

Development and Evaluation of an Operational Aerobraking Strategy for Mars Odyssey

Paul V. Tartabini,* Michelle M. Munk,[†] and Richard W. Powell*
NASA Langley Research Center, Hampton, Virginia 23681

The Mars 2001 Odyssey Orbiter successfully completed the aerobraking phase of its mission on 11 January 2002. The support provided by NASA's Langley Research Center to the navigation team at the Jet Propulsion Laboratory, California Institute of Technology, in the planning and operational support of Mars Odyssey aerobraking is discussed. Specifically, the development of a three-degree-of-freedom aerobraking trajectory simulation and its application to both preflight planning activities and operations is described. The importance of running the simulation in a Monte Carlo fashion to capture the effects of mission and atmospheric uncertainties is demonstrated, and the utility of including predictive logic within the simulation that could mimic operational maneuver decision making is shown. A description is also provided of how the simulation was adapted to support flight operations as both a validation and risk reduction tool and as a means of obtaining a statistical basis for maneuver strategy decisions. This latter application was the first use of Monte Carlo trajectory analysis in an aerobraking mission.

Nomenclature

C_A	=	axial force coefficient
C_N	=	normal force coefficient
C_Y	=	side force coefficient
h_a	=	apoapsis, altitude, km
h_p	=	periapsis altitude, km
h_s	=	scale height, km
\dot{Q}_{ref}	=	reference heat rate, W/cm ²
\dot{Q}_{per}	=	periapsis heat rate, W/cm ²
V_{atm}	=	velocity relative to atmosphere, m/s
α	=	angle of attack, deg
β	=	sideslip angle, deg
Δh	=	altitude change, km
Δh_p	=	periapsis altitude change, km
ΔV	=	velocity change, m/s
ρ	=	density, kg/km ³
ρ_0	=	reference density, kg/km ³
σ	=	standard deviation of density scale factor
τ	=	dust opacity factor

Introduction

ON 28 May 2002 a NASA press release heralded the 2001 Mars Odyssey Orbiter's discovery of vast quantities of water ice located beneath the surface of Mars in expansive areas surrounding the planet's south pole.¹ This important finding, which may help determine what happened to the abundant water that once flowed on the Martian surface, as well as to set the direction of future Mars exploration missions, marked NASA's successful return to Mars following the loss of the Mars Climate Orbiter and Mars Polar Lander missions in 1999.

After those well-publicized failures, NASA refocused its planetary exploration strategy and adopted a more extensive risk management style. The first mission to fly under this new approach was the 2001 Mars Odyssey Orbiter, which incorporated many risk-reduction measures that were derived from internal project assessments and external review boards. One of the actions taken by the

project management team at the Jet Propulsion Laboratory (JPL), California Institute of Technology, was to conduct an independent verification of Mars aerobraking, one of the highest risk segments of the mission.

This paper describes some of the independent analysis that was performed by NASA Langley Research Center (LaRC) and discusses the development of a three-degree-of-freedom (three-DOF) flight simulation of the aerobraking phase of the mission. In addition to verifying independently analysis that had already been conducted, the NASA LaRC simulation was used to perform numerous trade studies throughout the preflight planning stage of the mission, as well as to support the decision-making process during flight operations.

The Mars Odyssey Orbiter, which is shown in its aerobraking configuration in Fig. 1, arrived at Mars on 24 October 2001, following a seven-month interplanetary cruise, and was initially propulsively inserted into an 18.6-h elliptical orbit. After four orbits in which the spacecraft systems were verified, Odyssey commenced a 77-day aerobraking phase where the spacecraft made 332 successive drag passes through the Martian atmosphere to reduce nonpropulsively its orbit period to the roughly 1.8-h period required for the mapping phase of the mission. The propellant savings obtained from using aerobraking reduced the required mass of the spacecraft by more than 200 kg, which significantly lowered launch costs by enabling the use of a Delta II class launch vehicle. Aerobraking was completed in January 2002, and Odyssey's two-and-one-half-year mapping mission began just over a month later after four small maneuvers were performed to finalize the science orbit.

