Mars Reconnaissance Orbiter Navigation During the Primary Science Phase[1]

Dolan Highsmith[2], Tung-Han You[3], Stuart Demcak[4], Eric Graat[§], Earl Higa[§], Stacia Long[§], Ram Bhat[§], Neil Mottinger[§], Allen Halsell[5], and Fernando Peralta[§]
Jet Propulsion Laboratory / California Institute of Technology, Pasadena, CA, 91109

The Mars Reconnaissance Orbiter began science operations in November 2006, with a suite of seven instruments and investigations, some of which required navigation accuracies much better than previous Mars missions. This paper describes the driving performance requirements levied on Navigation and how well those requirements have been met thus far. Trending analyses that have a direct impact on the Navigation performance, such as atmospheric bias determination, are covered in detail, as well as dynamic models, estimation strategy, tracking data reduction techniques, and residual noise.

Nomenclature

a	=	semi-major axis
e	=	eccentricity
i	=	inclination
Ω	=	longitude of the ascending node
ω	=	argument of periapsis
M	=	mean anomaly
L_S	=	longitude of Mars around the Sun, an indicator of season (L_S = 0 deg is northern spring)
a_D	=	drag acceleration
S	=	density/drag scale factor
ρ	=	density
V	=	spacecraft velocity
C_D	=	drag coefficient
A	=	drag area
M	=	spacecraft mass

I. Introduction

THE Mars Reconnaissance Orbiter (MRO) launched August 12, 2005, headed for a low-altitude, sun-synchronous, frozen science orbit around the Red Planet. After a seven-month cruise, six-months of aerobraking, and a transition phase that spanned solar conjunction, MRO began its Primary Science Phase (PSP) on November 8, 2006, with its set of seven instruments and investigations ready to study Mars in greater detail than ever before. Henceforth, all the planning and preparation for science operations was to be put into action. Due to the stringent requirements levied by the instruments, Navigation performance was a key to the success of MRO science operations, in particular how well the actual performance matched the pre-launch analysis. Now, after twenty months and 8000 orbits in the PSP, the pre-launch assumptions can be evaluated against the achieved results, providing a gauge of the accuracy and effectiveness of the assumptions with corresponding recommendations for future missions.

---

[1] This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California under contract to National Aeronautics and Space Administration.
[2] Mars Reconnaissance Orbiter Navigation Team Chief (Acting), Jet Propulsion Laboratory.
[3] Former Mars Reconnaissance Orbiter Navigation Team Chief, Jet Propulsion Laboratory.
[4] Senior Member of Engineering Staff, Jet Propulsion Laboratory.
[5] Flight Path and Orbit Control Group Supervisor, Jet Propulsion Laboratory.

II. Mission Description

The MRO science objectives are a key component of the NASA Mars Program objective to "Follow the Water", a strategy that combines the unique perspectives of orbiters, landers, and rovers to identify and confirm the presence of water—a requirement for the existence of life.[1] To that end, the MRO science payload includes a high-resolution imager, context imager, hyperspectral spectrometer, climate sounder, weather camera, ground-penetrating radar, and gravity science investigation.[2] Each experiment has its own Navigation requirements, with the driving requirements being considerably more stringent than for previous Mars orbiters. Moreover, to increase the resolution of these experiments, MRO occupies a much lower-altitude orbit than its predecessors. Figure 1 shows the altitude versus latitude profile for MRO, Mars Global Surveyor (MGS), and Mars Odyssey. Given the approximate 120-km altitude difference, MRO is in a new orbit regime for a near-circular science orbit at Mars; for the first time outside of aerobraking, drag becomes a significant perturbation considering the highly variable atmosphere. Thus, the challenge facing MRO Navigation is to provide significantly increased accuracy flying in a much more *un*predictable orbit.

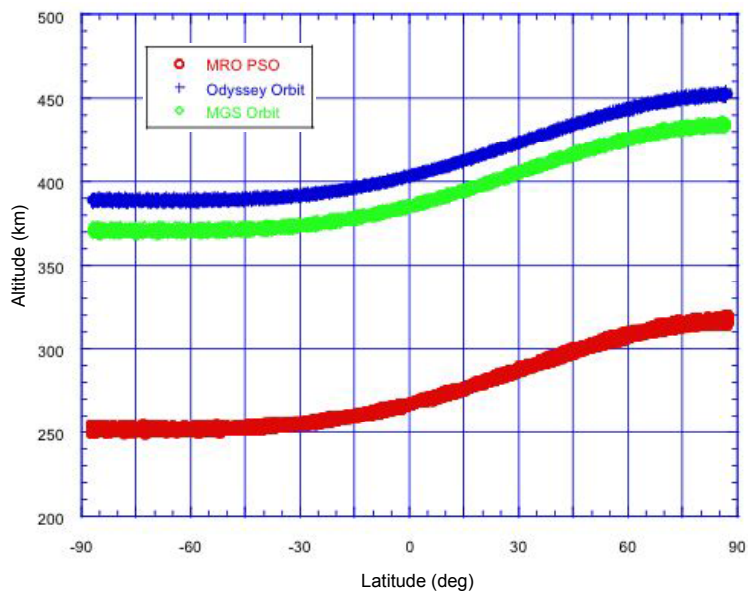

Figure 1: Altitude comparison of MRO, Odyssey, and MGS.

A. Science Orbit

Specifically, the MRO science orbit is a 252 x 317 km altitude, sun-synchronous orbit with periapsis frozen over the south pole and the ascending node at 3 pm. Mean orbital elements are given in Table 1. The ideal orbit is designed to exactly repeat after 4602 orbits in 349 sols (1 sol = 1.0275 days), providing sub-5 km coverage at the equator.[3] There are several shorter near-repeat cycles that are more useful for science planning, most notably a 211-orbit cycle that walks 0.5-deg in longitude (32.5 km on the ground) west of the previous cycle every 16 sols. This 211-orbit characteristic is the basis for the orbit maintenance strategy summarized in Section IV.E.

Table 1. MRO mean orbital elements.

Epoch: 05-May-2007 00:11:56.08 ET	
a (km)	3657.499
e	0.0088347
i (deg)	92.649
Ω (deg)	100.764
ω (deg)	270.368
M (deg)	-143.005

B. Spacecraft

The MRO spacecraft is comprised of a bus, two 12-m^2 solar panels, and a 3-m diameter high gain antenna (HGA). The solar panels and HGA are each pointed with two-axis gimbals. Figure 2 shows depictions of MRO in action: the bus is nadir-pointed, the HGA is tracking the Earth, and the solar panels are tracking the Sun. The top row in Figure 2 shows the view of the spacecraft from the drag or anti-velocity direction, and the bottom row shows the nadir view. Since the MRO orbit is sun-synchronous, the pattern of the solar panel gimbal motions is generally the same from orbit to orbit; however, the HGA per-orbit motion changes slowly as the Earth moves relative to the MRO orbit plane. Since drag is a significant perturbation, Navigation must accurately model these gimbal motions and the corresponding changes in effective drag area.

For attitude control, the spacecraft is equipped with 100 N-m-s reaction wheels and balanced, mono-propellant thrusters. Momentum desaturation events (desats) are spaced two days apart, with occasional low-torque periods with three-day spacing. However, even with balanced thrusters and relatively infrequent desats, the residual delta-v

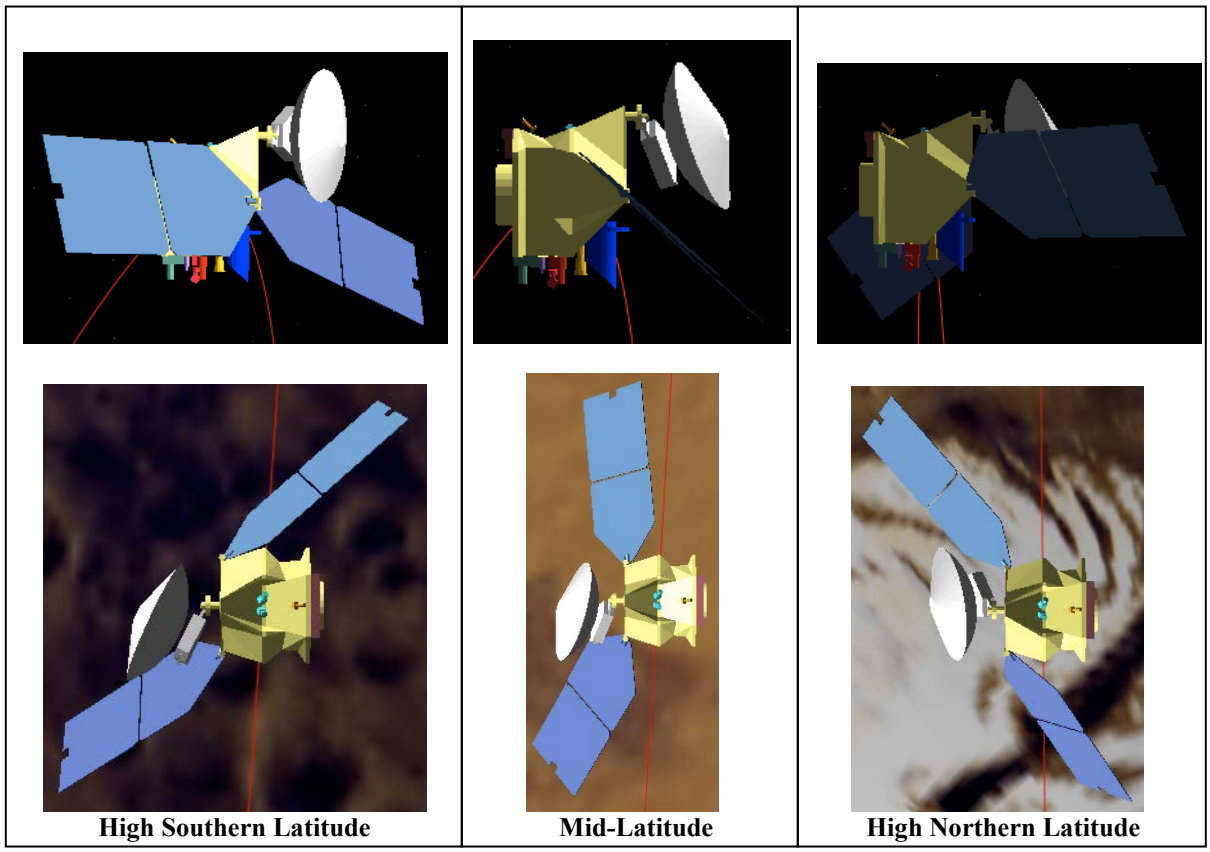

Figure 2. Depictions of MRO in various locations around an orbit on May 26, 2008. The top row shows the view along the drag/anti-velocity direction with Mars being down; the bottom row shows the view along nadir with the velocity direction being up. The solar panels are tracking the Sun, the high gain antenna is tracking Earth, and the bus is nadir pointed.

imparted due to misalignments and inconsistent thruster performance can be a driving error source for ephemeris prediction accuracy.

MRO is a very capable targeting spacecraft, with an on-board ephemeris and the ability to slew up to 30 deg off-nadir. The on-board ephemeris corrects for the most accurate target pointing and overflight time by referring to the latest predicted trajectory from Navigation. The off-nadir slew occurs in the crosstrack direction to view targets not directly under the groundtrack. Combined with the near-term groundtrack repeat patterns, the off-nadir capability gives MRO multiple opportunities to image the same site within several weeks. The implication for Navigation is that the predicted ephemerides must satisfy a long-term requirement for initial target selection in addition to a short-term requirement for the accuracy of the on-board ephemeris.

