

Mars Reconnaissance Orbiter Interplanetary Cruise Navigation[†]

Tung-Han You[‡], Eric Graat^{*}, Allen Halsell[§], Dolan Highsmith^{*}, Stacia Long^{*},
Ram Bhat^{*}, Stuart Demcak^{*}, Earl Higa^{*}, Neil Mottinger^{*}, Moriba Jah[#]

*Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-8099*

Carrying six science instruments and three engineering payloads, the Mars Reconnaissance Orbiter (MRO) is the first mission in a low Mars orbit to characterize the surface, subsurface, and atmospheric properties with unprecedented detail. After a seven-month interplanetary cruise, MRO arrived at Mars executing a 1.0 km/s Mars Orbit Insertion (MOI) maneuver. MRO achieved a 430 km periapsis altitude with the final orbit solution indicating that only 10 km was attributable to navigation prediction error. With the last interplanetary maneuver performed four months before MOI, this was a significant accomplishment. This paper describes the navigation analyses and results during the 210-day interplanetary cruise. As of August 2007 MRO has returned more than 18 Terabits of scientific data in support of the objectives set by the Mars Exploration Program (MEP). The robust and exceptional interplanetary navigation performance paved the way for a successful MRO mission.

I. Introduction

The 2180 kg Mars Reconnaissance Orbiter (MRO) was launched on 12 August 2005 from the Space Launch Complex 41 at Cape Canaveral Air Force Station (Figure 1). Carrying six science instruments and three engineering payloads, MRO is the first mission in a low Mars orbit to characterize the surface, subsurface, and atmospheric properties with unprecedented detail.

Similar to previous Mars missions, MRO consists of six mission phases – launch, cruise, approach and Mars orbit insertion (Approach-MOI), aerobraking and transition, primary science, and relay phases. A four-week launch period was selected with a 90-minute minimum launch window each day. The launch period extended from 10 August to 5 September 2005.

Following launch, MRO was first in-view from the supplemental acquisition station in Uchinoura, Japan in the Kagoshima prefecture. Three minutes after the separation, the Kagoshima station received the first angle tracking data. Twenty minutes after that, the Deep Space Network (DSN) at the Goldstone station acquired the first downlink signal. The navigation team delivered its first orbit solution three hours after injection. This solution enabled the successful signal acquisition of the follow-on DSN stations at Canberra, Australia and Madrid, Spain.

The MRO launch targets were biased due to the planetary protection requirements for Mars. The achieved injection parameters were within the 1σ level. The first interplanetary maneuver was scheduled at launch plus fifteen days to correct the separation errors, remove the injection bias, and adjust the Time-of-Flight (TOF).



**Figure 1: Atlas V-401
Launch of MRO**

[†] This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California under contract to National Aeronautics and Space Administration

[‡] Mars Reconnaissance Orbiter Navigation Team Chief, Jet Propulsion Laboratory

^{*} Senior Member of Engineering Staff, Jet Propulsion Laboratory

[§] Flight Path and Orbit Control Group Supervisor, Jet Propulsion Laboratory

[#] Senior Scientist, Oceanit Laboratories, Inc., 590 Lipoa Parkway, Suite 259, Kihei, HI, 96753

The subsequent maneuvers were designed to correct the navigation propagation uncertainty and maneuver execution errors.

Originally four interplanetary Trajectory Correction Maneuvers (TCMs) were planned, but only two were needed to meet the target requirements at Mars. Also, two contingency maneuvers were designed to divert from a Mars-impact trajectory, but were not needed. One contingency maneuver had two opportunities to execute prior to the Mars Orbit Insertion (MOI) while the second contingency maneuver had an execution opportunity after MOI.

After a seven-month cruise, MRO arrived at Mars. A real-time Doppler display unveiled a successful MOI maneuver of 1 km/s. MRO achieved a 430 km periapsis altitude with the final orbit solution indicating that only 10 km was attributable to navigation prediction error. With the last TCM performed four months prior to MOI, this was an unprecedented accomplishment.

This paper describes the MRO interplanetary cruise phase navigation. It summarizes the navigation requirements, spacecraft systems, dynamic models, measurement models, orbit determination strategy, and flight results. The implementation and performance of the TCM and MOI maneuvers will be discussed as well.

II. Navigation Requirements and Spacecraft Systems

A. Key Navigation Requirements

During the launch phase, Navigation's major responsibility was to provide trajectories or separation states to support initial acquisition and critical events monitoring. Continuous navigation support was critical through launch plus 24 hours. The orbit information updates in the first several hours required fast turn-around solutions. The updated knowledge was used to support the acquisition and handover at the subsequent ground-stations.

The MRO project required Navigation to place the Mars capture orbit within 50 km of the MOI targeted altitude, 491 km. Further, without executing additional maneuvers, Navigation needed to ensure that the post-MOI periapsis altitude was bounded within 300-500 km for the next eight consecutive orbits.

B. Spacecraft Systems

The MRO spacecraft consists of one structure subsystem (bus) and three main gimbal mechanisms. As shown in Figure 2, the skeleton of the bus supports all the science instruments and engineering subsystems such as




Figure 2: MRO Spacecraft Configuration

telecommunications, propulsion, command and data handling, guidance navigation and control, electrical power, and thermal systems. The gimbals, two each for the solar arrays (SA) and high gain antenna (HGA), have the capability of performing two-degree-of-freedom articulations. Each set of gimbals includes one inner and outer gimbal motor that allows independent orientations.

MRO utilizes a three-axis stabilized Attitude Control System (ACS) that relies on star trackers, Sun sensors, an Inertial Measurement Unit (IMU), and a Reaction Wheel Assembly (RWA) for attitude control. The spacecraft attitude measurements are provided via star trackers and Sun sensors. Between those measurements, it depends on the IMU for attitude estimation and propagation. The knowledge (i.e. on-board reconstruction) of the attitude information is temporarily stored on-board and played back through the engineering telemetry channels to the Ground Data System (GDS) when DSN communication is available. With the on-board reconstructed attitude quaternions and component gimbal angles, the navigation team is able to accurately model the spacecraft orientation.

The propulsion system operates in a blow-down mode for all thruster events except the MOI burn where a pressure regulator was used to improve the burn efficiency. A monopropellant system is used to reduce the complexity of the propellant management. As hydrazine was used during the cruise and approach phases, the tank pressure dropped to about 190 psi after Mars orbit insertion from a peak of 205 psi. To maintain a reliable thruster operation, a minimum of 100-psi tank pressure is necessary.

