Navigation Challenges of the Mars Phoenix Lander Mission

Brian M. Portock[1], Gerhard Kruizinga[2], Eugene Bonfiglio[†],
Behzad Raofi[†], Mark Ryne[†],
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 91101

The Mars Phoenix Lander mission was launched on August 4th, 2007. To land safely at the desired landing location on the Mars surface, the spacecraft trajectory had to be controlled to a set of stringent atmospheric entry and landing conditions. The landing location needed to be controlled to an elliptical area with dimensions of 100km by 20km. The two corresponding critical components of the atmospheric entry conditions are the entry flight path angle (target: -13.0 deg ±0.21 deg) and the entry time (within ±30 seconds). The purpose of this paper is to describe the navigation strategies used to overcome the challenges posed during spacecraft operations, which included an attitude control thruster calibration campaign, a trajectory control strategy, and a trajectory reconstruction strategy. Overcoming the navigation challenges resulted in final Mars atmospheric entry conditions just 0.007 deg off in entry flight path angle and 14.9 sec early in entry time. These entry dispersions in addition to the entry, descent, and landing trajectory dispersion through the atmosphere, lead to a final landing location just 7 km away from the desired landing target.

I. Introduction

THE Phoenix Lander is currently on the surface of Mars carrying out its exciting science mission, characterizing the chemistry of the materials at the local surface, subsurface and atmosphere. There exists the possibility of identifying organics amongst water ice that may point to the biological potential of Mars. In order to fulfill the science goals of the mission a landing region was chosen in the northern polar region of Mars where water ice was previously detected by the Mars Odyssey science instruments. Once a scientifically interesting region was selected, a specific landing area needed to be identified that met the strict safety requirements of the entry, descent, and landing (EDL) system. The key drivers for landing site safety were surface slopes and surface rocks. The landing site safety concerns required that Phoenix land in an area the shape of an ellipse with dimensions 100km x 20km centered at 234.292 degrees planetocentric east longitude and 68.109 deg. latitude.

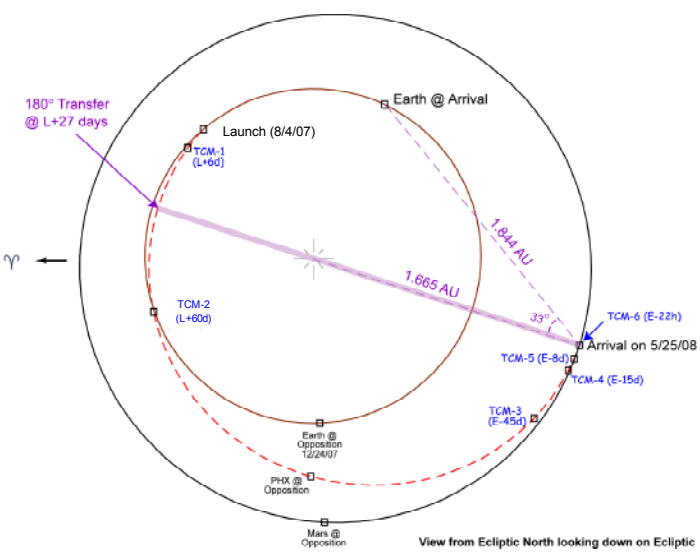

Figure 1. Phoenix Interplanetary Cruise Trajectory

The Phoenix spacecraft, including entry vehicle, cruise stage, and Lander, was launched from Kennedy Space Center on August 4, 2007. The launch targets and trajectory were designed to meet a specific launch energy or C3 of 29 km^2/s^2, while at the same time allowing for a 90-sol prime surface mission and the possibility of a 150-sol extended mission. The surface lifetime requirement constrained the Phoenix interplanetary

[1] Phoenix Navigation Team Chief, Guidance Navigation and Control Section, M/S 264-623.
[2] Senior Member of Engineering Staff, Guidance Navigation and Control Section , M/S 264-623.

trajectory and Mars arrival time to be on or before May 25, 2008 during the Mars northern summer for adequate solar power conditions and Lander temperature requirements. Any later Mars arrival time and the mission would risk the science activities being cut short due to the approach of northern winter. The Phoenix trajectory was a 10-month, type-two trajectory that reached 1.7 AU from the Sun and 1.8 AU from the Earth at Mars arrival (Figure 1).

There were two main navigation challenges that Phoenix faced on its interplanetary trajectory, or cruise, to Mars. Each challenge posed a significant risk to the overall landing safety and overall success of the science mission. The first challenge to overcome was to understand the trajectory effect of many small ΔVs, known as small forces, imparted to the trajectory by the unbalanced thruster controlled attitude control system (ACS). The spacecraft used thrusters to maintain spacecraft pointing in order to maintain communication with Earth and allow power to be collected by the solar arrays. Each time a thruster was pulsed to maintain attitude, a ΔV of, as much as 0.07 mm/s was imparted to the trajectory at a frequency as high as 160 thruster firings per day. These "small forces" needed to be understood in order to properly reconstruct and predict the spacecraft trajectory.

The second navigation challenge was also posed by the attitude control system. In order to perform trajectory control maneuvers (TCMs) the spacecraft would use its thruster to slew, or turn, the spacecraft to the proper direction for the ΔV correction. The turn itself caused a ΔV because of the unbalanced thrusters. For the TCMs near arrival the slew to and from the TCM attitude was a significant portion of the overall ΔV. In order to meet the stringent entry and landing site conditions, the slew ΔV needed to be understood and the TCM design strategy adjusted to minimize the overall effect of this error source.

The following sections of this paper describe the solutions to the two main navigation challenges as well as the final navigation results based on those solutions.

II. Navigation Plan and Requirements

The navigation design is made up of three main components: orbit determination, trajectory control, and error analysis. The orbit determination component includes the reconstruction of the traveled trajectory and the prediction of the trajectory into the future. The trajectory control component includes methods for changing or correcting the trajectory to meet certain navigation targets. The error analysis component is identified separately because it depends on the strategies of the other two components. The key areas of importance to Phoenix from each of these components will be described in this section and they will be compared to the Phoenix navigation requirements.

Table 1. Phoenix DSN Tracking Schedule

From	To	Type of Support
Launch	Launch + 14 days	Continuous coverage
Launch + 15 days	Entry - 60 days	Three 8 hour tracks/week
TCM - 4 days	TCM + 4 days	Continuous coverage
Entry - 60 days	Entry	Continuous coverage

a) Two-way Doppler and Ranging Tracking Schedule

Start Time	End Time	Measurement Frequency
Launch + 36 days	Launch + 42 days	1 measurement/day
Launch + 64 days	Entry - 60 days	1 measurement/week
Entry - 60 days	Entry - 18 days	3 measurements/week
Entry - 18 days	Entry	2- 3 measurements/day (DSN and ESA DDOR)

b) DDOR Tracking Schedule

A. Orbit Determination

There were two main types of orbit determination (OD) processes used on Phoenix. The first and most challenging type was the interplanetary or cruise phase OD. Determining the orbit from the Earth and predicting the Mars atmospheric entry conditions is extremely challenging. The second type is called surface positioning. Surface positioning is the process of determining the location of the Lander on the surface of Mars.

The three main data types used in the cruise phase were two-way Doppler, two-way Ranging and Delta Differential One-way Range (DDOR). The two-way Doppler measurement uses the Doppler effect of a signal from a DSN antenna to the spacecraft and back, to measure the line-of-sight

velocity of spacecraft[1]. Two-way Ranging uses the same signal to measure the distance between a DSN antenna the spacecraft[1]. DDOR is used to measure the position of the spacecraft in the geocentric plane-of-sky[1]. Each of the data types were processed using an X-band signal transmitted to the spacecraft and an X-band signal transmitted by the spacecraft in return. Each of these measurements was used on a regular basis according to the schedule in Table 1. The collection frequency of each data type increased the closer the spacecraft got to Mars arrival.

Surface positioning also used the two-way Doppler data type. In this case a UHF signal was transmitted from one of the Mars orbiters to Phoenix and back. The Doppler data was stored on board the orbiter and sent down in spacecraft telemetry to the DSN. The Doppler data was processed to determine the position of the Lander.

The tracking data was used to refine a number of different models. Some of the key models used in the estimation process were the solar radiation pressure model of spacecraft, the TCM or large delta-V model, and the small forces or small attitude control ΔV model. This paper will go into detail in later sections about the small forces model as that model was one of the major challenges of Phoenix navigation. The details of the other Phoenix navigation models can be found in Ref. 2.

B. Trajectory Control

The Phoenix spacecraft was designed to make trajectory corrections by pointing the spacecraft in a particular direction and firing the monopropellant hydrazine propulsion system. There are a number of different sources of trajectory error. The first trajectory error that needed to be corrected was the launch vehicle performance errors. Once that was complete the remaining error sources were OD reconstruction errors, OD prediction errors, and past TCM execution errors. All of these errors needed to be corrected, using TCMs, to the fidelity of the entry and landing accuracy requirements.

As described in the Section III below, the spacecraft has four TCM thrusters that are all pointed in the spacecraft –X axis. Firing these TCM thrusters causes a ΔV in the spacecraft +X axis. Therefore, in order to perform a TCM in a particular direction the spacecraft would turn to the desired attitude, and turn on the TCM thrusters until the desired ΔV was completed. The ΔV portion of the TCM, or burn, was terminated using onboard accelerometers to measure the achieved ΔV. This accelerometer cutoff capability allowed the TCM magnitude to be performed very accurately. The TCM direction was controlled by off-pulsing the TCM thrusters to keep the +X axis inertially pointed while the RCS thrusters controlled the spacecraft rotation about the X axis. The spacecraft rates were sensed by the gyroscope, which was included on the inertial measurement unit (IMU). The resulting attitude errors were corrected by commanding the TCM and RCS thrusters accordingly.

Table 2. TCM Summary

TCM #	TCM Location	TCM Date	Tracking Data Cutoff
1	Launch + 6 days	08/10/07	TCM – 5 days
2	Launch + 81 days	10/24/07	TCM - 5 days
3	Entry – 4 days	04/10/08	TCM - 5 days
4	Entry – 15 days	05/10/08	TCM- 22 hours
5	Entry – 7 days	05/18/08	TCM – 22 hours
6	Entry – 22 hours	05/25/08	TCM – 22 hours

Six TCMs were planned throughout the cruise phase in order to control the trajectory to the nominal target on the Mars surface. Each TCM had a specific purpose and scheduled time. Table 2 contains a summary of the TCM plan. The first TCM, TCM-1, was performed 6 days after launch in order to correct for the majority of the launch vehicle target bias and execution errors. The launch vehicle targets are biased in such a way to reduce the probability of a Mars impact of the spent third stage, which ends up in a similar orbit to the spacecraft[3]. TCM-2 was performed 81 days after launch in order to remove the remaining launch vehicle bias and errors. TCM-1 and TCM-2 were optimized, or designed together, in order to minimize their combined ΔV and propellant cost. At Mars arrival minus 45 days, TCM-3 was executed correcting errors that had built up in the trajectory in the over five month period since TCM-2. The purpose of TCM-4 at Mars atmospheric entry (referred to as entry) minus 15 days was to limit the ΔV size and error contributions from the TCM-5 and 6 burns. TCM-5 was strategically placed in order to meet the EDL system limits, by mitigating the risk of the failure of the next TCM. This allowed a reasonable chance for the EDL system to execute properly even if the landing site constraints were not met. Finally, TCM-6 was placed at entry minus 22 hours to guarantee that the landing site constraints and target were reached.

TCM execution errors posed a significant challenge in navigating the Phoenix spacecraft. TCM execution errors can be broken up into two major pieces. There are errors associated with the main ΔV execution and there are errors associated with the spacecraft attitude control actions before and after the main burn to prepare for and cleanup from the TCM. The main burn ΔV magnitude has errors due to controller latency and discreet processing. In general this error is small compared to the pointing errors, which are due to things like thruster imbalances, misalignments, and mass property uncertainty. These main burn errors did dominate the execution error for larger early burns of approximately 2 m/s or more, however they were not of concern for the smaller fine-tuning TCMs such as TCMs 4, 5 and 6. The more dominant error source for the smaller turns was the attitude control activities surrounding the main burn. In particular the slews to and from the main burn attitude caused the most ΔV error because of the unbalanced RCS thrusters. This was the second main challenge in navigating the Phoenix spacecraft. More details regarding trajectory control and TCM execution errors can be found in Ref. 3.

C. Requirements and Error Analysis

The most driving navigation requirement on Phoenix was the landing safety requirement. In order to avoid rocks that could damage the Lander upon touchdown or block a Lander solar array from the deploying, a specific landing area was chosen. Once a scientifically interesting landing region was chosen the Mars Reconnaissance Orbiter's HiRise camera was used to take images of the possible landing areas. With a resolution of 30 cm per pixel the HiRise images allowed the Landing Site Safety Team to identify and count rocks that were 1.5 m in height. The Lander was vulnerable to rocks taller than 40 cm, and therefore rock size distribution models were used to extrapolate from the HiRise images. The landing site selection process identified a suitable landing area that could accommodate a 100 km by 20 km landing 99% ellipse with a landing azimuth of 112 degrees centered at 68.151 deg north latitude and 233.975 deg east planetocentric longitude. More detail regarding the landing site safety considerations can be found in Ref 3.