C. Driving Requirements

Among other requirements, there are several drivers that define the weekly operational scenario for the science phase:

1) *Long-term prediction*: Off-nadir target angles must not change by more than 3 deg (3σ) twenty-eight days after OD data cutoff.
2) *Short-term prediction*: Predict the Mars-relative position to within 1.5 km (3σ) downtrack.
3) *Reconstruction*: Reconstruct the Mars-relative position to within 100 m (3σ) downtrack, 40 m (3σ) crosstrack, and 1.5 m (3σ) radial.
4) *Groundtrack Walk (GTW)*: Maintain the 211-orbit near-repeat cycle near the ideal westward walk of 32.5 km (0.5 deg in longitude).
5) *EDL Relay Phasing*: Perform maneuvers to phase the spacecraft in true anomaly to provide UHF communications to Phoenix during its entry, descent, and landing (EDL) phase. Specifically, cross a given latitude on an ascending track within 30 sec of the desired time.

The first three requirements listed above are paraphrased from the MRO Project Systems Requirements Document and Science Requirements Document, whereas the last two are de-facto requirements based upon operational desires. Covariance analyses are used to determine the degree to which these particular requirements could be met and the operational scenarios to address them.

Given a nominal trajectory, the covariance analysis is performed by simulating tracking observations at their expected frequency, computing partial derivatives of relevant quantities, and filtering the tracking and partials to evaluate the resulting formal uncertainties. During the development phase, the navigation performance of each mission phase is analyzed and compared to the relevant requirements, and the results are combined into the Navigation Plan released at the time of the Critical Design Review. Details of the MRO covariance analysis results and implications to the Navigation operations plan can be found in Ref. 4. The following sections will describe the science phase operational navigation and compare flight performance to the requirements and covariance analysis.

III. Operational Navigation

In operations, the Navigation process compares the observed two-way Doppler from Deep Space Network (DSN) tracking to computed Doppler values based on a nominal trajectory model, then generates *observed – computed* residuals. The residuals are then minimized with a batch least-squares filter by estimating changes to the spacecraft state and other quantities necessary to align the computed and observed Doppler. The key parts of this process are the accuracy and fidelity of the models used to generate the trajectory and the choice of estimated states and their a priori uncertainties used by the filter.

A. Models

Accurate trajectory modeling of the MRO science orbit relies on high-fidelity models of the atmospheric density, desat delta-v, Mars spherical harmonics, and MRO spacecraft attitude.

1. Atmospheric Density

By far the driver in ephemeris prediction uncertainty, it is important to have a density model that reflects reality in a mean sense over the span of weeks and months. Divergence of the actual from the predicted atmosphere into biased behavior quickly reduces the prediction accuracy due to the exponential behavior of the resultant accumulating timing offset.

Navigation currently uses the Mars Global Reference Atmosphere Model 2000 Version[5], with modifications made specifically for MRO science orbit altitudes (MarsGRAM 2000 MRO Special Edition). Above 170 km, this version uses a modified Stewart thermospheric model, including diurnal and seasonal variations. Additional MarsGRAM models include 2001 MRO Special Edition (SE) and 2005[6,7]. MarsGRAM 2001 MRO SE was used in the pre-launch covariance analysis; MarsGRAM 2005 was used during MRO aerobraking and the first few months of the science phase. In February 2007, Navigation switched to MarsGRAM 2000 after it became apparent that its underlying seasonal model was a much better predictor than the MarsGRAM 2005 Map Year 1 model used for aerobraking. Comparisons are provided in Section V.

2. Angular Momentum Desaturations

After atmospheric density, the unbalanced portion of the desats causes the next largest trajectory perturbation. In fact, during the low-drag season, desats can be the driving error source. For trajectory reconstruction, Navigation directly queries the spacecraft telemetry for the thruster pulse information to include in the trajectory integration. For prediction, the Attitude Control System (ACS) team at Lockheed Martin models the momentum and provides a file of predicted desat thruster pulses for the upcoming two to three weeks. After that period, a constant acceleration model is used based on the average delta-v per two days from the latest ACS prediction.

3. Mars Oblateness

Because the MRO science orbit is approximately 130 km below MGS and Odyssey, it is sensitive to some different spherical harmonics – terms that are not well-defined based on the MGS and Odyssey gravity field solutions. Through the first few months of the PSP, Navigation used MGS95J, an 95th degree and order field based on MGS and Odyssey tracking data through December 6, 2004[8]. Periodically throughout the science phase, Navigation received updated, preliminary solutions from the Gravity Investigation team that included MRO tracking; the version currently in use is the final MRO95A[6], which shows 10-50% improvement in the MRO residual RMS over MGS95J, depending on the specific fit span.

[6] Submitted to the Planetary Data System Geosciences Node.

4. Spacecraft and Appendage Attitude

Figure 2 shows the challenge for modeling the independent attitudes of the bus, HGA, and solar panels. First, the trajectory inputs include models for the shape and size of the components: two double-sided flat plates for the solar panels, a parabolic dish for the HGA, and three double-sided flat plates to represent the faces of the bus. Each component can be pointed individually, either using geometric relationships (e.g., Sun, Earth, nadir) or by directly inputting predicted attitude quaternions delivered by ACS.

For trajectory integration, the geometric relationships are sufficient, where the driver is the computation of the changing drag area around the orbit. This is not precise because of the variable component shadowing from the drag flow (top row of Fig. 2), but the difference simply manifests in the per-orbit drag scale factor estimates described in the next sections. The attitude quaternions are used along with the locations of the spacecraft center of mass and HGA gimbals to model the phase center offset of the HGA for Doppler observation modeling. Since the trajectory integration is based at the center of mass, but the Doppler is produced from the phase center of the HGA, the position difference can induce a signature in the Doppler, particularly when MRO slews to an off-nadir target. Modeling the HGA phase center significantly improves the residual noise.

B. Filter Strategy

After integrating the trajectory with all the necessary models, residuals are computed and filter iterations are performed to converge the state estimates to a trajectory that not only matches the observed tracking, but can be used to accurately predict forward. Key aspects of the filter strategy include weighting of the Doppler observables and filter setup, including choice of state parameters and a priori uncertainties.

1. Data Weighting

In orbit around Mars, Doppler is the only measurement required because of its strong signature from the orbit being tied to the planet. For the covariance analysis, the Doppler weight was set to the conventional 0.1 mm/sec at a count time (T_C) of 60 sec. In operations, Navigation uses a tool that automatically computes a data weight for each arc of Doppler, with an arc generally defined to be the span between Earth occultations. This technique is referred to as *weight by pass*. The automatic weight is computed to be the residual RMS of the arc scaled up by a factor that depends on the data type: for two-way Doppler: 1.0; for three-way Doppler: 1.2; and for one-way Doppler: 1.5. As an example, for one-way and two-way Doppler passes with the same residual RMS, the one-way is deweighted by 50% relative to two-way.

In addition, there is another scale factor applied to all the Doppler weights to effectively provide the "whitened" data noise equivalent. The theory, described in Ref. 9, is based upon the fundamental frequency exhibited by the Doppler signature. In the case of interplanetary cruise, that frequency is 24 hours because of the motion of the DSN ground station; however, in the case of MRO, the frequency is the orbit period of 112 min because of the large Doppler variation as MRO completes a circuit around Mars. The relevant equation is $\sigma_W = 0.470\left(t_D / T_C\right)^{1/3}$, where t_D is the time scale of interest and T_C is the Doppler count time[9]. MRO Doppler has a count time of 10 sec, so $\sigma_W\,(T_C{=}10) = 4.12$.

2. Filter Setup

Table 2 shows the estimated parameters and their corresponding 1σ a priori uncertainties. The parameters are limited to those quantities that drive the dynamics of the orbit or, in the case of the one-way Doppler terms, correct for the offset of the on-board frequency reference relative to the DSN hydrogen masers (the frequency reference for two- and three-way Doppler). Estimating per-rev drag scale factors is necessary due to the large orbit-to-orbit variability in the density; estimating a constant factor over multiple orbits leaves significant signatures in the Doppler. These per-rev estimates provide invaluable insight into the current atmospheric trend, which is then used to formulate the prediction model.

The desat delta-v components are estimated as stochastic white noise biases, with scale factor estimates for each spacecraft axis and stochastic batches for each wheel being desaturated. To desaturate a given wheel (X, Y, or Z), the balanced thruster system fires two pairs of thrusters. By solving for a scale factor for each set of thruster pairs, the performance of each pair can be resolved and trended to produce a de-facto thruster calibration. Unfortunately, the cruise calibration results[10] did not persist after Mars orbit insertion, so updated results had to be generated from the estimates of routine desats performed as part of nominal science operations.

Since the gravity model has been tuned with MRO95A, only the near-resonant zonals of degree 12 and 13 are included in the nominal filter. Note that the MRO orbit is near resonant with the degree and order 13 terms because there are approximately 13.2 orbits per sol. The a priori is set to 10 times the formal uncertainties obtained in the

Table 2. Estimated filter parameters and a priori uncertainties.

Estimated Parameters	Estimate Type	*A priori* uncertainty (1σ)	Remark
Position	-	100 km	
Velocity	-	10 m/s	
Angular Momentum Desaturation Scale Factors	Stochastic	0.5	Est. spacecraft X,Y,Z scale factors separately. White noise batches for each wheel.
Gravity: J12, J13	Bias	10x MRO95A formal uncertainties	Near-resonant terms.
Orbit Trim Maneuver	Bias	➢ < 10 cm/s: 1.67% of magnitude. ➢ > 10 cm/s, < 50 cm/s: 1.67 mm/s fixed plus 0.67% proportional. ➢ > 50 cm/s: 6.7 mm/s fixed plus 0.67% proportional	Occurs no more frequently than every 28 days.
Solar Pressure Scale Factor	Bias	10%	
Drag Scale Factors	Bias	25% per rev	Apoapsis to apoapsis.
One-Way Doppler Bias	Bias	10 Hz	On-board frequency reference initial bias.
One-Way Doppler Rate	Bias	1e-6 Hz/sec	On-board frequency reference drift rate.

MRO95A fit—a reasonable scaling given the optimistic values at the lower degrees[7]. Gravity terms are still included in the filter because small signatures remain that appear to be due to gravity; this is evidenced by the fact that the estimated changes to the nominal values have decreased with the successive deliveries of interim MRO95A solutions as more and more MRO tracking was added.

The one-way Doppler bias and rate terms are required due to the fact that the frequency reference on board MRO is generated by an oven-controlled crystal oscillator with an approximate stability of 10^{-12} over the Doppler count time. This is good enough for Navigation purposes, as was expected based on the analysis in Ref. 11. The key, however, is to have two-way Doppler surrounding the one-way in order to resolve the bias and rate terms. The one-way Doppler enables Navigation to fill-in gaps between two-way passes, resulting in drag scale factor estimates on orbits that otherwise would have had no observability.