The aerobraking phase was considered the most hazardous part of Odyssey's mission. To initiate aerobraking, the periapsis altitude of the initial capture orbit was lowered from 292 km to a point within the Martian atmosphere. Thus, with each periapsis, energy was removed from the spacecraft's trajectory for a brief segment of each orbit, and the orbit was incrementally shrunk to the proper size required for the mapping mission. Without aerobraking, additional fuel would have been needed to provide a 1.08-km/s ΔV to slow the spacecraft down propulsively. The risk enters into aerobraking because the friction from the atmosphere that depletes orbital energy also causes surface temperatures to rise on the spacecraft. This problem is compounded by the difficulty in accurately predicting the atmospheric conditions on Mars due to the uncertainty in the atmospheric dynamics that produce orbit-to-orbit variability that is not predicted well by current atmospheric models. To accommodate this uncertainty, Odyssey's trajectory had to be closely monitored on each orbit and occasionally adjusted to control how deeply Odyssey penetrated into the Martian atmosphere. If Odyssey penetrated too

Received 20 December 2002; accepted for publication 4 November 2004. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0022-4650/05 \$10.00 in correspondence with the CCC.

*Senior Research Engineer. Member AIAA.

[†]Lead Systems Engineer for Aerocapture In-Space Propulsion Technology Project Office. Senior Member AIAA.




Fig. 1 Mars 2001 Odyssey Orbiter in aerobraking orientation.

deeply, temperatures on the spacecraft could exceed the thermal limits of the solar arrays. If Odyssey penetrated too shallowly, the aerobraking mission element could drag on for months, keeping the orbit from being aligned properly for the science mission. The periodic adjustments were made by performing small, propulsive aerobraking trim maneuvers (ABMs) at apoapsis to raise or lower the periapsis altitude depending on the predicted behavior of the spacecraft over the next day's orbits.

A significant amount of planning was required to establish the operational strategies and guidelines used to monitor and control Odyssey during aerobraking. The purpose of this paper is to describe the aerobraking simulation that was developed for Mars Odyssey mission planning, to demonstrate how it was used to provide data to the project before launch to aid in decision making, to discuss its function as a testbed for screening operational strategies during preflight planning, and to show how it was adapted to assist the project team in making daily operational decisions while the spacecraft was aerobraking. Because much of the workload during operations involved deciding whether or not ABMs were needed to adjust Odyssey's closest approach altitude, it was imperative that a representative form of the operational decision-making logic be included in the simulation to understand truly how changes in strategy would affect the various trajectory parameters. Also, in the preflight planning stage of the mission, it was discovered that small perturbations due to modeling uncertainties prevented the accurate prediction of many of these parameters based on only a single aerobraking simulation. Consequently, it was necessary to run the simulation in a Monte Carlo fashion by using a range of input uncertainties to deduce the mean and standard deviation of the variables of interest. Not only was this approach used extensively during mission planning before launch, but it also marked the first time that Monte Carlo analysis was used to support flight operations for an aerobraking mission. By having a completely independent simulation with these capabilities, it was possible to provide the JPL-led navigation team with a continuous validation of their results, as well as an additional analysis tool that gave a probabilistic basis for making maneuver decisions and a means of assessing how operational strategy affected aerobraking progress.

Approach

Simulation Development

During the preflight planning stage of the 2001 Mars Odyssey Mission, a three-DOF simulation of the entire aerobraking phase

Table 1 Orbital elements used to initialize aerobraking simulation, Mars-centered Mars mean equator and IAU vector of J2000 coordinate system

Initial condition	Nominal value
Apoapsis, km	25,000–30,000
Inclination, deg	93.410
Longitude of ascending nodes, deg	28.0973
Argument of periapsis, deg	110.6018
True anomaly, deg	180

of the mission was developed by using the program to optimize simulated trajectories (POST).² This program is a widely used trajectory simulation and optimization tool that has been employed in the design, analysis, and operational support of other planetary exploration missions, including Mars Pathfinder.³ POST was the ideal tool for performing this analysis because it could be easily modified by an experienced user to include nonstandard trajectory models that made the simulation more accurate. In addition, the program's adaptability proved invaluable for modeling logic within the mission simulation that mimicked operational decision making, a capability that enabled rapid screening of various maneuver control strategies that were considered for use during aerobraking operations.

