

FILTER STRATEGIES FOR MARS SCIENCE LABORATORY ORBIT DETERMINATION*

**Paul F. Thompson,[†] Eric D. Gustafson,[‡] Gerhard L. Kruizinga,[§]
and Tomas J. Martin-Mur^{**}**

The Mars Science Laboratory (MSL) spacecraft had ambitious navigation delivery and knowledge accuracy requirements for landing inside Gale Crater. Confidence in the orbit determination (OD) solutions was increased by investigating numerous filter strategies for solving the orbit determination problem. We will discuss the strategy for the different types of variations: for example, data types, data weights, solar pressure model covariance, and estimating versus considering model parameters. This process generated a set of plausible OD solutions that were compared to the baseline OD strategy. Even implausible or unrealistic results were helpful in isolating sensitivities in the OD solutions to certain model parameterizations or data types.

INTRODUCTION

The Mars Science Laboratory (MSL) spacecraft successfully delivered the Curiosity rover to Gale Crater, Mars on August 06, 2012. A crucial part of this success was orbit determination (OD) during the cruise from Earth to Mars. The primary goal of this process was to determine the spacecraft state and accelerations acting on the spacecraft along with their associated uncertainties. This allowed the navigation team to predict the future trajectory and to characterize the expected error in the propagated trajectory. Understanding and limiting the uncertainties was necessary in order to meet the delivery and knowledge requirements at the time of atmospheric entry. Differently from previous Mars landers, the MSL entry, descent, and landing (EDL) system had active guidance control. Knowledge requirements needed to be met in order to initialize EDL. The baseline OD process and results details have been discussed elsewhere¹. Herein, we discuss the OD solution filter variations that were used to increase our confidence in the OD and helped guide changes made to the baseline OD solution strategy in order to meet the ambitious requirements for EDL.

* © 2013 California Institute of Technology. Government sponsorship acknowledged.

[†] MSL Orbit Determination Analyst; Mission Design and Navigation Section; Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California 91109.

[‡] MSL Orbit Determination Analyst; Mission Design and Navigation Section; Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California 91109.

[§] MSL Orbit Determination Lead; Mission Design and Navigation Section; Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California 91109.

^{**} MSL Navigation Team Chief; Mission Design and Navigation Section; Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, California 91109.

MSL Mission Overview

Mars Science Laboratory was launched from Cape Canaveral on November 26, 2011. The cruise phase ended successfully with a precision landing on Mars in Gale Crater on August 06, 2012. The MSL spacecraft was the first mission to use entry guidance at Mars to meet the precise landing requirements for Gale Crater. In previous Mars missions, e.g., Phoenix and the Mars Exploration Rovers (MER), the navigation objective was to control the entry state such that the desired landing location would be achieved assuming ballistic entry. The guidance system of MSL allowed a reduction of the landing ellipse from 80×10 km for the MER rovers² to one approximately 20×7 km for MSL. While the guidance system used active control to land at a desired location by processing inertial measurement unit (IMU) data, it still needed to be initialized with an entry state. Therefore, the emphasis of the MSL navigation team at the end of the approach phase was not on precisely controlling the location of atmospheric entry, but in improving the trajectory knowledge in order to initialize the EDL system.

Based on pre-launch analysis of navigation performance, the project had originally planned for a number of trajectory correction maneuvers (TCMs) and entry parameter update (EPU) opportunities (Table 1). The first of these updates, EPU-1, was based on tracking data up to Entry-6.5 days, a time soon after TCM-4 execution. However, due to the stability of the OD solutions for the remainder of the approach phase up through atmospheric entry, the remaining TCM and EPU opportunities were all canceled. The final OD reconstruction shows that MSL hit the top of the atmosphere entry point only 200 m from the EPU-1 state³. Further details on these key events along with the navigation strategy and OD analysis in support of them can be found in References 1 and 3.

Table 1. MSL Major Navigation Cruise Events.

Mission Event	Event Date	Relative Mission Time
Launch	26-Nov-2011	L+0 days
TCM-1	11-Jan-2012	L+45 days
TCM-2	26-Mar-2012	L+120 days
TCM-3	26-Jun-2012	E-40 days
TCM-4	29-Jul-2012	E-8 days
EPU-1	30-Jul-2012	E-6.5 days
TCM-5*	04-Aug-2012	E-2.0 days
EPU-2*	04-Aug-2012	E-33 hrs
EPU-3*	05-Aug-2012	E-15 hrs
TCM-6*	05-Aug-2012	E-9 hrs
EPU-4*	05-Aug-2012	E-6 hrs
Entry	06-Aug-2012	E-0 hrs

**cancelled activity*

MEASUREMENT MODELING AND DYNAMICS

An interplanetary tracking schedule was developed for range, Doppler, and delta difference one-way range (Δ DOR) measurements for cruise. The baseline OD strategy was to include the available range, Doppler, Δ DOR data. We assumed a per-pass data weight based on the post-fit RMS of each tracking pass, estimation of charged particle delay parameters, and a data elevation cutoff of 20 deg. In addition, there was pre-processing of data to remove the sinusoidal signature and Doppler bias introduced by the spinning spacecraft⁴. Once the spin signature had been removed and the Δ DOR measurements processed, we used the Mission analysis and Operation

Navigation Toolkit Environment (MONTE) program developed at the Jet Propulsion Laboratory (JPL)⁵ to process data and produce an OD solution. For additional background on the fundamentals of statistical orbit determination (e.g., linearization, measurement processing, covariance propagation, and consider parameters to name a few topics), please see Reference 6. The specific process used for MSL OD and the details of the baseline OD filtering strategy are discussed in Reference 1.

Force models included TCMs, delta-V introduced by turns, solar radiation pressure, and thermal radiation pressure due to the radio-isotope thermoelectric generator (RTG) on-board the Curiosity rover. The measurement model included Deep Space Network (DSN) station coordinates and the quasar catalog. The MSL OD process used dedicated planetary ephemerides that were focused on improving the Mars ephemerides with the latest observations. There were two planetary ephemerides used during cruise: (1) DE424⁷, which was delivered two months prior to launch, and (2) DE425⁸, which was delivered three months prior to Mars entry. The change introduced by the updated ephemeris was on the order of tens of meters – contributing very little error to the total error from other sources at entry. Note that the planetary ephemerides covariance from DE423⁹ was used during the entirety of MSL cruise and EDL. The spherical harmonic expansions of the Mars gravity field was taken from mgs95j¹⁰ and truncated to degree and order eight.

SRP modeling

The MSL OD team used a novel approach for modeling the Solar Radiation Pressure (SRP) in interplanetary navigation³. Most interplanetary missions model the SRP using a geometric model of the spacecraft and assign to each the necessary surface optical properties to compute the SRP force. In Figure 1, the MSL spacecraft is depicted showing all the SRP forces and surfaces involved: solar panels, the launch adaptor, the parachute cone, antennas, radiators, and sensors. Modeling each surface and SRP force becomes rather complex with a complicated spacecraft geometry. For MSL the SRP force was computed by expanding the net total SRP force into Fourier series as a function of the solar colatitude. The Fourier coefficients represent the effective surface areas. Only the average effect was important as the MSL was a spin stabilized spacecraft rotating approximately about the Z-axis.

