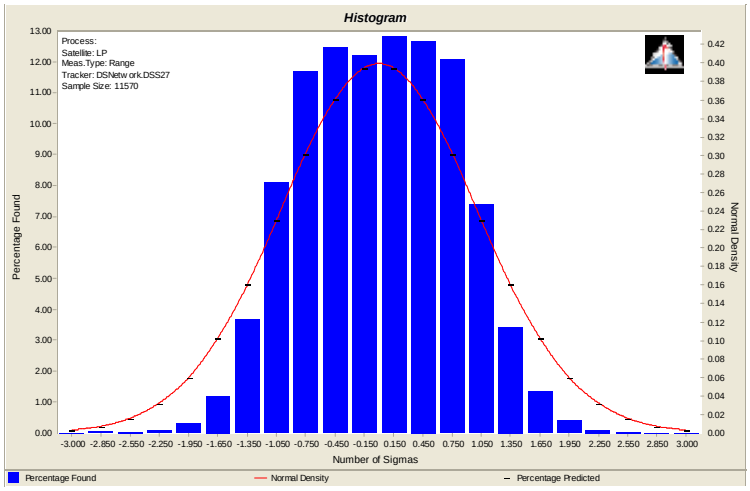

Figure 4a Range residual histogram

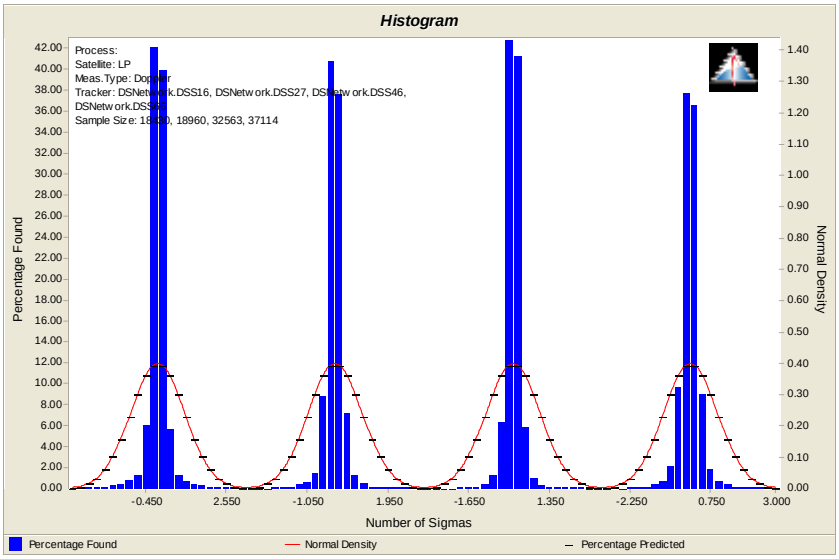

Figure 4b Doppler residual histograms

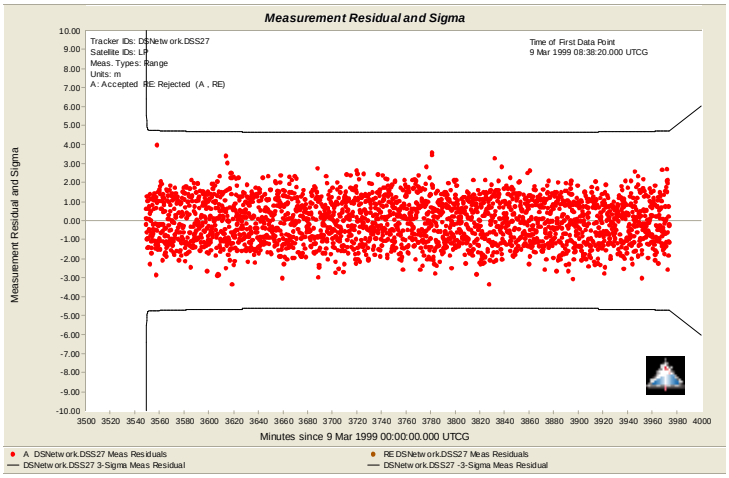

Figure 5a Range pass residuals

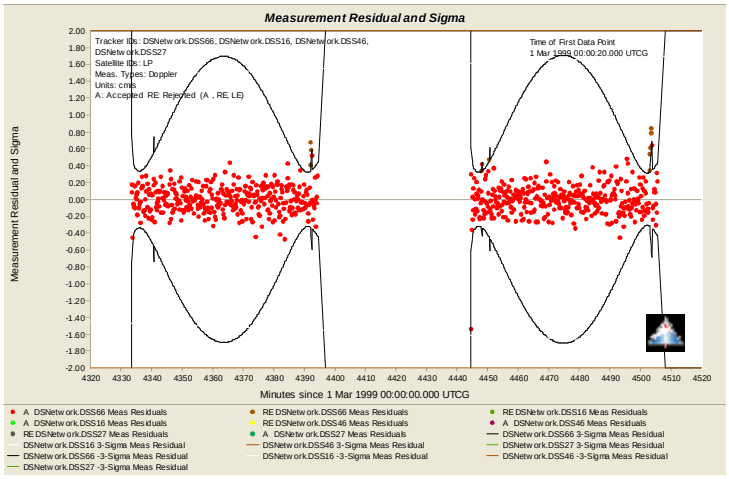

Figure 5b Doppler pass residuals – Earth in LP orbit plane

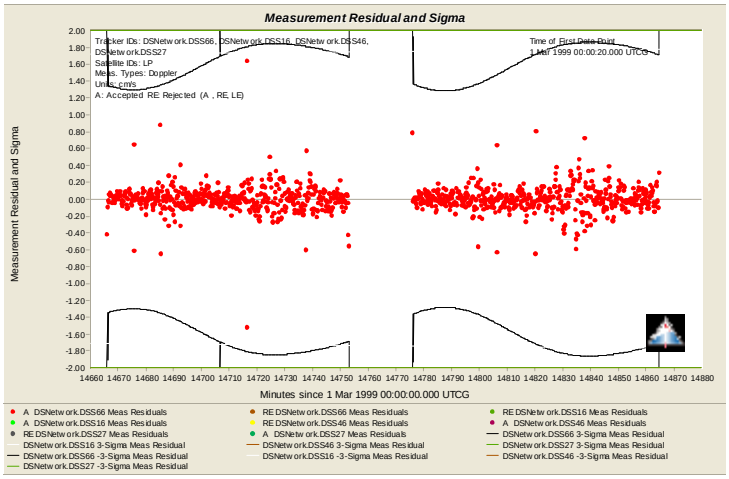

Figure 5c Doppler pass residuals – Earth direction perpendicular to LP orbit plane