Table 1: Thruster Types and Specification

Specification	Thrust (N)	Number	Use
MR-107N	170.0	6	Main Engine Thrusters: TCM1 and MOI
MR-106E	22.0	6	TCM Thrusters: TCMs, MOI thrust vector control
MR-103D	0.9	8	ACS Thrusters: ACS and AMDs

There are three types of the thrusters used in operations. Table 1 summarizes their usages and configuration. The main engine thrusters were specifically designed for the MOI burn. To minimize the risk of first-time use of these powerful 170-Newton thrusters during orbit insertion, TCM1 was executed via the Main Engine as part of the risk reduction management.

Besides use on the flight path control, TCM thrusters were also employed to perform MOI thrust vector control.




Figure 3: Thruster Configuration

Figure 3 shows the locations of the thrusters. A_1 through A_8 are the ACS thrusters, T_1 through T_6 are the TCM thrusters, and M_1 through M_6 are the Main Engine thrusters.

In addition to performing the Angular Momentum Desaturations (AMD), the ACS thrusters help maintain the aerobraking drag attitude. Arranged in couples, the ACS thrusters are fired in pairs, so that ideally the resulting net ΔV is zero. Although it is designed as a balanced thruster system, a small amount of the residual ΔV for each thruster activity was anticipated during operation.

MRO relies on a X-band radio system to communicate with the Earth. The on-board antenna system includes a 3-meter diameter High Gain Antenna (HGA), and two Low Gain Antennas (LGA). Although the LGAs have the capability of transmitting and receiving data, Navigation primarily depends on the HGA to obtain its radiometric tracking data. The Small Deep Space Transponder (SDST) is capable of

supporting two-way and one-way radio operations. It either gets the reference frequency from an uplink signal source for a two-way link or obtains the reference frequency from the on-board Ultra Stable Oscillator (USO) or the

built-in auxiliary oscillator (AUX/OS) for a one-way link. In addition to the conventional tracking data capability, the SDST can generate the Differenced One-way Range (DOR) tones for VLBI observations.

C. Spacecraft Coordinate System and Orientation

As illustrated in Figure 2, the spacecraft body fixed coordinate system is defined such that +Z is along the normal direction of the nadir deck, +Y is along the center line of propellant tank and MOI thrusters, and +X is determined by the cross product of +Y and +Z. The origin is located at the geometric center of the Launch Vehicle Separation Adapter plate (i.e. at the end point of the propellant tank along the center line of the MOI-thruster module).

To ensure the continuous communication ability, the HGA is capable of tracking the Earth via its gimbaling motions. At zero-degree gimbaling angles (both inner and outer gimbals), the normal directions of the cell-side panels are along -Y axis and the long sides of the panels are along the $\pm X$ axes with each panel canted toward the +Z axis. Both the solar-array gimbals are capable of rotating along I_G (inner gimbaling rotation axis) and O_G (outer gimbaling rotation axis) to track the Sun. Illustration of the gimbaling rotation axes (i.e. I_G and O_G) is shown in Figure 2 (the directions denote positive rotations). The gimbaling locations are fixed relative to the spacecraft mechanical frame.

After the orbiter separated from the launch vehicle, through a series of appendage-deployment activities, MRO was settled in a preset inertial-fixed initial acquisition attitude. The orbiter +X was determined by the cross product of Sun and V_{inf} , -Y axis was pointed to the Sun with a bias of 15-degree rotation about -X axis, and the gimbaling positions of solar arrays and HGA, (I_0, O_0) , were $(0,0)$ and $(180, -45)$ degrees, respectively.

After spacecraft system checkout and the appendage deployments, approximately 3 days after launch, MRO transitioned from the initial acquisition attitude to the Sun-point cruise attitude. The attitude was configured such

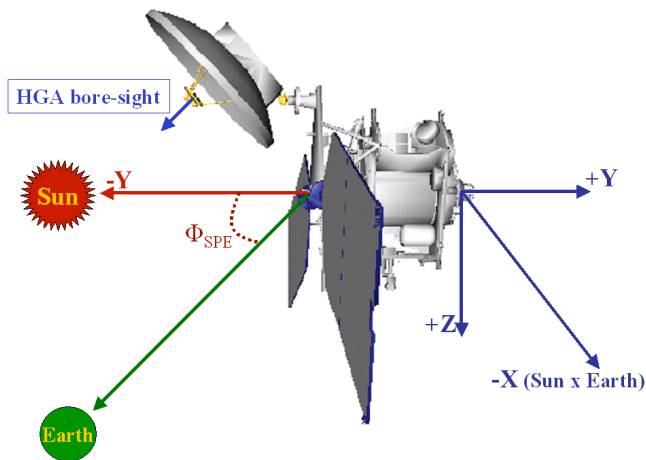


Figure 4: MRO Cruise Attitude

that the spacecraft -Y axis tracked the Sun, -X axis was determined by the cross product of the Sun and Earth, the solar gimbaling angles, (I_0, O_0) , were set to $(0,0)$, the HGA inner gimbaling angle, I_0 , was fixed at 180 degrees, and the HGA outer gimbaling angle, O_0 , was the result of the Sun-Probe-Earth angle, Φ_{SPE} , subtracting 90 degrees. This ensured the Earth always lied on the -YZ plane. In this configuration, the HGA could easily track the Sun via the outer gimbaling rotation and maintain the inner gimbaling angle constant. Figure 4 demonstrates the Sun-point cruise attitude.

Two months prior to the Mars Orbit Insertion, the spacecraft was articulated to a “spread-eagle” configuration as illustrated in Figure 2. Both of the solar-array gimbaling positions, (I_0, O_0) , were set to $(0,0)$ degrees. The HGA inner gimbaling angle, I_0 , and outer gimbaling angle, O_0 , were fixed at 180 degrees and -90 degrees respectively. The nominal

spacecraft attitude in this phase was defined as: -Y-axis was Earth pointed (i.e. HGA to Earth), the pointing direction of +Z was the vector-cross product of Earth X Sun, and +X was determined by +Y X +Z. This “spread eagle” configuration was used in approach and MOI phase but also during TCMS and AMD events.

D. Spacecraft Dynamic Models

MRO’s cruise trajectory was influenced by solar radiation pressure; by weekly AMD thruster firings; by small, short duration outgassing events due to the evaporation of residual substances; and by gravity. To characterize the AMD thruster behavior, in addition to the passive trending analysis, two active calibrations were conducted during the mid-cruise phase.

The gravitational model consisted of Newtonian point masses for the Sun, Earth, Moon, Mars, Phobos, Deimos and the remaining planets with an additional relativistic influence calculated for the Sun. JPL DE410¹, which contains the planetary ephemerides and constants, has been used for all mission phases. The associated Mars position uncertainty is less than a kilometer, which provides the best forecasted position of Mars at the time of MRO arrival. The ephemerides of Phobos and Deimos are important for considering close approaches and collision avoidance of these bodies. The total position uncertainties of the delivered Phobos and Deimos ephemerides are 8 and 35 km (1σ) respectively².