There are two main contributors to the size of the landing ellipse. The first contributor is the uncertainty in the trajectory at the Mars atmospheric interface point. This uncertainty can essentially be reduced to the one-dimensional quantity known as the entry flight path angle (EFPA), which is the angle between the spacecraft velocity vector and the local horizontal at the atmospheric interface point (Mars radius of 3522.2 km). The second major contributor to the landing ellipse size is the trajectory dispersion as a result of the EDL system's flight through the atmosphere. The EDL dispersions are determined by factors such as center-of-mass knowledge, atmospheric density knowledge, and the aerodynamic properties of the entry vehicle. The EDL dispersions combined with the selected landing area required that the navigation entry flight path angle be accurate to ±0.2 deg (3-sigma).

Table 3. TCM-5 and TCM-6 Entry Flight Path Angle Errors Budget.

Error Source	Entry Error (3σ) TCM-5 EFPA (deg)	Entry Error (3σ) TCM-6 EFPA (deg)
Small Forces Errors	0.53	0.10
TCM Slew Errors	1.10	0.15
TCM Execution Errors	0.32	0.05
Other Orbit Determination Errors	0.26	0.11
TOTAL	1.29	0.20

The absolute entry flight path angle target is a critical parameter to the EDL system. If the entry flight path angle is too shallow there is a risk that the spacecraft may skip out of the atmosphere. If the entry flight path angle is too steep then there is a risk of overheating the entry vehicle. Both steep and shallow entry flight path angles also have an adverse impact on EDL timing and the terminal descent/touchdown phase. An entry flight path angle target of -13.0 degrees was deemed optimal for the EDL system. The flight path angle accuracy of ±0.2 deg (3-sigma) was well within the limits of the EDL system. EDL robustness analysis showed that the EDL system could reasonably survive EFPA accuracies of ±1.0 deg (3-sigma) from the target.

A third and final requirement on the Mars entry conditions was the time of entry. The Phoenix EDL configuration was designed to communicate to the Mars orbiters (Mars Odyssey,

and Mars Reconnaissance Orbiter) using a UHF signal as no reliable direct-to-Earth link was available[5]. The orbiters would store the data sent by Phoenix and would in turn send that data to the Earth using their respective X-band systems. One of the important requirements on Phoenix was to have a system that allowed the diagnosis of failures during EDL, if such a thing were to occur. The UHF link to the orbiters was the main pathway to store the necessary data to meet that requirement. The orbiters needed to be in the correct location, orientation, and configuration at the time of EDL, and so the entry time was an important parameter. The navigation requirement was for the Phoenix entry time to be within 30 sec of the time targeted at TCM-3.

Table 3 shows the final predicted navigation error budget for the entry delivery accuracy from the final two TCMs. This table illustrates the relative importance of the error sources and points to the fact that the two major challenges of navigating Phoenix would be to understand the small forces errors and the TCM slew errors.

III. Spacecraft Description

Understanding certain systems of the Phoenix spacecraft was a critical element in navigating it successfully. This section describes the parts of the spacecraft that were the most important to navigation.

The spacecraft as a whole can be broken up into 3 major components: the Lander, the entry system, and the cruise stage (see Figure 3). The Lander contains all of the spacecraft avionics, the propulsion subsystem, the surface power components, the surface telecom components, and the surface payload. The entry system encapsulates the Lander with the backshell and the heatshield. The cruise stage contains the cruise power and telecom components.

Figure 2 shows the spacecraft in the cruise configuration. The spacecraft cruise coordinate system is given in the figure by the included coordinate axes. As it is shown in Figure 2**,** the solar array normal points in the direction of the –X axis. The cruise low gain antenna (LGA) boresight is also oriented in the –X axis direction. The boresight of the medium gain antenna (MGA) is oriented 33 deg off of the –X axis towards the +Z axis. Also shown in the Fig. 3 at the top of the backshell are two of the cruise rocket engine modules (REM). Each of these modules contains a single TCM and a single reaction control system (RCS) thruster. The two other REMs are located, in the figure, at the bottom side of the backshell.

The cruise stage contained two X-band transponders that were used to communicate telemetry and tracking data with the DSN antennas. The LGA was used to transmit the spacecraft signal for the first 157 days of the mission when the range to Earth was adequate. The LGA allowed the spacecraft to be off-pointed from the Earth in order to point the solar arrays closer to the Sun. During the second half of cruise the telecom system was configured to use the MGA. The MGA boresight was offset from the solar array normal allowing the solar array to point directly at the Sun while the Earth was in the field of view of

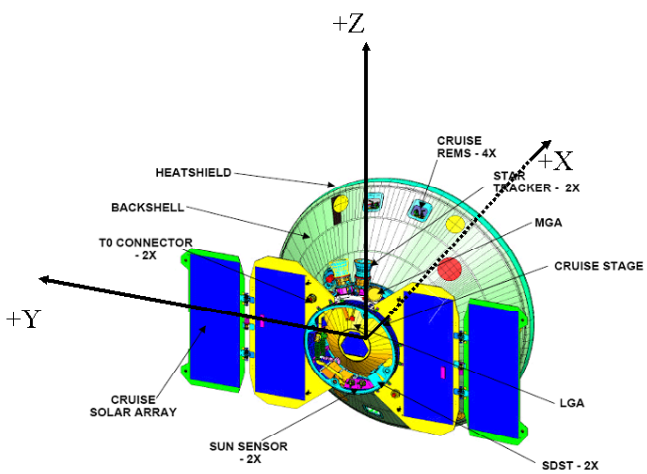

Figure 2. Spacecraft in Cruise Configuration

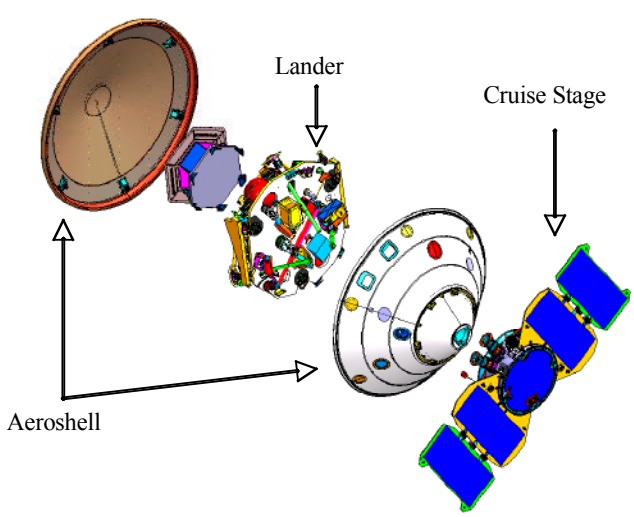

Figure 3. Spacecraft in Cruise Configuration Exploded View

the antenna. This late cruise attitude, being mostly solar torque balanced, was a key ingredient in keeping the attitude control effect on the trajectory consistent and predictable. The telecom and power subsystems were the main drivers for the two spacecraft attitudes used during cruise. However, there were other constraints levied by other subsystems. Most notably, the thermal subsystem constrained the amount of time away from the nominal cruise attitude to 20 min or less. This constraint had an effect on the design of TCMs, where the statistical ΔV could be in any direction[3].

The main navigation focus of the attitude control system was the RCS thruster system and its ΔV effect on the trajectory. The TCM thrusters were also important for performing TCMs, however their use and ΔV effect were reasonably straightforward. The RCS thrusters were a different story and posed a significant challenge to navigation. The RCS thrusters were the attitude control mechanism used to keep the spacecraft pointed correctly. Using only four thrusters to maintain attitude requires that each thruster have a torque component in each spacecraft axis. The thrusters were mostly fired in pairs to control the attitude and rate about each axis inside a set of control deadbands. When the attitude control system detected that it was at the edge of one of the deadbands the controller would command a set of 15 ms thruster firings to avoid crossing the constraint. The control was not explicitly limited to firing in pairs of thrusters, but odd numbers of thruster firings were rare during normal cruise deadbanding operations. Figure 4 shows the thruster locations and the components of the thruster vectors. Table 4 shows the as-measured nominal thruster direction per thruster. As you can see from the table, every thruster has a thrust component in the +X axis direction. Therefore any thruster firing would impart ΔV in the +X axis direction. The other thrust components in the Y and Z directions are designed to cancel, in an average sense, given equal number of firings on each thruster. However, solar torque imbalances caused certain thrusters to fire more than others, which broke the cancellation symmetry and caused accumulation of ΔV. Furthermore any thrust magnitude or thrust direction asymmetries in the Y or Z directions caused additional net ΔV.

A number of deadband sizes were planned for uses throughout cruise in order to maintain spacecraft pointing and minimize fuel usage. Tight (±1 degrees in each axis) deadbands were used just after launch for communication and power robustness. Once the spacecraft was checked out after launch, the deadbands were opened up to 11 degrees, 8 degrees, and 4 deg about X, Y, and Z respectively. At 70 days after the launch the deadbands were opened up further to 44 degrees in X, 12 degrees in Y, and 14 degrees in Z, saving fuel to be used during EDL. Half-way through cruise the original plan was to step down the deadbands to meet pointing constraints until they reached the 10, 2, and 2 degrees set in X, Y, and Z. The final deadbands were then planned to be constant from 60 days prior to entry until just minutes prior to EDL. Once the deadband fuel performance was understood, it was deemed that the final deadbands could be started a great deal sooner than planned. The increased fuel usage was small versus the margin that was available in the propellant tanks at the time. The propellant margin increase was due to the fact that the launch trajectory was within a sigma of the target, while the TCM propellant budget accounted for the 99% case. The Phoenix project took advantage of the opportunity, tightened the deadbands to their final values of 10, 2 and 2 degrees in X, Y and Z respectively, which allowed the operations teams to characterize the small forces at those deadbands for an additional 130 days.

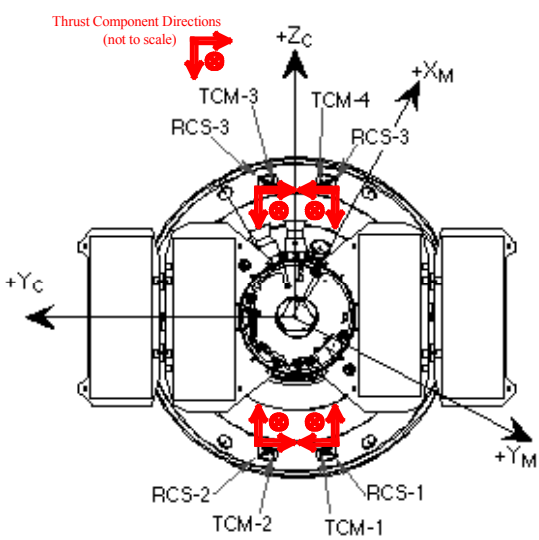

Figure 4. RCS Thrust Component Direction

IV. Attitude Control Thruster Calibration and Verification

The total attitude control ΔV effect on the trajectory depended upon a number of factors. The main factors include the total ΔV per pulse and the number of pulses fired by each thruster to maintain attitude. The ΔV per pulse depends on the thrust magnitude for each thruster, which is a function of propellant tank pressure, thruster catalyst bed (aka "catbed) temperature, minimum on-time control, valve dynamics, nozzle direction, nozzle shape, etc. Each of these nozzles was designed to match the contour of the backshell, or scarfed, which means that the nozzle exit plane was not perpendicular to the nozzle axis of symmetry. Scarfed nozzles affect the thrust direction and therefore ΔV direction. Some prelaunch testing was completed using the thrusters, but it was mainly geared to attitude control performance issues and not focused on ΔV characterization.

Table 4. Prelaunch Nominal RCS Thrust Directions

	Scarfed RCS Thrust Vectors, Cruise Frame		
Thrusters	X_C	Y_C	Z_C
RCS-1	0.2874	0.8840	0.3737
RCS-2	0.2874	-0.8840	0.3737
RCS-3	0.2874	-0.8840	-0.3737
RCS-4	0.2874	0.8840	-0.3737

The pulse frequency of each thruster, in turn, depended on another set of factors. The thrust magnitude and direction of each thruster determined the change in angular rate of the spacecraft after a firing. The deadband size determines how often a thruster would fire to keep from crossing the constraints. The spacecraft attitude determines the size and direction of the solar torque on the spacecraft, which biases the attitude and thruster firings to a particular area of the deadband space.

All of the factors that govern the ΔV effects of the attitude control system combined with the lack of full system test data, drove the need to perform an in-flight calibration campaign in order to properly reconstruct and predict the spacecraft trajectory. The small forces calibration campaign was made up of three specific activities. The first was the active thruster calibration where the thrusters were pulsed in a designed sequence at two specific spacecraft attitudes while two-way Doppler tracking data and attitude control gyroscope data was collected. The data was used to estimate the thrust magnitude and direction for each thruster. The second calibration activity was the passive thruster calibration. In this phase, no special thruster firing sequence was developed. The two-way Doppler and gyroscope data that was collected during normal cruise deadbanding was used to trend and evaluate the results of the active thruster calibration. The third activity used the standard operations OD process to evaluate the active thruster calibration model adjustments. This method could not isolate the trajectory effects of the thruster activity as well as the other methods, but it was a good way to verify the results in a bounding sense.