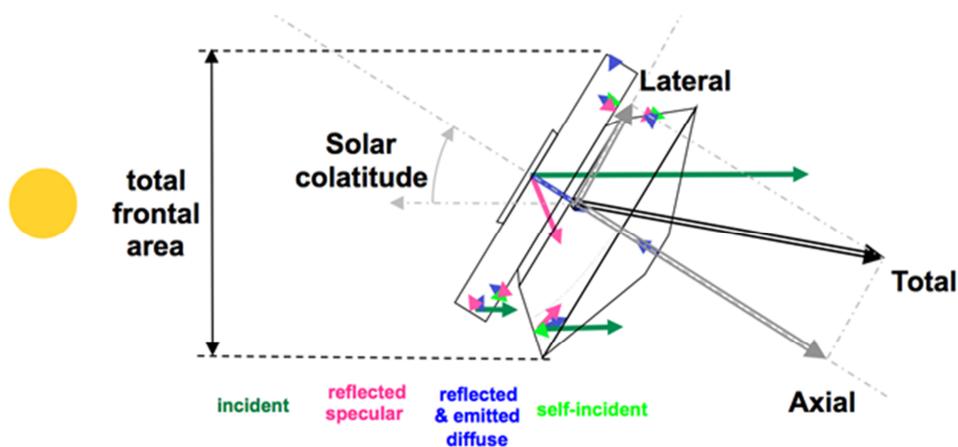


Figure 1. SRP Model Illustration.

The sums of the Fourier coefficients were multiplied by an appropriate scale factor to account for the spacecraft mass, solar distance and solar flux. The SRP Z-axis coefficients represented the SRP force in approximately the rotation axis (Z-axis) direction and the SRP X-axis coefficients represented the SRP force perpendicular to the rotation axis in the plane formed by the Sun vector and the spacecraft Z-axis. The Y-axis coefficients represented the SRP force perpendicular to the plane formed by the Sun vector and the spacecraft Z-axis. In general this force in the Y-direction is expected to be small because of the spacecraft symmetry, though thermal imbalance may cause a force in this direction.

Prior to launch, a MSL geometric model was constructed with appropriate surface optical modeling. A Fourier expansion of the SRP was estimated up to degree two for only the X and Z coefficients. This was used as the a priori model during launch and early cruise during the OD process. The number of parameters estimated, a priori values, and uncertainties were refined and modified as our experience with the model increased. In addition, there was a net acceleration along the Z-axis due to the thermal effects of the radioisotope thermal generator that was difficult to separate from the effects of the SRP model. As such, we generally analyzed its effects in combination with the acceleration caused by SRP.

FILTER STRATEGY AND RESULTS

The OD process and strategy evolved with each phase of the cruise mission because of unique aspects of that phase or lessons learned in previous phases. For discussion purposes, the cruise mission was divided up into the following phases: launch, early cruise, mid-cruise, late cruise and final approach (Table 2). A key characteristic of any OD solution is the timespan of data included in the solution. The last time at which measurements are included in the data set is called the data cutoff (DCO). For each DCO, several OD solutions were computed by using a small number of data start times. One of these arcs would then be chosen to be the officially delivered, baseline OD solution for a DCO.

Table 2. Orbit Determination Phase Definitions

Phase	Start Date	End Date	Comment
Launch	26-Nov-2011	26-Nov-2011	First tracking pass
Early Cruise	26-Nov-2011	11-Jan-2012	Launch to TCM-1
Mid-cruise	11-Jan-2012	05-Mar-2012	Starts at TCM-1
Late cruise	05-Mar-2012	29-Jul-2012	Ends at TCM-4
Approach	29-Jul-2012	06-Aug-2012	TCM-4 to Entry

Once a baseline OD solution had been created, a large number of solution variations of the baseline were computed. These filter variations were collectively known as “filterloop” – multiple OD solution variations calculated by automatically looping through many different versions of the filter setup. This was used to calculate a new solution to the model parameters along with a new trajectory and error predicted at Mars entry. The details of these different filter cases changed and evolved significantly during the cruise to Mars.

Data variations that we examined, depending on the mission phase, could include: a single weight for each data type instead of per pass, tighter or looser weights for Δ DOR, selected data types (e.g., Doppler and range only, Doppler and Δ DOR only), or Δ DOR observation pairs that include only those from the North-South baseline or the East-West baseline.

Another filter variation we discuss is that of varying data arc length. This is not identical to having independent OD trajectory arcs that start at different epochs. However, they are compara-

ble in that these variations of the data arc alone keep the same nominal trajectory model (i.e., the same initial epoch and state) but remove data from the start of the arc

Dynamic model variations included varying SRP model a priori covariances, older SRP models, and TCM and other DV event covariance scaling. Particularly during the approach phase, the filterloop focus was on the effect of estimating or removing consider parameters such as Mars and Earth ephemerides, or media corrections (troposphere and ionosphere). While not all variations were realistic, they all helped to determine what was a plausible set of OD solutions for a given data cutoff. Even less realistic assumptions helped highlight sensitivities to data or specific model parameters.

There were as many as 40 different filterloop cases computed for each baseline OD delivery as the mission progressed, taking as long as 90-120 minutes to complete if done in sequence. This was computationally prohibitive during the final days of approach when baseline ODs were computed for new DCOs every one to two hours. A parallelization process was developed to allow multiple workstations to work on one to several cases at once. This cut the computation time for the complete set of filterloop case to as little as 10 min per baseline OD.

In all following figures, the OD solutions are shown mapped to the Mars B-plane. This plane is one which passes through the center of gravity of Mars and is perpendicular to the incoming asymptote of the spacecraft. The position in this plane is approximately the point that the spacecraft would pass through if Mars had no mass. For a detailed discussion of the definition of the B-plane and its application to spacecraft navigation, see Reference 11.

The three-sigma requirements for OD were to provide an entry state with an accuracy of 2.8 km in position, 2.0 m/sec in velocity, and a flight path uncertainty of 0.2 deg. In the Mars B-plane figures that follow below, the 0.2 deg flight path angle uncertainty corridor is shown in black. After launch, a set of TCM-5 decision criteria were defined with the EDL team. These TCM-5 decision criteria are visible in the following figures as +/- 0.05 deg in flight path angle, as well as in the cross track direction 3.0 km to the south and 4.0 km to the north. If the OD solution fell within the green box, then no TCM-5 maneuver would be executed. The yellow box TCM-5 decision criteria was defined by +/- 0.1 deg in flight path angle and in cross track 3.0 km to the south and 6.0 km to the north. If the OD solution were within the yellow box but outside the green box then a TCM-5 would only be executed if the spacecraft was healthy. Outside the yellow box a TCM-5 would be executed regardless of spacecraft health.

Early Cruise

Immediately after launch the challenge was to isolate the outgassing signature from the other dynamics. Once this decayed to an imperceptible level, the largest source of uncertainty contributing to the propagated trajectory was in SRP model. At this time, data variations or estimating some of the other consider parameters had no significant effect on the OD solution. The filterloop variations were primarily focused on variations in the solar radiation pressure (SRP) model or dynamical models that could be incorrectly attributed to SRP (Figure 2). And the primary source of error mapped to the Mars B-plane was due to a reset of the assumed error at the time of the DCO. This was accomplished by modeling the SRP coefficients as stochastic parameters. Instead of a single parameter defined and estimated for the entire arc, a stochastic parameter is a time varying one with predefined model for the noise, e.g., random noise⁶. In the case of the SRP model parameters, it allowed for an estimate of that parameter during the data arc to be separate from the value and uncertainty of that parameter after the DCO. While the parameter was relatively well determined during the data arc (i.e., the uncertainty was small), the assumed error for the future propagation was reset to a larger value after the DCO.

Table 3. Postfit Residuals for Early Cruise Filterloop Cases (RMS).