POST was used to integrate the three-DOF equations of motion for the complete aerobraking trajectory. The nominal orbital elements used to initialize the simulation are shown in Table 1. An initial periapsis altitude was not specified in Table 1 because it varied from case to case and was generally chosen so that the peak heat rate of the first drag pass was within a predetermined range. A range of apoapsis altitudes for the initial orbit were considered in the aerobraking trade studies before flight because the precise period of the capture orbit following the Mars orbits insertion (MOI) maneuver was not known. The uncertainty arose because the insertion maneuver was terminated when the oxidizer used to power the main engine was depleted, and the uncertainty in the performance for this type of burn is generally larger than that with a more traditional accelerometer derived cutoff.⁴

The aerobraking phase of the mission was divided into four segments: walk-in, main phase, end game, and walk-out. Each of these segments had a different constraint that limited how far into the Martian atmosphere the spacecraft could penetrate. During walk-in, the periapsis altitude of the spacecraft was gradually lowered in

seven predetermined steps over the course of 12 orbits from the initial capture orbit value of from 292 to 111 km, which was well within the sensible region of the upper Martian atmosphere. This careful transition to the nominal aerobraking altitudes allowed the spacecraft team to evaluate the spacecraft response to the aerobraking loads while the atmospheric modeling teams analyzed the effective drag that was observed relative to the predictive models. No atmospheric data had ever been taken during winter in the northern latitudes (the season and location where *Odyssey* would be performing most of its aerobraking passes) and so evaluating the models in this early phase was particularly important. It was difficult to model the walk-in phase of the mission because the maneuver decisions were so subjective and depended on the latest atmospheric observations. Consequently, simulations used for mission planning purposes were often initiated at main phase, immediately after walk-in.

A typical simulation began at the apoapsis state before the start of main phase aerobraking and included several hundred passes through the atmosphere, each gradually removing orbital energy until the apoapsis altitude was reduced to 400 km. During each orbit, the effects of gravity (from both Mars and the sun), atmospheric drag, and ABMs were modeled. The equations of motion were integrated using a fixed-step fourth-order Runge–Kutta (RK4) integrator inside the atmosphere with a 1-s time step and a variable-step Krogh integrator outside the atmosphere.⁵ When compared to an independent simulation generated with the JPL code DPTRAJ, the Krogh/RK4 integration scheme used in POST was shown to produce differences of less than 0.1 m in position and 0.1 mm/s in velocity for a test simulation of a completely exoatmospheric orbit integrated for a duration of nine days. The close agreement between the two test simulations suggested that the numerical errors in the POST aerobraking simulation were well within the expected level of error from other, less certain models, for example, atmosphere and aerodynamics.

During the main phase segment of aerobraking, the limiting constraint on the spacecraft's periapsis altitude was the peak heat rate. That is, for each drag pass the maximum heat rate encountered by the spacecraft had to be within a predetermined heating control corridor that was related to the material temperature limits of the vehicle. For trajectory calculations, the corridor was parameterized in terms of a reference-heating indicator,

$$\dot{Q}_{\text{ref}} = \frac{1}{2} \rho V_{\text{atm}}^3 \quad (1)$$

A representative corridor is shown in Fig. 2. The upper corridor limit was chosen to ensure that there would be sufficient margin to prevent the spacecraft thermal limits from being violated in the event of a worst-case spike in density. Typically, the amount of margin was set at 80–100% of the flight allowable limit, for example, an 80% margin was a factor of 1.8 less than the flight allowable limit. This safety factor was based on the amount of orbit-to-orbit variability that was observed during the aerobraking phase of the Mars Global Surveyor mission.⁶ The placement of the upper corridor was also important because it affected the duration of aerobraking. When a




Fig. 2 Representative heating control corridor used in the POST simulation.

more conservative approach to aerobraking was taken by adding margin and lowering the upper corridor limit, larger atmospheric density perturbations could be tolerated, but less energy was removed during each drag pass, and the total duration of aerobraking increased.

The other defining feature of the heating control corridor was its width, which established its lower bound. The corridor width affected aerobraking duration because a narrower corridor kept the peak heat rate on each orbit closer to the top of the heating control corridor where the aerobraking altitudes were lowest and the most energy could be safely removed from each orbit. However, decreasing the width of the corridor tended to increase the ABM frequency because more ABMs were needed to adjust the trajectory to stay within the more constraining (narrower) corridor.