A. Active Thruster Calibration Objective

The objective of the active RCS thruster calibration was to measure the average RCS thrust vectors with an accuracy of 10 % or better. To achieve this objective a sequence of RCS thruster firings was executed and high-rate Doppler tracking data plus gyroscope data from the spacecraft's IMU were collected. The Doppler tracking data and gyroscope data were combined in a sigma-point consider filter to estimate the average RCS thrust vectors.[8]

1. Design and Plan

The active RCS thruster calibration design had to take into account attitude, thermal, and telecom constraints, which were unique to the Phoenix mission. For the Odyssey and MRO[6,7] active thruster calibrations, the spacecraft maintained attitude using reaction wheels and the selected attitudes allowed for complete observability of the RCS thrust vectors.

Phoenix does not have reaction wheels and therefore cannot maintain a selected attitude during the active RCS thruster calibration. Analysis by the spacecraft team showed that only 18 RCS thruster pairs could be fired, before the attitude would drift outside the thermal and telecom attitude constraints. Therefore after each 18 RCS thruster pair firings an attitude reset or "clamp-down" is needed. Nominally the RCS thrusters fire for 0.015 s during deadbanding and the RCS thruster calibration. The RCS thruster firings during the clamp-down can use larger on-times and cannot be used to estimate the average thrust vectors for a 0.015 s on-time.

The thermal and telecom attitude constraints limit complete observability of the RCS thrust vector using Doppler tracking data only, contrary to the Odyssey and MRO thruster calibration campaigns[6,7]. This resulted in using the gyroscope data in conjunction with the Doppler tracking data. The attitudes chosen for the Phoenix active calibration were chosen such that maximum observability for the RCS thrust vectors could be achieved by complementing the gyroscope and Doppler tracking data observabilities. The gyroscope data can only observe the component of the

RCS thrust vector that is perpendicular to the position vector from the center of mass (COM) to the throat location of the RCS thruster (torque component). From Figure 4 it can be derived that the z-component of the RCS thrust vector is nearly unobserved in the gyroscope data because the position vectors from the COM to the RCS thruster throat location are nearly parallel to the z axis and hence the torque component is small. Furthermore, the z-component of the RCS thrust vector is not observable in the Doppler data in the inner cruise attitude because the X and Y spacecraft coordinate axes are located in the Sun-Phoenix-Earth plane. This results in the Z axis being perpendicular to the line of sight to earth, which is unobservable by Doppler tracking.

The attitudes chosen for the active RCS thruster calibration were rotations of +/- 37 deg about the sun vector as seen from the Phoenix spacecraft. These two attitudes maximize the Doppler observability of the z-axis component of the RCS thrust vectors. These attitudes satisfied the thermal constraints by not changing the illuminated area of the spacecraft. The telecom attitude constraint to provide continuous telecom during the active RCS thruster calibration was also met.

During the RCS thruster calibration the following thruster pair firing sequence was repeated twenty times. The thruster-pair firing sequence consists of three identical subsets. Each subset of thruster-pair firings was planned as follows: -X(2,4); +X(1,3); -Y(1,2); +Y(3,4); +Z(1,4); -Z(2,3). The sign and letter refer to the control axis plus direction and the numbers in parenthesis identify the RCS thruster numbers in Figure 4 for the thruster pair. A total of 18 thruster pair firings were executed per sequence. The thruster pair firings were separated by 1 minute to allow cooling of the catbed and the collection of sufficient 1 sec Doppler tracking data for data noise reduction. The firing sequence order was selected to minimize attitude deviations based on observed angular rates prior to the active RCS thruster calibration. Once a sequence executed it is followed by a clamp down to restore the active RCS thruster calibration attitude. In total 10 sequences were executed in the +37 deg rotation about the sun line followed by 10 sequences in the -37 deg rotation about the sun line .

Covariance studies[1] showed that the selected attitudes and the combination of Doppler tracking data and gyroscope data would result in the estimation of the average RCS thrust vectors with an accuracy of 10 % or better.

2. Active RCS thruster calibration execution

The RCS thruster calibration started on 14 September 2007 17:40:00 UTC and lasted for 8 hrs. All 60 thruster pair firing sequences executed nominally. However, some extra thruster firings occurred after a clamp-down within the 1 minute prior the start of the sequence. All high-rate Doppler tracking data and gyroscope data were successfully collected. The Radio Science Receiver (RSR) at the Canberra ground station also provided Doppler tracking in addition to the standard Doppler tracking data from the DSN antennas.

During the execution, the RCS catbed temperatures were observed to be 100 to 150 deg C higher than the nominal deadbanding thruster firings. The higher catbed temperatures were caused by the clamp-down thruster firings and insufficient time between sequences preventing a cool down to nominal catbed temperatures. The catbed temperatures, during the RCS active thruster calibration are shown in Figure 5. In Figure 5 it can be seen that after each clamp-down, the catbed temperature increases and a cool down occurs between clamp-downs but the catbed temperatures do not cool back to the nominal catbed temperature range of 120 to 130 deg C.

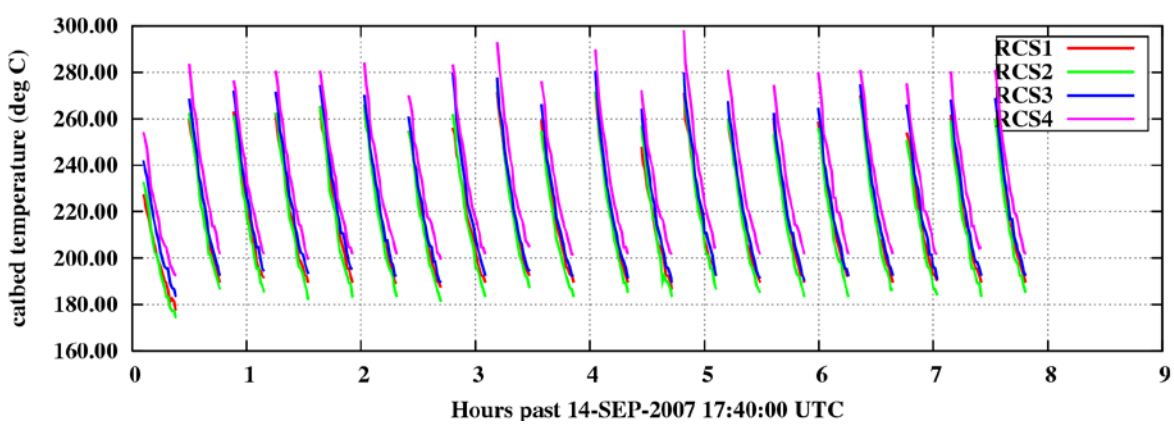

Figure 5. Catbed Temperatures for RCS thrusters during active thruster calibration

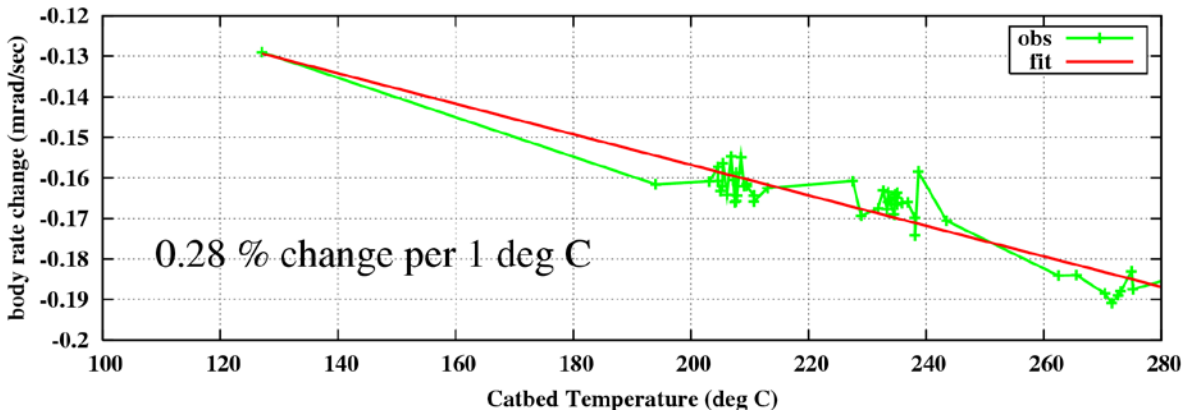

Figure 6. Empirically derived linear relationship between body rate change and RCS catbed temperature for –X axis thruster pair firings

The increased catbed temperatures magnify the thrust of the RCS thrusters because less energy is needed to warm up the catbed[10]. The objective of the active RCS thruster calibration is to estimate the average thrust vectors during nominal deadbanding thruster firings, therefore a model was needed to remove this effect from the data. An empirical linear relation between catbed temperature and thrust magnitude change was derived from angular spacecraft body rate changes and the observed catbed temperature data. In Figure 6 the angular body rate change for the –X axis control is plotted as a function of the catbed temperature. Also shown is the linear fit to this data. Based on this fit, the percent change in magnitude thrust per 1 deg C is 0.28 %. For all other axes the results were within 0.02 % compared to the –X axis result. Based on this model and the observed catbed temperature increase of 100 to 150 deg C, the thrust magnitude was 28-42% higher during the active RCS thruster calibration compared to nominal deadbanding thruster firings.

3. Active RCS thruster calibration results

The average RCS thrust vectors were estimated using a sigma point consider filter (SPCF)[9]. The SPCF combines the gyroscope angular body rate change data and Doppler tracking data to estimate:

- RCS thruster vector components at an RCS catbed temperature of 120 deg C
- center of mass vector
- Doppler bias (stochastic reset at each clamp-down)

and the uncertainty of the following parameters is taken into account as consider parameters:

- moments of inertia
- spacecraft mass
- RCS thruster throat location vector (assumed point of force imparted on spacecraft)
- low gain antenna phase center vector.

The results for the RCS thrust vector magnitudes are shown in Table 5. The first result is that all the estimated magnitudes are 23-56% higher than the magnitude predicted from the preflight impulse (Ibit) equation[11]. Furthermore, the magnitudes for RCS thrusters 3 and 4 are 13-23% higher than the magnitude for RCS thrusters 1 and 2. The solution covariance was used to assess the accuracy of the estimate by computing the Root Sum Square

Table 5. RCS thruster magnitude results

Thruster	Magnitude (N)	Solution RSS Sigma (N)	Solution RSS Sigma/Magnitude ratio (%)	Magnitude Sigma (N)	RSS direction Sigma (deg)
RCS1	1.92421	0.117	6.09	0.077	2.651
RCS2	1.88243	0.115	6.11	0.076	2.665
RCS3	2.39787	0.154	6.43	0.098	2.861
RCS4	2.16982	0.090	4.16	0.064	1.708
Pre-flight Ibit	1.53000				

(RSS) of the three Cartesian thrust component sigma's and by mapping the covariance into magnitude and direction angles. It should be noted that the uncertainties reported in Table 5 are a function of the data weights used for the Doppler and gyroscope data. In this filter solution the average Doppler weight was 0.032 mm/sec and the gyroscope data weight was 0.0035 mrad/sec. The spherical mapping shows that the 1-sigma uncertainty in the magnitude is < 0.1 N and the 1-sigma uncertainty in thrust direction is < 3.0 deg for all RCS thrusters. Finally, the ratio of the RSS Cartesian sigmas and the estimated magnitude showed that the total error was less than 7% for all RCS thrusters, and thus the goal of determining the average thrust vectors to better than 10% was achieved according to the filter covariance..

In Table 6 the estimated thrust direction and pre-flight scarfed thrust vectors are shown. From Table 6 it is clear that the estimated x and z components are much larger than the pre-flight x and z components. The increase in the x-component is consistent with the empirically derived thrust[8] vector for the Mars Polar Lander (MPL) spacecraft, which has the same RCS configuration and spacecraft design as Phoenix. The likely causes for the change in thrust direction are the thruster plume impingement with the Phoenix back shell and minimum impulse effects. Also notable is that RCS thrust vectors 3 and 4 appear out-of-family with RCS thrust vectors 1 and 2.

Table 6. Estimated and pre-flight RCS thrust directions

	Estimated Unit Thrust Direction Vector Cruise Frame (-)			Pre-flight Scarfed Unit Thrust Vector. In Cruise Frame (-)		
Thruster	X	Y	Z	X	Y	Z
RCS-1	0.4035	0.7493	0.5251	0.2874	0.8804	0.3773
RCS-2	0.3828	-0.7654	0.5174	0.2874	-0.8804	0.3773
RCS-3	0.3694	-0.7907	-0.4882	0.2874	-0.8804	-0.3773
RCS-4	0.3423	0.7458	-0.5716	0.2874	0.8804	-0.3773

Figure 7a shows the observed and SPCF-modeled Doppler tracking data, relative to a predicted orbit without modeling the active RCS thruster calibration activities. From Fig. 7a is it clear that most of the ΔV in the line of sight to Earth was caused by the clamp-downs. The clamp-down ΔV can not be used in the SPCF because the thrusters did not fire with a 0.015 sec on-time, and therefore only piece-wise continuous segments of Doppler data could be used in the SPCF. This significantly weakens the Doppler data because after every 18 thruster pair firings a new Doppler bias needs to be estimated. This is different than the Odyssey and MRO active RCS thruster calibrations where many hours of continuous Doppler data was available to constrain the thrust vector solution.