Case	Doppler (mHz)	Range (RU)	Δ DOR (ps)
Baseline OD	0.845	1.429	23.267
Data: Doppler only	0.839	N/A	26.564
Data: range only	N/A	1.426	22.078
Lateral calibration scale sigma X 0.5	0.846	1.429	23.333
Lateral calibration scale sigma X 2.0	0.845	1.429	23.249
Lateral calibration scale sigma X 5.0	0.845	1.429	23.243
SRP cannonball	14.382	5.351	444.129
SRP baseline OD with scale factor	0.846	1.429	22.962
SRP fixed nominal model, estimate scale factor only	1.672	1.493	175.465
SRP fixed nominal model, estimate stochastic accelerations	0.802	1.427	24.771
Gasleak: sigmas loose in Z-direction	0.840	1.428	25.722

Consider error cases

SRP future error deleted
SRP & ACS consider error deleted

The DCO for the filterloop cases shown in Table 3 and Figure 2 was 05-Jan-2012, which was just prior to TCM-1. The OD solution with the largest ellipse is the baseline OD. It has 12 terms of the Fourier series used for the SRP to be estimated – 9 as bias terms and 3 as stochastics. However the stochastic batches were defined to be equivalent to a single bias term during the data arc and another stochastic batch after the DCO. The noise model was a random walk model, where the value estimated during the data arc is propagated after the DCO but with the covariance increased to reflect our increased uncertainty in the SRP model at that time.

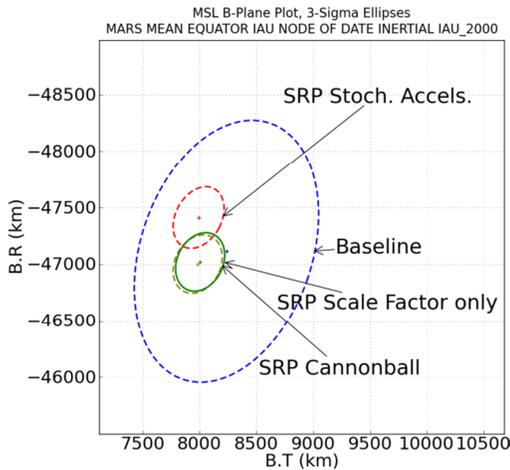


Figure 2. Early-cruise SRP cases.

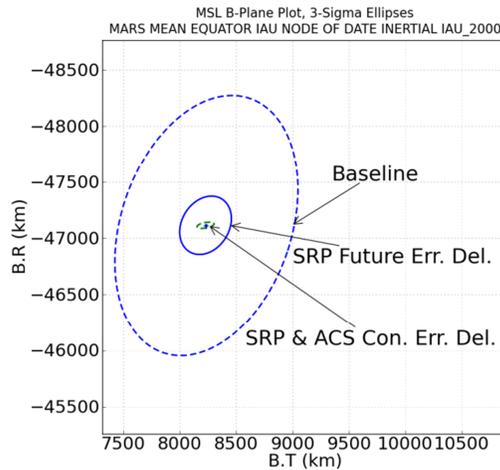


Figure 3. Early-cruise Consider Error Cases.

The other filterloop variations were attempts to significantly alter the SRP model. One variation was to fix all of these parameters to their a priori values and instead estimate a single parameter, solar pressure scale factor. First, to the lower left is the SRP model which fixed the nominal Fourier series coefficients and instead used a single scale factor estimate for the model – replacing 12 parameters with only one. The result was an OD solution giving a relatively poorer fit to the Doppler data and not really fitting the Δ DOR measurements at all. This result says in part that there needs to be some separability in the X and Z-components of the solar pressure model.

The coordinate system is defined such that the X-Z plane is in the Sun-Earth plane; therefore, in order to fit the Earth-line information provided by the Doppler data both components need to adjust. While the Δ DOR measurements provide more information about the out-of-plane direction, that is, the Y-components of the SRP model. A single scale factor is unable to fit this data. This illustrates the sensitivity of coupling between the X and Z terms of the SRP coefficients (the two directions which defined the Sun-Earth plane). Another example is to completely remove the SRP model parameters from the set of estimated parameters and introduce frequent, empirical, stochastic accelerations to fit the data arc

Another variation shown is simplifying the model even further and replacing it with a sphere or ‘cannonball’. This resulted in a trajectory similar to the single scale factor, but the postfit residuals indicate that this is a very poor model of the spacecraft. Finally, another filterloop variation fixed the SRP coefficients to their nominal values and instead allowed stochastic accelerations to fit the data. As is shown in the postfit residuals, this is designed to give a very good fit to the data. However, this is at the expense of learning anything about the SRP model. The trajectory is equivalent to allowing the unadjusted SRP model to propagate all the way to Mars; this is the solution up and to the left of the baseline. A key conclusion to be made is even given these extreme variations, it was clear that there was sufficient uncertainty prescribed to the SRP model after the DCO. It is this error source that is responsible for the large 3-sigma ellipse in the Mars B-plane for the baseline OD.

Figure 3 illustrates the main sources of error in filterloop variations for different consider error assumptions. The largest ellipse shown is the baseline OD. This was dominated by the error in the SRP model assumed after the DCO. Once this error source is removed, the second largest ellipse shows what is left, which was found to be the error considered for future Attitude Control System (ACS) activities, i.e., planned thrusting other than due to TCMs. Once the ACS thrusting error source is removed, the smaller error ellipse is left. This is the error remaining from to all other error sources, e.g., media effects, ephemeris errors, and spacecraft state error.

Mid-cruise

Mid to late cruise, the goal of the OD was to reconstruct TCM-1 and to prepare for the design and execution of TCM-2. Here we discuss the filterloop variations for a DCO of 22-Mar-2012, just prior to TCM-2. Table 4 is the postfit residuals for mid-cruise filterloop cases. In Figure 4, all the B-plane filterloop variations are shown. All cases but a couple of the poorly fitting SRP cases are well grouped around the baseline OD. In this instance, the largest ellipses are not due to the baseline OD but instead are SRP model variations. The baseline OD is among the tight group of ellipses in the center.

As in the early cruise cases, the error in the OD was dominated by the SRP model uncertainties. However, one of the key changes in the baseline OD assumptions was that the SRP model was updated. In addition to updating the values for all of the components, the filter setup was significantly simplified. Only three bias parameters were estimated: one for each direction, X,Y, and Z, with an a priori uncertainty of 5%, 1%, and 5% respectively. There was also a stochastic acceleration model in the Z-direction introduced to account for the heat generated by the RTG on the rover. Because of these changes, the error propagated beyond the DCO due to the SRP estimates had decreased considerably.

Table 4. Postfit Residuals for Mid-cruise Filterloop Cases (RMS).

Case	Doppler (mHz)	Range (RU)	Δ DOR (ps)
Baseline OD	0.829	6.163	30.276
Data: charged particle delay removed	1.113	6.148	40.808
Data: Doppler and Δ DOR	0.812	N/A	22.59
Data: Doppler and range	0.811	6.16	N/A
Data: fixed weights	1.504	7.82	26.087
SRP cannonball	2.133	7.377	1132.38
SRP fixed nominal model, estimate stochastic accelerations	0.772	6.111	14.38
SRP early cruise model	0.834	6.171	95.705
SRP a priori sigmas x 3	0.829	6.161	30.198
Thermal accelerations estimated bias with a looser a prior sigma	0.829	6.162	29.386
Thermal accelerations removed	0.86	6.198	128.249

Consider error cases

SRP future consider error removed
ACS future consider error removed
SRP & ACS future consider error removed

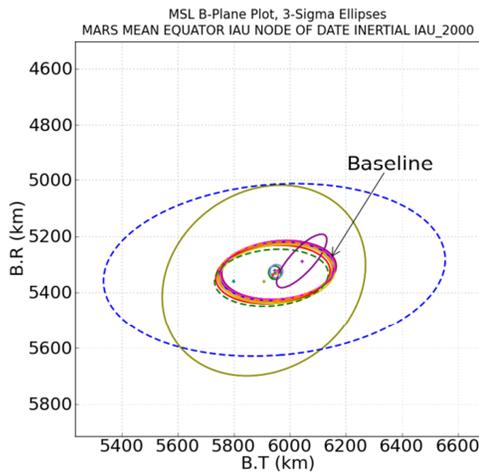


Figure 4. Mid-cruise All Cases.