Figure 6 shows the filter-smoother consistency statistic for the satellite position components expressed in the radial, in-track, and cross-track directions. The consistency metrics are very small over most of the month long tracking data interval. This is evidence that the filter and smoother position variances are too large, and is primarily due pessimistic gravity error process noise. There are short spans of time when the Earth-Moon line lies in the spacecraft orbit plane and the gravity process noise becomes optimistic. This is reflected in the filter-smoother consistency metrics as well. Inconsistency between the filter and smoother state on the 24th of March is due to a maneuver which is further analyzed in the sequel. There are a small number of additional spikes in the filter-smoother consistency metrics that are not fully understood, but it may not be possible to determine their cause without detailed information on spacecraft activities during the period in question.

Finally, we wish to discover the stability and expected accuracy of the orbit resulting from the filter-smoother process. During normal operations, the sequential filter process would be run to process observations and produce a predicted trajectory to be used for mission planning. Each filter run begins with the final state and covariance from the prior run. If the filter is stable, this process can continue for the life of the mission without the need to re-initialize the filter. The smoother would be run, using output from the filter, to produce a high precision, time-lagged ephemeris for mission analyses.

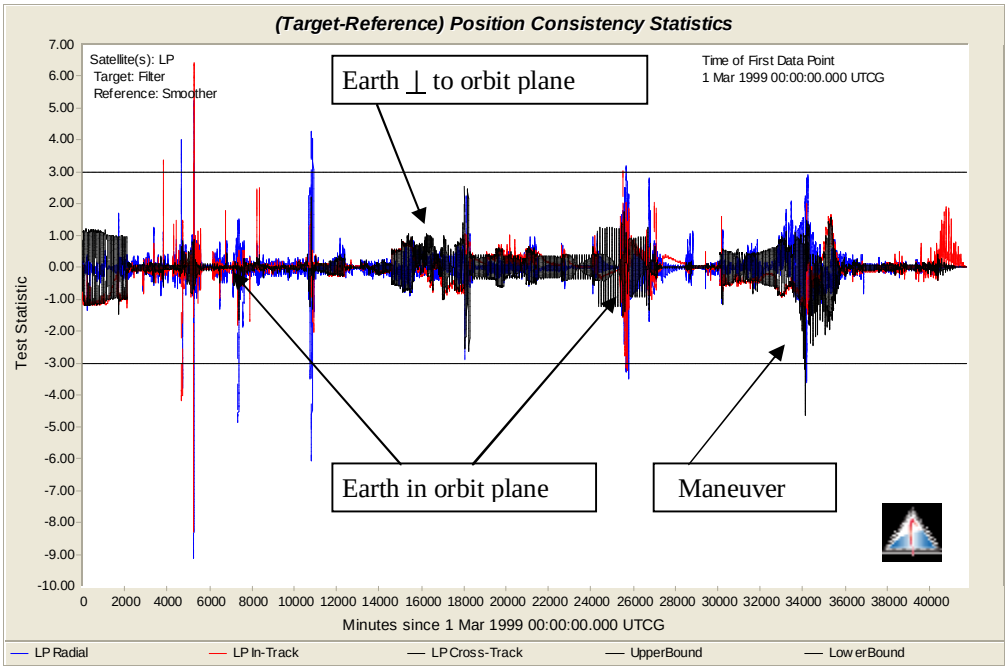

Figure 6 Filter-smoother consistency statistics for LP position

Each state estimate from the filter contains information about all tracking data taken prior to the current time in the filter. Smoother estimates contain information from all tracking taken prior to the current time in the smoother plus all tracking data after the current time in the filter up to the time of the smoother initial state. The smoother runs backwards in time. Since smoother estimates are conditioned by more information at each time that their corresponding filter estimates, the uncertainty in the smoother estimates are smaller than those from the filter.