IV. Performance vs. Requirements

A. Post-Fit Residual Noise

Before evaluating any results of the estimated quantities, it is important to examine the post-fit Doppler noise as an indication of how well the filter was able to match the observations with the given set of estimated parameters. Navigation reconstructs the MRO trajectory in 20-24 orbit batches (1.5-2 days), then merges four to five batches together to form our weekly reconstruction delivery. In general, the passes of 10-second Doppler points are fit down to the noise. Figure 3 shows the standard deviation of the combined post-fit residuals for each merged reconstruction batch, such that there is one point per week. The noise for all three Doppler types is shown in the units of mm/sec at a count time of 60 sec. The plotted data start at the beginning of the science phase as MRO emerged from solar conjunction. Clearly, the two- and three-way Doppler noise values are comparable, with the one-way noise exhibiting the deleterious effect of the less-stable frequency from the on-board oscillator. Note also the reduction in noise as Mars passed through opposition in December 2007.

[7] Alex Konopliv, MRO Gravity Science Team, personal communication.

With the weight-by-pass technique, the equivalent data weight is obtained by scaling all the values by 1.68 ($= 4.12/\sqrt{6}$ for $T_C = 60$ sec), then scaling the different data types as discussed in Section III.B.1. Doing so with the two-way Doppler produces equivalent data weights in the range of 0.034 to 0.067 mm/sec. This compares to the 0.10 mm/sec data weight used in the covariance anaysis, indicating that the two-way Doppler data noise is better than expected. It also indicates that, except for occasional minor adjustments, the filter strategy adopted for the science phase appropriately determines the key factors affecting the trajectory.

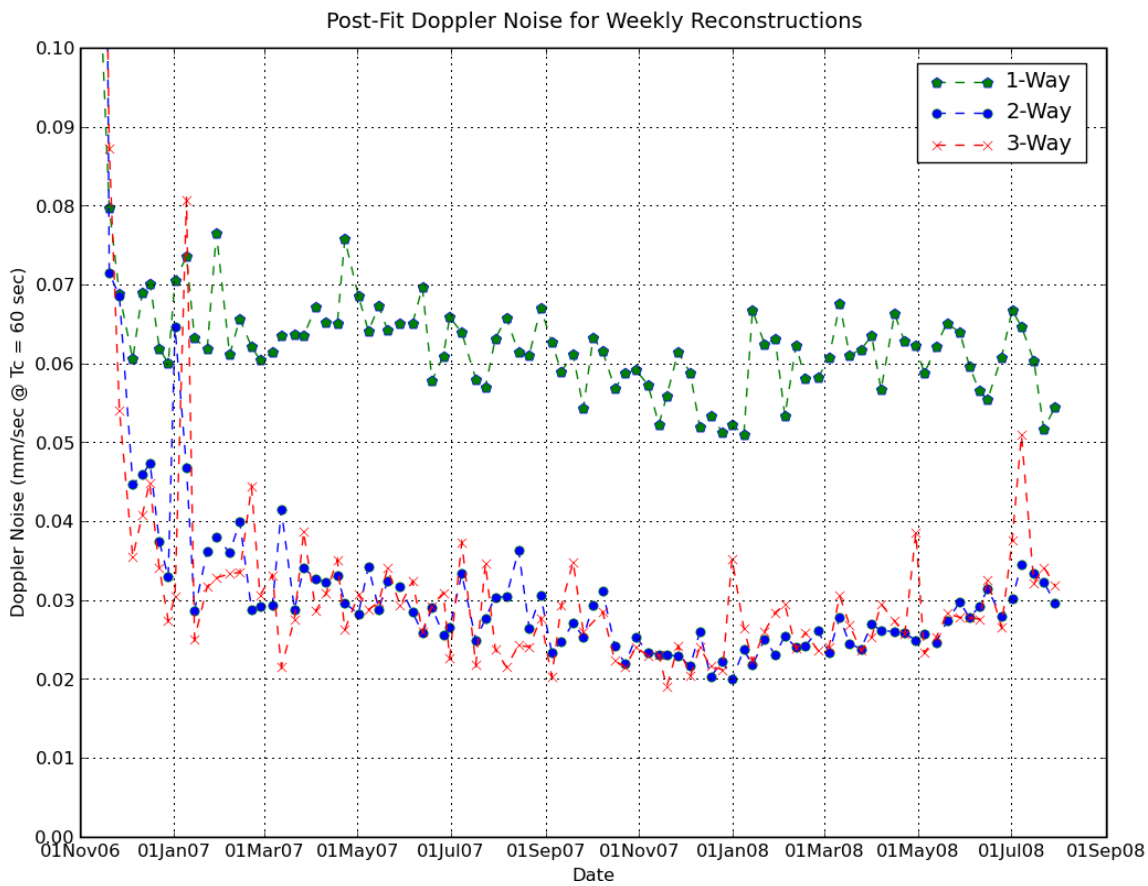

Figure 3. Post-fit Doppler noise for weekly reconstruction batches. One point represents the average noise for the specified data type over the span of one week. The data begin as MRO emerged from solar conjunction at the start of the PSP.

B. Long-term Prediction

The long-term requirement listed in Section II.C states that a target on the ground must not drift by more than three degrees (3σ) off-nadir (or crosstrack) in 28 days after data cutoff (DCO). This is to allow the scientists to choose off-nadir targets during the four-week science planning period with confidence that the targets will not rotate outside of the spacecraft roll capability of 30 degrees. In other words, it limits the margin that must be carried in the process of selecting targets. In the pre-launch covariance analysis, Navigation determined that the stressing case for this requirement is the downtrack timing error caused by drag variability. The timing error maps to ground target pointing in the crosstrack direction through MRO altitude and the rotation rate of Mars. Details are given in Ref. 4, with the result being that approximately 60 seconds of downtrack timing error would cause a three-degree off-nadir pointing error at the equator. Based on the covariance, the high drag season was the stressing case, during which the ephemeris could last only 24 days at the 3σ level. As a consequence, the science teams adopted a 24-day planning

cycle for choosing targets. For reference, the low-drag season reached the requirement level after 47 days, so the expected performance varied by a factor of two depending on season.

As MRO approaches its first full Mars year in orbit, the sample size of timing errors starts to become large enough to determine meaningful performance statistics for comparison to expectations. Navigation has delivered over 200 predicted ephemerides and reconstructed 8000 orbits since early November 2006. Figures 4 and 5 show

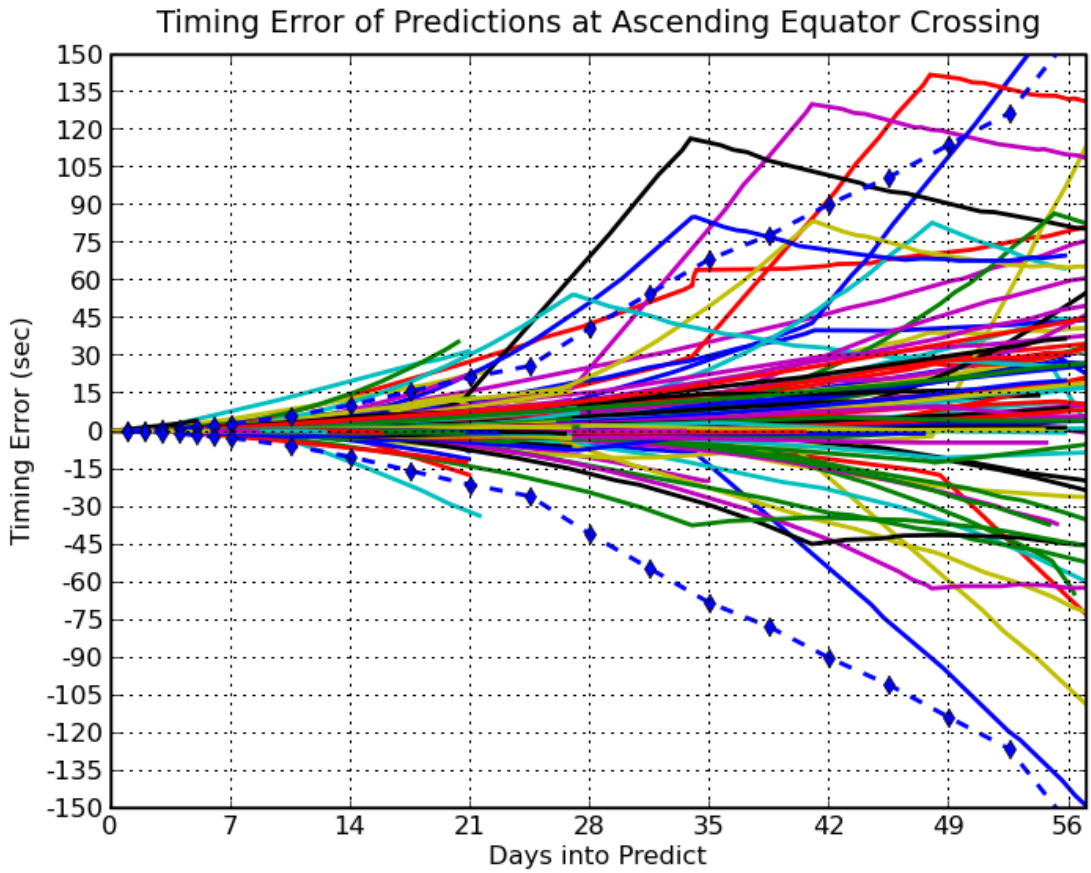

Figure 4. Long-term predicted ephemeris timing errors. Dashed line shows three times the standard deviation of the samples.

different perspectives of the timing error of the predicted ephemerides. The plot in Fig. 4 shows the timing errors relative to days into the prediction, with the dashed line indicating three times the standard deviation of the samples at 3.5-day intervals through the prediction. This statistic gives an indication of the 3σ noise across the entire times span, assuming a Gaussian process. The dashed line crosses 60 sec about 32 days into the prediction. Since the sample set includes results from all seasons, this statistic should represent more of an average performance. Given the covariance results that varied from 24-47 days, 32 days falls well within the range of expectations, if slightly worse than a pure average.

Taking samples from specific seasons does not provide a large data set since Navigation delivers only one long-term prediction per week. However, Fig. 5 shows the total timing error with respect to date, which more clearly shows correlations to certain events and seasons. Many of the large excursions are due to maneuver execution errors, as evidenced by the linear nature of the error growth. An example that stands out is the error caused by the delay of OTM-2 by two weeks in April 2007; those predictions that included OTM-2 on the original date show fast error growth for two weeks, but then turn around when the actual maneuver is performed. By contrast, drag bias errors accumulate on every orbit, so they exhibit a curved signature in the timing errors. An example of this behavior is the nearly global dust storm that occurred in July 2007; Navigation generally underpredicted the drag as the storm grew (positive timing errors), but overpredicted as the storm waned (negative timing errors). This storm also occurred right at the peak of the high-drag season – southern summer at L_S = 180 deg. Clearly, there is increased noise during southern summer around July 2007, and less noise during southern winter (L_S = 90 deg) around June 2008, as expected from the covariance analysis.