During mission planning the POST simulation was used extensively to determine how various corridor control strategies affected the duration of the aerobraking phase, which was important for two main reasons. First, it was desirable to complete aerobraking in the fewest possible orbits because there was a certain amount of risk associated with each drag pass. Reducing aerobraking orbits became a prime consideration during flight, when the POST simulation was used to estimate how many aerobraking orbits could be eliminated by adopting more aggressive maneuver strategies. In addition, aerobraking duration was also important because the alignment of the spacecraft's orbit became less favorable for power collection the longer that aerobraking continued. As time passed, the local true solar time (LTST) of the orbit's descending equator crossing decreased by 1–2 min/day due to the motion of Mars about the sun. As the LTST fell, the time that the spacecraft's solar arrays were shadowed by Mars on each orbit grew longer. If the LTST of the descending node decreased below 1400 hrs, the power degradation due to the increasing duration of the solar occultations would be unacceptable. Thus, the final LTST was an important factor in establishing the upper bound of the corridor. For mission planning, the final LTST limit was set at 1418 hrs to provide an additional margin in the event of a major dust storm or other unplanned interruption of aerobraking.

The desire to reduce aerobraking duration and maximize the final LTST by narrowing the corridor had to be balanced with the need to avoid an excessively high ABM frequency. Consequently, the corridor width was ultimately selected to meet the project goal of performing no more than 20–30 ABMs over the course of aerobraking.

In the two remaining phases of aerobraking following the main phase, endgame and walk-out, different constraints limited the depth of the drag passes, but the same corridor philosophy was used to control the spacecraft's trajectory. During endgame, the duration of the drag passes were longer, and, consequently, the total integrated heat load of the vehicle over a drag pass became the limiting constraint. The corridor was lowered during this part of aerobraking to ensure that the heat load constraint was satisfied (Fig. 2). Walk-out, as the name suggests, was the part of the mission where *Odyssey* was gradually raised out of the atmosphere before aerobraking termination. During walk-out, the spacecraft's orbit lifetime was the constraint that limited how far the spacecraft could dip into the atmosphere. The project mandated that the spacecraft always have at least a 24-h lifetime, which was defined as the time it took for the apoapsis altitude to decay to 300 km if no ABMs could be performed (due to a loss of ground contact or other spacecraft failure). The lifetime constraint was parameterized as a heat rate limit and was reflected in the heating control corridor (Fig. 2). The transition from endgame to walk-out occurred when the apoapsis altitude was roughly 1500 km, where the 24-h lifetime limit became the dominant constraint. Near the end of walk-out when the orbit period was less than 2 h, an orbit with a 24-h lifetime had to be capable of making as many as 13 drag passes without a periapsis adjustment. Thus, it was typical during this part of aerobraking for an ABM to be performed each day.

Corridor Control

In the three-DOF POST simulation when a violation of the upper corridor limit was predicted, a small impulsive ABM would

be performed at apoapsis to raise the spacecraft's periapsis altitude, resulting in lower heat rates on subsequent drag passes and a lower position within the corridor. Likewise, when a lower corridor violation was predicted, an ABM was performed to lower the periapsis altitude, thus raising the position within the corridor and increasing the rate of aerobraking because more energy would be removed from each orbit. During flight operations, actual maneuver decisions were more subjective and often considered observed short-term changes in atmospheric variability as well as other operational concerns. Consequently, the heating control corridor was used more as a guide for the project management during operations, and predicted violations did not always result in ABMs being performed.

Keeping the spacecraft within the corridor required continuous monitoring and prediction of future behavior because the peak reference heat rate [Eq. (1)] often fluctuated significantly from one orbit to the next. These fluctuations were almost entirely due to variations in atmospheric density caused by changes in periapsis altitude and inherent atmospheric variability. Changes of as much as 1.5 km in periapsis altitude could occur between successive orbits because of gravitational perturbations. Because the Mars gravitational field is known to a high level of precision, these altitude changes could be accurately predicted in advance. Much harder to predict were the large variations in density that were independent of altitude and have been known to occur throughout the atmosphere of Mars, especially above 75 km (Ref. 7). For atmospheric modeling within the three-DOF POST simulation, this density variability was divided into two components, a random part representing variation due to weather processes that can not be predicted in advance and a longitudinally dependent term that accounted for the existence of density waves, which had been observed at aerobraking altitudes during the Mars Global Surveyor (MGS) mission.⁸ The contribution of this longitudinally dependent component to the overall density variability was frequently substantial because Odyssey's aerobraking orbit was nearly polar, and the longitude of periapsis often varied significantly from one orbit to the next.