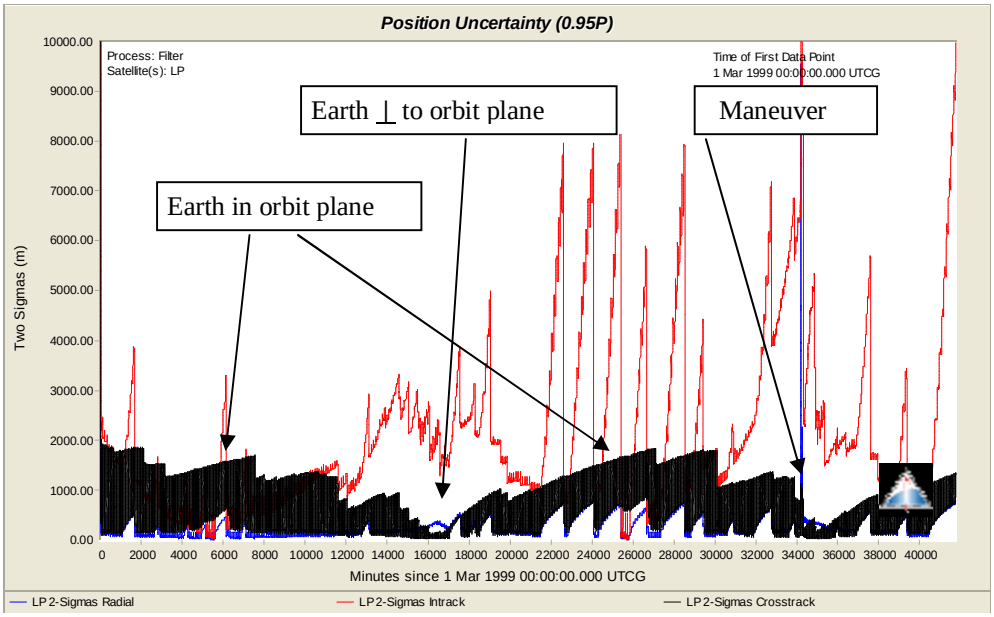

Figure 7a LP position uncertainty from filter

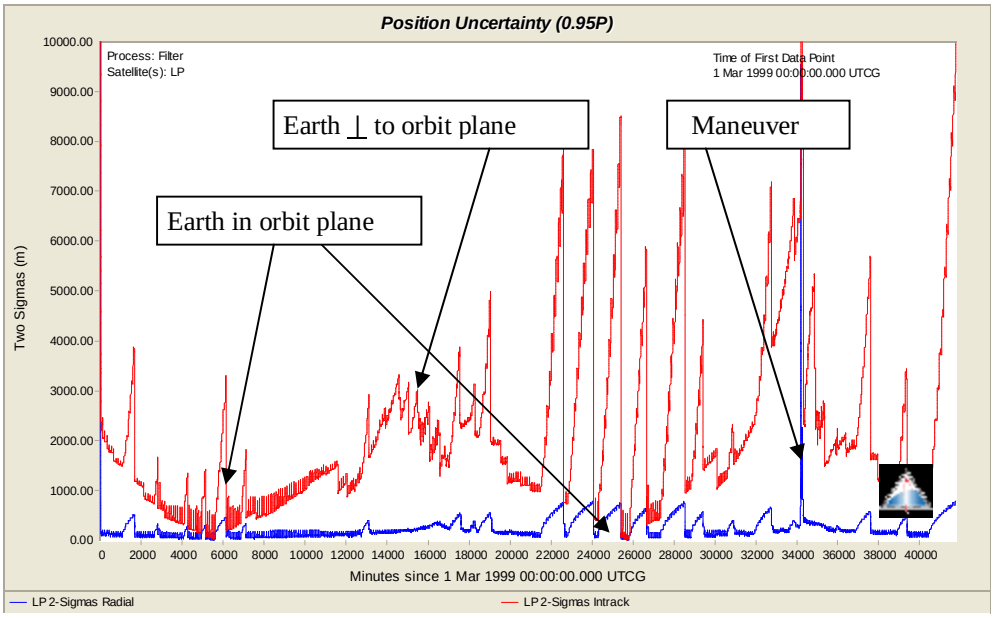

Figure 7b LP radial and in-track position uncertainty from filter

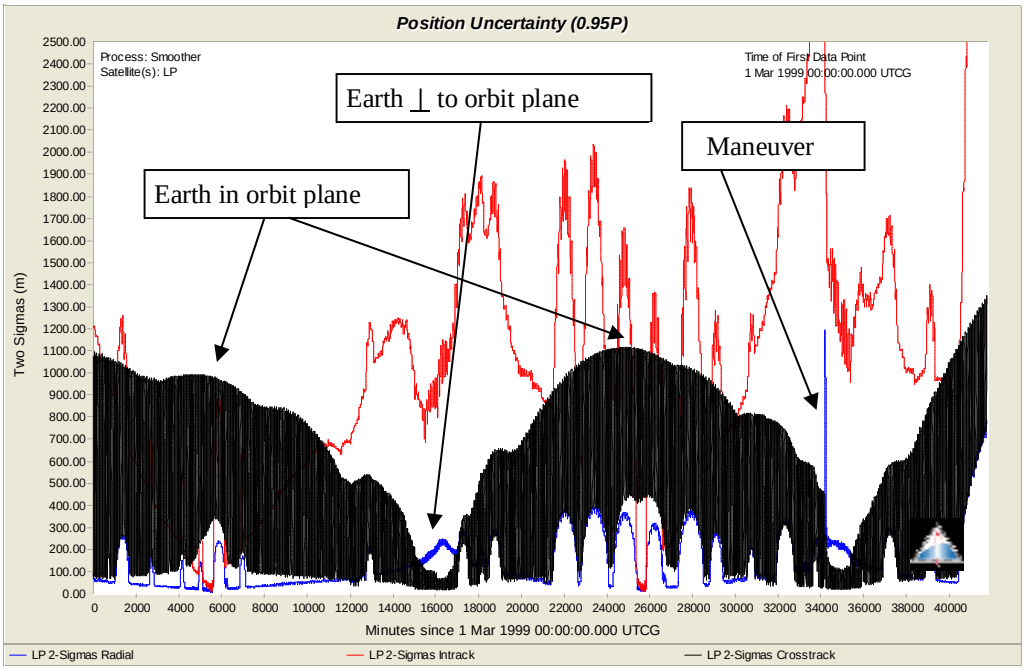

Figure 8a LP position uncertainty from smoother

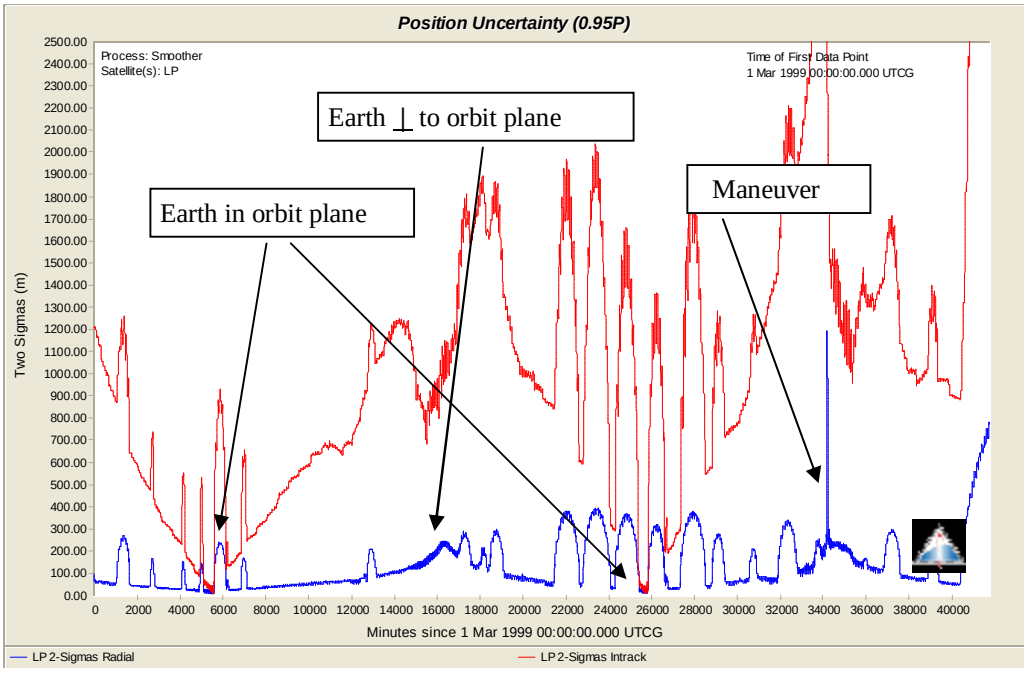

Figure 8b LP radial and in-track position uncertainty from smoother

The position uncertainty from the filter is shown in Figure 7a for the radial, in-track, and cross-track components. To provide better insight into the behavior of the radial and in-track uncertainties, the cross-track uncertainty is removed in Figure 7b. The important feature of these plots is that the uncertainty, while showing the effects of the tracking geometry, is neither growing nor shrinking in a secular fashion. We also did not experience any increase in the rate of tracking-data rejection by the filter over the duration of the run. The filter, therefore, appears to be stable without indicating the need for re-initialization. The position uncertainty from the smoother is shown in Figure 8a and Figure 8b. Note that the effects of the orbit-Earth geometry can be clearly seen in the gross behavior of the lines.

Lunar prospector trajectories are available from the Planetary Data System, making a comparison of results possible.[11] These trajectories originated from the Goddard Trajectory Determination System (GTDS). The trajectory differences for March 3 are shown in Figure 9.

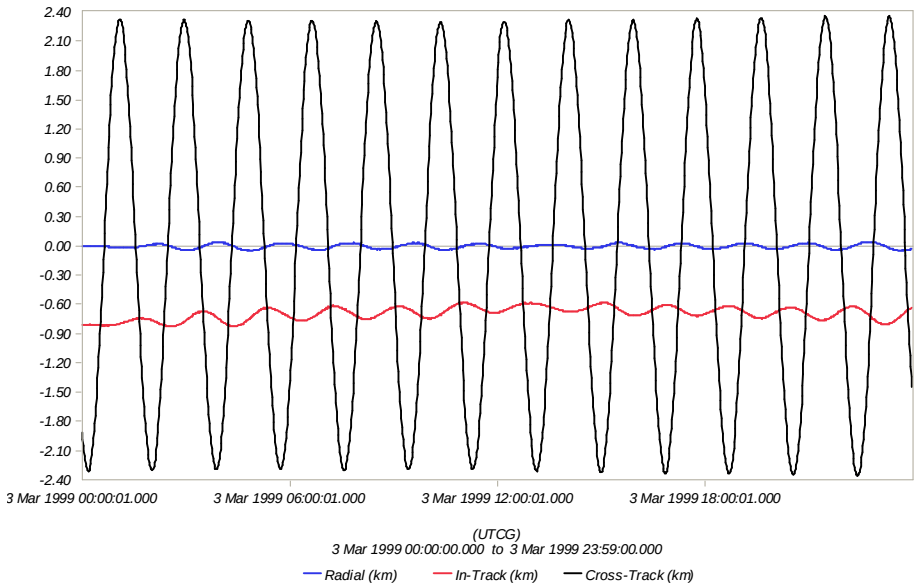

Figure 9 Trajectory difference between smoother and GTDS solutions

Maneuver