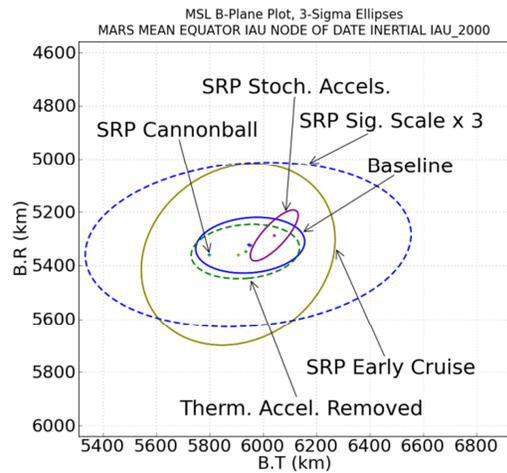


Figure 5. Mid-cruise SRP Cases.

The ellipses in Figure 5 show the filterloop SRP model variations. The largest ellipse (light blue) is the baseline OD with the a priori uncertainties for the SRP model increased by a factor of three. This was to test if the SRP model uncertainty was too tight. It resulted in almost a negligible improvement in the residuals and change in the trajectory. The next smaller ellipse (yellow) is the older, early cruise SRP model discussed previously. The resulting trajectory is within the 3-sigma uncertainty of the baseline OD. However, the fit to the data was significantly worse for the Δ DOR data, indicating the out-of-plane components were improved by the mid-cruise SRP model. The smaller ellipse comparable in size to the baseline by slightly offset is the case where the thermal model for the RTG was removed. While resulting in only a minor change in the trajectory, the postfit residuals showed that accounting for this effect was important for fitting the Δ DOR data. The two outliers from other are the cannonball model (green ellipse) and the empirical, stochastic accelerations (pink). The filter indicates that this model is well estimated – the error ellipse only shows up as a dot on this scale. This is a very poor model as indicated by the

data fit and it was soon after this OD that we were unable to converge the solution using such a model. At the other extreme, the model using frequent stochastics with 6-hr batches and 24-hr time constant in place of estimating the SRP coefficients resulted in a very good fit to the data. This is to be expected. The trajectory for this case is being propagated from the DCO to Mars entry using the uncorrected, baseline SRP model.

Only minor differences were observed in the trajectory of error ellipses for the cases that used different combinations of the data or data weights (Figure 6). The solutions shown are where we removed charged particle delay parameters, used Doppler and range data only, Doppler and Δ DOR data only, or fixed data weights instead of weights per pass. The primary conclusion being that the charged particle delay parameters were helping to improve the fit for some noisy data received during a period of high solar activity.

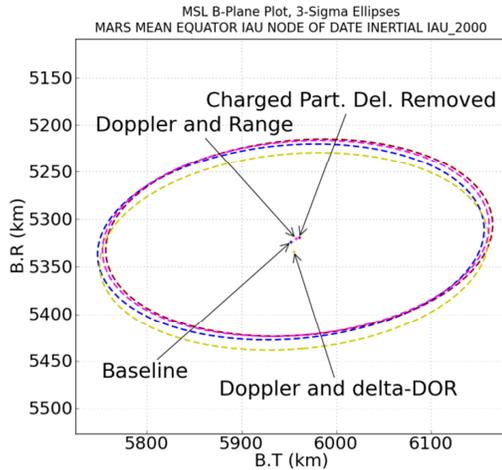


Figure 6. Mid-cruise Data Variation Cases.

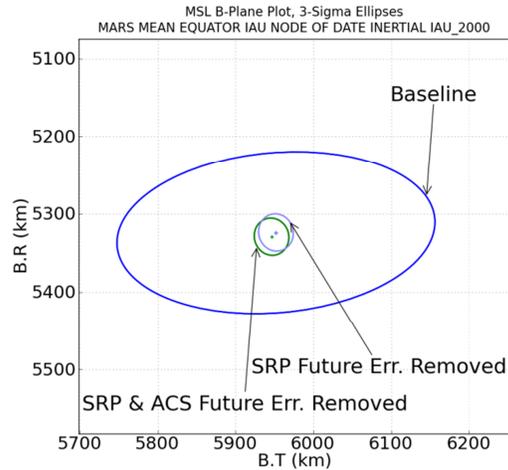


Figure 7. Mid-cruise Consider Error Cases.

The filterloop cases shown in Figure 7 were designed to isolate the main contributions to the propagated error. The largest ellipse is the baseline OD and it was dominated by the error in the SRP model. The next smaller ellipse is from the removal of the future error in the SRP model. The remaining ellipse came from removing the future ACS events, both their delta-V effect and their uncertainty. This illustrated that during this phase of the mission not only the uncertainties in the events, but the full magnitude of the delta-V due to non-TCM thrusting events had a very small effect relative to the much larger uncertainty in the SRP model.

Late-cruise: TCM-3

The late cruise phase covers all the maneuvers from TCM-2 through TCM-4. Here, we discuss the results of the filterloop cases with a DCO on 19-Jun-2012, just prior to TCM-3 (Table 5). We began to include variations on Earth orientation parameter (EOP) modeling at this phase of the mission. During late-cruise, the SRP model was also refined again. The basic model setup was similar in that it had the same number and type of estimated parameters. However, the nominal values were updated based on flight experience to date and the uncertainties were reduced relative to the mid-cruise model. Compared to the mid-cruise model discussed previously, the uncertainty was reduced by a factor of 5 in the X and Z-directions, and a factor of 2 in the Y-direction. As expected, this resulted in a much smaller error ellipse for the SRP filterloop cases (Figure 8). The largest ellipse (near-circular) is using the early cruise SRP model, the next small-

er ellipse is the mid-cruise model, and the smallest ellipse is the baseline model. At first it may appear that the models are in fairly good agreement; however, this is only after using tracking data to allow the model parameters to adjust. The late-cruise model was tuned to give better trajectory predictions, which was verified by passing new tracking data through the late-cruise nominal model. The late-cruise SRP model required less correction during the estimation process than the early and mid-cruise SRP models. This stability in the model is also illustrated by the other case in Figure 8. The other ellipse which is slightly bigger than the baseline is the case where the SRP model was fixed to the baseline model, that is, no estimated SRP parameters. Instead, frequent stochastic accelerations were used to fit the data: one-hour batch size, eight-hour time constant, with error equivalent to 10% of the total SRP acceleration. This shows that the late-cruise nominal model and the estimated model gave roughly the equivalent trajectories to Mars.

Table 5. Postfit Residuals for Late Cruise Filterloop Cases (RMS).