Figure 7b shows that the Doppler residuals are colored-noise instead of Gaussian distributed. After noise reduction by averaging over 60 seconds, the Phoenix project had two Doppler tracking data sets available; 1) The Doppler data from the standard DSN receiver with a noise level of 0.03 mm/sec and 2) the Doppler data from the RSR with a noise level of 0.01 mm/sec. For the filter solution the RSR Doppler data was used, but no significant change in the colored-noise Doppler residuals was observed between the two data sets. Thus the Doppler noise is not the limiting error source. Most likely the colored-noise residuals are caused by pulse-to-pulse variations, which can be significant for 0.015 sec on-times. The filter model assumes a stationary process and can not accommodate these variations, which result in colored-noise residuals. In order to mitigate the colored-noise effect, the Doppler data was weighted separately for each segment between clamp-downs based on the residual noise for that segment. This is reflected by the filter residual covariance also plotted in Fig. 7b. The filter residual covariance represents the uncertainty of the estimated parameters mapped back into the measurement space. Part of this mapping includes the data weight used for the measurements. It can be seen in Fig. 2b that the filter residual covariance changes with the size of the residuals. The main purpose of the filter residual covariance is to check the tuning of the filter. The values in the plots show +/- 3-sigma values. If the filter is tuned, then the filter residual covariance should envelope the residuals, instead of the residuals being too small or too big.

Figure 8a shows the observed and SPCF-modeled body rate change for the X-axis. It is clear that the model and observations are in good agreement. Furthermore, the effect of the decreasing catbed temperature can be seen in the decreasing amplitude of delta body rate when the +X or –X thruster pairs are fired. Figure 8a also shows that for 8 sequences the –X pair firing (first in sequence) was removed because an extra thruster pair fired within the 60 seconds prior to the start of the sequence. This gave insufficient time for Doppler data noise reduction. The observed and model plots for the Y and Z axes show similar agreement

Finally, Figure 8b shows the body rate change residuals for the X-axis, which are more Gaussian-distributed than the Doppler residuals. It is also shown that the data weight selected for the gyroscope data is correctly chosen because the filter residual covariance is nicely enveloping the residuals. Analysis of the filter solution revealed that the gyroscope data acts as a strong constraint on the solution because of the low noise characteristics of the data. Many filter strategies were investigated and the gyroscope data residuals were always nearly Gaussian-distributed while the Doppler residuals varied significantly (being the weaker data type). The solution presented in this paper (using the per-segment Doppler weight) was chosen because it represented the median of all the investigated filter solutions. The filter solutions investigated included using subsets of the data, data type weight variations, and Doppler media correction variations.

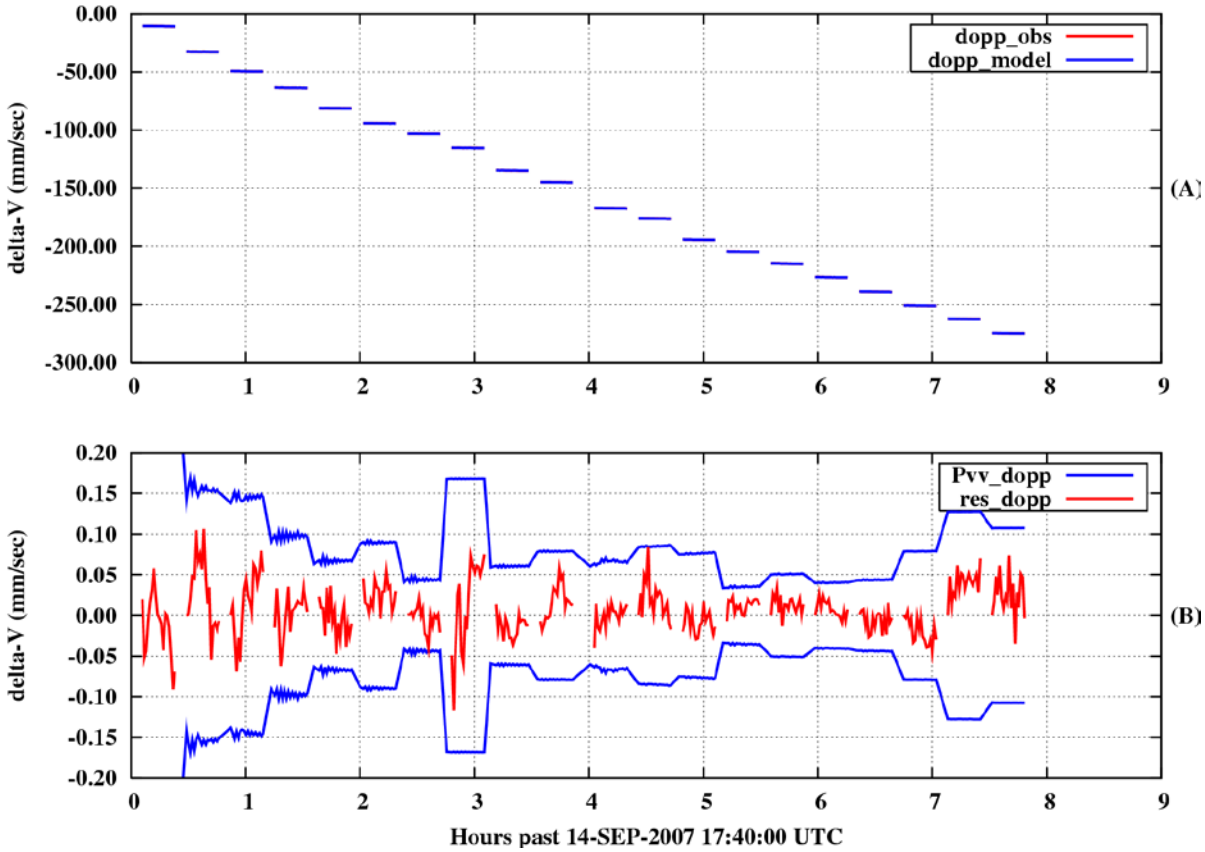

Figure 7. Doppler data and residuals during active thruster calibration:
A) Modeled and observed Doppler data
B) Doppler residuals and filter residual covariance

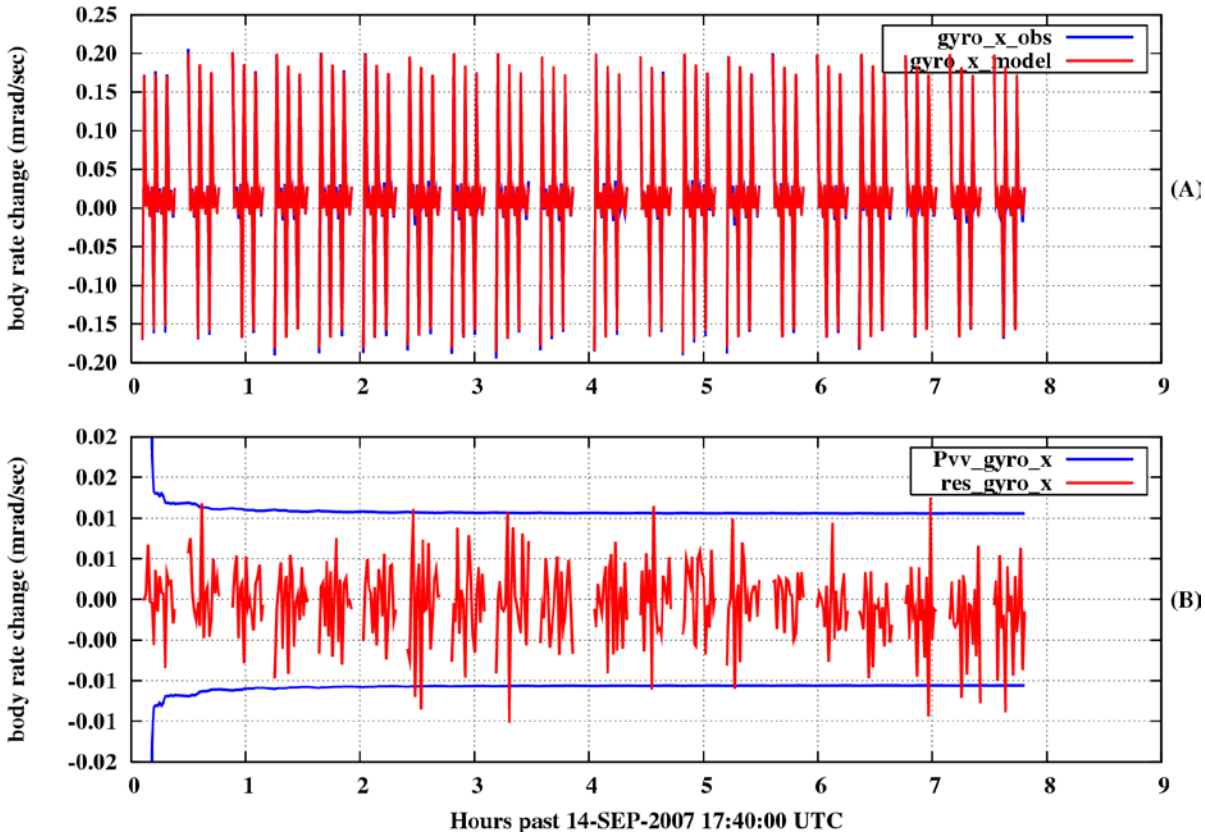

Figure 8. Gyroscope spacecraft X-axis body rate change data during active thruster calibration:
A)Modeled and observed gyroscope spacecraft X-axis body rate change data
B)Gyroscope spacecraft X-axis body rate change residuals and filter residual covariance

B. Passive Thruster Calibration

The objective of the passive RCS thruster calibration was twofold. The first objective was to characterize the thruster activity in the late cruise attitude and associated tight dead bands. The second objective was to assess the accuracy of the active RCS thruster calibration result (SPCF model) by comparing the accumulated ΔV of all thruster firings from the Doppler tracking data to the SPCF model. The passive thruster calibration was extended to entry. The TCM-1 ΔV was relatively small, which increased the propellant margin and allowed additional propellant usage to extend the passive RCS thruster calibration. This extension enabled the estimation of a constant acceleration error in the orbit determination process, which was used to bound the error in the SPCF model.

1. Design and plan

The original plan for the passive RCS thruster calibration was to put the spacecraft into the late cruise attitude with deadbands of 10 deg in X, 2 deg in Y and 2 deg in Z (10,2,2) for one week to characterize the thruster firings and deadbanding characteristics. After one week the deadbands would be increased to conserve fuel. Sixty days before entry the 10,2,2 dead bands would be reinstated and the characterization of the thruster firings would be continued until entry. During the passive thruster calibration it was decided to remain in the late attitude deadbands because the increase in propellant usage was small compared to the propellant saved due to a relative small TCM-1 maneuver. It was deemed more advantageous to continue trending the deadbanding in this configuration.

The accumulated ΔV analysis has the advantage that no preplanned sequence is needed and that it does not depend on the deadband. Therefore, all acceptable RCS thruster firings between the TCM-1 maneuver and entry were included in the analysis. The variation in viewing geometry of the Phoenix coordinate axes determines how well the accumulated ΔV analysis can assess the errors in the SPCF model, because that determines how much of the SPCF mode error will be visible in the line of sight to Earth. The viewing geometry for Phoenix is shown in Figure 10d. It can be seen that for long periods of time the viewing geometry is nearly constant and that certain axes

are unobservable. Another limiting factor of the accumulated ΔV analysis is the DSN tracking schedule. For Phoenix only a few passes per week were available during the period where most of the viewing geometry change occurred.

2. Execution

The passive RCS thruster calibration was started on 15 November 2007 00:00:00 UTC by tightening the dead bands to (10,2,2). A few days prior the attitude was changed to the late cruise attitude, which caused out-gassing to occur. Out-gassing caused torques resulting in an increase of thruster firings and one-sided deadbanding. The increased thruster one-sided firings showed an exponential decay signature over time which is the main symptom of out-gassing. If the passive calibration had not been extended this would have added an extra complication to the analysis.

3. Passive RCS thruster calibration results

The results for trending the RCS thruster firings are shown in Fig. 9. Fig. 9 shows the daily averaged acceleration based on the reported RCS thruster firing ΔV for the spacecraft X,Y and Z axis. The reported RCS thruster firing ΔV is computed onboard Phoenix using the active calibration result (SPCF model) corrected for the propellant tank pressure difference after TCM-2. In Figure 9 it can be seen that the acceleration is changing until the middle of December 2007, due to out-gassing caused by the attitude change to the late cruise attitude a few days prior to the start of the passive RCS thruster calibration. After the middle of December 2007, the RCS thruster firings remained stable until entry. The stability is illustrated in Figure 9 by the constant predict values used in the orbit determination process for propagating the trajectory forward to entry. It should be noted that the reported accelerations are not the actual accelerations experienced by the spacecraft, but rather the accelerations predicted by the active RCS thruster calibration result (SPCF model).