1. Solar Radiation Pressure

The solar radiation pressure (SRP) model developed for MRO consisted of a parabolic antenna, six flat plates to represent the spacecraft bus and two flat plates to represent the solar power arrays. The six plates were aligned to the spacecraft X, Y, Z axes, a pair representing the front and back. The antenna component could be either aligned with respect to the spacecraft X, Y, Z axes or could track a specified body (i.e. the Earth) while the solar array plates could be oriented with respect to the spacecraft-Sun vector. The spacecraft attitude was reconstructed from flight data with an attitude prediction modeled to Mars encounter. Areas for each component were calculated from spacecraft schematics provided by Lockheed Martin Space System (LMSS). Each component had a diffuse and specular reflectivity coefficient. Diffuse reflectivity imparts acceleration normal to the reflecting surface, while the specular term imparts accelerations both normal to and in the plane of the reflecting surface.

Early in flight three refinements were made to improve the solar radiation pressure model:

- The solar array diffuse reflectivity was recalculated using an “energy balance” derivation of the emissivity coefficient. An anomalous acceleration normal to the solar arrays was first thought to be due to outgassing. However, the decay signature did not match the expected pattern. Further investigation uncovered the source as thermal re-radiation from the solar array backside, which was roughly 27% of the total solar pressure acceleration. The front solar array emissivity was reduced from 0.85 to 0.12 and the associated diffusivity from 0.27 to 0.05 to account for this thermal re-radiation effect.
- Component areas were scaled to account for self-shadowing. Figure 5 illustrates the Sun-to-MRO view of the spacecraft, highlighting the shadowing of the spacecraft bus by the solar arrays during the Mars approach attitude. Software developed for the primary science phase was used to compute the “true” effective area. The corrections are expressed in terms of area scale factor and applied in the trajectory integration process.
- A small, 7th flat plate was added to better model the irregular shape of the spacecraft bus.

Figure 6 demonstrates the changes in overall SRP acceleration, 3.0×10^{-11} km/s² in total, due to each refinement. Comparison of the pre-launch baseline model (shown in red) to the post-launch improved baseline model (in blue) shows that the pre-launch baseline is about 34% over-estimated. Component self-shadowing accounted for about 11% over-estimated.




Figure 5: MRO Component Self-Shadowing




Figure 6: Cruise SRP Accelerations Comparison

2. Outgassing

Pre-launch analysis indicated that the maximum possible outgassing acceleration of 10^{-8} km/s² and would most likely occur soon after launch.

Assuming an exponential decay, it would decrease to 10^{-12} km/s² by mid-cruise. After launch, the maximum observed acceleration was near 10^{-10} km/s² at the MRO separation, decaying to 10^{-12} km/s² in less than a week (two orders of magnitude smaller than the pre-launch prediction). Figure 7 shows the post-launch navigation reconstructed accelerations due to outgassing.




Figure 7: Outgassing Accelerations during Launch Phase

However, these events did not have a long-term impact on the trajectory due to their short durations.

3. Thruster Calibrations

The MRO thrusters are coupled by design so in theory, no net translational ΔV was expected to be imparted to the spacecraft. However, since the thrusters could not be mounted and oriented perfectly, it was realistic to expect a net translational ΔV .

In order to be able to predict the motion of the spacecraft, well enough to meet prediction requirements, it was necessary to calibrate this ΔV . This was done by performing a thruster calibration early in the cruise phase of the mission. This calibration involved slewing the spacecraft to three mutually orthogonal attitudes, where typical AMD-type burns were executed and the resultant ΔV was measured. The orientation pointed each of the three spacecraft fixed axes (X, Y, Z in Figure 8) toward the Earth.




Figure 8: Thruster Calibration Attitude Configuration

$\pm X$). The RWA wheels were spun up and down in order to produce a positive and negative desaturation of each axis. Since the net ΔV vector was three-dimensional, three orthogonal attitudes were required in order to reconstruct it. The desaturation logic used in the thruster calibration was the same as that for a typical AMD. This means that the thruster pulse-width was fixed at 0.4 seconds per pulse and the time between pulses was fixed at 10 seconds. An average of 75 pulses were performed for each AMD.

During the calibration, the AMD ΔV s were one of several perturbing forces that MRO experienced. To distinguish the ΔV effect from the other forces (e.g. solar radiation pressure) a 10-minute quiescent time was placed between each burn and attitude change. Including these quiescent times, the total thruster calibration duration grew to approximately eight hours. High-rate Doppler data were collected throughout the thruster calibration activities via the LGA. RWA wheel speeds, obtained from telemetry, were used as an independent source of force reconstruction and were compared against the results obtained from the Doppler data. The other data type used in the reconstruction of the burns was the small force telemetry. Thus, the calibration event served the dual purpose of both determining the residual ΔV and validating the onboard computation of the small force events.




Figure 9: Thruster Calibration in Real-time Display (SC +Z at Earth)

Figure 9 shows the real-time Doppler display during thruster Calibration #1 (spacecraft +Z axis points to Earth). The residual should have been zero if the thrusters were balanced. In this case, the largest unbalanced ΔV occurred when desaturating the “-Z-axis”.

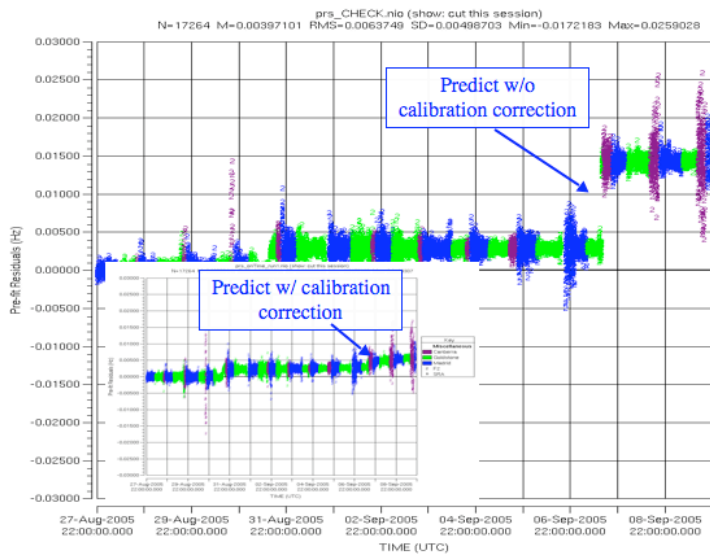


Figure 10: Doppler Residual With and Without Thruster Calibration Correction

Figure 10 shows the comparison of the predicted two-way Doppler residual based on small force predictions generated with and without the calibration data. It is clear that the residual improved significantly with the calibrated small force prediction (shown in the foreground insert).