Case	Doppler (mHz)	Range (RU)	Δ DOR (ps)
Baseline OD	1.021	6.1	26.43
Frequent stochastic accelerations estimated	0.998	6.091	24.444
Arc epoch 01-Jun-2012	1.131	6.574	27.151
Arc epoch 06-May-2012	1.065	6.558	22.075
EOP Estimated as a bias – looser sigmas	1.019	6.099	26.203
EOP Estimated as stochastics	1.020	6.100	26.187
Mars ephemeris estimated as bias	1.021	6.100	26.43
Mars ephemeris estimated as stochastics	1.021	6.100	26.429
Media calibrations estimated as biases	1.009	6.099	25.018
Media calibrations estimated as stochastics	1.001	6.096	23.452
Thermal acceleration estimated as a bias	1.025	6.102	28.014
Thermal accelerations estimated as a bias – looser sigma	1.021	6.100	26.449
Thermal accelerations removed	1.025	6.102	28.325
Data: charged particle delay removed	1.175	6.105	28.472
Data: Δ DOR from E-W baseline only	1.012	6.099	21.322
Data: Δ DOR from N-S baseline only	1.019	6.100	28.410
Data: Doppler and Δ DOR	1.014	N/A	24.859
Data: Doppler and range	1.011	6.099	N/A
Data: elevation cutoff of 10 deg	1.289	6.125	29.200
Data: fixed weights	1.825	7.650	27.677
SRP early cruise model	1.022	6.169	27.776
SRP fixed nominal model, estimate stochastic accelerations	0.994	5.931	17.514
SRP mid-cruise model	1.021	6.108	26.592
SRP a priori sigmas x 3	1.021	6.100	26.231
Consider error cases			
SRP future consider error removed			
ACS future consider error removed			
SRP & ACS future consider error removed			

Figure 9 shows the filterloop variations using different epochs for the start of the data arc, but the same DCO: epochs of 01-Jun-2012, 06-May-2012, and the baseline of 10-April-2012. The shortest arc contained just over two weeks of tracking data and was statistically consistent with the longer data arcs. This gave us more confidence in the baseline data arc, but it was clear that a longer arc was needed to reduce uncertainties in initial state such that they were not the dominant error source.

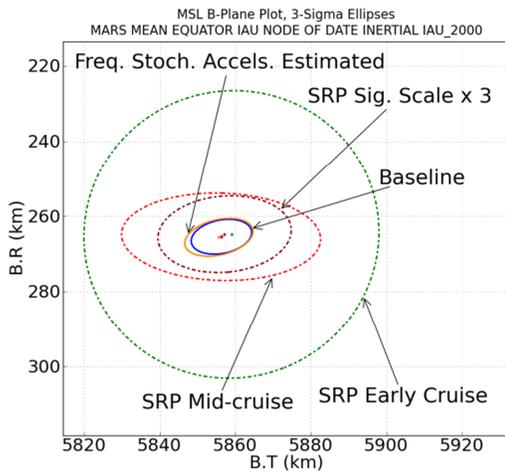


Figure 8. Late Cruise SRP Cases.

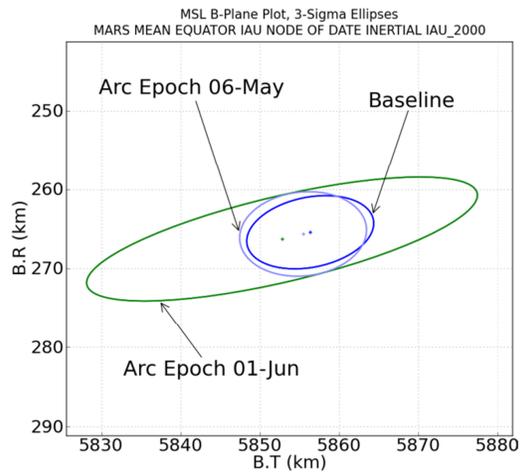


Figure 9. Late Cruise Data-arc Cases.

Now that the SRP errors were reduced significantly (relative to early and mid-cruise assumptions), we were able to observe the differences in the filterloop cases for varying data types and combinations. In Figure 10 are shown selected data variations that we used during late-cruise. The smallest ellipse is the baseline that used Doppler, range, and Δ DOR data. The next largest ellipse is Doppler and Δ DOR (range deleted) and the largest is Doppler and range (Δ DOR deleted). It was clear that at this phase of the mission the Δ DORs were a key measurement needed to reduce the OD uncertainties.

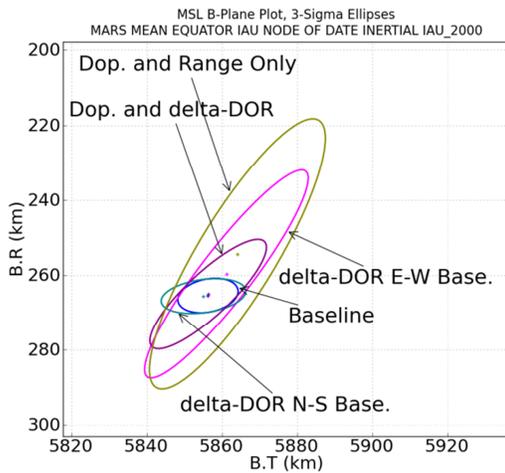


Figure 10. Late Cruise Data Type Variations.

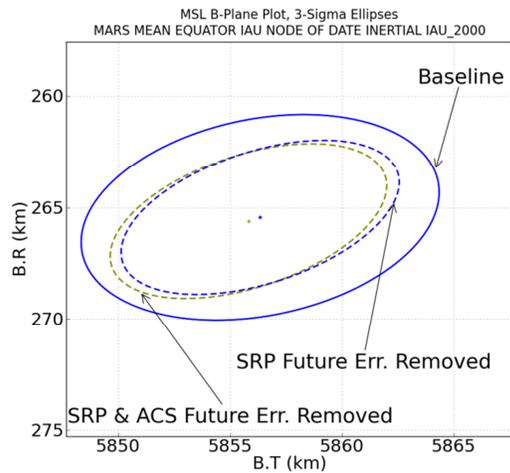


Figure 11. Late Cruise Consider Error Cases.

In Figure 11 are a couple of cases showing the dominating error sources, other than SRP model error, which contributed to the propagated error. The largest ellipse is the baseline OD. The smaller ellipse with the same midpoint as the baseline OD has the consider error of the SRP mod-

el after the DCO removed. The ellipse of similar size but shifted also has the same SRP model error removed. However, it also has future non-TCM thrusting activities removed – not just the associated uncertainties but the delta-V as well. This is to illustrate that the SRP model error was still the dominant error source in the dynamics, when ignoring the error introduced by future TCMs.

Late cruise: TCM-4

At the tail end of the late cruise phase was the design and execution of TCM-4. While there were placeholders in the schedule for TCM-5, TCM-6, and other contingency maneuvers, this was to be the final TCM for the mission. Here we discuss the filterloop variations of the OD with the DCO of 28-Jul-2012, the final OD for the design of TCM-4. The postfit residuals for the late-cruise filterloop cases are shown in Table 6. Along with cases similar to those used for the earlier phases, additional cases were added to study data variations as well as effects of consider parameters.

Table 6. Postfit Residuals for TCM-4 DCO Filterloop Cases (RMS).

Case	Doppler (mHz)	Range (RU)	Δ DOR (ps)
Baseline OD	1.359	5.058	33.273
Frequent stochastic accelerations estimated	1.321	5.047	32.627
Arc epoch 07-Jul-2012	1.399	5.128	31.472
Arc epoch 18-Jul-2012	1.455	5.376	28.309
EOP Estimated as a bias – looser sigmas	1.355	5.052	33.039
EOP Estimated as stochastics	1.357	5.056	33.128
Mars ephemeris estimated as bias	1.359	5.058	33.273
Mars ephemeris estimated as stochastics	1.359	5.058	33.272
Media calibrations estimated as biases	1.347	5.056	31.274
Media calibrations estimated as stochastics	1.340	5.050	29.091
Thermal acceleration estimated as a bias	1.360	5.058	33.684
Thermal accelerations estimated as a bias – looser sigma	1.359	5.057	32.464
Thermal accelerations removed	1.360	5.058	33.755
Data: charged particle delay removed	1.422	5.073	32.960
Data: Δ DOR sigma 30 ps	1.359	5.057	35.085
Data: Δ DOR sigma 60 ps	1.359	5.057	35.898
Data: Δ DOR sigma 240 ps	1.359	5.057	38.605
Data: Δ DOR from E-W baseline only	1.358	5.057	36.461
Data: Δ DOR from N-S baseline only	1.360	5.055	32.042
Data: Doppler and Δ DOR	1.353	N/A	32.781
Data: Doppler and range	1.359	5.055	N/A
Data: elevation cutoff of 10 deg	1.911	5.063	33.157
Data: fixed weights	3.105	5.506	33.355
SRP early cruise model	1.360	5.056	32.938
SRP mid-cruise model	1.359	5.059	32.356
SRP a priori sigmas x 3	1.359	5.057	32.175
Consider error cases			
SRP future consider error removed			
ACS future consider error removed			
SRP & ACS future consider error removed			