None of these orbit-to-orbit changes in density were biased in any direction, and the density was just as likely to rise on the next orbit as it was to fall. The lack of bias was due to the random element of atmospheric variability, as well as to the orbit period decreasing throughout aerobraking and the regions of the atmosphere and gravity field through which the spacecraft passed continuously changing. Although changes in heat rate due to density fluctuations between orbits were unbiased, there were trends in heat rate over multiple orbits that were caused by a natural drift in periapsis altitude due to the oblateness of Mars. Even if the periapsis radius remained constant over a series of orbits, the oblateness-induced precession of the argument of periapsis caused the latitude of periapsis to shift and the altitude to vary because of the oblate shape of Mars. This effect tended to cause the heat rate to decrease naturally during the first part of aerobraking when the periapsis point moved northward toward the pole and to increase after passage of the pole, when periapsis began to move southward.

Each time a predicted corridor violation triggered an ABM in the POST simulation, the magnitude of the maneuver was determined to take advantage of the natural drift in periapsis altitude. That meant performing maneuvers that moved the adjusted trajectory high in the corridor early in the mission before periapsis moved over the pole (so that it would naturally drift lower in the corridor) and low in the corridor for the rest of the mission (knowing it would drift higher). This same strategy was employed during operations to keep the adjusted trajectory within the corridor for as long as possible to reduce the ABM frequency.

Simulation Models

A number of mission-specific models were incorporated into the three-DOF simulation to increase its accuracy and realism for use during mission planning and flight operations. Outside the atmosphere, gravity was the principal force that affected the three-DOF trajectory. Effects due to solar radiation pressure were found to have a negligible effect on the three-DOF aerobraking trajectory

and were, therefore, ignored. Mars' gravity was modeled by using the JPL MGS75E gravity field.⁹ Even though this model included terms up to a degree and order of 75, a truncated version that only made use of the terms up to a degree and order of 20 was employed in the POST simulation because the computational time required to perform a simulation decreased significantly without a noticeable change in accuracy. The third-body effect of the sun was also modeled because it made a noticeable effect during the first 50–60 orbits of aerobraking when the orbit period was fairly large.

For altitude calculations, Mars was modeled as a biaxial ellipsoid with the same ellipsoid radii that were used to parameterize the atmospheric properties in the Mars-GRAM atmospheric model (discussed hereafter). In addition, the planetary rotation rate was set as 350.891983 deg/Earth day, which was taken from Ref. 10. Finally, solar position vectors, used for LTST calculations and determination of the third-body effect of the sun, were computed by using an analytic ephemeris that employs Van Flandern and Pulkkinen's low-precision formulas.¹¹

For simulation of flight within the atmosphere, an aerodynamic database was developed to provide force and moment coefficients for a range of vehicle attitudes and atmospheric densities that were expected to be encountered during aerobraking. The database utilized in the POST three-DOF simulation was created by applying free molecular and direct simulation Monte Carlo computational fluid dynamic techniques to a high-fidelity geometry model of the spacecraft.¹² These methods are ideal for aerothermodynamic analysis of flight vehicles in rarefied transitional flow regimes such as the upper atmosphere of Mars, where all of Odyssey's drag passes were performed. The database that was developed provided predictions of the force and moment coefficients for angles of attack and sideslip between ± 60 deg. The moment coefficients were used in the three-DOF simulation to determine the angle of attack and sideslip at each time step that resulted in a trimmed vehicle orientation, with the static pitching and yawing moments about the center of gravity equal to zero. The functional dependence of the axial coefficient on density is shown in Fig. 3 for $\alpha = \beta = 0$, which was close to a typical trimmed attitude of the spacecraft at periapsis. The reference area associated with the LaRC aerodynamic database was 11.14 m² (based on the projected area of the solar arrays), and the moment reference length used to dimensionalize the pitching and yawing moments was 5.43 m. Uncertainties in the aerodynamic modeling were also determined for use in the Monte Carlo analysis, which will be discussed in a later section.