Figure 10 shows the comparison of the predicted two-way Doppler residual based on small force predictions generated with and without the calibration data. It is clear that the residual improved significantly with the calibrated small force prediction (shown in the foreground insert).

4. Angular Momentum Desaturation Events

MRO’s attitude is maintained via the RWA. The RWA consists of three 100 Nms reaction wheels mounted perpendicular to each other with a fourth skew wheel for redundancy. Since the spacecraft is non-symmetric about at least one axis, a net torque was impressed on the orbiter due to solar radiation pressure. This solar torque was absorbed by the RWA. Because the reaction wheels have a limited allowable spin rate, the ACS 0.9 N thrusters are periodically fired to unload the accumulated angular momentum. These AMD events were accomplished by firing pairs of ACS thrusters when the angular momentum reaches ~40 Nms.

The small force file is used to model any resulting net ΔV from the ACS thruster activities. The navigation team queries the on-board reconstructed small force packets via the telemetry system. Key information included in the packet is pulse width, ΔV s (constructed by an on-board algorithm), pulse count, and quaternion. By incorporating the resulting ΔV s (per pulse) from the thruster calibrations, navigation produced updated on-board small forces.




Figure 11: Desat Pulse Count History




Figure 12: Reconstructed Average ΔV /Desat

These ΔV s are then used as nominal values in orbit determination. Through the estimation process, reconstructions of small forces were produced by navigation weekly for trending analysis. As shown in Figure 11, the average pulse count decreased from 35 to 10 a few weeks after launch. The high initial pulse count (i.e. 35) was mainly due to the post-launch activities and outgassing events. Figure 11 also shows that the AMD direction reversed from +roll (rotation about +X) direction to -roll due to a 180-degree spacecraft flip when the Sun-Probe-Mars angle approached zero in October 2005. Trending analysis, illustrated in Figure 12, shows that the average ΔV per AMD stabilized at 1.0 mm/sec by mid-cruise. During the interplanetary phase, the predicted AMD ΔV s were modeled as discrete events based on this trending information.

III. Flight Path Control And Orbit Determination

A. Maneuver Strategy

The cruise navigation operations strategy was based on the pre-flight OD covariance analyses, maneuver analyses, and dynamic model sensitivity studies. Based on these studies, four trajectory correction maneuvers were planned for the interplanetary phase to best condition the Mars orbit insertion trajectory. Additionally, a contingency maneuver (TCM5) was designed for risk reduction. TCM1 was designed specifically to remove the targeting bias and orbiter launch injection error. The target bias was introduced in the design to satisfy the planetary protection (i.e. Mars impact avoidance) requirement and to reduce the risk of main engine first time use. TCMs2-4

Table 2: Maneuver Epochs

Maneuver	Epoch	Purpose
TCM1	Launch + 15 days	Correct injection errors, remove injection bias, target to final aimpoint
TCM2	Launch + 99 days	Correct for OD and maneuver execution errors
TCM3	MOI - 40 days	Correct for OD and maneuver execution errors
TCM4	MOI - 10 days	Correct for OD and maneuver execution errors
TCM5	a	MOI - 24hrs
	b	MOI-6hrs
		Orbiter Safety only; no statistical components

were intended to correct for orbit determination uncertainties. TCM5 was an off-the-shelf contingency maneuver with two execution opportunities at MOI-24 hours and MOI-6 hours (TCM5a & TCM5b). The relative timing of the cruise maneuvers took into account the dynamical capability to change the Mars encounter at various times. The Type-I interplanetary trajectories for MRO did not have transfer-angle constraints or other singularities, and so the

TCM timing was based primarily on operational considerations and standard practice. The cruise TCM epochs are shown in Table 2. The relative timing of the TCMs did not change throughout the launch period.

B. Orbit Determination Strategy

This section briefly describes the radiometric tracking data types and the filter strategy used in the orbit determination (OD) process during MRO's 210-day transfer from Earth to Mars.

1. Radiometric Tracking Data Types

MRO utilized X-band Doppler, range and Δ DOR radiometric tracking data from NASA's Deep Space Network (DSN). The DSN is a global network of tracking stations located in Goldstone, California, Madrid, Spain, and Canberra, Australia. The tracking data coverage provided by the DSN is summarized in Table 3. The information content of each tracking data type is described next.

Table 3: DSN Tracking Data Coverage

Mission Phase	Key Events	Begin	End	Doppler Range	Δ DOR	Remark
Launch		L+000d	L+030d	Continuous		
	TCM1	L+015d		Continuous		Dual Track
Cruise		L+030d	M-060d	1pass/day	L+040d to M-060d: 1/wk	8 hrs/pass
	TCM2	L+099d		Continuous		\pm 3days, Dual Track
Approach & MOI		M-060d	M+007d	Continuous, X	M-060d to M-40d: 1-2/wk	
					M-040d to M-005d: 3/wk	
	TCM3	M-040d		Continuous, X		Dual Track
	TCM4	M-010d		Continuous, X		Dual Track
	TCM5ab	M-001d	M-006h	Continuous, X		Dual Track
	MOI	M+000d		Continuous, X		Dual Track

Doppler data directly measures the line-of-sight component of the spacecraft velocity while the spacecraft right ascension and declination can be inferred when collected over the length of a typical tracking pass (i.e. several hours). Doppler data units are expressed in hertz (Hz). One-way Doppler measurements are made when the spacecraft is the source of the signal's reference frequency. The frequency reference is usually provided by the AUX/OS or the USO. Unfortunately, the one-way Doppler data are prone to irregular biases and ramps. Two-way Doppler measurements are generated when a single tracking station radiates a signal to the spacecraft that in turn multiplies the received signal by a constant (turn-around ratio) and sends the signal back to the transmitting station. Since the two-way Doppler data are based on the station's extremely stable frequency reference, the biases and ramps that plague the one-way Doppler are not observed, making the two-way Doppler a "bedrock" data type for many interplanetary missions. Three-way Doppler data are formed in the same manner as the two-way Doppler but with a second station simultaneously recording the downlink signal.

Range measurements are the round-trip light time for a signal to propagate between a tracking station and spacecraft and measure the station-spacecraft distance. DSN range data are often expressed in so called "range units", or RUs, that at MRO's X-band frequency are equal to 0.142 meter.