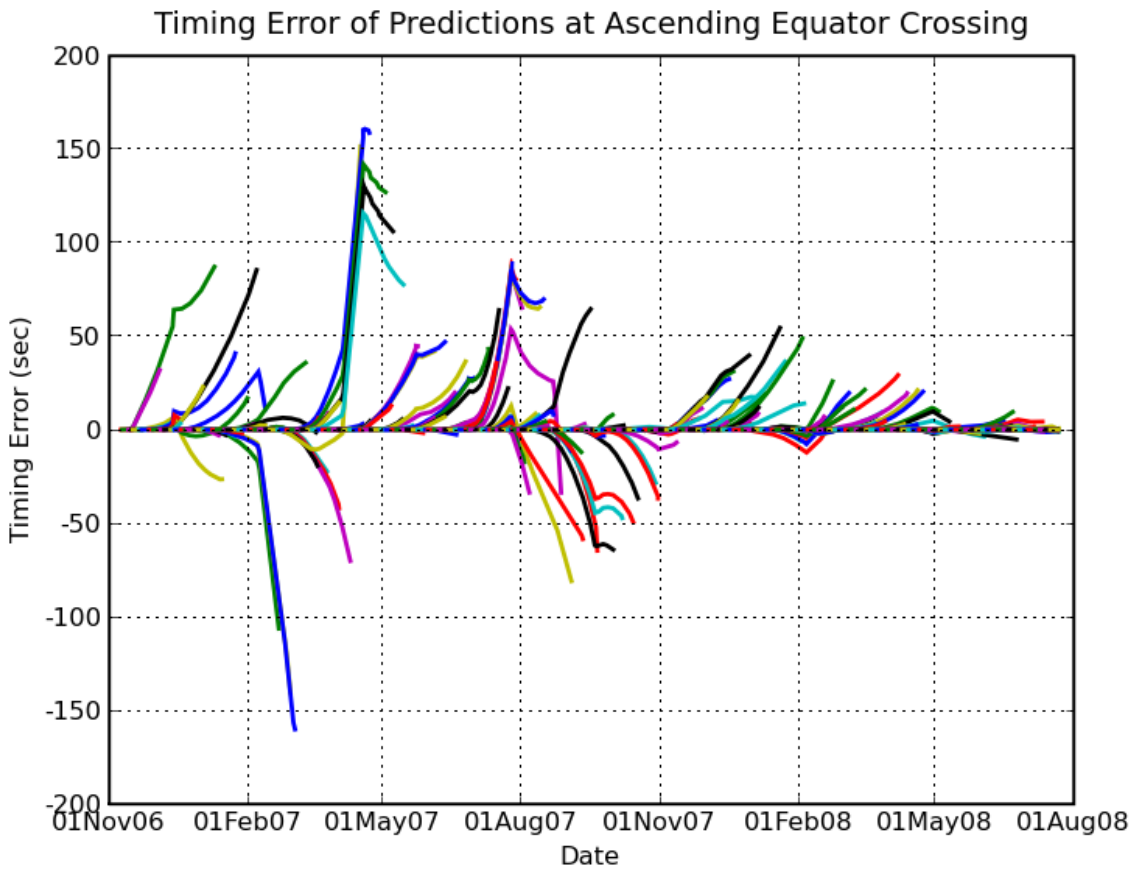

Figure 5. Timing errors of predicted ephemerides. Abrupt slope changes are caused by maneuvers. A dust storm caused the timing error runoff around July-August 2007.

C. Short-term Prediction

The driving short-term requirement, listed in Section II.C, states that the downtrack position must be maintained to 1.5 km (3σ). Converted to the equivalent timing value using the orbital velocity, the requirement becomes to maintain the on-board timing to 0.44 sec (3σ). The covariance analysis showed that the expected lifetime of a predicted ephemeris to the 3σ level would be 2 to 3.5 days from data cutoff[4]. With a day of latency before the ephemeris becomes active on-board, this meant a minimum of three on-board ephemeris updates per week, with the potential for updates virtually every day. The plan became to deliver a short-term predicted ephemeris three days per week on Monday, Wednesday, and Friday, in addition to the long-term ephemeris to be delivered on Thursday. However, after aerobraking while MRO was transitioning to the PSP, Navigation established a twice-per-week delivery schedule with a short-term ephemeris on Monday and a long-term on Thursday; the Flight Engineering Team (FET) at Lockheed Martin used these two deliveries to update the on-board ephemeris. This schedule persisted into the science phase, with the understanding that Navigation could increase the delivery frequency if the short-term requirement became at risk to be exceeded. Given the twice-per-week delivery schedule, the predictions would need to stay within the requirement for 3.5 to 4.5 days before the following update could take effect on board.

Because this ephemeris lifetime was longer than the 3σ lifetime from the covariance analysis, Navigation instituted a daily quicklook orbit determination (OD) solution to compare to the current on-board ephemeris. The quicklook process is a completely automated OD analysis that runs every day, resulting in an email to the Navigation Team around 8:00 am with relevant plots and timing error metrics. In this way, Navigation monitors the timing performance of the current on-board ephemeris daily, including predictions of the error in the near-term, to enable rapid identification of and response to unexpected perturbations. Just such an occasion occurred during late June 2007, when the quicklook OD showed a rapid increase in the mean drag over the span of a few days. This was a consequence of the onset of a dust storm, which severely reduced the predictability of the atmosphere. Navigation

adapted by increasing the prediction delivery frequency to three times per week until the dust storm had begun to settle out a month later.

The timing performance of the ephemeris predictions is shown in Fig. 6 and 7. Figure 6 is the same as Fig. 4, but focused in to the first seven days; Figure 7 shows the timing error between predictions with a 15-hour latency at the end to mimic the delay in processing and uplinking an on-board ephemeris. Figure 6 shows the first seven days of prediction for the 200+ ephemerides delivered since early November 2006. The blue dashed line with diamonds indicates the 3σ envelope based on the standard deviation of all the samples, including several outliers. The red dashed line at +/-0.44 sec highlights the short-term requirement. Obviously, many of the samples cross the requirement line before reaching 4.5 days into the predict, as indicated by the 3σ statistic at 4.5 days: 1.5 sec. This implies that this distribution does not meet the requirement even at the 1σ level; however, this may be misleading due to the large outliers and the fact that about half the samples (those delivered on Mondays) need only last for 3.5 days.

Figure 7 may be a better gauge of the short-term performance. The plot shows the error growth in each predict until 15 hours (8 orbits) after the DCO of the next predict. Most of the lines that exceed the requirement are due to maneuver execution errors, which were not considered in the short-term covariance analysis. Since routine maneuvers always occur on Wednesday, their execution errors show up on the Monday prediction at two days into the predict (see some of the large slopes that start at 2 days in Fig. 6). These cases are driving the statistics shown in Fig. 6 more so than the variability of the atmosphere or desat delta-v. Note also the seasonal dependence on the performance in Fig. 7, in particular the larger errors in July 2007 (high drag season) and smaller errors in June 2008 (low drag season). The low-drag season also requires fewer maneuvers, so there are not as many opportunities for predictions to be driven off by maneuver execution errors. Overall, Fig. 7 shows that the short-term requirement level is approached and occasionally exceeded, but often not by much. Thus, the twice-per-week delivery frequency works well most of the time, with the expectation that more frequent deliveries may be needed during the high drag season and immediately after maneuvers.

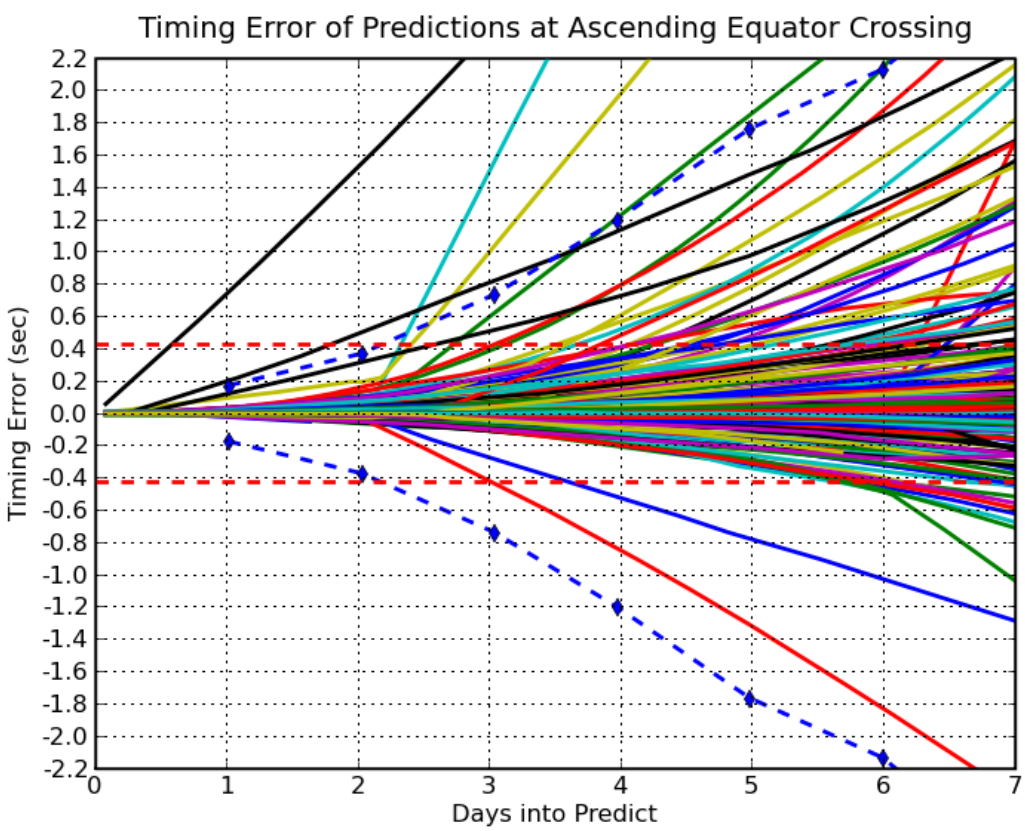

Figure 6. Timing errors of predictions relative to OD data cutoff. The short-term prediction requirement of 0.44 sec is shown as the dashed red line. The dashed blue line with diamonds shows three times the standard deviation of the errors.

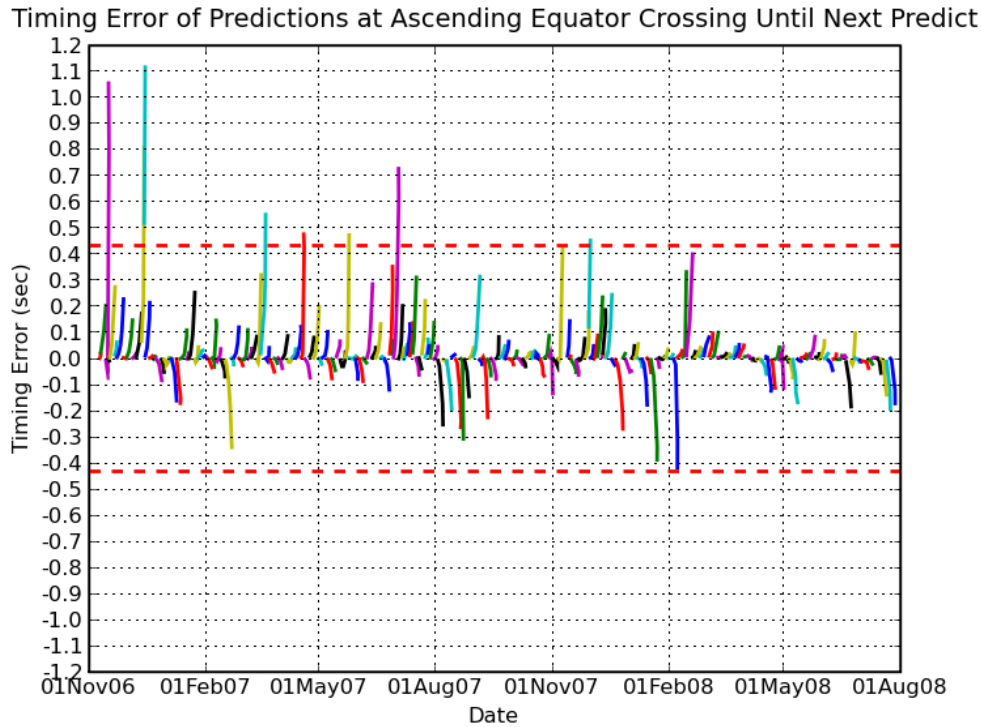

Figure 7. Predicted timing errors between ephemeris updates to mimic on-board timing errors. Each line extends 15 hours past the DCO of the next predict to account for latency. The dashed line represents the short-term requirement of 0.43 sec.