No maneuver information was known during the time span of the available tracking data. Initial processing of the data, however, indicated some type of modification to the trajectory during a significant part of one orbit period on 24 March 1999. The start time of the unknown activity, at approximately 17:30 UTC, was detected via the sudden and unexpected rejection of Doppler tracking data (Figure 10). The end time of the unknown activity was determined by adding a zero-magnitude finite maneuver of large uncertainty, and adjusting the end time of the maneuver until data were processed normally by the filter. Once the approximate end time of the activity was determined to be 18:25 UTC, the large uncertainty in the unknown finite maneuver was reduced in an iterative manner to determine the approximate smallest amount of process noise which was required to allow the filter to resume nominal processing at the end of the activity period. It was possible to stop at this point, and let the filter span the maneuver period courtesy of the process noise injected by the unknown maneuver model; however, we continued to investigate the event.

Further refinement of the unknown maneuver was accomplished by turning on maneuver estimation. The filter was able to lock onto an initial maneuver estimate, but began rejecting tracking data suddenly about three-quarters of the way through the maneuver span. It appeared that the net thrusting direction changed suddenly at this time from mainly radial to mainly cross-track. To enable the filter to continue processing data, we ended the first modeled maneuver at the apparent time of change of thrust direction, 18:11 UTC, and added a second maneuver of unknown nominal thrust, and having significant cross-track uncertainty, over the remainder of the maneuver time span. The result of our maneuver modeling selections is shown in terms of the Doppler residuals across the thrusting interval in Figure 11.