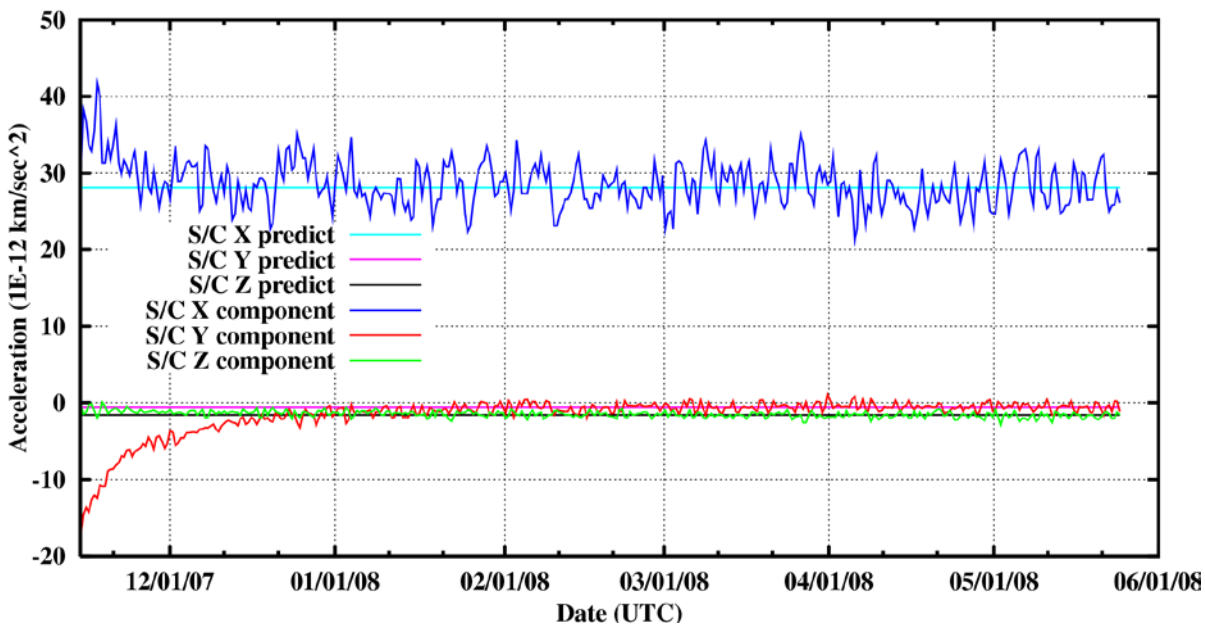

Figure 9. Reported Small Force Acceleration (24 Hour Averages)

One way of validating the active RCS thruster calibration results (SPCF model) was to evaluate the accumulation of ΔV over a long period. For this analysis the ΔV for each suitable RCS thruster pair firing is computed from the SPCF model and from high-rate Doppler tracking data. For both ΔVs, an accumulated sum is computed and the sums are compared. The advantage of this validation method is that no sophisticated spacecraft dynamic modeling is needed and it can be applied over the complete mission. Furthermore, as the number of RCS thruster pair firings increases, the accumulated ΔV becomes more accurate compared to the noisy ΔV observations derived from Doppler tracking data. The Doppler tracking-data-derived ΔV was noisy because the size of the individual ΔV per thruster pair firing was about 0.03 mm/sec in the line of sight to Earth. The Doppler tracking data

for this analysis has a 1-sigma noise of 0.03 mm/sec, making it difficult to detect the individual thruster ΔV in the data directly.

Since the accumulated ΔV analysis spans nearly the complete mission, it is necessary to add an additional correction to SPCF model. As Phoenix consumes propellant, the tank pressure dropped, resulting in lower RCS thrust magnitudes as the mission progressed. Therefore, a blow-down ratio[11] was applied to the RCS thrust magnitude based on the tank pressure difference with respect to the tank pressure during the active RCS thruster calibration. Analysis revealed that the empirical linear catbed temperature correction used for the active RCS thruster calibration gave significantly poorer results in the accumulated ΔV analysis. This suggests that the catbed temperature correction is not that linear and is only valid over the temperature range during the active RCS thruster calibration.

Figure 10a shows the accumulated ΔV based on Doppler tracking data and the SPCF model corrected for blow-down. A close agreement between the two accumulated ΔV curves can be seen. In the beginning of the mission, not much ΔV is accumulated because the spacecraft X-direction has low observability, as can be seen in Figure 10d. From Table 6 it can be seen that the net forces for all RCS thruster pair firings nearly cancel in the Y and Z-axis, but not in the X-axis. As the X-axis becomes more observable during the mission, the accumulated ΔV grows much faster.

Figure 10b shows the individual ΔV derived from Doppler tracking data and the SPCF model. This plot reveals the noisy characteristic of the ΔV derived from Doppler tracking data. This is also clear in the residuals of the two ΔVs. Also shown are the 5-day boxcar smoothing of the two ΔV-data-sets and their residual. The 5-day boxcar-smoothed residual curve is nearly zero, which indicates that the model is providing an unbiased ΔV compared to the Doppler tracking ΔV in a mean sense. The 5-day boxcar-smoothed Doppler-derived ΔV and SPCF model ΔV can be used to quantify the SPCF model error by computing their ratio. If the ratio is one, then the SPCF model and the observed ΔV are in agreement. If the ratio is less than 1, then the model is over-predicting the ΔV, and if the ratio is higher than 1, then the model is under-predicting the ΔV. Figure 10c shows this ratio for the complete mission.

The ratio in Fig. 10c shows very large values in the beginning because very little ΔV is accumulated, causing the large variations. After the active RCS thruster calibration the ratio value varies between 90 and 110%. The average ratio value from the active RCS thruster calibration to entry is 0.99 with a sigma of 7%. In summary, the passive RCS thruster calibration showed that the RCS thruster performance was stable from the start of the passive calibration campaign and validated the active RCS thruster calibration result. The accumulated ΔV analysis determined that the average ratio between the observations and the model was 0.99, with a 1-sigma variation of 7%. These results were consistent with the error estimated based on the active calibration SPCF covariance, which was less than 7% for each RCS thruster.

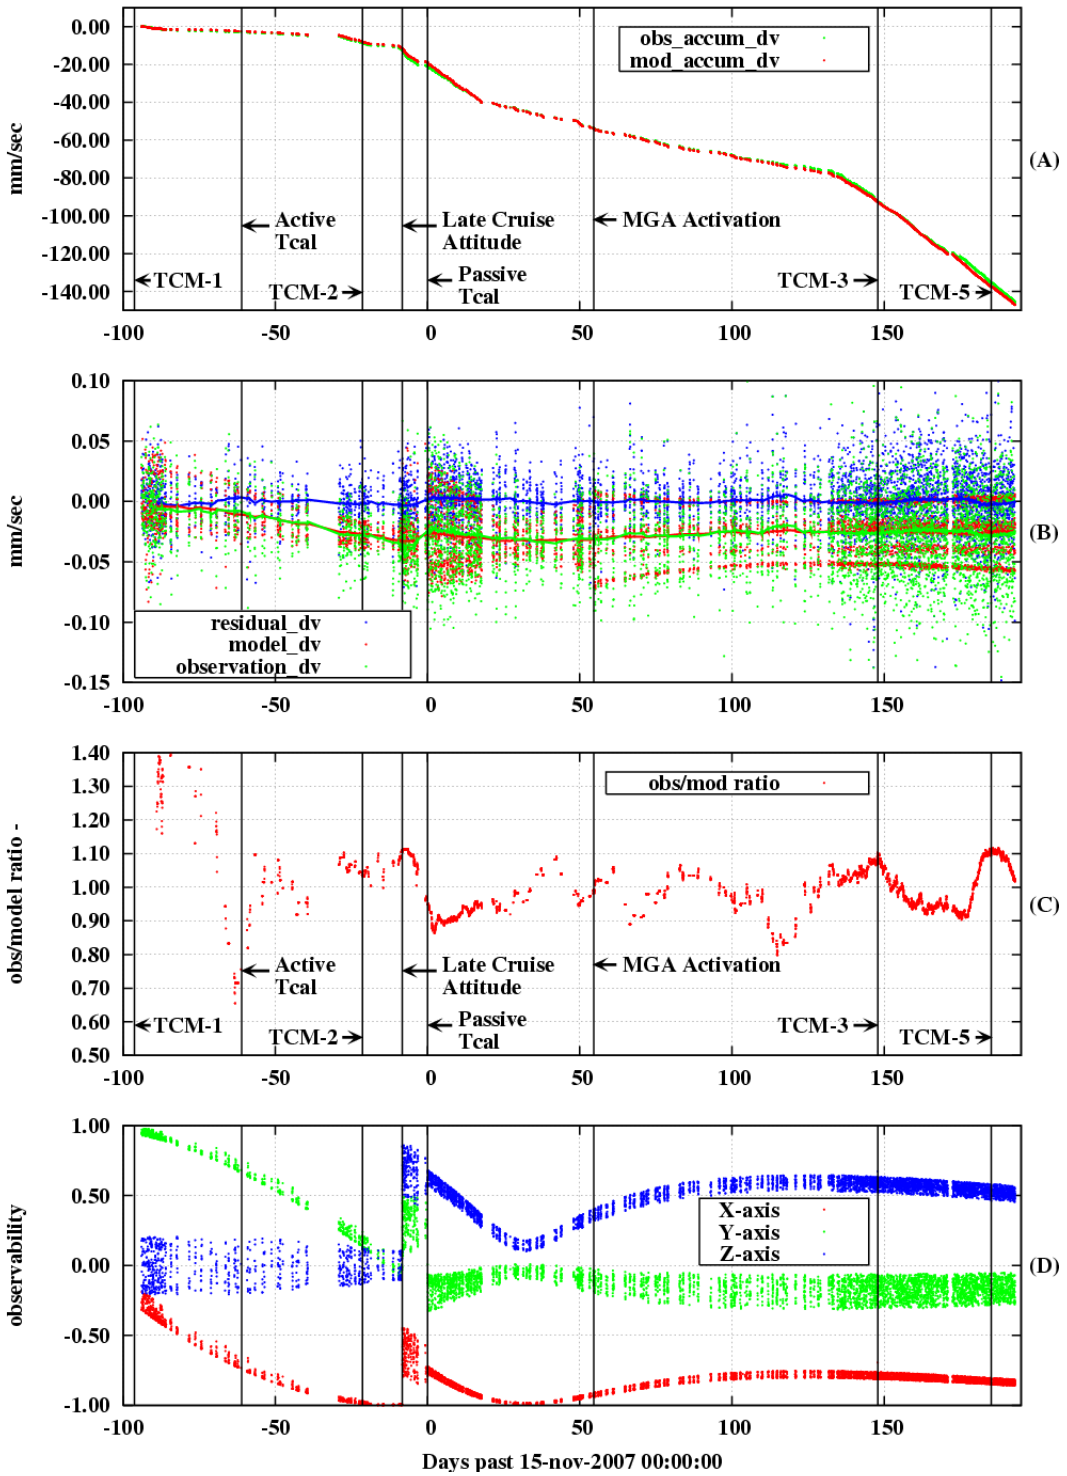

Figure 10. Phoenix accumulated ΔV analysis without catbed temperature correction:

A) Active Thruster Cal Model and Doppler observed accumulated ΔV

B) Active Thruster Cal Model and Doppler observed per thrust ΔV

C) Doppler observed/ Active Thruster Cal Model ΔV ratio

D) Space Craft frame coordinate axes observability in Earth line of sight

C. Orbit Determination Verification

Finally, a constant acceleration was estimated in the orbit determination process, which was used to bound the error for the SPCF model. The constant acceleration is an upper bound of the error because it estimates any unmodeled acceleration acting on the spacecraft. The two non-gravitational forces acting on the spacecraft were the RCS thruster firings and the solar radiation pressure. By assuming that the estimated acceleration is all due to SPCF model error, an upper-bound SPCF model error can be calculated and compared to the active and passive thruster calibration results. The constant acceleration was estimated after the late cruise attitude out-gassing subsided[2].

A constant acceleration could be estimated because the spacecraft remained in the late cruise attitude and (10,2,2) deadbands for the remainder of the mission after 15 November 2007. In the orbit determination process the reported RCS thruster-firing accelerations are applied when integrating the orbit without modification. Any constant errors in the reported RCS thruster-firing accelerations will cause a drift in the propagated orbit solutions as more tracking data becomes available. This behavior was observed during Phoenix navigation operations and estimating a constant spacecraft per axis acceleration made the propagated orbit solutions consistent. The estimated constant spacecraft X-, Y-, and Z-axis accelerations were found to be small and nearly constant as more tracking data became available.

To quantify the error in the SPCF per axis, the ratio of the ΔV due to the estimated constant acceleration over one day and the sum of all the ΔVs in that axis in an absolute sense during that day was computed. In essence, for each axis this ratio gives the ΔV error of the SPCF model compared to the total ΔV expended during that day. The ratio is calculated as follows:

$$percent_error = \frac{86400 * accel_axis}{\sum_{j}^{Nthr}\left(\sum_{i=1}^{4}\left|\Delta Vaxis_i\right|\right)_j} * 100\%$$

where *accel_axis* is the estimated constant acceleration for a given axis and $\Delta Vaxis_i$ is the ΔV (RCS thruster i) component for a given axis. This ratio was calculated for every day since 1 January 2008 and then the average ratio was calculated. The average ratio results are for X (6.6%), Y (6.5%) and Z (1.7%). These error estimates are consistent with the active RCS thruster calibration and the passive calibration results where the error in the SPCF was estimated to be < 7%.