D. Reconstructions

Given the trajectory reconstruction requirement to determine the Mars-relative position to within 100 m (3σ) downtrack, 40 m (3σ) crosstrack, and 1.5 m (3σ) radial, the covariance analysis indicated no problem reaching those levels with the expected 16 hours per day of DSN tracking[4]. Because Doppler is the only tracking data used, the driving sensitivity for reconstruction covariance was determined to be the angle of the MRO orbit plane relative to the Earth line-of-sight: for face-on geometries, the downtrack position is less observable, whereas the crosstrack position observability deteriorates during edge-on geometries.

To evaluate the flight performance, examining the difference in the overlap region of reconstruct batches provides the best metric of the noise level. The trajectory is reconstructed in 20- to 24-orbit batches, with overlaps of approximately one orbit covered by two-way Doppler. Figure 8 shows the maximum magnitude difference in each overlap region, expressed in radial, downtrack, and crosstrack coordinates, along with the RMS of the points for each axis. There are several spikes in the radial direction that exceed the 1.5-km requirement, but the overall noise is well below that. Spikes in the overlap plot are sometimes caused by reconstructing a span that includes entry into safe mode, such as in November 2007 and February 2008. As MRO enters safe mode, it transitions to thruster attitude control and produces delta-v during the attitude changes. Unfortunately, the telemetry from these types of maneuvers is sparse, so the portion of the reconstruction batch that follows the safe mode entry may be less accurate.

In addition to the spikes, the downtrack and crosstrack directions in Fig. 8 show increased noise during certain periods. As expected from the covariance analysis, these periods correspond to edge-on and face-on geometries of the MRO orbit plane with respect to the Earth line-of-sight. Figure 9 shows the Earth beta angle, the elevation of the Earth as viewed from the MRO orbit plane. On the plot, a beta angle of zero indicates the edge-on condition, and a beta angle of 90 deg indicates the face-on condition. The crosstrack noise increased around July and November 2007 when the beta angle passed through zero; and the downtrack noise increased as the beta angle persisted above 60 deg in the first half of 2008. During these periods, the downtrack difference never exceeded the 100 m requirement, but the crosstrack difference shows a few batches that approach and exceed the more stringent 40 m requirement right at the zero crossing of the beta angle. In both directions, the RMS is well below the requirements.

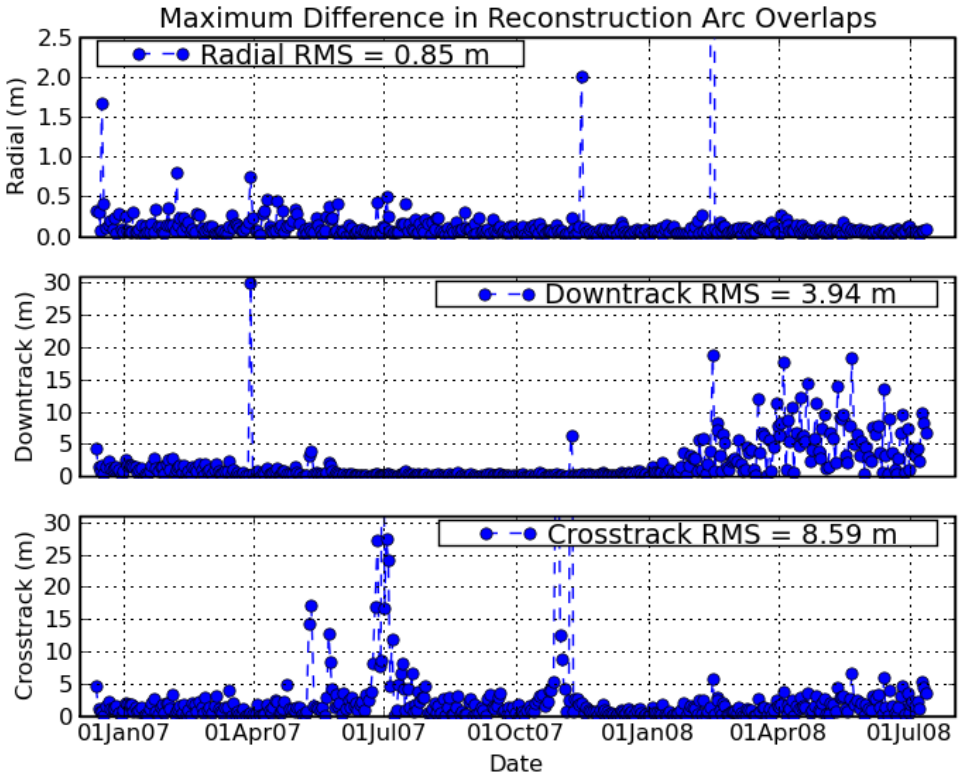

Figure 8. Maximum difference in each overlap region between successive reconstruction arcs. The 3σ requirements are 1.5 m radial, 100 m downtrack, and 40 m crosstrack.

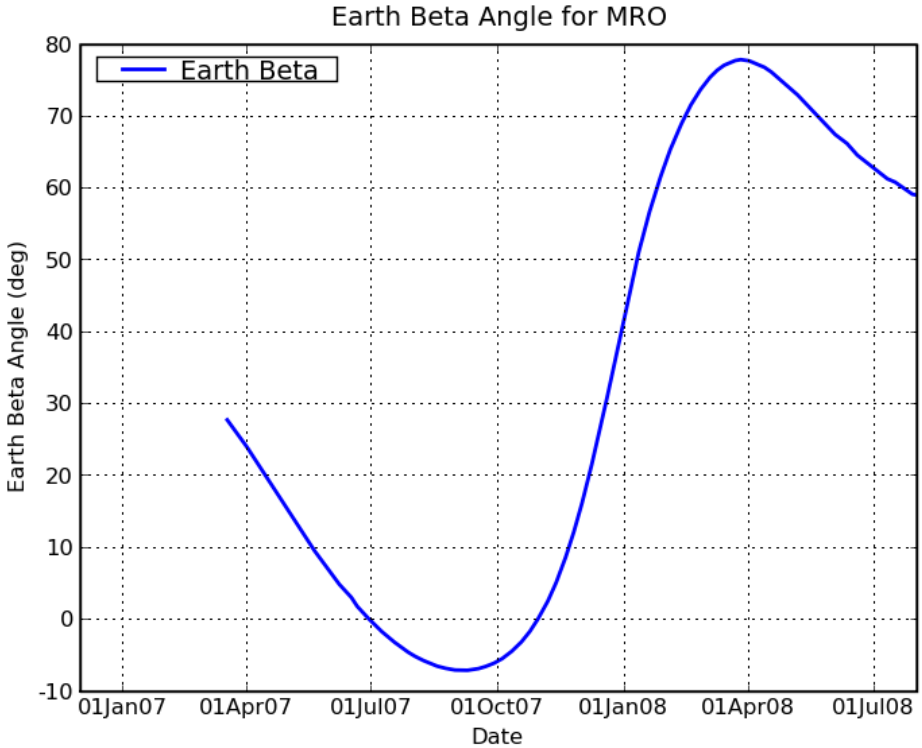

Figure 9. Earth beta angle with respect to MRO – the elevation of Earth above the MRO orbit plane. Edge-on to Earth = 0 deg, face-on = 90 deg.

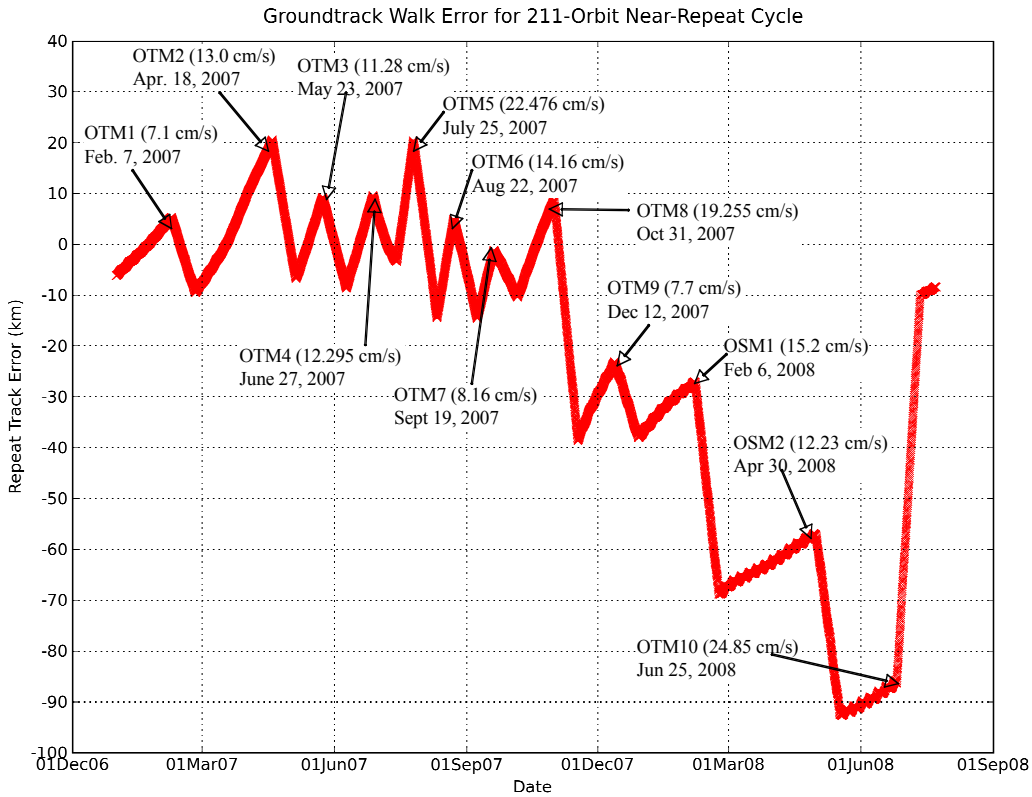

Figure 10. Groundtrack walk error of 211-orbit cycle with respect to the desired separation of -32.5 km (signified by a repeat track error of 0 on the plot). The goal was to maintain the error to +/-10 km, except during the Phoenix EDL phasing period.

E. Groundtrack Walk Error

As mentioned in Section II.C, the requirement for maintaining the nominal groundtrack walk pattern of 32.5 km west after 211 orbits is a de-facto requirement based upon the operational desire to maintain the general characteristics of the orbit. Because of atmospheric drag, the semimajor axis decays, causing the period to reduce and the groundtrack to drift eastward. The drift rate varies due to seasonal fluctuations in the atmospheric density, particularly at periapsis over the South Pole. Pre-launch analysis of various strategies resulted in a plan to perform orbit trim maneuvers (OTMs) at 28-day intervals, where some OTM opportunities during the low-drag season could be skipped[4]. The 28-day frequency fit into a consistent slot in the science and on-board sequence planning processes.