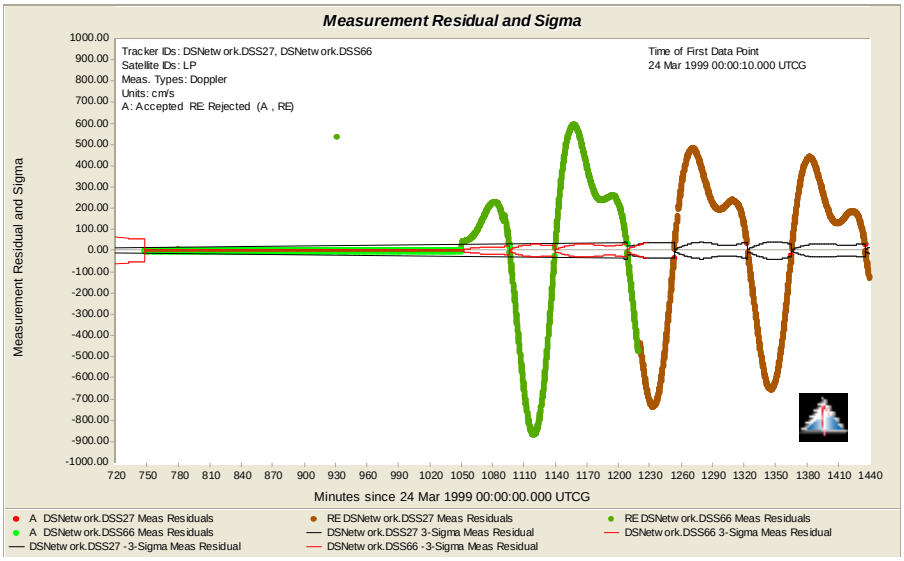

Figure 10 Rejection of tracking data indicates unknown satellite activity

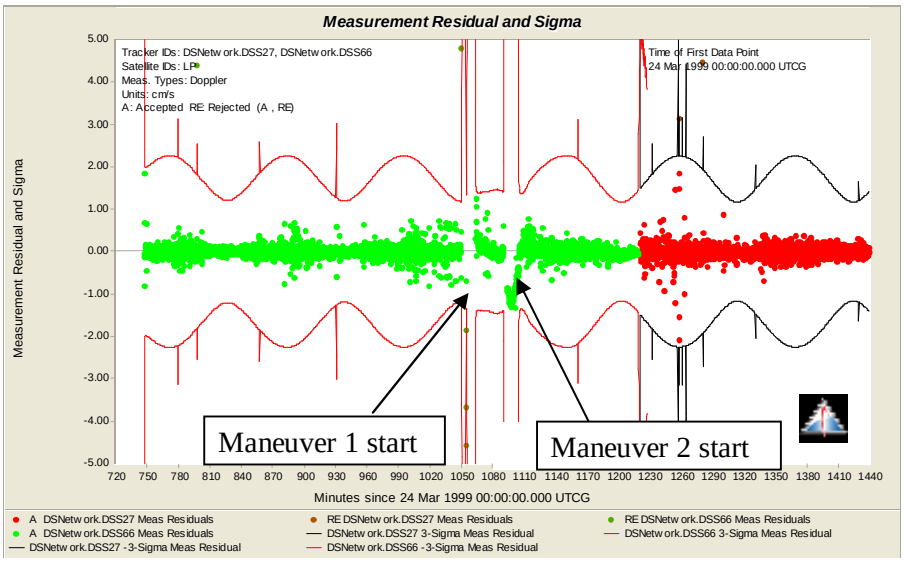

Figure 11 Residuals over satellite activity interval with unknown finite maneuver modeled

Table 5a Force modeling for STEREO, selenocentric

Lunar Gravity			
	Field	LP100K	8 x 8
	Solid Tides	Off	
3rd Body Gravity			
	Sun	Mercury	Venus
	Earth	Mars	Jupiter
	Saturn	Neptune	Uranus
Solar Pressure			
	Model	Spherical	Cp = 1.2
	Shadow	Dual Cone	Earth,Moon

Table 5b Force modeling for STEREO, geocentric

Earth Gravity			
	Field	GGM02C	8 x 8
	Solid Tides	Off	
3rd Body Gravity			
	Sun	Mercury	Venus
	Moon	Mars	Jupiter
	Saturn	Neptune	Uranus
Solar Pressure			
	Model	Spherical	Cp = 1.2
	Shadow	Dual Cone	Earth,Moon

Table 5c Force modeling for STEREO, heliocentric

Sun Gravity			
	Field	Point mass	
3rd Body Gravity			
	Moon	Mercury	Venus
	Earth	Mars	Jupiter
	Saturn	Neptune	Uranus
Solar Pressure			
	Model	Spherical	Cp = 1.2
	Shadow	Dual Cone	Earth,Moon

STEREO

DSN tracking data for the STEREO-B spacecraft were obtained for the lunar fly-by portion of the STEREO-B trajectory. The STEREO-A and STEREO-B spacecraft utilized lunar gravity assists to direct their trajectories out of the Earth-Moon system towards their heliocentric mission orbits ahead (STEREO-A) and behind (STEREO-B) the Earth. "Definitive" ephemerides were also provided to allow for post-fit trajectory comparisons; these ephemerides were in binary SPICE format and originated from the Goddard Trajectory Determination System (GTDS). Because the fly-by - the selenocentric phase of the trajectory - represented a transition from the geocentric phase to the heliocentric phase, we took the opportunity to use each of the three bodies as the central-body origin for spacecraft orbit determination, and then compared the estimation results. Tables 5a-5c provide the nominal force modeling selections used for each run. In ODTK, the specification of the primary central body for a satellite implies that it is to be the origin for trajectory integration. It is also the only celestial body for which the full gravitational field is evaluated during trajectory integration. It is also the only celestial body for which gravity process noise is available.