Delta-Differential One-way Range, Δ DOR, is a Very Long Baseline Interferometric (VLBI) measurement technique that uses two widely separated tracking stations to simultaneously view a spacecraft and then simultaneously view an angularly nearby natural radio source (e.g. a quasar) to provide an angular position determination. For both the spacecraft and natural radio source, the difference in signal arrival time between the stations is measured. This time delay, coupled with the knowledge of the baseline joining the stations, provides a direct determination of the angle between the baseline and signal source. The Δ DOR observable is formed by differencing the time delay measurements from the spacecraft and natural radio source. This reduces the effects of common error sources, including station clock offsets, instrumental phase shifts and media path delays. Using the Goldstone-Madrid and Goldstone-Canberra baselines, the spacecraft's angular position in the plane-of-sky can be determined to better than 5 nanoradians.

Table 4: Tracking Data Observed Accuracy and Applied Weights

Data Type	Accuracy (1σ)	Weight (1σ)	Remark
Two-way Doppler	0.02 mm/s	0.10 mm/s	60 second compression
Two-way range	0.16 m	2.00 m	1.00 m = 7.04 RU
Δ DOR	1.17 nrad	3.37 nrad	37.45 nrad = 1.00 ns

During the launch phase Differenced One-way Doppler (DOD) measurements were generated from the one-way Doppler received by two different tracking stations. The DOD data remove the oscillator frequency drift along with other common errors and provide plane-of-sky velocity information on the spacecraft trajectory. Table 4 summarizes the observed tracking data accuracy and applied weights during interplanetary phase.

2. Filter Configuration

The JPL Orbit Determination Program’s (ODP) pseudo-epoch state weighted least-squares filter was used in the determination of MRO’s cruise trajectory. The baseline filter configuration estimated for the spacecraft epoch state, solar radiation pressure, thermal re-radiation along the spacecraft X and Z axes, AMD Δ V components, and per

Table 5: OD Filter Configuration

Parameter	Estimate Type	<i>A priori</i> uncertainty (1σ)	Remark
Epoch State Position	-	100,000 km	
Epoch State Velocity	-	10 km/s	
AMD Events	Constant Bias	30%	Per axis Δ V scale factor
Solar Radiation Pressure	Constant Bias	10%	Overall scale factor
	Stochastic	3%	Uncorrelated 2 day batches
Thermal Reradiation	Constant Bias	3.0×10^{-12} km/s ²	Along X and Z axes
Outgassing Events	Constant Bias	3.0×10^{-6} km/s	Estimation type and <i>a priori</i> would vary with the instrument calibration
	Stochastic	3.0×10^{-11} km/s ²	
Range Bias	Stochastic	2.0 m	Uncorrelated per pass
Earth, Moon & Mars GM	Consider	0.05 / 0.001 / 0.05 km ³ /s ²	
Station Locations	Consider	Full Covariance	
Quasar Locations	Consider	2 nrad	
Earth Pole X, Y and UT1	Consider	2 cm / 2 cm / 3 cm	
Ionosphere (day/night)	Consider	55 cm / 15 cm	S-band units
Troposphere (dry/wet)	Consider	1 cm / 1 cm	

tracking pass range data biases. The OD covariance considered the contributions of the following error sources: future AMD events, station locations, Earth and Mars ephemerides and gravity, troposphere, ionosphere, polar motion and UT1, and quasar locations (for Δ DOR data). Table 5 summarizes the standard filter configuration.

IV. Flight Results

A. Launch

Orbit determination during the launch phase of any mission is required to support the DSN in acquiring the initial signals from the spacecraft after launch vehicle separation. Prior to launch, the navigation team provides the DSN with predicted nominal and $\pm 3\sigma$ trajectories for the available launch dates and times. The DSN uses these trajectories to generate pointing and frequency predicts for the tracking stations supporting the launch. Though

these predicts may be sufficient for the first station signal acquisition, the angular position of the spacecraft must be known within ± 0.032 deg (the 34 m station -3 dB half-cone beamwidth) to ensure the second tracking station acquisition. It is the responsibility of the orbit determination analysts to provide the DSN with an updated trajectory based on the available tracking data.

For MRO launch on August 12, 2005 11:43:00 UTC, the available tracking data types included angles, range and Doppler. The Doppler data included one-way, two-way, three-way and differenced one-way (DOD). In addition to the DSN data, one-way Doppler tracking data was collected at the Uchinoura station that was then differenced with the DSN Goldstone station one-way Doppler data to form long baseline DOD data. The tracking data was updated every 30 minutes (or when requested) during the first 12 hours after launch. This wonderful menu of tracking data types allowed the OD analysts the opportunity to execute many solution variations to aid in the selection of the best trajectory for delivery to the DSN. For each tracking data update, solutions fitting one, two, three or all tracking data types were generated and compared to each other and to earlier solutions. The filter strategy was simple: estimate the spacecraft state (position, velocity) and any angle, range and DOD data biases while considering errors due to the media, station locations, solar radiation pressure and other non-gravitational accelerations (e.g. out gassing). At launch + 6.5 hours, a solution fitting DSN two-way Doppler, three-way Doppler, angles and range data was delivered to the DSN to update the Madrid station predicts. The pointing requirements at Madrid were easily met.

Table 6: Launch Vehicle Target Performance

Parameter	L.V. Target	Error in σ
C_3 (km^2/s^2)	16	-0.8
DLA (deg)	39	-0.9
RLA (deg)	29	0.3

The post-launch trajectory estimate was based on 11-hour two-way Doppler and range data collected after spacecraft separation from the upper stage. By the time of orbit determination, the trajectory estimates were stabilized to the point that the injection state was known to better than 20m in position and 0.7 mm/sec in velocity. The launch vehicle performance is tabulated in Table 6 in terms of orbiter energy (C_3), Declination

(DLA), and Right Ascension (RA). The error is presented as a sigma-level miss from the target.

B. TCM1

The final tracking data cut-off for TCM1 was on August 18, 2005. Based on six days of flight, the OD analysts had already started to refine the dynamic models. First, any long term, large (i.e. $> 10^{-11}$ km/s^2) non-gravitational accelerations due to out-gassing were ruled out. Next, the AMD prediction model and error assumptions could be calibrated against the two AMD events that had occurred, but significant refinements would have to wait until Thruster Calibrations 1 and 2. Lastly, the SRP model was reviewed. It was decided that a better understanding of the thermal balance of the spacecraft was needed, but this information was not available before the final TCM1 design tracking data cut-off.

The radiometric tracking data used in the TCM1 design solutions consisted of two-way Doppler, three-way Doppler and range. The filter strategies estimated the spacecraft state, the ΔV components of the two AMD events, a constant 3-axis non-gravitational acceleration and per pass range biases.