In operations, the specific OTM implementation strategy focused on targeting the maneuvers to a relative groundtrack error of -10 km and performing the next one when the relative error reaches +10 km. In other words, the OTM magnitude would be chosen to always push the 211-orbit repeat cycle to -42.5 km (measured positive eastward), then let the groundtrack drift back to the east to -22.5 km before performing the next OTM. The benefit of targeting the relative error of -10 km comes from a timing prediction standpoint. Since Navigation delivers a long-term predicted ephemeris that spans eight weeks, there could be up to two OTMs modeled in the trajectory. From week to week, these maneuver designs are updated with the latest OD and prediction models. By adjusting the OTMs to always go to the same relative walk, the maneuvers compensate for errors in the modeling of earlier predicts. Thus, this plan allows the long-term timing to be somewhat self-correcting, as evidenced by the lines in Fig. 5 that reverse the direction of timing error growth after OTMs.

Figure 10 shows the reconstructed 211-orbit groundtrack walk error over the science phase to date. A repeat track error of 0 km means that the pattern exhibits the desired -32 km eastward (+32 km westward) pattern. The positive slope segments are eastward drifts, whereas the negative slope segments are westward drifts. All the negative slopes and the last large positive slope are the result of OTMs. The steep positive slopes around July 2007,

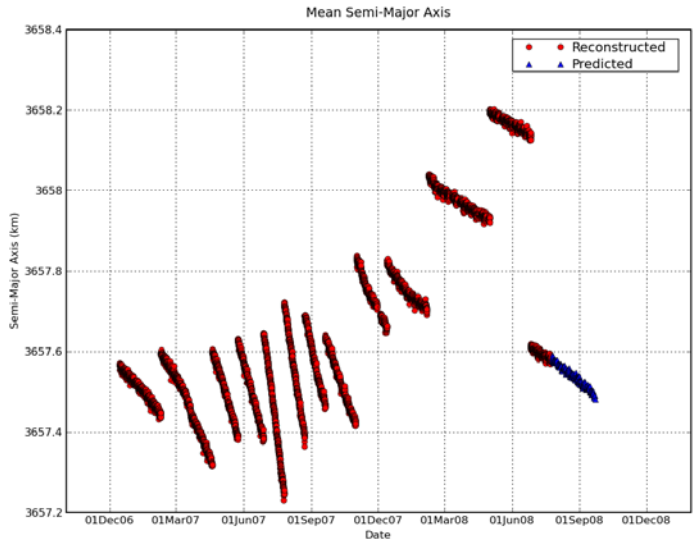

Figure 11. Mean semimajor axis.

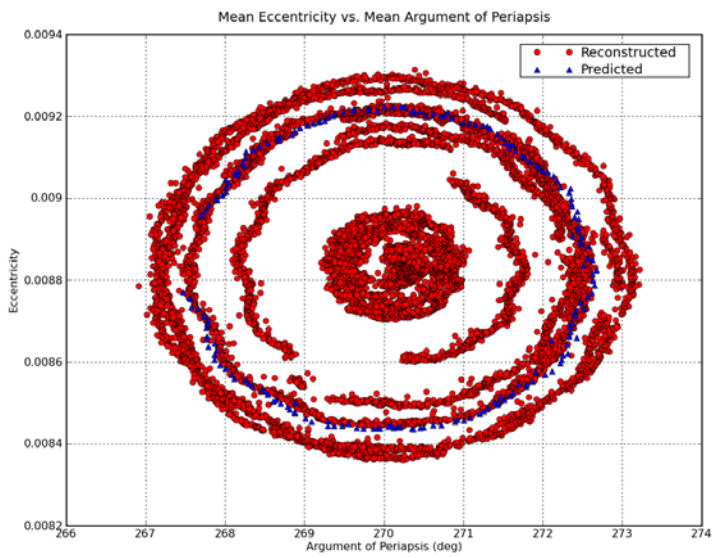

Figure 12. Mean eccentricity versus argument of periapsis.

along with the more frequent maneuvers, occurred during the high-drag season; the shallower positive slopes in 2008 are indicative of the low-drag season.

Until OTM-8 in October 2007, the groundtrack walk error was kept within -14 to +20 km, acceptably close to the +/- 10 km goal. However, MRO needed to phase within its orbit by approximately 45 min by May 25, 2008, to be in position to communicate with the Phoenix lander during its entry, descent, and landing (EDL). With their long lead times, OTM-8 and OTM-9 were biased to target -40 km, thereby significantly reducing the amount of phasing remaining for the two orbit synchronization maneuvers (OSMs). During the orbit synchronization period with Phoenix, the groundtrack walk error was ignored. One month after Phoenix EDL, however, OTM-10 returned the groundtrack error to within the +/- 10 km bounds.

Finally, it is interesting to note the effect of this groundtrack maintenance strategy on the mean elements. Figure 11 shows the mean semimajor axis, and Figure 12 shows the mean eccentricity versus mean argument of periapsis (e-ω plot). Clearly, controlling the groundtrack walk pattern close to a certain value requires maintaining a nearly constant period or semimajor axis. Figure 11 is analogous to Fig. 10, showing the corresponding semimajor axis control corridor and the departures due to the Phoenix phasing maneuvers. The e-ω plot in Fig. 12 indicates the degree to which the MRO orbit is frozen, and how that condition changes with the orbit maintenance maneuvers. Jumps in this plot occur at the maneuvers, which are usually alternated between apoapsis and periapsis to attempt to minimize the departures from the frozen condition. Without a formal requirement for *how* frozen the orbit must be, the goal is to stay within +/-3 deg in mean argument of periapsis. Recently, OTM-10 was executed approximately 40 deg away from apoapsis to reduce the ω variation from 3 deg to 2.5 deg.

F. Phoenix Phasing

The Phoenix phasing requirement states that MRO must ascend through the specified target latitude within 30 sec of the specified time. The only maneuver that matters with regard to the 30-sec tolerance is the final synchronization maneuver, OSM-2, at 28-days prior to Phoenix EDL. Based on the long-term covariance analysis shown in Ref. 4, the predictive capability during the low-drag season (when Phoenix would arrive) is 30 sec (3σ). This is the value at 35 days into the prediction to allow for a one-week maneuver design process. When determining the phasing strategy, the covariance analysis was updated with more accurate models and uncertainties, including variations of desat delta-v biases. Given current atmospheric scaling and a nominal or conservative desat delta-v bias, the resulting covariance indicated OSM-2 delivery performance in the range of 15-20 sec (3σ).

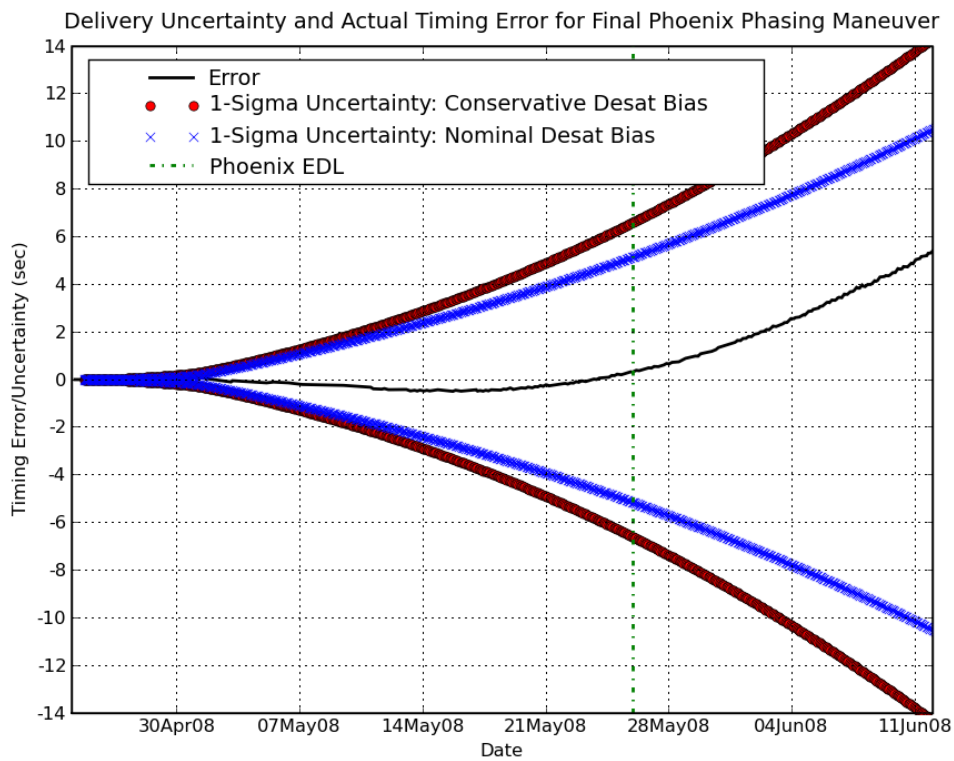

Figure 13. Phoenix EDL trajectory delivery uncertainty compared to actual performance. Achieved error represents < 0.05σ.

Figure 13 shows these updated covariance bounds at the 1σ-level compared to the actual *error* of the OSM-2 delivery ephemeris. The delivery ephemeris is the trajectory with OD DCO one week prior to OSM-2 execution. It includes the final maneuver design that takes MRO exactly to the EDL target. As Fig. 13 shows, MRO reached the target latitude only 0.25 sec early, better than 0.05σ performance. Part of this success, however, must be attributed to serendipity: the OSM-2 execution error offset the OD predict model error until just after EDL, at which point the plot clearly shows an acceleration of the timing error. Nonetheless, the MRO phasing performance far exceeded expectations.

V. Atmospheric Drag Trending

The most critical aspect of MRO trajectory modeling is the atmosphere. Its variability drives the prediction performance and the necessity of the reconstructions to include per-orbit drag scale factors. Similar to aerobraking, the key metric that defines the state of the atmosphere is per-rev drag delta-v, the measure of the integrated effect of the drag acceleration between apoapses. The drag equation implemented in the Navigation modeling is

$$a_D = -\frac{1}{2}(S\rho)V^2\frac{C_D A}{m} \tag{1}$$

where a_D is the drag acceleration, S is the density scale factor, ρ is the density, V is the spacecraft velocity, C_D is the drag coefficient, A is the drag area, and m is the spacecraft mass. Integrating this equation over one revolution produces the drag delta-v for a given orbit. The density scale factor is the parameter estimated in the filter as a constant over each revolution. In reality, this parameter scales the entire drag delta-v, accounting for variations in both density and drag area. Recall that the drag area is not modeled precisely since the computation does not compensate for shadowing from the drag direction among the components (see Fig. 2). Thus, in this case, S is more appropriately referred to as the *drag* scale factor. By examining trends in the drag scale factor estimates, Navigation determines how to properly scale the drag for the prediction model. In addition, results from various atmosphere models can be compared to the nearly two years of reconstructed drag delta-v values to determine which model would have been the best predictor.

A. Reconstructed and Predicted Drag Scale Factors

Trending begins by evaluating the estimates of the per-rev drag delta-v from the reconstructed trajectory. Figure 14 shows all the estimates since November 2006. The random, short-term noise is apparent, as is a seasonal trend that peaks during the southern summer (L_S = 270 deg in July 2007). The sudden increase in drag delta-v in late June to early July 2007 is attributable to two nearly simultaneous events: the spacecraft transitioned from fixed solar panels to Sun-tracking solar panels, significantly increasing the drag; and a dust storm began that became nearly global in extent and dramatically increased the density at the MRO altitude. The solar panels gimbals were fixed from March 2007 through June 2007 such that the panels were edge-on to the drag direction; at all other times, they tracked the Sun as illustrated in Fig. 2.