The preliminary estimation runs showed reasonable agreement for most estimates but did show a small difference in the transponder delay. An investigation into the source of the difference revealed two primary causes: the value of the gravitational parameter for the Moon differed depending on if the Moon was the primary body or not, and the number of lunar gravitational harmonics in the selenocentric case differed. To remedy the first cause, the source code was updated to allow the option to specify the gravitational parameter used for third-body perturbations. This option will be included in a future release of ODTK. To remedy the second cause, an input setting was changed and the pass of tracking data closest to the lunar periapsis passage was excluded. The estimation process was repeated and the estimates from all three cases were then nearly identical (Figure 12). The agreement between estimation results using different centers for trajectory integration provides a strong validation of the measurement model implementations across very different frames of reference. While simplifying the gravity modeling for the primary body allowed us to achieve consistency, for best accuracy we still recommend using a more complete gravitational potential for the Moon. Filter-smoother consistency test results are shown in Figure 13.

Observations were provided from DSN stations 25, 26, 45 and 54 during this time period with all tracking being done in X-band. Common measurement statistics were used for all stations and are provided in Table 6. One pass of tracking data from station 45 during the actual lunar flyby exhibited larger than expected residual scatter resulting in the rejection of a significant number of total-count-phase observations, see Figure 14. The cause of the difficulty in processing the total-count-phase measurements during the flyby has not yet been determined. We also note the brief violation of the filter-smoother consistency test during this time period in Figure 13.

Table 6 Measurement statistics

Measurement Type	Constant Bias	Bias Est.	Bias Sigma	White noise
Sequential Range	0.0 m	Yes	3 m	1.0 m
Total Count Phase	0.0 Cycles	No	N/A	0.003 Cycles
Transponder delay	425 m			