TCM1 was implemented as planned on August 27, 2005. The earth look angle with respect to ΔV was 58.7° , so 4.05 m/s of velocity change was directly observed through Doppler data. The total expected 2-way Doppler shift was 227.6 ± 22.1 (1σ) Hz. The observed Doppler shift (Figure 13) was approximately 223.2 Hz. The difference between observed and expected Doppler shifts (-4.5 Hz) indicated that the TCM1 performance was clearly within expectations. The change in slope around the loss of lock reflects the settling and main burns.

Table 7: TCM1 performance

Parameter	Design	Actual	Error in σ
ΔV (m/s)	7.79	7.80	0.16
Right Ascension (deg, EME2000)	128.10	130.31	0.83
Declination (deg, EME2000)	65.20	65.43	0.21
MOI Target Altitude (km)	395km	3800km	0.85

The achieved result with its covariance is mapped to the Mars B-plane to visualize the target uncertainties. It is clear that from Figure 14 the difference between the designed and achieved TCM1 target is somewhat less than 1 sigma. The B-Plane system is a convenient way of expressing errors at the target planet as a linear function of errors in the orbit. This is known as R-S-T coordinate system. B is the target parameter, S is the direction of the incoming nominal trajectory asymptote, T is the vector parallel to the ecliptic plane and perpendicular to S, and R is the vector normal to S-T plane directed south.

TCM1 performance was within 1 sigma of the expected value (Table 7). The reconstructed delta-V magnitude is 8.4 mm/sec greater than the design, the reconstructed attitude differs by 0.96°, and the target altitude is off by 3800 km.

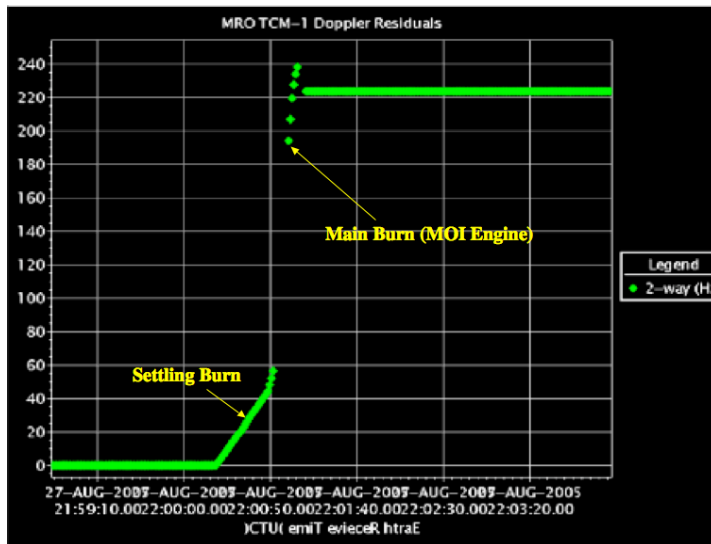


Figure 13: TCM1 in a Real-Time Doppler Display

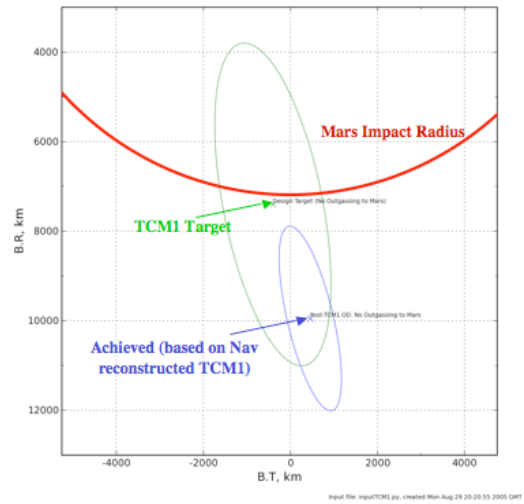


Figure 14: Encounter B-Plane Error Ellipses (1-6)

C. TCM2

The final tracking data cut-off for TCM2 was on November 11, 2005. Starting from September 20, 2005, Δ DOR measurements augmented the two-way Doppler and range tracking data. Also by this date, two significant improvements to the dynamic models were incorporated. First, the solar radiation pressure model was updated: the solar power array diffuse reflectivity was recalculated using an “energy balance” derivation of the emissivity coefficient and spacecraft self-shadowing was now accounted for. Second, the AMD ΔV reconstruction and prediction models were updated based on the results of Thruster Calibrations 1 and 2. These dynamic model improvements would allow for the cancellation of TCM3 and TCM4.

The filter strategy at the TCM2 design had transitioned to using the automated solution generation methods developed during the Mars Exploration Rover (MER) mission³. The method of automating the generation of multiple solutions based on a single tracking data arc was achieved via a Perl script christened “*filter loop*”. By the TCM2 final design, three tracking data arc lengths had been selected and the OD analysts had defined 100 solution variations (cases) for execution in *filter loop*. There were eight tracking data type combinations, or series, as summarized in Table 8. For each of the listed series, solutions varied the tracking data weights, editing criteria, and parameter estimation strategies in the filter.

The three tracking data arcs analyzed were: 1) a long arc with epoch on 16-SEP-2005, 2) a medium arc with epoch on 12-OCT-2005 and 3) a short arc with epoch on 03-NOV-2005. Excluding the prediction cases, 261

solutions were generated on the date of the TCM2 tracking data cut-off. When mapped to the Mars B-plane, the solutions displayed consistency across data arcs, estimation strategies and data type variations. The final TCM2 design was based on a medium arc solution that fit two-way Doppler, range and Δ DOR data.

Table 8: Filter Loop Strategy

Data Type Series (E-W: East-West Baseline, N-S: North-South Baseline)										
	Series100	Series200	Series300	Series400	Series500	Series600	Series700	Series800		
Doppler	√	√	√	√	√	√	√	X		
Range	√	√	X	X	√	√	X	√		
Δ DOR	√	X	√	X	√ (E-W)	√ (N-S)	√	X		
Data Weight and Editing Variations for Series100 to Series800 (ns: nano-sec)										
Data Weight Variations					Data Editing Variations					
	Standard	Tight	Loose	Open	Case-1	Case-2	Case-3	Case-4	Case-5	Case-6
Doppler	5.63 mHz	2.81	56.3	563	Minimum 20deg elevation	Ignore data after a specified epoch	Remove Goldstone two-way data	Remove Canberra two-way data	Remove Madrid two-way data	20m/pass <i>a priori</i> uncertainty for range data
Range	2.0 m	1.0	20.0	200						
Δ DOR	0.09 ns	0.065	0.03	1						
Dynamic Parameter Filtering Variations										
Case-1	Case-2	Case-3	Case-4	Case-5	Case-6					
Add SC Y-axis acceleration to the baseline	Loose acceleration <i>a priori</i> Uncertainties (x5)	Tight Acceleration <i>a priori</i> uncertainties (/5)	No stochastic solar coefficient estimated	150% uncertainty on the AMD scale factors	300% uncertainty on the AMD scale factor					

Table 9: TCM2 performance

Parameter	Design	Actual	Error in σ
Δ V (m/s)	0.751	0.754	0.14
Right Ascension (deg, EME2000)	37.0	37.1	0.14
Declination (deg, EME2000)	12.3	12.1	0.23
MOI Target Altitude (km)	491	507	0.27

TCM2 was implemented on 18-Nov-2005 at 22:00:00 UTC. Six TCM thrusters and eight RCS thrusters were used for attitude control during the burn. It was expected that most of the effect of the maneuver would be visible through Doppler shift, as the earth look angle was 160.7 degrees. The expected Doppler shift was 39.94 Hz. The observed Doppler shift was 40.0 Hz. The difference between observed and expected Doppler shift, +0.105 Hz, indicated that the TCM2 performance was nominal.