V. TCM Slew Errors and Landing Site Targeting

The TCM slew errors played a large role in the navigation of Phoenix. The thruster calibrations performed throughout cruise were designed to calibrate minimum impulse (15 ms) isolated pulses. Slews are mostly made up of pulse trains designed to set a specific turn rate. Therefore the deadbanding calibrations would not be adequate to characterize slews. A slew calibration was discussed, bit it was decided that the reasonably nominal performance of the early TCMs did not warrant such an activity.

A. Slew Description

TCM slews consist of turns both to and from the TCM ΔV attitude. The slews were completed using the RCS thrusters which when fired imparted ΔV to the trajectory. This ΔV is predicted and accounted for as a part of the TCM implementation process, but there were significant errors associated with the prediction. One error source that could not be accounted for was the randomness of the attitude in the deadbands just prior to the start of the slew. Therefore there was an error in the number of pulses needed to impart the desired slew rate and there was uncertainty in the inertial direction of the pulsing. Other error sources include thruster misalignments and pulse-to-pulse variations.

There are other activities associated with slews and TCMs that cause ΔV error. The spacecraft attitude control system makes transitions through certain modes as part of a slew or TCM execution sequence. Each of these modes has different characteristics. One of the typical modes is rate damping. For certain mode transitions the spacecraft makes sure that the rates are below some limit. If they are not, the spacecraft fires thrusters to force the rates down to an acceptable level. This mode is used multiple times during the TCM sequence, most notably just after the TCM thrusters shutdown following the main burn. Another mode is called inertial hold where the spacecraft holds a specific attitude with 1 deg deadbands on each axis. This can cause thruster firings because the deadband is so tight. This mode is used just prior to the main burn of a TCM and as the spacecraft waits to make a mode transitions following rate damping. Finally, there are a number of different slew modes that specify how a slew or turn is completed. There are slew modes that cause the spacecraft to seek a specific attitude and there are slew modes that

Table 7. Prelaunch TCM Non-Main TCM ΔV Error Budget

Error Component	Units	There/Back 1.0 deg/sec slew	Full Circle Slew	There/Back 0.4 deg/sec slew
Slew	**mm/s**	**13**	**10**	**8**
inertial hold	mm/s	2	1.5	1.5
rate damping	mm/s	2	2	2
Ops tool prediction error	mm/s	5	5	5
Analysis Total (1σ)	mm/s	**14**	**11**	**10**

are turns specified relative to the current attitude. Even though these slew modes are different they all have a common profile. The controller starts a slew in a certain direction by firing a pulse train to reach a specific slew rate, typically ranging from 0.25 deg/s to 1 deg/s. Once the correct direction and rate are achieved the spacecraft continues to use its RCS thrusters to maintain the rate and the slew profile during the coast period. Once the spacecraft reaches the end of the slew profile attitude, it fires thrusters to stop the turn. The start and stop portions of the slew impart the most ΔV, but the coast-period thruster firing is difficult to predict because it is a partially random process. The desired slew rate and the initial attitude rate in the deadband both determine the amount of thrusting needed for the slew start and stop.

The prelaunch slew Monte Carlo analysis provided some insight into the dominant error sources involved with slews. The slew error budget is shown in Table 7. The prelaunch error budget showed that the errors from the actual turns dominate the total non-main-TCM burn error budget.

B. Slew Error Mitigations

In order to minimize the ΔV errors from slews, different slew strategies were selected for particular TCM directions. The slew Monte-Carlo analysis showed a significant portion of the slew ΔV error came from the thruster pulsing used to start and stop the slew. Slew strategies were designed to minimize the ΔV error due to slews.

1. Slew Strategies

The first strategy was designed to use lower slew rates, which caused fewer thruster pulses during the slew starts and stops. This strategy could not be used to its fullest extent because of the thermal constraint limiting the amount of time off of the nominal attitude to 20 minutes. Therefore the slower slew rates could only be used for smaller angles. For large slews of greater than 100 deg the standard slew rate of 1.0 deg/s was employed. For slews less than 100 deg/s a slew rate of 0.4 deg/s was used. The reduction in overall error for the 0.4 deg/s slews can be seen in Table 7. Slews as low as 0.25 deg/s were considered for TCM slews less than about 30 degrees.

Another strategy that was used to minimize slew errors was called a full-circle slew. This strategy would slew out to the TCM attitude, using the minimum distance slew. Once the TCM ΔV was over, instead of slewing back in the opposite direction of the first slew (as in a there-and-back slew), the spacecraft would continue in the direction of first slew until it was back at the nominal cruise attitude. In essence the spacecraft completed a 360 deg "full circle" turn during the activity. In the example where the TCM attitude was 180 deg away from the nominal attitude, and in the absence of other error sources, all of the slew ΔV would cancel. In reality not all of the ΔV would cancel, however most of it and any regular ΔV biases would cancel. For slews smaller than 180 deg the amount of cancellation is less. The statistical benefit for using full-circle slews for slews greater than 100 deg can be seen in Table 7.

2. Landing Site Targeting Strategies

In addition to adjustments made to the way slews were executed, TCMs were designed to limit the size of slews. The standard way to design TCMs for landing site targeting was to target EFPA, latitude, and longitude. This process was completed using two separate trajectory propagator programs. One propagator program completed a ΔV search and propagated the interplanetary trajectory to the entry point. The other propagator took the state at entry and propagated it through the EDL events to the surface. This process was iterated until the landing site target was met. By allowing EFPA to be a free parameter, the TCM ΔV direction could be constrained to minimize the slew angle. Minimizing the slew angle allowed slower rate slews to be implemented, which minimized slew errors.

Allowing the EFPA parameter to be free did not come without some cost. First of all, there was an EFPA target (-13.0 deg) and tolerance (±0.2 deg) that the spacecraft had to stay within. Secondly, one of the entry conditions affected by this targeting strategy was entry time. The entry time requirement was to be within 30 sec of the TCM-3 target. The basic physics that allowed this targeting strategy is illustrated in Figure 11. There exists a ΔV direction at any given time that changes the entry time and EFPA together in a such a way as to not change the resulting landing point on the surface. By adding this alternate ΔV direction to the standard ΔV design it is possible to change the resultant ΔV direction to minimize the slew angle. Figure 11 shows a nominal EFPA with the green solid arrow that is steepened by a ΔV in the alternate direction, which results in the EFPA with the blue solid arrow. When the EFPA is steepened, the down range travel of the spacecraft through the atmosphere is shortened shown by the difference

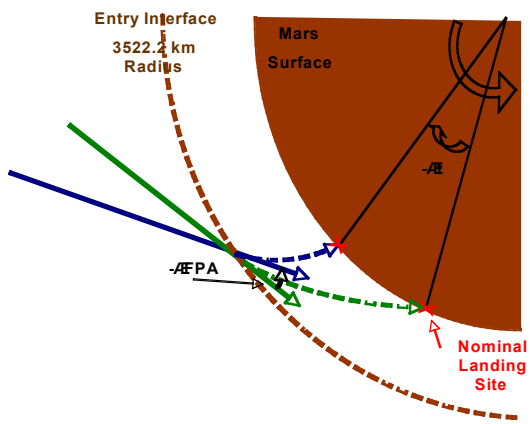

Figure 11. Landing Site Targeting Diagram

between the two dotted lines. When the atmospheric flight time is shortened the landing time is earlier and Mars has not had the same amount of time to rotate underneath the spacecraft. Therefore in order for the Lander to touchdown at the same surface location, the entry and landing time need to be earlier, which is specified by the -Δt in Figure 11. More details regarding Phoenix landing site targeting can be found in Ref. 3 and Ref. 4.

C. In Flight Slew Results

During actual flight there were two there-and-back slews and 2 full-circle slews. TCM-1 and TCM-2 used there-and-back slews and TCMs 3 and 5 used full-circle slews. This section will go into the details of the performance of those slews and how the TCM-1, 2, and 3 slew performances affected the decision for the TCM-5 slew type.

TCM-1, executed just 6 days after launch, was the largest ΔV, at 18.5 m/s, performed by Phoenix. The slew angle to the desired ΔV direction was 58.5º and therefore a 0.4 deg/s there-and-back slew could be used to execute the TCM and stay within the twenty-minute off-nominal attitude thermal constraint. Even though a slower slew rate was used it was inconsequential because the size of the main burn ΔV errors swamped the slew errors**.** Table 8 shows the difference between the slew ΔV reported by the spacecraft telemetry and pre-TCM predicted ΔV. It can be seen that the TCM-1 slew had a 1-sigma error based on the prelaunch assumption from Table 7.

TCM-2, at launch plus 81 days, was also fortunate to only require a small slew to reach the TCM attitude. The TCM-2 slew angle was only 14.8º. Again, the there-and-back slew at 0.4 deg/s was used to execute TCM-2. As can be seen in Table 8, the TCM-2 slew was well below the prelaunch expected 1-sigma error.

The TCM-3 slew was significantly different than the previous two. The slew angle for TCM-3 was 145.6 degrees and the ΔV was 1.4 m/s. In order to minimize slew errors for this large slew the full-circle slew mode was used to execute TCM-3. The TCM-3 slew executed very precisely and the results are shown in Table 8. It can be seen that even though TCM-3 was larger in magnitude and was executed in a different mode, its errors were very similar to the TCM-2 slew errors.

Table 8. In Flight Slew Performance

Activity	Slew Type	Slew Rate (deg/s)	Slew Angle (deg)	Slew Magmitude Error (mm/s)
TCM-1	There/Back	0.4	58.5	**11.0**
TCM-2	There/Back	0.4	14.8	**3.6**
TCM-3	Full Circle	1	145.6	**3.8**

It is easy to see that the slew error from TCM-1 was significantly larger than the slews errors for the other two TCMs. An investigation was undertaken to see if the cause could be identified. Analysis showed that each slew coast phase (the period between a slew start and slew stop) contained significantly more thrusting than the slew prediction accounted for. During a slew coast phase the controller is attempting to follow a slew profile made up of a slew axis and a rate. The coast phase thrusting was not determined by an attitude error deadband. The control was determined by a pulse-width threshold. The minimum thruster pulse width was 15 ms. The controller was set to only command a pulse if the required pulse width, to stay on the slew profile, was greater than 7 ms. After an attitude update cycle, if the controller needed to pulse a thruster for 5 ms to stay on the profile, the pulse would be skipped. If on the next cycle the controller needed an 8 ms pulse, a 15 ms pulse was commanded. In some cases this was a rather inefficient controller behavior, which would cause a larger number of pulses. This behavior was difficult to model because it was very dependent on the initial conditions at the start of the slew, which was a statistical event within a set of deadbands.

Given the inability to predict this behavior well, there were two ways to mitigate the overall impact. One way to mitigate the impact was to limit the amount of time the spacecraft spent in the coast phase of any slew. This was essentially the difference between the TCM-1 and TCM-2 slews. This mitigation is consistent with the overall strategy to minimize the total slew angle to any TCM attitude. However, since this behavior is really a function of time, limiting the time in the slew coast phase went against the strategy of using slow slew rates in order to mitigate the amount of pulsing required to start and stop a slew. The second way to mitigate this error source was to use full-circle slews and allow the coast ΔV prediction error to cancel itself out. This was the main difference between the TCM-1 and TCM-3 slews. For the remaining TCMs it was decided that the best way to reduced overall slew errors

was to limit the slow slew rate strategies to slews of up to 60 degrees instead of 100 degrees and to use full-circle slews for all slews greater than 60 degrees The results for the last three TCMs are described in the next section.

VI. Final Entry and Landing Results

This section summarizes the final entry and landing results achieved by the Phoenix mission. The results from the final four TCMs, as well as entry and landing will be described. The results will be described using a number of different parameters: B-plane parameters, entry parameters, and planetocentric longitude and latitude. The B-plane coordinates used in this section are made up of two orthogonal axes in a plane that is centered at Mars and is orthogonal to the incoming V-infinity vector. The X axis of the B-plane is called B•T and is centered at Mars parallel with the Mars mean equator of date. The Y axis is called B•R and lies in the plane perpendicular to B•T. The B-plane is used in this section because it is a convenient way to show two-dimensional error ellipses that describe the relative uncertainties of arrival conditions[12].

D. TCM-3 and TCM-4

TCM-3 was a relatively straightforward maneuver. There were more than 5 months of trajectory and tracking data between TCM-2 and TCM-3 where the only spacecraft ΔV events were due to regular ACS deadbanding small forces. The sun-relative spacecraft attitude was constant with the –X axis of the spacecraft pointed at the Sun for 5 months starting in early November. The ACS deadbands were constant and in their final state since about the same time. The active thruster calibration results were complete and verified using with the passive calibration (accumulated ΔV method) and the standard orbit determination process. All of those factors contributed to the stability of the OD solutions leading up to the final TCM design shown in Figure 12. Figure 12 also shows the correction TCM-3 made in the B-plane along with the post TCM-3 result.