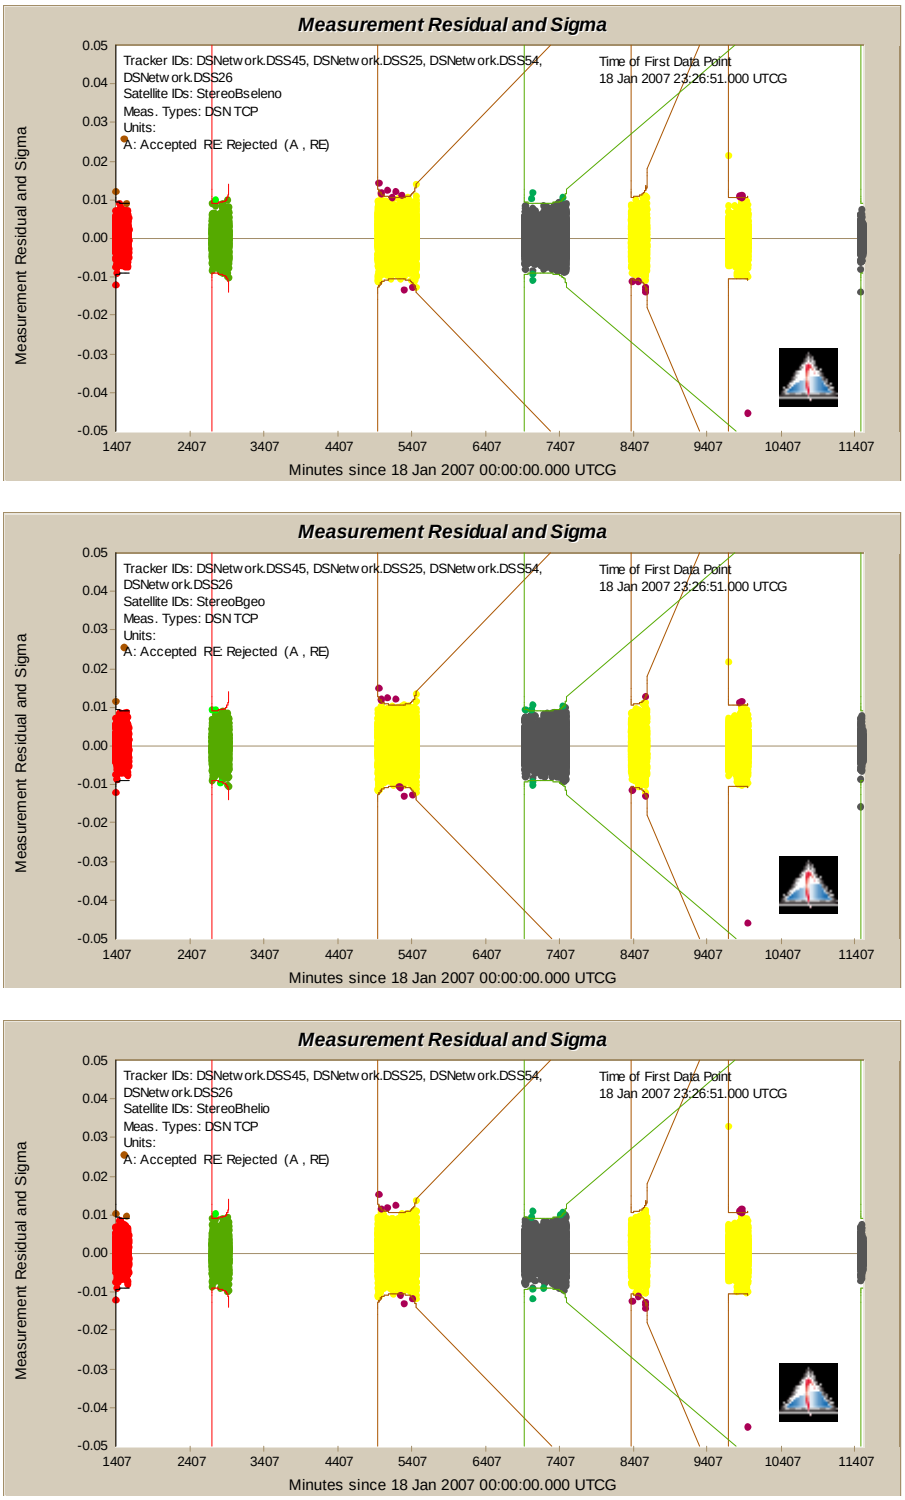

Figure 12 Total phase count residuals using different integration centers

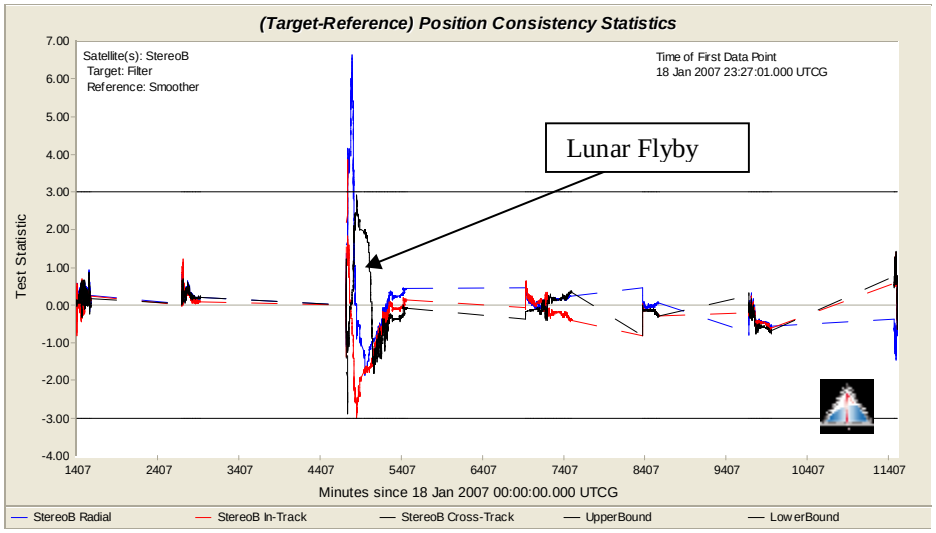

Figure 13 Filter-smoother consistency test results for STEREO

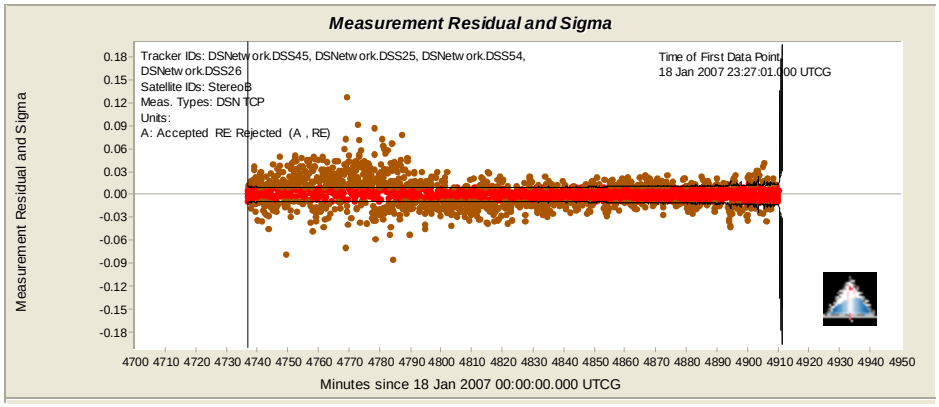

Figure 14 Residuals during tracking interval when STEREO was closest to the Moon

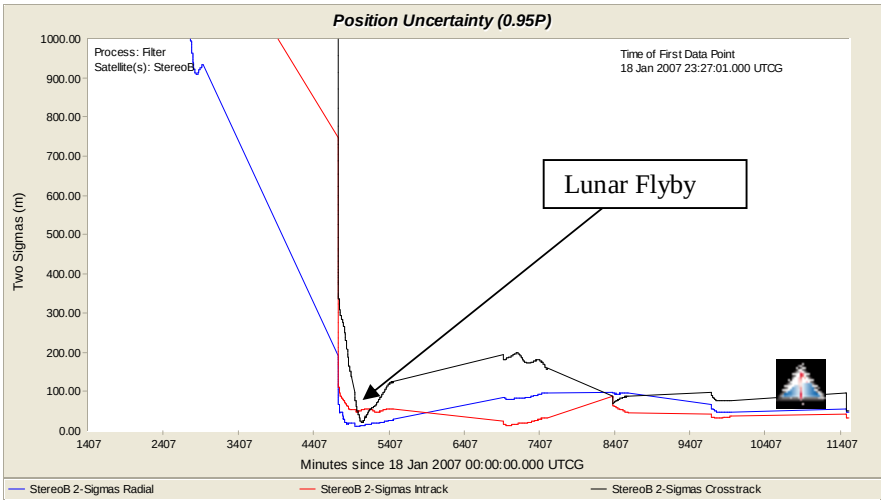

Figure 15 STEREO selenocentric position uncertainty from filter

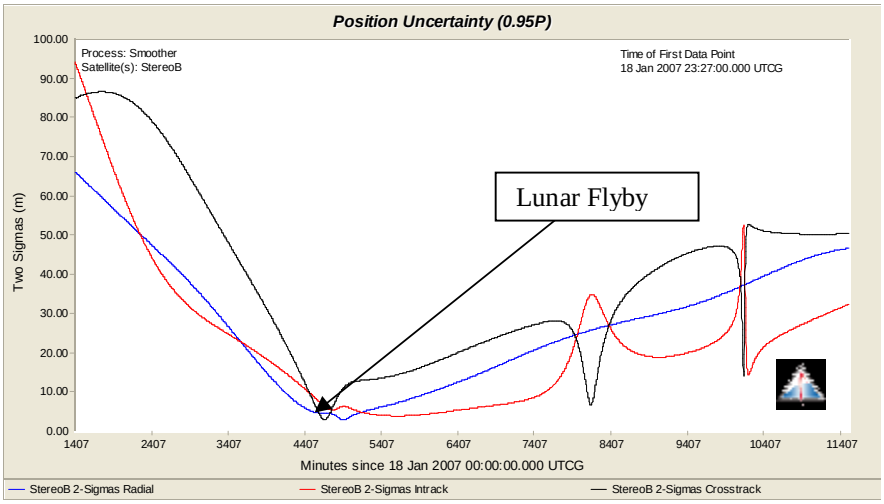

Figure 16 STEREO selenocentric position uncertainty from smoother

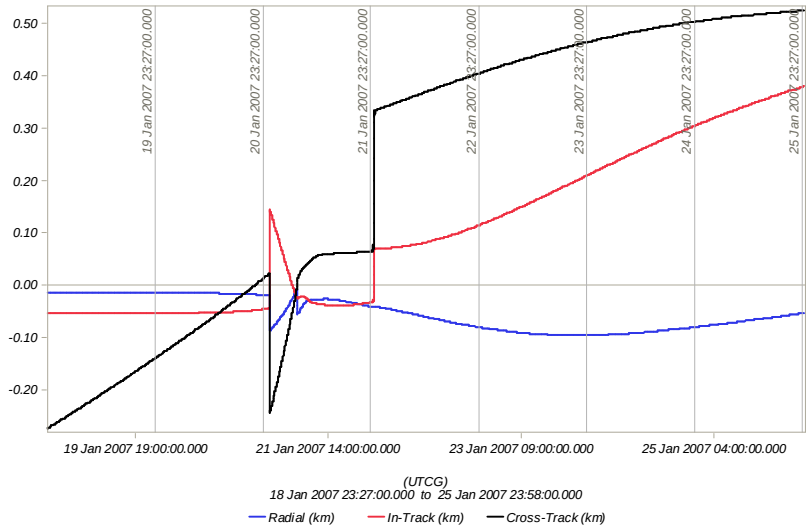

Figure 17 Trajectory difference – selenocentric smoother result vs GTDS result

STEREO position uncertainty from the filter in a Moon-centered frame is presented in Figure 15. The corresponding smoother result is given in Figure 16. We note increased observability of the orbit during the lunar fly-by. In fact, while the estimation process is fairly insensitive to the transponder delay for the majority of the trajectory, the ability to process data during the fly-by is greatly degraded by errors in the transponder delay. Orbit determination results from GTDS were provided with the STEREO tracking data making a comparison of results possible. Figure 17 shows the ephemeris differences between the selenocentric filter-smoother trajectory and the GTDS result.

CONCLUSION

We have demonstrated the ability to process approximately one month of tracking data on a low altitude lunar-orbiting satellite with a single run of the sequential filter. During this time period, the filter showed no signs of divergence or degradation. While no ad-hoc process noise

was required to achieve this goal, the existing gravity process-noise model is overly conservative most of the time and optimistic during short periods of time. The non-optimal nature of the gravity error covariance function for the lunar gravity field is mainly a result of the large variation in gravity field uncertainty over the lunar globe. Lunar gravitational potentials developed based on SELENE should be both more accurate and have more homogeneous uncertainty. This should result in more accurate orbit determination and more realistic orbit covariance.