For scale, Figure 12 shows all of the filterloop variations for this OD. While difficult to see at this scale, all of plausible OD solutions from the filterloop variations are tightly clustered in the center of all those error ellipses. As opposed to the other phases in the mission, the data-type var-

iations were significantly larger than the dynamical model variations. What is primarily removed in Figure 13 shows the filterloop cases remaining after removing all of the data variations. This lack of sensitivity to dynamics models was due to a few factors at this phase of cruise: (1) an improvement in the dynamic models, in particular the SRP model, (2) with less time to propagate from the DCO to Mars, the uncertainties in the dynamic force models had less time to grow, and (3) frequent Δ DORs.

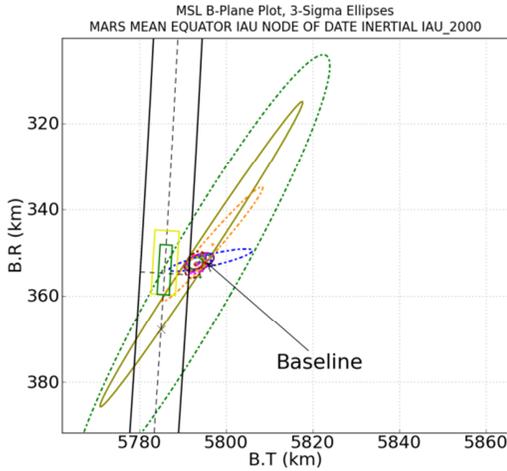


Figure 12. TCM-4 DCO All Cases.

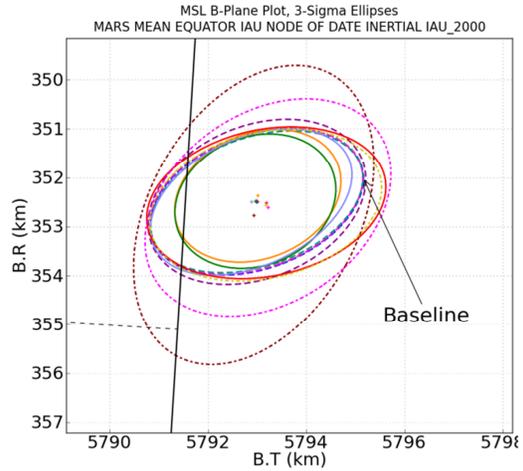


Figure 13. TCM-4 DCO Dynamic Model Cases.

The importance of the Δ DORs can be seen in Figure 14. These are compared to the baseline OD which used all available Doppler, range, and Δ DOR. Particularly important was to get more than one Δ DOR baseline. Also shown are selected filterloop cases using only East-West baseline Δ DORs or North-South Δ DORs. Other filterloop cases not specifically shown in this set of figures for late cruise, but noted in Table 6, are variations in Δ DOR weights, fixed weights for Doppler and range, and lower elevation cutoff.

The baseline OD at this time used a starting epoch of 26-Jun-2012, giving a total data-arc of approximately one month. Shown in Figure 15 are some of the shorter data-arc length variations with starting epochs of 07-Jul-2012 and 18-Jul-2012. Even a relatively short data arc gave results well within 3-sigma of the baseline. Not shown are some of the longer data-arcs which were being independently run at this time. They were also statistically consistent with these cases. However, those longer arcs did not use the same model or data assumptions as the baseline OD and were not strictly speaking simple variations in the filter relative to the baseline OD.

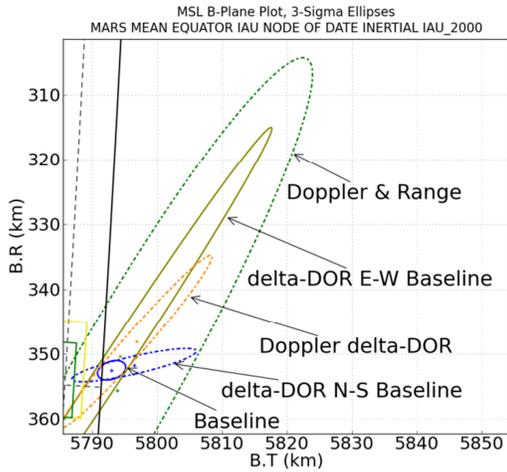


Figure 14. TCM-4 DCO Selected Data Type Variations.

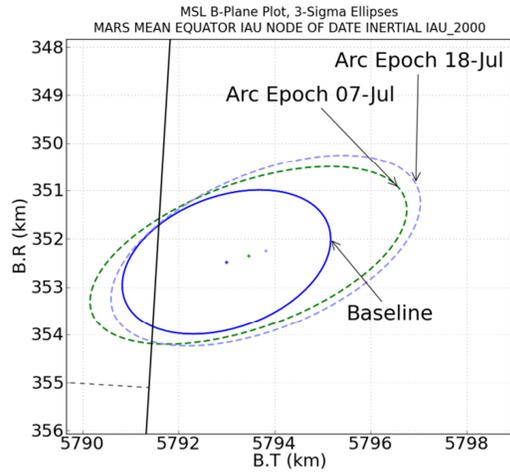


Figure 15. TCM-4 DCO Data-arc Cases.

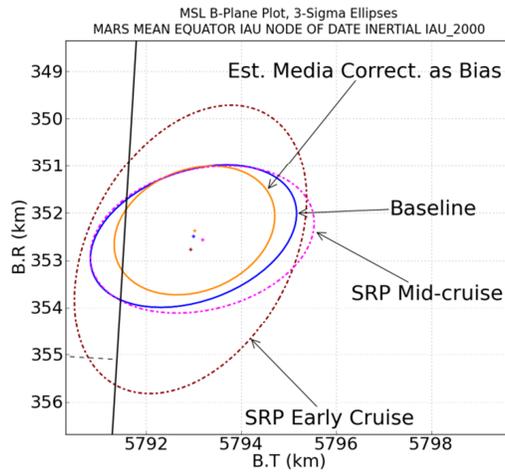


Figure 16. TCM-4 DCO SRP Model Cases.

As mentioned above, given the shorter time to propagate from the DCO to Mars, the dynamical force model uncertainties were becoming less important. In particular, the SRP model cases using early-cruise and mid-cruise models were all giving very similar trajectories (Figure 16). Only the uncertainty in the early-cruise SRP model case was significantly larger but that was a function of the uncertainty forced to be larger at the DCO. The SRP models during the data arc were relatively well determined and the consistency of the trajectories indicated that the OD was becoming relatively insensitive to which version of the SRP model was used (i.e., early, mid, or late cruise version). It was no longer the baseline assumptions for the SRP model which dominated the error in the error ellipse at Mars entry. It was at this time that the consider errors were now larger than any of the uncertainties in the dynamic force models, that is, the SRP, future thrusting events, or thermal radiations. Also in Figure 16, a filterloop case is shown which estimated the tropospheric and ionospheric media corrections relative to the baseline (the baseline

OD having these as consider errors). This does not suggest that the better solution would have been to estimate those parameters. The errors still needed to be treated as consider parameters in order to account for potential systematic errors in the media corrections. What this case and others like it showed was that the OD errors were no longer dominated by dynamics or state errors, but instead by consider errors that we were not estimating.