Figure 15 shows the reconstructed drag scale factors compared to the predicted scale factors for all the ephemeris deliveries to date during the PSP. The discontinuity on February 23, 2007, signifies the change from the aerobraking baseline of MarsGRAM 2005 Map Year 1 to MarsGRAM 2000 MRO Special Edition, which showed much more stable and predictable behavior at the time. The flat lines of the predicted scale factors represent the strategy of predicting with a constant average scale factor. The exception to this occurred as the dust storm waned in August 2007, at which point a negative scale factor rate term was implemented until the scale factor reduced to the pre-dust-storm level. A departure from the underlying seasonal trend in MarsGRAM 2000 caused the scale factor to continue to decline into October 2007. Since then, the bias has not varied by more than about 10 percent.

Figure 15 shows significant short-term noise, as well as periods of long-term trends in the bias. A plot of the standard deviation of the scale factor estimates for non-overlapping three-day (39-orbit) spans is shown in Fig. 16. The standard deviation is scaled by the three-day average bias over the same span, so the metric becomes a percentage of noise relative to the mean. The covariance analysis assumption was 35% per rev (1σ) with a two-day correlation. Figure 16 shows that after changing the MarsGRAM model in late Feburary 2007, the scale factor exhibited random noise at approximately 25% or less. However, there are clearly periods when the bias behavior exhibits correlated trends at time scales much longer than two days. A more appropriate covariance model would include per-rev random white noise at 25% (1σ) along with a 10% (1σ) bias.

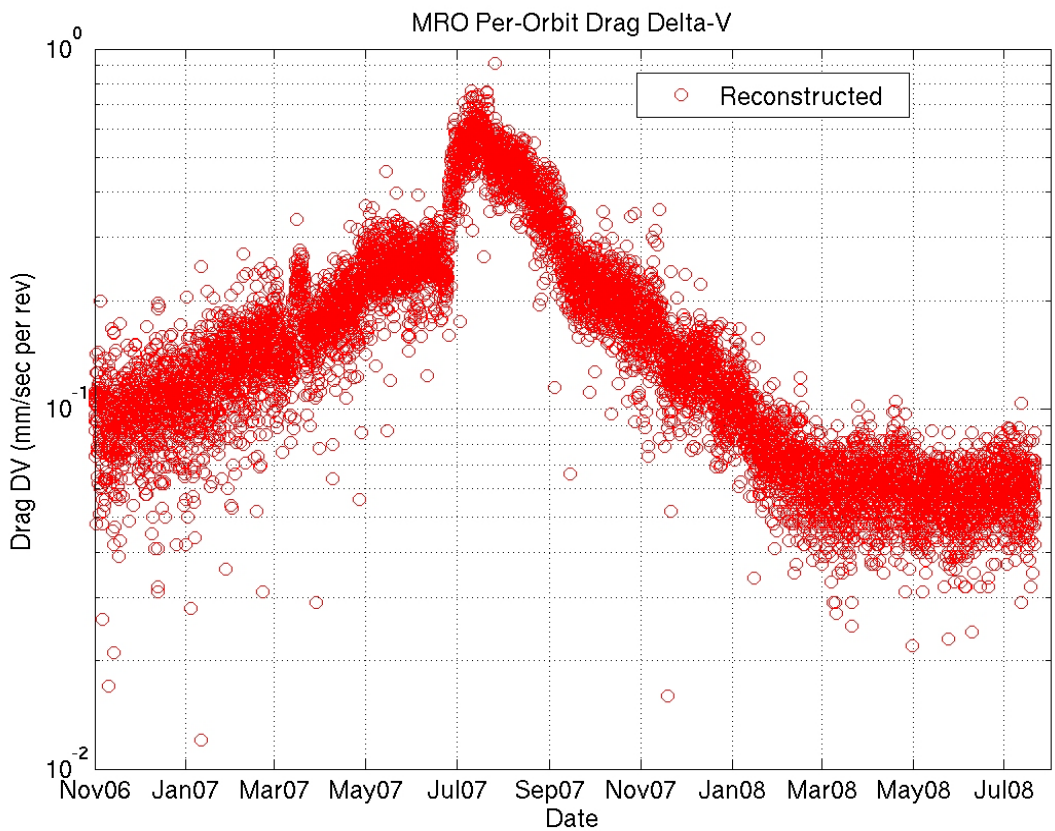

Figure 14. Reconstructed per-orbit drag delta-v.

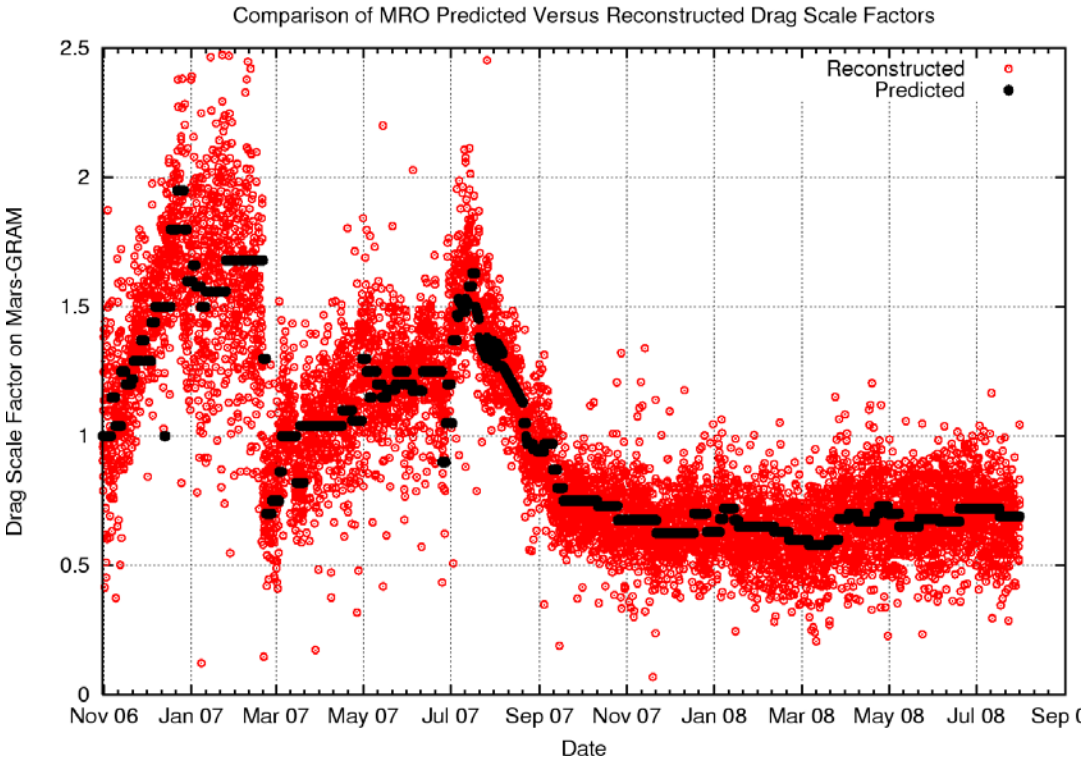

Figure 15. Reconstructed and predicted per-orbit drag scale factor estimates for densities from MarsGRAM 2005 Map Year 1 (prior to Feb. 23, 2007) and MarsGRAM 2000 MRO Special Edition (after Feb. 23, 2007).

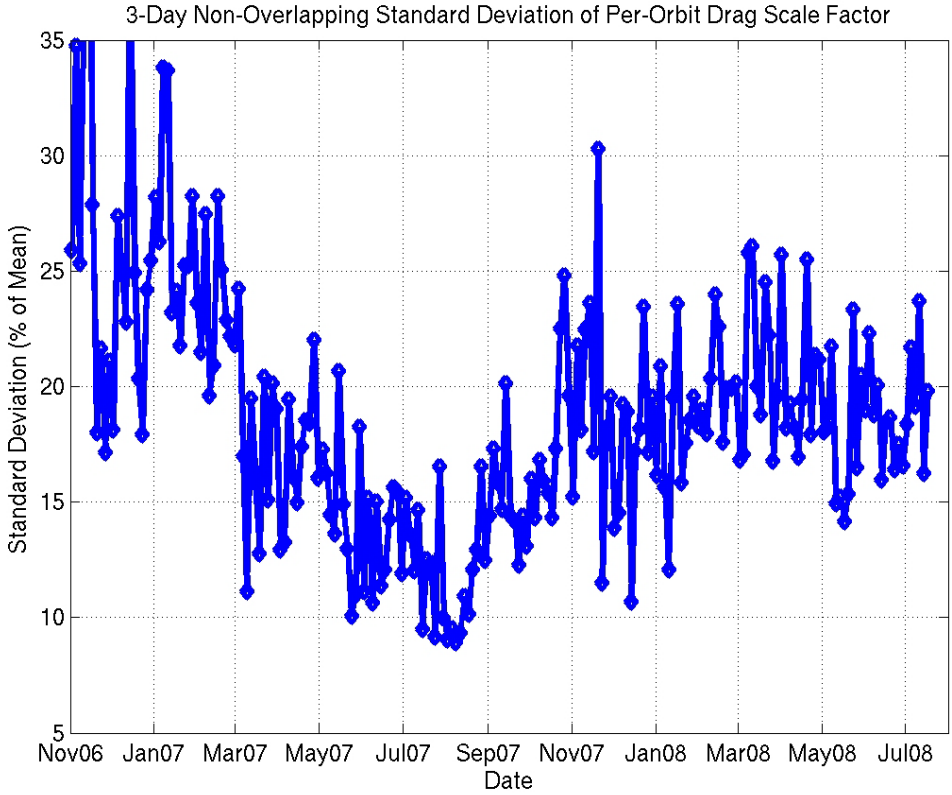

Figure 16. Standard deviation of non-overlaping three-day (39-orbit) samples of drag scale factor as a percentage of the three-day mean.

B. Model Comparisons

The dramatic difference in drag scale factor estimates between MarsGRAM models in Fig. 15 highlights the issue of determining the most appropriate density model for the MRO orbit. The latest MarsGRAM models are empirically tuned based on retrievals of density from accelerometer measurements during MGS and Odyssey aerobraking. However, the accelerometers are not sensitive enough to measure the density at 250 km and above, and there have not been any previous Mars orbiters with as sustained a presence in the 250 to 300 km region. The closest recent comparison is Mars Express, which inhabits an 6.5-hour-period elliptical orbit with a periapsis altitude of approximately 265 km when precessing over the south pole. In Ref. 12 several months of Mars Express orbits were reconstructed as its periapsis moved across the south pole, where the MRO periapsis would be frozen. Using MarsGRAM 2001 MRO Special Edition, the reconstructed drag scale factors showed a mean of approximately 0.5. Since the same density model had been used in the MRO covariance analysis, this result set the expectation that the MRO covariance may be conservative.

For MRO aerobraking operations, however, the baseline density model became MarsGRAM 2005 Map Year 1[13]. The primary updates in the 2005 model were the addition of data sets from the MGS Thermal Emission Spectrometer (TES) instrument of the lower atmosphere[7]. Two Mars years of data were implemented into separate models: Map Year 1 and Map Year 2. There remained a Map Year 0, as well, that is based on MarsGRAM 2001 models.