We have demonstrated the capability to process DSN sequential range and total-count-phase measurements during a lunar flyby. Estimation was performed using Earth, Moon and Sun centered frames for trajectory integration with similar results.

ACKNOWLEDGMENT

The authors wish to thank Mark Beckman of NASA Goddard for providing tracking data. We also wish to acknowledge the invaluable assistance provided by Tomas Martin-Mur of JPL in helping us understand the DSN tracking data format and measurement models.

REFERENCES

[1] Hujsak, Richard S., Woodburn, James W., Seago, John H., "The Orbit Determination Tool Kit (ODTK) – Version 5," AAS 07-125, AAS/AIAA Space Flight Mechanics Meeting, Sedona, AZ, January 2007.

[2] Roncoli,Ralph B., "Lunar Constants and Models Document", JPL D-32296, Sept 2005.

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

[4] Wright, James R., Woodburn, James, Truong, Son, Chuba, William, "Sample Orbit Covariance Function and Filter-Smoother Consistency Tests," AAS 08-159, AAS/AIAA Space Flight Mechanics Meeting, Galveston, TX, January 2008.

[5] Wright, James R., Woodburn, James, Truong, Son, Chuba, William, "Orbit Covariance Inner Integrals with Polynomials," AAS 08-161, AAS/AIAA Space Flight Mechanics Meeting, Galveston, TX, January 2008.

[6] Moyer, Theodore D., "Formulation for Observed and Computed Values of Deep Space Network Data Types for Navigation", Monograph 2, Deep Space Communications and Navigation Series.

[7] Martin-Mur, Tomas (2007), Personal Communication.

[8] "TRK-2-34 DSMS Tracking System Data Archival Format", 820-013, Deep Space Mission System External Interface Specification JPL D-16765, Revision I.

[9] Konopliv, A.S., S.W. Asmar, E. Carranza, W.L. Sjogren, and D.N. Yuan (2001), "Recent Gravity Models as a Result of the Lunar Prospector Mission", *Icarus*, Volume 150, Issue 1, March 2001, Pages 1-18.

[10] McReynolds, S.R. (1984), "Editing Data Using Sequential Smoothing Techniques for Discrete Systems." Paper AIAA-1984-2053, Proceedings of the AIAA/AAS Astrodynamics Conference, Seattle, WA, Aug 20-22, 1984.

[11] http://pds.nasa.gov/