A 67-day tracking data arc was used to reconstruct the TCM2 maneuver thrust and direction. The tracking data included post-TCM2 Δ DOR measurements. This Δ V estimate was consistent with the observed Doppler shift and telemetry data. The achieved results were mapped to the Mars B-plane to compare the achieved with the designed target parameters. Figure 15 shows the difference between the designed and achieved TCM2 target. TCM2 performance was well within 1σ of the expected value. In Table 9, the reconstructed Δ V was 0.754 m/s, which is 3 mm/s higher than the designed Δ V of 0.751 m/s, the burn attitude is off by 0.14 degrees, and the MOI aimpoint differs by 16 km.

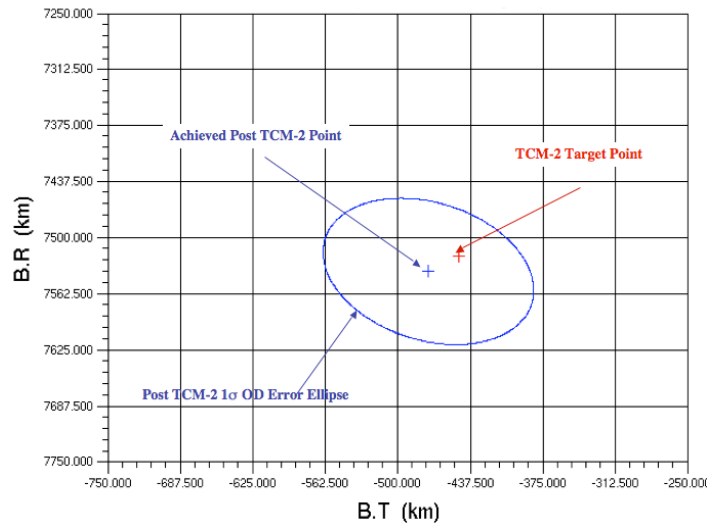


Figure 15: TCM2 Performance (1- σ)

D. TCM3 and TCM4 Cancellation

Since the performance of TCM2 placed MRO on a trajectory within 1σ of the Mars encounter target, the need for TCM3 depended on the subsequent movement of the OD solutions in the B-plane. If the solutions should drift outside of a prescribed region in the B-plane, a TCM would be required to place MRO back on target. This region was defined by ± 50 km about the target altitude and ± 0.3 degrees about the target inclination.

The tracking data cut-off for the final design of TCM3 and TCM4 were January 19 and February 16, 2006

respectively. Five tracking data arcs were analyzed for TCM3: 1) a 5-month long arc, 2) a medium arc with 4-month data arc, 3) a short arc with 2.5 month data arc, 4) a post-TCM2 data arc, and 5) a 1-week Approach phase data arc. Four tracking data arcs were examined for TCM4: 1) a nearly 3-month long data arc, 2) a 1-month medium data arc, 3) a 2-week short data arc, and 4) a 1-week very short arc. On the dates of the tracking data cut-off, 432 TCM3 and 129 TCM4 solutions were generated and analyzed respectively. The solutions displayed consistency across data arcs, estimation strategies and data type variations. Figure 16 shows the B-plane history of post-TCM2 solutions with the altitude and inclination bounds marked.




Figure 16: Post-TCM2 Mars B-Plane 1σ Error Ellipses

MOI aim-point were 28 km in B.R, 17km in B.T, and 9 second in Time of Closest Approach (TCA). The errors were on the level of 1-sigma. This resulted the cancellation of TCM3. If the TCM3 had been designed to hit the aimpoint, it would have been 15 mm/s in ΔV .

Table 10: Orbit Determination Solutions Since OTM2 (1σ)

Solution	B•R (km)	B•T (km)	TCA: March 10, 2006 utc	MOI Target Altitude
MOI Target	7519.3	-448.4	21:29:31.0	491 km
Post-TCM2	7536.5 \pm 78.6	-474.5 \pm 86.6	21:29:14.2 \pm 44.8 sec	507 km
Final TCM3	7547.2 \pm 23.3	-464.0 \pm 24.0	21:29:22.3 \pm 9.1 sec	514 \pm 19 km
Final TCM4	7550.7 \pm 5.3	-470.8 \pm 4.9	21:29:24.9 \pm 1.8 sec	518 \pm 4.3 km
Final TCM5a	7551.3 \pm 0.5	-470.6 \pm 0.4	21:29:24.8 \pm 0.1 sec	518 \pm 0.4 km
Final TCM5b	7551.0 \pm 0.4	-470.7 \pm 0.3	21:29:24.8 \pm 0.1 sec	518 \pm 0.4 km

TCM4 was also cancelled, because the 3-sigma OD error ellipse was within the target corridor. The decision to cancel TCM4 was based on the medium arc solution that fit two-way Doppler, range and ΔDOR data. Table 10 shows that the 3-sigma MOI target altitude for the final TCM4 solution was bounded within 505 - 531 km (i.e. 518 \pm 13 km, 3-sigma), which was well within the MOI target requirement (i.e. 491 \pm 50 km).

E. TCM5 Go/No-Go Decision and Mars Orbit Insertion

The need to perform the TCM5 contingency maneuver was based on the predicted altitude of the second Mars periapsis (P2) after the MOI burn. It was designed solely to divert from a Mars impact trajectory. Since the orbiter’s safety and health was the only concern, there was no intent to use the TCM5 maneuver for trajectory correction or optimization. To reduce the workloads of implementation and sequence tests, “off-the-shelf” (i.e. pre-designed and verified) maneuvers shared one burn attitude and one ΔV file. Opportunities for TCM5 were scheduled at MOI–24 hr (TCM5a) with a 4-hour execution window and MOI–6 hr (TCM5b) with a 2-hour execution window. Each of the opportunities had three ΔV selections and they were capable of raising P2 altitude from 80 km to 350 km⁴.