With the extensive amount of data from MRO, a comparison can now be made between the various MarsGRAM models to evaluate, in hindsight, which model would have provided the most consistent results. Figure 17 shows the reconstructed three-day average per-orbit drag delta-v versus the drag delta-v that would have been produced by the models listed in Table 3 with the specified relevant inputs. In other words, given the MRO trajectory and densities obtained by the models listed in Table 3, Figure 17 shows the resulting drag delta-v from each for a scale factor of 1.0.

Table 3. MarsGRAM model variations for comparison to drag delta-v reconstructions.

Model	Full Version Name	*Dusttau*	*zoffset*	*ibougher*	Scale Factor STD	Comment
MG2k MRO	2000 MRO Special Edition	0.5	-	-	30%	Current baseline model for science operations
MG01 MRO SE	2001 MRO Special Edition	0.5	+7 km	1	33%	Baseline for covariance analyses
MG05 MY0	2005 Map Year 0	0.1 – 0.5	+11 km	1	17%	Tuned to reconstructions
MG05 MY1	2005 Map Year 1	0.3	-2 km	0	47%	Aerobraking baseline
MG05 MY2	2005 Map Year 2	0.3	-2 km	0	50%	Use MY1 inputs

The input parameters listed in Table 3 allow MarsGRAM to be tuned for specific output. *Dusttau* controls the optical depth of the background dust, with the option to have a seasonally varying range; *zoffset* shifts the atmosphere up or down by the specified amount, where a positive value increases density; and *ibougher* is a flag that determines whether a seasonal variation is applied to zoffset (*ibougher*=1) or not (*ibougher*=0). For details refer to Ref. 7.

Given the results from Fig. 17, Fig. 18 shows by how much these models must be scaled to match the reconstructions. For reference, the *Baseline MG2k MRO* line is the three-day running mean of the per-rev estimates shown in Fig. 15 after February 23, 2007. The standard deviations (STD) of the scale factor profiles in Fig. 18 are also given in Table 3. As a percentage of the mean, this metric provides a measure of the long-term variability of the respective models. The very large variability in the 2005 Map Year 1 scale factor bears out the decision to switch to a different model. The 2005 Map Year 2 results are even worse, while the 2000 and 2001 scale factors show reasonable stability. The 2001 scale factors for most of the time are at a very similar level to the Mars Express results at approximately 0.5, indicating an overall consistency in the model predictions between 2004 and 2008. The relatively flat profile from the 2005 Map Year 0 scale factors shows a significant improvement over the current baseline. Note, however, that this model was specifically tuned to attempt to match the reconstructions, whereas the other models use inputs consistent with current or former usage. Therefore, it is plausible that the other models could be tuned for improved performance, particularly 2001, on which 2005 Map Year 0 is based. Nonetheless, the conclusion of this analysis would be to use a similarly-tuned MarsGRAM 2005 Map Year 0 model for the analysis of future missions in MRO-like orbits.

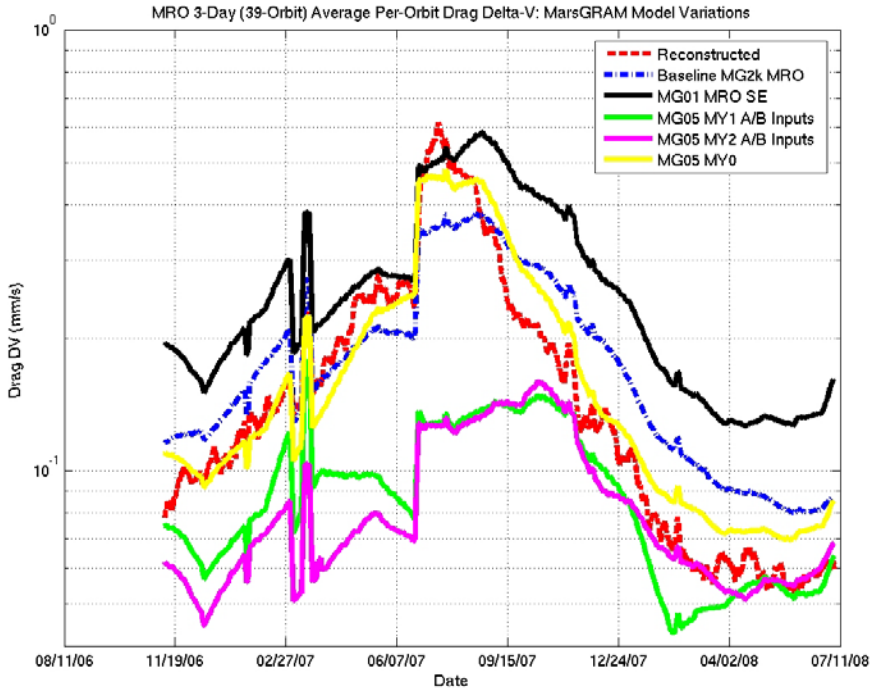

Figure 17. Comparison of reconstructed three-day average drag delta-v to various MarsGRAM models, including the 2000 MRO Edition (MG2k MRO), 2001 MRO Special Edition (MG01 MRO SE), 2005 Map Year 1 and 2 with aerobraking inputs (MG05 MY1/MY2 A/B Inputs), and 2005 Map Year 0 (MG05 MY0).

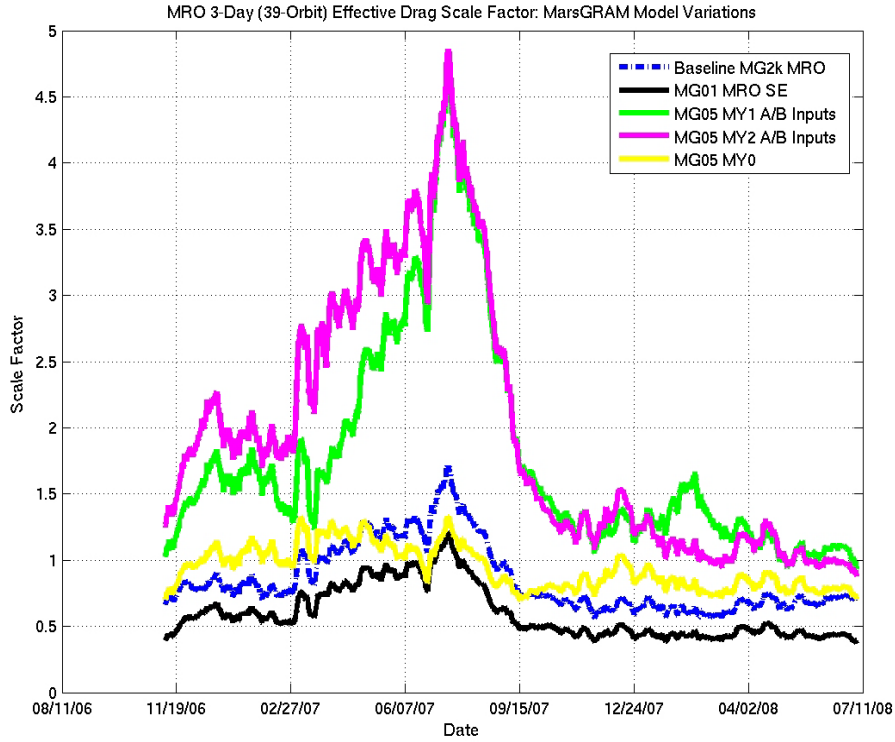

Figure 18. Comparison of three-day running mean drag scale factors that would have been obtained with the different MarsGRAM models based on the reconstructed drag delta-v. The Baseline MG2k MRO line is the three-day running mean of the data plotted in Fig. 15. Note that the MarsGRAM 2005 Map Year 0 scale factor exhibits the smallest variance.

VI. Conclusion

Prior to science phase operations, the Mars Reconnaissance Orbiter Navigation team set the expectations of performance through extensive covariance analysis and requirements verification. Analysis of the flight performance after the first 8000 science orbits has demonstrated the expected results, further emphasizing the challenge of meeting the stringent navigation requirements in the presence of a highly-variable atmosphere. Maintaining accurate navigation requires continual monitoring and trending, particularly of the behavior of the density model. MRO Navigation has demonstrated a robust strategy of performance evaluation and trending analysis—keys to the successful navigation of the next 8000 orbits and beyond.

Acknowledgments

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California under contract to the National Aeronautics and Space Administration. The authors would like to thank the Mars Reconnaissance Orbiter Project, Lockheed-Martin Space Systems, and the JPL Guidance/Control & Navigation Section for their support.

References

[1]Zurek, R.W., and Smrekar, S.E., "An overview of the Mars Reconnaissance Orbiter (MRO) science mission", *J. Geophys. Res*., Vol. 112, E05S01, doi:10.1029/2006JE002701, 2007.

[2]Johnston, M.D., et al., "The Mars Reconnaissance Orbiter Mission", 2004 IEEE Aerospace Conference, Big Sky, WY, 6-13 March 2004.

[3]Bhat, R., et al. "Mars Reconnaissance Orbiter Primary Science Orbit Acquisition", AAS/AIAA Astrodynamics Specialists Conference, Mackinac Island, Michigan, 19-23 August 2007.

[4]You, T., et al., "Mars Reconnaissance Orbiter Navigation", AIAA/AAS Astrodynamics Specialist Conference, Providence, RI, 16-19 August 2004.

[5]Justus, C.G., and James, B.F., "Mars Global Reference Atmospheric Model 2000 Version (Mars-GRAM 2000): Users Guide", NASA/TM-2000-210279, May 2000.

[6]Justus, C.G., and Johnson, D.L., "Mars Global Reference Atmospheric Model 2001 Version (Mars-GRAM 2001): Users Guide", NASA/TM-2001-210961, April 2001.

[7]Justus, C., and Duvall, A., "Atmospheric Models for Mars Aerocapture", 41st AIAA/ASME/SAE/ASEE Joint Propulsion Converence and Exhibit, Tucson, AZ, 10-13 July 2005.

[8]Konopliv, A.S., Yoder, C.F., Standish, F.M., Yuan, D., and Sjogren, W.L., "A global solution for the Mars static and seasonal gravity, Mars orientation, Phobos and Deimos masses, and Mars ephemeris", *Icarus*, Vol. 182, 2006, pp. 23-50.

[9]Mackenzie, R., and Folkner, W.M., "Applying Appropriate Weights to Doppler data", JPL IOM-343J-06-034, 24 October 2006.

[10]You, T., et al., "Mars Reconnaissance Orbiter Interplanetary Cruise Navigation", 20th International Symposium on Space Flight Dynamics, Annapolis, MD, 24-28 September 2007.

[11]Highsmith, D., "The Effect of USO Stability on One-Way Doppler Navigation of the Mars Reconnaissance Orbiter", AAS/AIAA Space Flight Mechanics Meeting, San Antonio, TX, 27-30 January 2002.

[12]Highsmith, D., et al., "Atmosphere Variability at Mars Reconnaissance Orbiter Science Orbit Altitudes Based on Mars Express Reconstructions", AAS/AIAA Space Flight Mechanics Meeting, Copper Mountain, CO, 23-27 January 2005.

[13]Demcak, S.D., et al., "Mars Reconnaissance Orbiter Orbit Determination During Aerobraking", AAS/AIAA Space Flight Mechanics Meeting, Sedona, AZ, 28 January – 1 February 2007.