Approach

During the final approach, the OD solutions were very stable – the result of a highly accurate execution of TCM-4 as well as significant reduction in the uncertainties due to all predicted forces acting on the spacecraft. There were 40 different filterloop variations for the DCO approximately nine hours before entry – all listed in Table 7 and in shown Figure 17. All of the larger error ellipses are due to different data variations; the largest ellipse being the case using Doppler and range only. A couple of the cases with larger error ellipses highlighted the importance of having both North-South and East-West Δ DOR measurements (Figure 18).

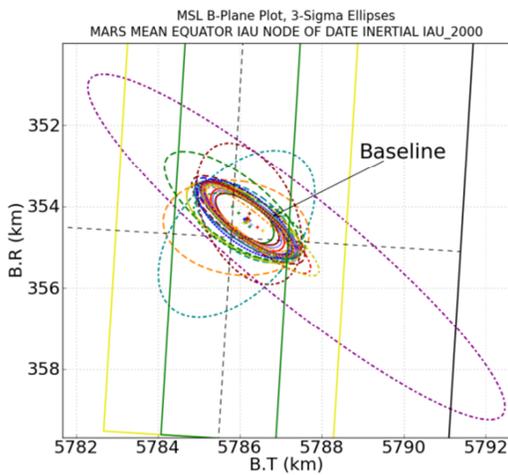


Figure 17. Approach All Cases.

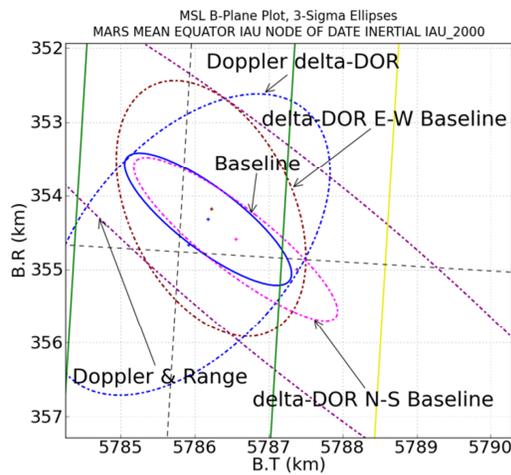


Figure 18. Approach Δ DOR Cases.

The predicted OD trajectories were very stable relative to all of the cases varying data weight assumptions (Figure 19). For context, note that the baseline OD was \sim 200 m from the trajectory used to help define EPU-1 – the onboard entry state to be used by the EDL system. The fixed data weight case, which gave the worst postfits, was only \sim 100 m off from the baseline OD trajectory. The weights in this case were tighter on average than the baseline OD using the weights defined per pass. The most significant variations from the baseline trajectory only occurred as the Δ DORs were increasingly downweighted. Even relatively short data arcs gave OD results comparable to the baseline data arc starting on 29-Jun-2012. As shown in Figure 20, the only outlier was for an extreme case using a data arc starting epoch of 01-Aug-2012, or approximately four days of tracking data. This was not enough data to resolve the large a priori errors assumed for the initial state and TCM-4 execution.

Table 7. Postfit Residuals for Approach Filterloop Cases.

Case	Doppler (mHz)	Range (RU)	ΔDOR (ps)
Baseline OD	1.381	4.839	35.523
Frequent stochastic accelerations estimated	1.343	4.83	34.592
Arc epoch 01-Aug-2012	1.753	3.175	38.321
Arc epoch 07-Jul-2012	1.419	4.811	34.626
Arc epoch 18-Jul-2012	1.463	4.775	34.523
Arc epoch 29-Jul-2012	1.587	4.085	42.575
EOP Estimated as a bias – looser sigmas	1.374	4.839	35.310
EOP Estimated as stochastics	1.379	4.839	35.406
Mars ephemeris estimated as a bias	1.381	4.839	35.493
Mars ephemeris estimated as a bias – looser sigmas	1.381	4.839	35.313
Media calibrations estimated as biases	1.366	4.839	34.507
Media calibrations estimated as stochastics	1.361	4.833	30.041
Thermal accelerations est. as stoch: 1-hr batches, 1-day time const.	1.379	4.839	35.178
Thermal accelerations est. as stoch: 12-hr batches, 7-day time const.	1.380	4.839	35.368
Thermal accelerations est. as stoch: 1-day batches, 7-day time const.	1.380	4.839	35.376
Thermal acceleration estimated as a bias	1.381	4.84	35.946
Thermal accelerations estimated as a bias – looser sigma	1.381	4.839	34.979
Thermal accelerations removed	1.381	4.840	35.967
Data: charged particle delay removed	1.451	4.855	35.130
Data: Δ DOR sigma 30 ps	1.381	4.84	38.142
Data: Δ DOR sigma 60 ps	1.381	4.839	38.749
Data: Δ DOR sigma 240 ps	1.381	4.839	40.396
Data: Δ DOR from E-W baseline only	1.380	4.839	39.507
Data: Δ DOR from N-S baseline only	1.381	4.839	31.822
Data: Δ DOR using secondary pairs	1.381	4.839	35.401
Data: Doppler and Δ DOR	1.376	N/A	35.121
Data: Doppler and range	1.381	4.838	N/A
Data: elevation cutoff of 10 deg	1.885	4.844	35.352
Data: fixed weights	2.933	5.268	35.451
SRP early cruise model	1.381	4.844	35.445
SRP mid-cruise model	1.381	4.842	34.906
SRP a priori sigmas x 3	1.381	4.839	34.811
TCM-4 looser a priori sigmas	1.381	4.839	35.501
TCM-4 tighter a priori sigmas	1.381	4.840	35.645

Consider error cases

All consider error removed
Media consider error removed
SRP future consider error removed
ACS future consider error removed
SRP & ACS future consider error removed

With such little time to go to before entry and all delta-V events after TCM-4 having been cancelled, the error in spacecraft dynamics was no longer a dominant error source as they were in the previous cruise phases. For example, the SRP model cases no longer showed any significant variation (Figure 21). This was due in part to an improved understanding of the SRP effects, but also due to the fact that the angle of the spacecraft relative to Sun was relatively stable about a mean of ~ 40 deg from TCM-2 onwards. The primary sources for error were now due to non-dynamical, consider error effects. This included the errors in Earth ephemeris, Mars ephemeris,

Earth polar motion and length of day, and tropospheric and ionospheric (media) calibrations, illustrated in Figure 22. The largest ellipse is the baseline OD. The slightly smaller and offset ellipse is the case with media corrections included in the estimated bias parameters. The two remaining ellipses are the cases removing the consider error of the media effects and removing all consider error effects. Either removing the error completely or estimating it had the same effect on the covariance: a reduction in the size of the error ellipse. Though we did not believe it would have been appropriate to estimate these parameters, only that any undetected bias in the media calibrations could affect the solution in this way. In hindsight, a more realistic estimate of the mean tropospheric error was 25% and the ionospheric error was 50% of what we used for operations³.

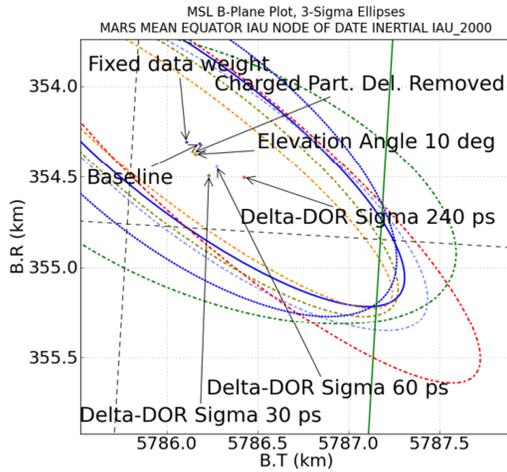


Figure 19. Approach Data Weight Variations.