To derive the P2 target altitude, a baseline minimum P2 altitude was identified at 200 km, outside of the sensible Martian atmosphere. Other factors such as ΔV s from fault protection, excessive or insufficient ΔV s from the MOI timer cutoff strategy, 3-sigma engine over-burn, and MOI interruption and restart analysis brought the minimum P2 altitude needed prior to MOI to 340 km. Further, taking into account the navigation margin of 50 km, the P2 target altitude was rounded up to 400km. Due to the drop in altitude caused by the pitch-over MOI burn, the MOI altitude was re-targeted to 491 km at TCM2, which was an increased of 96 km from the original TCM1 target of 395 km.

Prior to the data cutoffs, the navigation team provided daily assessments to the MRO project starting from TCM4. Key elements in these daily assessments included trends of the dynamics models, filter variations, and solution consistency. The decisions on whether to perform the TCM5a/b maneuvers were based on OD solution updates having tracking data cut-offs of MOI-32 hr and MOI-13 hrs, respectively. Figure 17 shows the sensitivity of TCM5a solutions due to the data-type variations. Figure 18 shows the data-arc sensitivity. As seen in the figures, the solutions were robust and consistent. Through the navigation review process, one of the TCM5 solutions was

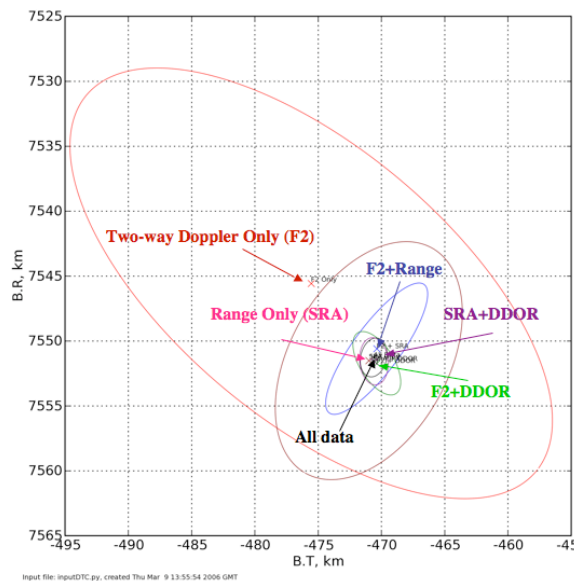


Figure 17: TCM5a Data-Type Solution Sensitivity

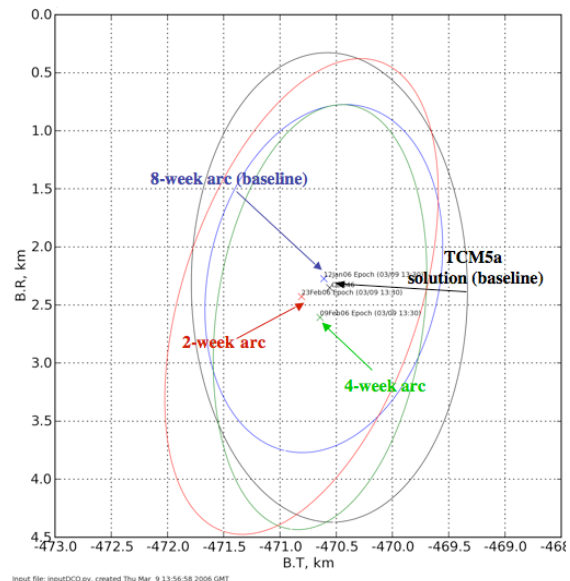


Figure 18: TCM5a Data-Arc Solution Sensitivity

selected to generate the “minimum solution” for “Go/No Go” process. To produce the “minimum solution”, the navigation 3-sigma uncertainty was subtracted from the approved solution. One condition needed to be satisfied in order to waive the contingency maneuver: TCM5 would be executed if and only if the “minimum solution” indicated the periapsis altitude was less than 340 km at P2. As Figure 19 illustrated, the approved TCM5a solution at P2 was 426 km, therefore the minimum solution at P2 was 425 km (since the 3-sigma navigation uncertainty was 1.2 km). This, with the 85 km margin, satisfied the TCM5a cancellation criteria. A similar process was performed for TCM5b and resulted in its cancellation as well.

Table 11 compares the MOI final design with the post-MOI reconstruction. The final design was based on post-TCM2 solution, which was implemented 3.5 months before MOI. The achieved MOI ΔV was only 0.01%

percentage short from the design, the P2 altitude was off by 1 km, the burn duration disagreed by merely 35 sec, and the inclination did not change from the prediction. The classic performance laid an excellent foundation for the aerobraking and the subsequent mission phases.




Figure 19: TCM5a P2 Altitude Margin

Table 11: Mars Orbit Insertion Maneuver

MOI Solution	ΔV (m/s)	Duration (sec)	Orbit Period (hr)	Inclination (deg)	MOI Target Altitude (km)	P2 Altitude (km)
Final Design	1000.4	1606	35	93.5	491	427
Achieved	1000.5	1641	35.5	93.5	518	426

V. Conclusion

Since the completion of the Mars Orbit Insertion, MRO has successfully implemented the aerobraking operation and established the primary science orbit. As of August 2007, it has returned over 18 Terabits of scientific data in support of the objectives set by the Mars Exploration Program. This paper describes the orbit determination and flight-path control strategy used to meet the project requirements. No other planetary missions nailed the MOI target at TCM2, thereby providing the MRO project much-needed time to get ready on the Mars Orbit Insertion operation.

The MRO navigation team has gained significant experience in improving the dynamic modeling of the component self-shadowing, the thermal imbalance, and the small force trending during the interplanetary phase. The sophisticated filter configuration and strategy further increased our confidence in evaluating navigation solutions and decision-making. The textbook performance of the interplanetary navigation not only paves the way for a successful MRO mission, but also extends its experience to future interplanetary missions.

References

- ¹ E. M. Standish, JPL Planetary Ephemeris DE410, JPL IOM 312.N-03-009, April 24, 2003.
- ² Jacobson, R.A. (1989) "The orbits of the satellites of Mars from spacecraft and Earth based observations", *Astronomy & Astrophysics* 225, 548.
- ³ McElrath, T. P., et al., "Mars Exploration Rovers Orbit Determination Filter Strategy", AIAA/AAS Astrodynamics Specialist Conference, AIAA-2004-4981, Washington DC, 2004.
- ⁴ T. H. You, Allen Halsell, et al., "Navigating Mars Reconnaissance Orbiter: Launch Through Primary Science Orbit", Space 2007 Conference & Exposition, AIAA-2007-6093, Long Beach, California, 18-20 September 2007.