TCM-3 performed so well that TCM-4 was cancelled. The ΔV needed to correct for the TCM-3 errors was below the minimum that the spacecraft could handle relative to the errors involved. The minimum ΔV that was executable on the spacecraft was 0.05 m/s. For more information regarding TCM-4 and the minimum ΔV constraint see Ref 3.

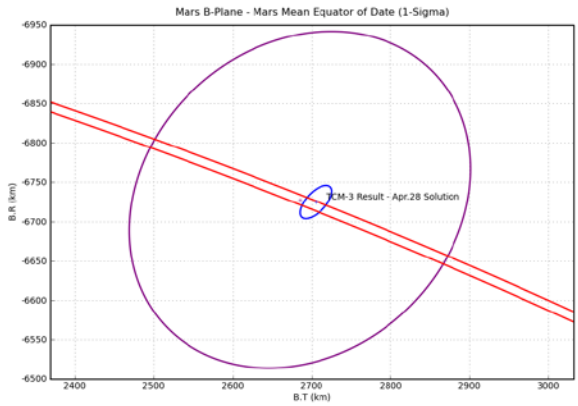

a) TCM-3 OD Solution History

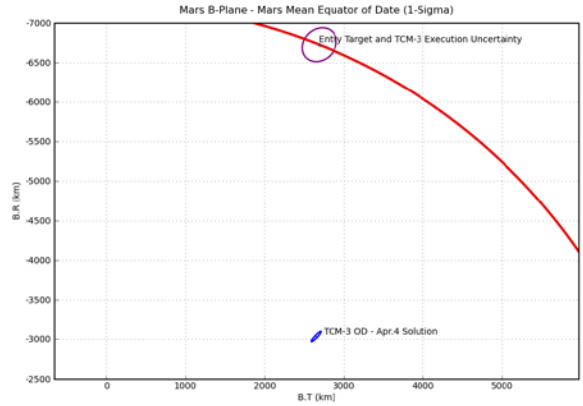

b) TCM-3 Designed B-plane Correction

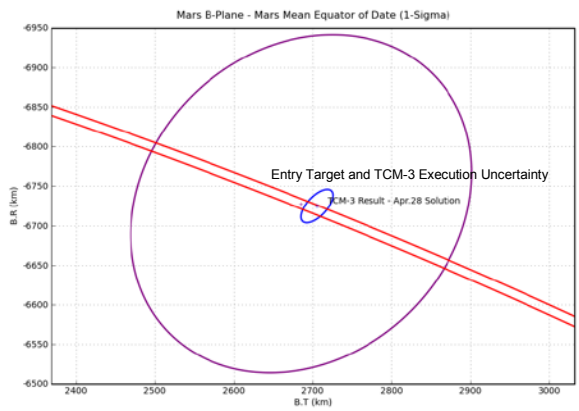

c) Post TCM-3 Results

Figure 12. TCM-3 B-plane Results

E. TCM-5 and TCM-6

The purpose of TCM-5 and TCM-6 was to make fine adjustments to the trajectory in order to meet the EDL system limits and target the nominal landing site. The OD errors prior to the TCM-4 opportunity (Entry minus 15 days) were relatively large compared to the EFPA and landing site requirements. It was not until about entry minus 12 days (TCM-5 minus 5 days) that the EFPA and landing ellipse size started to approach the requirement level. At that point the navigation target focus shifted from the B-plane to the predicted landing location.

1. Landing Safety Requirements

There were four major landing safety requirements that needed to be met. The requirements served as the main criteria for making TCM decisions. Table 9 summarizes the landing safety criteria. The least-driving landing safety requirement (#1 in Table 9) was the surface slope requirement. The requirement was that 99% of the 1-sigma ellipse, by ellipse area, needed to fall in an area with slopes that were safe for the Lander. See Ref. 4 for the definition of safe slopes. Criterion 2 was relative to the total probability of landing in what was known as a "certified safe zone" (CSZ) given the probability of landing at any particular point in a 99% Gaussian ellipse. The CSZ areas, labeled green in the surface map of Figure 13, were areas deemed safe by the landing site safety team. These areas had to meet certain requirements with regard to rock distributions or slopes. Rock distributions were calculated by either counting rocks in a HiRise image or by extrapolation of a particular geologic terrain type. By the time the MRO landing-site imaging campaign was complete over, 99% of the nominal landing ellipse was imaged. Slopes were determined from MOLA elevation data.

Table 9. Landing Safety Criteria

Number	Criterion	Value
1	Percent of 1-sigma landing ellipse (by area) with acceptable slopes	99%
2	Probability landing in the Certified Safe Zone	0.8
3	Percent of the 1-sigma ellipse (by area) in the CSZ	95%
4	Percent of the 2-sigma ellipse (by area) in the CSZ	97%

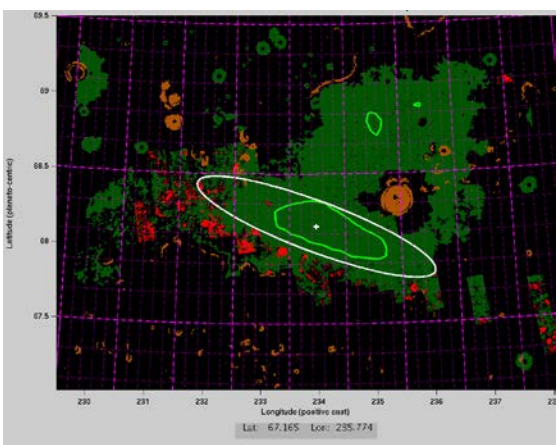

Figure 13. Certified Safe Zone Landing Area Map

The next two landing safety requirements were developed to avoid situations where the CSZ probability criterion was met with a known landing hazard very close to the ellipse center (the statistically most probable landing area). These criteria also guard against overly-conservative landing ellipses due to the stack-up of margins from different systems. Criterion 3 required that 95% of the 1-sigma ellipse, by area, be identified as in the CSZ. Criterion 4 stated that 97% of the 2-sigma landing ellipse, by area, be in the CSZ The green solid contour in Figure 13 is the contour that meets all of the landing site criteria. If the center of the landing ellipse (for the 100x20 km nominal ellipse) fell on or inside the green contour, then all of the criteria were met. If the ellipse center fell outside of that line then the criteria were not met and a TCM was required to correct the landing point.

2. Additional TCM Criteria

There were a number of decision factors that governed whether or not the next TCM was executed in addition to the landing site criterion. The first major factor was the spacecraft capability and health at the time of the TCM execution. The next factor was the OD knowledge and stability. If the size of the OD uncertainty was significantly larger than the correction then it is not statistically correct to perform a maneuver. The EDL system criteria needed to be met, and if not, then a TCM was required. Lastly, a secondary TCM decision factor was EDL communications. Certain trajectories had an impact on the timing of communicating with the Mars Orbiters during EDL. The EDL communications and EDL system criteria are not described in detail in this paper[6]. The flight circumstances were such that both EDL criteria-sets were not factors in the decision making process that came to fruition. OD knowledge criteria combined with the landing site safety criteria were the driving decision factors.

3. TCM-5

Figure 14 contains a time history of the OD solutions on the surface and in the B-plane for the four days leading up to TCM-5. There is a single solution labeled for each day. The relationship between the B-plane and the surface location can be seen in Figure 14 by comparing the two plots. The red constraint lines on the B-plane plot are EFPA contours representing the EFPA constraints of ±0.2 degrees As the solution moves towards Mars the EFPA is getting steeper or more negative. As the EFPA steepens the downrange travel of the spacecraft shortens. The downrange travel direction of the spacecraft is shown on the surface plot from the northwest to the southeast. Steeper EFPAs land short of he target and shallower EFPAs land long.

As the OD solution uncertainties were converging with added tracking data the solution center was “moving” on the surface toward the northwest, getting closer to the nominal landing point (steepening EFPA). Figure 14 shows that the landing ellipses had a relatively small uncertainty in the crosstrack direction relative to the downrtack direction. This is because the downtrack flight is much more sensitive to B-plane errors than the crosstrack direction. It was clear that even if the solution continued to converge in the direction of the landing point, the crosstrack miss would remain and the landing safety constraints would not be met. However the last few TCM-5 OD solutions settled down and the apparent drift ended leaving the predicted landing point outside the landing site safety constraint line, thus requiring a TCM. Figure 14 also shows the same solutions plotted in the B-plane where the solution convergence is evident from the fact that the lower left ends of the 1-sigma error ellipses are not intersecting.

The landing site targeting strategies in Section IV were used in the design of TCM-5. The standard landing site targeting process yielded a TCM-5 magnitude of 0.04 m/s with a slew angle of 83 degrees. The TCM ΔV was modified, by the EFPA and entry time targets as described in Section IV, primarily so that the ΔV magnitude could be increased to the minimum size of 0.05 m/s. Consequentially this strategy also decreased the slew angle to 68 degrees. The acceptable impact of this strategy was an EFPA target miss of 0.006 deg and an entry time target miss of 13 seconds. More details regarding the TCM-5 design strategy can be found in Ref. 3. TCM-5 was executed successfully on May 18, 2008. The real-time line-of-sight Doppler data showed less than 5 mm/s of prediction error when compared to the TCM-5 model.

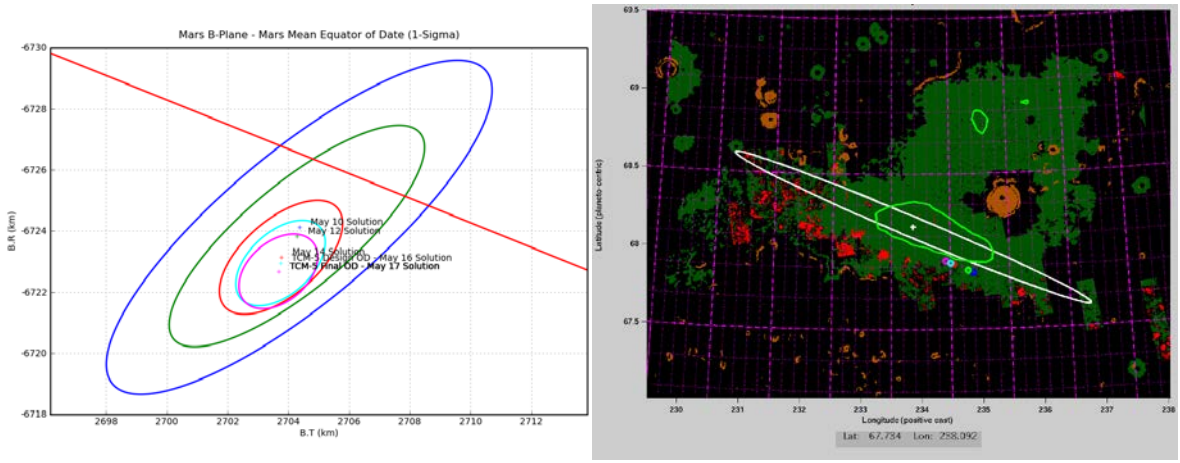

a) TCM-5 OD B-plane Plot **b) TCM-5 OD Surface CSZ Plot**

Figure 14. OD Solutions Leading Up To TCM-5 - solution colors are coordinated between plots a) and b)

4. TCM-6

The excellent TCM-5 performance decreased the probability of needing TCM-6 in order to meet the entry and landing site safety constraint. The limiting factor at this point was the OD. Figure 15 contains a time history of the OD solutions on the surface and in the B-plane leading up to TCM-6. There is a single solution labeled for each day. It can be seen that the early solutions just after TCM-5 were very close to the target. As with the solutions prior to TCM-5, the TCM-6 solutions were converging on entry states with steeper flight path angles. Again, this would result in a landing site point short of the nominal target. As the solutions converged the last OD moved closer and closer to the landing site safety constraint line. The final TCM-6 OD solution predicted that the center of the landing ellipse was just inside of the criteria line. This landing location met all of the landing site constraints, as well as the EDL system constraints. Based on those criteria TCM-6 was not executed. The TCM-6 landing criteria values for

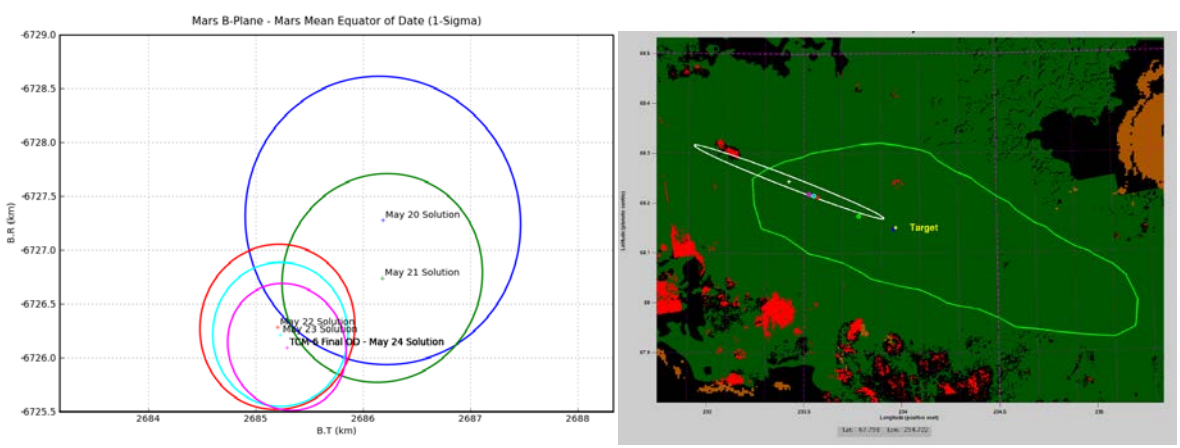

a) TCM-6 OD B-plane Plot **b) TCM-6 OD Surface CSZ Plot**

Figure 15. OD Solutions Leading Up To TCM-6 - solution colors are coordinated between plots a) and b)

Table 10. TCM-6 OD Landing Safety Criteria

Number	Criterion	TCM-6 OD Value	Decision Value
1	Percent of 1-sigma landing ellipse (by area) with acceptable slopes	100.0%	99%
2	Probability landing in the Certified Safe Zone	0.980	0.8
3	Percent of the 1-sigma ellipse (by area) in the CSZ	97.1%	95%
4	Percent of the 2-sigma ellipse (by area) in the CSZ	99.0%	97%

the final TCM-6 OD solution can be seen in Table 10.