The Mars Global Reference Atmospheric Model (Mars-GRAM) was used to model the atmospheric density in the POST simulation. Mars-GRAM is an engineering model of the Martian atmosphere that synthesizes the results of state-of-the-art global circulation models, which have been calibrated by using data obtained from previous missions including MGS, Viking, and Pathfinder.¹³

Within the POST simulation, the density was updated by Mars-GRAM at each time step and was a function of the altitude, date, latitude, and longitude of spacecraft. The atmospheric interface altitude (above which vacuum conditions are assumed and below which




Fig. 3 Variation of axial force coefficient with atmospheric density for $\alpha = \beta = 0$ deg.

atmospheric properties are computed) was defined as 170 km. Another significant input was the dust opacity τ , which is a measure of the background dust level of the Mars atmosphere. (Dust level increases with larger values of τ .) As the background dust level goes up, the atmosphere expands in response to dust-induced atmospheric heating, and the density increases at aerobraking altitudes. In the simulation, τ was assumed to remain constant throughout aerobraking.

The POST simulation was used extensively for understanding the differences between different versions of Mars-GRAM. The Mars-GRAM 2000 (MG2K) version of the model was selected as the official atmospheric model for use on the Mars Odyssey Project and was employed in the POST simulation for much of the pre-flight planning activities and all of flight operations. One reason that MG2K was chosen was the ease by which it could be altered to reflect observed atmospheric phenomena, through the use of a seasonally dependent height offset and a longitude dependent density wave model. For Odyssey mission planning, a longitude-dependent wave (LDW) model was constructed that provided a least-squares fit to the data obtained during MGS aerobraking.¹⁴ Engineering judgment was used to extend the MGS data, which were taken primarily at southern latitudes, to the northern latitudes where Odyssey aerobraking was to take place by assuming that wave activity was symmetric at the poles. The resulting wave-3 (three wave peaks over a 0–360 longitude range) LDW model had amplitude and phase coefficients that were a function of latitude. In the POST simulation, these coefficients were updated at the beginning of each drag pass and were assumed to remain constant for the duration of the pass. The coefficients were used to calculate a multiplier that was applied to the density computed by MG2K. The variation of the LDW-derived wave multiplier in the region of Mars Odyssey aerobraking is shown in Fig. 4.

A final model that was instrumental in making the POST simulation an effective tool for aerobraking design and planning was the ABM decision-making logic. The logic implemented in the simulation was intended to mimic the method in which ABM decisions would be made during operations, as well as the timing constraints that would restrict when maneuvers could be performed. During operations, the decision to perform an ABM was typically based on predictions of future spacecraft behavior. Also, only one ABM could be performed per day because the decision involved a significant amount of analysis that was presented and discussed at daily programwide planning meetings. In addition, ABMs could only be performed on the last apoapsis of a command sequence (called a slot) so that the maneuver could be factored into orbit determination predictions, which were only generated once per sequence. The number of orbits in the command sequence was a function of how many drag passes occurred during the sequence and increased throughout aerobraking due to the declining orbit period. Early in aerobraking when the orbit period was greater than 12 h, there was typically only one

slot per planning period, but by the end of aerobraking there were as many as five slots in a 24-h interval.

To model this process within an aerobraking mission simulation, a predictive technique was developed and employed within POST that determined when an ABM had to be performed to prevent future corridor violations. This decision-making logic was implemented in a way that ensured compliance with the various timing constraints that would be active during operations. The process began when the simulation reached the first ABM opportunity (slot) in a 24-h planning period. To determine how the trajectory would behave if an ABM were not performed, the mission simulation was paused, and a separate, secondary simulation was spawned and propagated forward, without a maneuver, to the first slot in the next planning period. The peak heat rates at each drag pass in the spawned case were compared to both the upper and lower bounds of the heating control corridor, and, if one was exceeded, a maneuver was scheduled in the main simulation at the first slot before the predicted violation. In the event that both an upper and lower corridor violation were detected, a periapsis-raise maneuver was implemented to avoid the breach in the upper corridor limit. When an ABM was required, the information from the spawned case was also used to size the maneuver magnitude so that the maximum heat rate over the 24-h prediction interval was at an appropriate location within the corridor. As in operations, this location was chosen to take advantage of the natural altitude drift and was, thus, placed 80% of the way between the bottom and top of the corridor early in the mission when the periapsis was moving northward, as well as at the 40% level when it was moving southward after passage over the pole. When the simulation reached walk-out, this procedure had to be modified because the required maneuver magnitude could not be accurately determined from the single spawned simulation that was performed without an ABM because, at this point in the mission, there were 12 or more orbits in a planning period, and the effect that a maneuver performed early in the planning interval had on drag passes near the end of the interval was nonlinear. Consequently, the required maneuver magnitude had to be determined iteratively by spawning additional simulations that included the ABM.