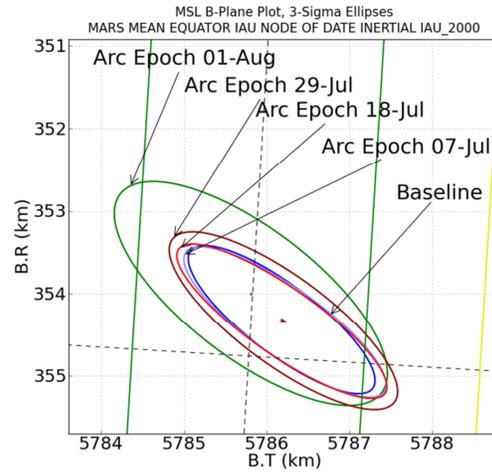


Figure 20. Data-arc Variations.

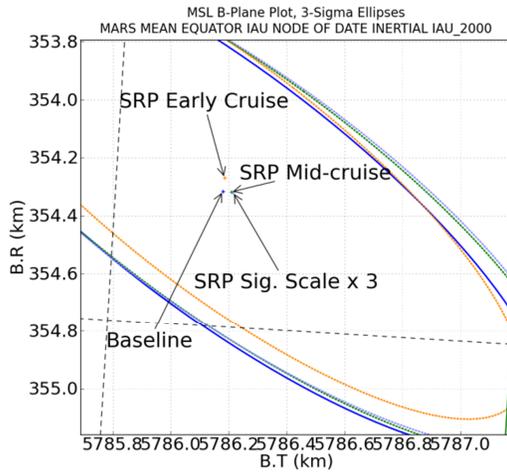


Figure 21. Approach SRP Model Variations.

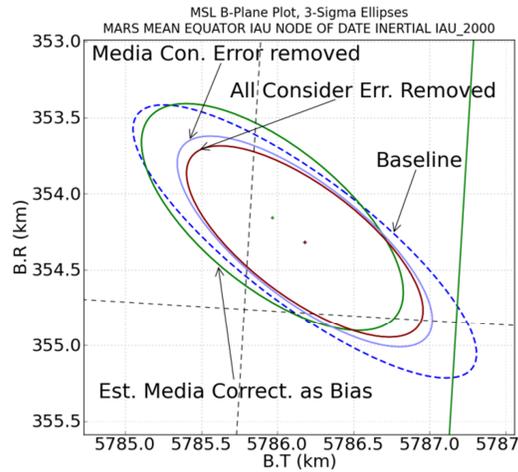


Figure 22. Approach Media Error Cases.

CONCLUSION

The MSL orbit determination was very successful and enabled the very precise landing of the Curiosity rover in Gale Crater on Mars. Throughout cruise, the orbit determination met all requirements with a considerable margin. The MSL OD team developed a spin signature removal tool, which successfully removed the spin signature and bias from the Doppler and range data. A novel approach was successfully used for modeling the solar radiation pressure via a Fourier expansion of the net solar radiation force. The solar radiation pressure model evolved into a highly accurate model by estimation of only three Fourier coefficients. Additionally, a stochastic acceleration in the direction of the rotation axis was used to model thermal radiation from the RTG. All trajectory correction maneuvers were successfully reconstructed and execution errors were found to be less than the assumed execution errors pre-flight. The delta-Vs associated with ACS turns were successfully calibrated in the beginning of the mission, but a decrease in delta-V magnitude was observed as cruise progressed. A stochastic scale factor for the ASC turn delta-V was estimated, which significantly improved the delta-V prediction accuracy for future ACS turns. For each cruise phase, the OD solutions showed statistical consistency as more tracking was included in the solution. Small systematic differences could be seen between solutions from different OD arc lengths but these differences were well within the 3-sigma B-plane uncertainties.

During final approach, the OD team provided one update (EPU-1) for the onboard entry state to be used by the entry guidance system. The OD solutions after EPU-1 were stable enough for the remaining the planned EPUs to be cancelled. The final reconstructed trajectory differed from the EPU-1 trajectory by only about 200 m in the B-plane and 0.11 m/s in velocity.

ACKNOWLEDGMENTS

The orbit determination team could not have done our work without the other members of the JPL MSL navigation team. We would like to acknowledge the work of the rest of the MSL navigation team, and in particular the other members of the MSL orbit determination team responsible for maintaining the baseline OD as well as running the filterloop scripts for the baseline OD: David C. Jefferson, Neil A. Mottinger, Frederic J. Pelletier, and Mark S. Ryne. We would also like to thank the Navigation Advisory Group, led by Joe Guinn, for discussions and suggestions often leading to the creation of additional, helpful filterloop cases.

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

- ¹ Kruizinga, G.L., E.D. Gustafson, P.F. Thompson, D.C. Jefferson, T.J. Martin-Mur, N.A. Mottinger, F.J. Pelletier, and M.S. Ryne, "Mars Science Laboratory Orbit Determination", 23rd International Symposium on Space Flight Dynamics. Pasadena, CA, USA, 2012.
- ² D'Amario, L.A., "Mars Exploration Rovers Navigation Results", AIAA 2004-4980, AIAA/AAS Astrodynamics Specialist Conference, Providence, RI, USA, 16-19 August 2004.
- ³ Martin-Mur, T.J., G.L. Kruizinga, and M. Wong, "Mars Science Laboratory Interplanetary Navigation Performance", AAS 13-232, 23rd AAS/AIAA Space Flight Mechanics Meeting. Kauai, Hawaii, USA, 10-14 February 2013.

- ⁴ Gustafson, E.D., G.L. Kruiizinga, and T.J. Martin-Mur, “Mars Science Laboratory Orbit Determination Data Preprocessing”, AAS 13-231, 23rd AAS/AIAA Space Flight Mechanics Meeting. Kauai, Hawaii, USA, 10-14 February 2013.
- ⁵ Sunseri, R.F., H.-C.Wu, S.E. Evans, J.R. Evans, T.R. Drain, and M.M. Guevara, “Mission Analysis, Operations, and Navigation Toolkit Environment (Monte) Version 040.” NASA Tech Briefs, Vol. 45, September 2012.
- ⁶ Tapley, B.D., B.E. Schutz, and G.H. Born, *Statistical Orbit Determination*, Elsevier Academic Press, Burlington, MA, 2004.
- ⁷ Folkner, W.M. “Planetary Ephemeris DE424 for MSL Early Cruise Navigation, IOM 343R-11-003.” JPL Inter-Office Memorandum, 2011.
- ⁸ Folkner, W.M. “Planetary Ephemeris DE425 for Mars Science Laboratory Arrival, IOM 343R-12-002.” JPL Inter-Office Memorandum, 2012.
- ⁹ Folkner, W.M. “Planetary Ephemeris DE423 Fit to MESSENGER Encounters with Mercury, IOM 343R-10-001.” JPL Inter-Office Memorandum, 2010.
- ¹⁰ Konopliv, A.S., C.F. Yoder, E.M. Standish, D.-N. Yuan, and W.L. Sjogren, “A Global Solution for the Mars Static and Seasonal Gravity, Mars Orientation, Phobos and Deimos Masses, and Mars Ephemeris.” *Icarus*, Vol. 182, pp. 23–50, May 2006. doi:10.1016/j.icarus.2005.12.025.
- ¹¹ Kinzer, W., “A Method of Describing Miss Distances for Lunar and Interplanetary Trajectories”, JPL External Publication No. 674, 1959.