F. Entry and Landing Results

The final approach navigation plan contained a contingency TCM, called TCM-6x. This contingency was developed to mitigate a number of different possible but unlikely TCM-6 failures. These failures included everything from a problem with the TCM-6 ground design and uplink process to a spacecraft execution failure. With the decision to skip TCM-6 and with a nominal spacecraft the only credible failure mechanism was a significant change to the OD as the Mars gravity became the dominant force on the trajectory. OD solutions were completed every few hours to monitor for such a scenario.

In addition to the TCM-6x contingency monitoring, the OD solutions were being evaluated by the EDL team in order to update certain key parameters used by the EDL system to land the spacecraft. These parameters included such things as entry time, parachute deploy triggers, backup timers and a host of other possibilities[13].

1. Entry

Figure 16 shows the final OD leading up to the Mars atmospheric entry point. As more tracking data was processed the solutions continued to get ever so slightly steeper in EFPA. The predicted landing point consequently moved to the northwest. However the solutions were moving around in the 1-sigma ellipse in the B-plane and on the surface as expected. The TCM-6x contingency opportunity was not used as all the constraints continued to be satisfied and the final EDL parameter update was uplinked to the spacecraft. The spacecraft was in its final preparations for a landing on Mars. The final OD solution, in Figure 17 contained tracking data up to about 15 minutes prior to entry, where the two-way Doppler data is very strong, resulting in a very well-known entry state. The center of the predicted landing ellipse was 22 km from the target to the northwest. The EFPA was very close to the optimal value of -13.0 deg and entry time was well within the EDL communications 30-second limit. The final entry state is shown in Table 11.

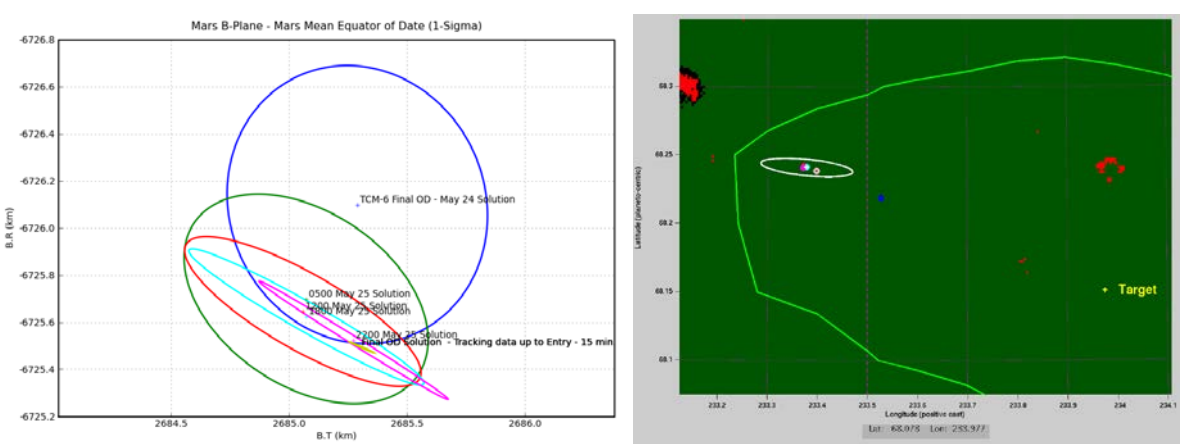

a) Entry OD B-plane Plot **b) Entry OD Surface CSZ Plot**

Figure 16. OD Solutions Leading Up To Entry - solution colors are coordinated between plots a) and b)

Table 11. Final Entry Conditions

Parameter	TCM-6 OD Value	Target
EFPA	-13.007 deg	-13.0 deg
Entry Time	25-May-2008 23:30:57.7 UTC	25-May-2008 23:31:12.6 UTC
Radius	3522.2 km	3522.2 km

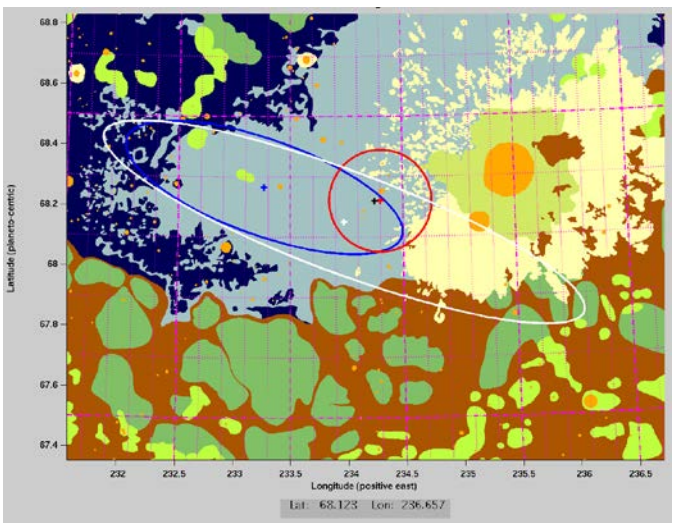

Figure 17. Final Surface Location Estimates

2. Landing Site Position

The EDL system performed nominally and the Phoenix Lander touched down safely on the Mars surface. The Lander used its UHF communication system to communicate with the Mars orbiters, MRO and Odyssey, which in turn relayed the data to the DSN here on Earth. Certain UHF passes contained two-way coherent UHF Doppler tracking data. This data was used to estimate the position of the Phoenix Lander on the surface. In addition to the two-way Doppler landing location solution, the spacecraft integrated its IMU accelerometer data during EDL as a part of the EDL navigation system. Once on the surface, the spacecraft sent its calculation of the surface location landing location in telemetry. Finally, each of these solutions was positively verified by taking an image of the Phoenix Lander on the surface using the HiRise instrument on MRO.

Figure 17 shows the final Lander surface position estimates using each method relative to the predicted landing point based on the final entry state. Figure 17 shows that the final landing point was 22 km (2.5 sigma) away from the pre-entry ellipse center and only 7 km (1.5 sigma) away from the nominal TCM-5 landing target. Given the knowledge of the trajectory at entry the 22 km offset from the predicted landing point was due to dispersions in the EDL trajectory. Post landing reconstruction analysis has shown that an offset in the trim angle of attack of the vehicle during the hypersonic phase of EDL caused more lift on the vehicle resulting in a landing point downtrack of the prediction[14].

The uncertainty of the spacecraft EDL navigation system was a 10 km circle on the surface shown in red. The final UHF based solution uncertainty, using eight Odyssey and three MRO UHF passes, was an ellipse 23 m by 5 m shown in green centered on 68. 218834 deg north latitude and 234.24832 deg east longitude (planetocentric). See the JPL website at http://marsprogram.jpl.nasa.gov/mro for the astounding MRO images of the Phoenix Lander on the surface of Mars.

VII. Conclusion

In conclusion, overcoming the navigation challenges posed by the Phoenix Mars Lander mission was a necessary factor in the successful landing on the surface of Mars. The active thruster calibration allowed accurate reconstruction and prediction of the spacecraft trajectory. The passive calibration validated the results of the active calibration, improving the confidence in the navigation solutions. Characterizing and mitigating the slews errors associated with TCMs using varying slew strategies and alternate TCM design strategies produced very accurate TCM executions. This total effort resulted in the successful navigation of the Phoenix spacecraft, using just 4 of 6 TCMs, both to the Mars atmospheric entry conditions and to within the landing site safety constraints, ultimately landing just 7 km from the desired landing point.

VIII. Acknowledgments

The successful navigation of the Phoenix Lander was made possible by the excellent support of the following groups. The team at Lockheed Martin in Littleton, Colorado not only provided a spacecraft that performed beyond expectations, but the overall operations support was also excellent. The Phoenix GNC team in particular was a major

contributor in the understanding of the major navigation challenges. The navigation team also received excellent support at JPL including DSN operations, tracking data support, and tracking data calibrations. In particular navigating this mission would have been extremely challenging without the use and support of the ΔDOR data type and operations team. Finally, the incredible support and dedication of the full Phoenix Navigation Team made up of the team leads (Tomas Martin-Mur and Tim Mcelrath), the orbit determination analysts (Eric Graat, Robert Haw, Gerhard Kruizinga, Eunice Lau, Sumita Nandi, Mark Ryne), the maneuver analysts (Ramachandra Bhat, Cliff Helfrich, Behzad Raofi), the entry descent and landing analysts (Eugene Bonfiglio, Lynn craig, William Strauss), the navigation software team (Theodore Drain, Scott Evans, Richard Sunseri, Joe Wu), and the system administrators (Dimitrios Gerasimatos, Katherine Nakazono). The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

IX. References

[1]Border, J.S., Lanyi, E.G., Shin D.K., "Radiometric Tracking Data for Deep Space Navigation," AAS Paper 08-52, Proceedings of the *31st Annual AAS Guidance and Control Conference*, Breckenridge, CO, Feb. 2008.

[2]Ryne, M.S., Graat, E.J., Lau E., et al. "Orbit Determination for the 2007 Mars Phoenix Lander," AIAA Paper 2008-7215, Proceedings of the *AIAA Astrodynamic Specialist Conference*, Honolulu, HI, Aug. 2008.

[3]Raofi, B., Bhat R. Helfrich C., et al. "Propulsive Maneuver Design for the 2007 Mars Phoenix Lander Mission," AIAA Paper 2008-7216, Proceedings of the *Astrodynamic Specialist Conference*, Honolulu, HI, Aug. 2008.

[4]Bonfiglio, E., Spencer D., Arvidson, R., et al., "Landing Dispersion Analysis and Landing Site Statistical Assessment for the Mars Phoenix Lander," AIAA Paper 2008-7348, Proceedings of the *AIAA Astrodynamic Specialist Conference*, Honolulu, HI, Aug. 2008.

[5] Kornfeld, R.P, Garcia, M.D., Craig, L.E., Butman S., Signori, G.M, " Entry, Descent, and Landing Communications for the 2007 Phoenix Mars Lander", Journal of Spacecraft and Rockets, Volume 45, Number 3, Pages 534-547 ,

[6].Mcelrath, T., "Odyssey Thruster Calibration Report", Internal JPL Memorandum, IOM-MARS01:NAV-01-012, 06 September 2001

[7].Jah, M., Graat E., Highsmith D., You T., "Mars Reconnaissance Orbiter Non-gravitational Force Modeling During Interplanetary Cruise", Internal JPL Memorandum, MRO-NAV-2005-07, 4 January 2006

[8] Mcelrath, T., "Mars Polar Lander (MPL) Small Forces File Modeling Development", Internal JPL Memorandum, 312/00.A-005, 30 March 2000

[9] Lisano, M.E., Kruizinga G.L., Portock B., "The NASA Phoenix Lander Thruster Calibration Estimator: Design and Validation", AIAA Paper 2008-7217, Proceedings of the *AIAA Guidance, Navigation and Control Conference, Honolulu, HI, Aug. 2008*

[10] McAllister, Greg, Lockheed Martin Space Systems, Denver, Personal communication

[11] Fisher,T., Internal Lockheed Martin Memorandum, PHX-PR-06-0131 Rev A, 19 April 2007

[12] Portock, B.M., "Mars Approach Navigation Challenges", AIAA Paper 2000-4425, Proceedings of *Astrodynamic Specialist Conference*, Denver, CO, Aug. 2000.

[13] Prince, J. L., Desai, P. N., Queen, E. M., and Grover, M. R., "Lander Entry, Descent, and Landing Operations Analysis for the Mars Phoenix Lander," AIAA Paper 2008-7349, Proceedings of *AIAA Guidance, Navigation, and Control Conference,* Honolulu, HI, Aug. 2008.

[14] Grover, M., "Overview of the Phoenix Entry, Descent and Landing System Architecture" AIAA Paper 2008-7218, Proceedings of *AIAA Astrodynamic Specialist Conference*, Honolulu, HI, Aug. 2008.