Monte Carlo Analysis

The POST three-DOF simulation discussed in this paper was originally developed as an independent validation of an earlier simulation that was used to perform initial Odyssey aerobraking studies. Generally, results from both simulations compared well, but there were significant differences in the predicted amount of ΔV required for ABMs. In the examination of differences, it was discovered that running a single simulation was not adequate for obtaining an accurate estimate of ABM ΔV usage. The reason had to do with the nature of corridor-based ABM logic because small differences in modeling combined with strict adherence to the constraints could cause maneuvers to be triggered differently in different simulations. These differences propagate forward, often resulting in significant differences in total ΔV , number of ABMs, and duration.

The large disparity observed in runs starting with similar initial conditions during preliminary analyses indicated that, for long-term aerobraking simulations, a single simulation was not adequate to predict certain parameters such as ABM ΔV usage. Furthermore, if only a single-trajectory simulation was performed, the effect of known modeling uncertainties could not be ascertained. Therefore, the POST simulation was modified so that it could be implemented in a Monte Carlo analysis. Running the simulation in this fashion captured the coupling of the various uncertainties, thus enabling the true variance of the key aerobraking parameters to be determined. The Monte Carlo simulation that was ultimately developed became a valuable tool that was used during preflight planning to evaluate a range of aerobraking design options and, throughout operations, to help make maneuver decisions and assess the progress made toward meeting the final LTST constraint.

Each Monte Carlo analysis generally included 1000 simulations. Computational constraints prevented the use of more than 1000 samples because a single aerobraking trajectory, which typically encompassed 300–500 drag passes, took approximately 45 min of




Fig. 4 Variation of MGS LDW-derived wave multiplier in the region of Mars Odyssey aerobraking.

CPU time on a 400-MHz Silicon Graphics Origin 2000 processor. Running only 1000 samples produced reasonable statistics because fluctuations in the mean and standard deviation of the various trajectory parameters of interest had stabilized by that point.

The only quantities that were dispersed in the Monte Carlo analysis were those that influenced the atmospheric flight of the vehicle (aerodynamic force coefficients and the atmospheric density) because the effect of these uncertainties overwhelmed any uncertainty in the initial state or the gravity field, which is well known. A description of the dispersed quantities is listed in Table 2. The three aerodynamic force coefficients were assumed uncorrelated and normally distributed about their nominal values with a standard deviation of 15%. Uncertainties in atmospheric density were partitioned into both large- and small-scale components. The large-scale component was the background dust level (opacity) of the atmosphere, which tended to increase the density uniformly at all altitudes as the atmosphere became dustier. For each simulation in the Monte Carlo analysis, a different value of τ was selected from a uniform distribution ranging from 0.3 (low dust) to 1.6 (high dust) and was kept constant for the duration of the simulation. The nominal value of τ was assumed to be 1.0. Short-scale dispersions were also applied to model the orbit-to-orbit fluctuations that had been encountered during MGS aerobraking. These variations were implemented by varying the constant term in the MGS-derived LDW model. A new constant was selected at the beginning of each drag pass and acted as a multiplier on the density derived from the baseline Mars-GRAM. The standard deviation of this density multiplier (nominally set to 21%) was varied throughout mission planning to assess the affect of different levels of orbit-to-orbit density variations on the aerobraking design. Also, during operations, this parameter was modified to reflect the observed level of atmospheric variability.

Results and Discussion

Preflight Planning

In the months before arrival at Mars, the POST three-DOF simulation was used to perform numerous trade studies and analyses that impacted many aspects of the mission design and flight preparation. Issues relating to aerobraking strategy such as the size and placement of the heating control corridor, the effect of changing the thermal limit lines, and the efficacy of various ABM decision-making schemes were frequently assessed by modeling the option within the POST simulation and performing a Monte Carlo analysis to evaluate its effect on a number of trajectory parameters, including the final LTST, number of maneuvers, and number of drag passes. In addition, similar analyses were performed to understand how

atmospheric characteristics such as scale height and orbit-to-orbit variability affected the aerobraking trajectory. Results from these studies facilitated the selection of MG-2K as the official atmospheric model for the project. Some highlights of these preflight analyses are presented here to demonstrate the capabilities of the simulation and to show some results that were generated in the months leading up to aerobraking.

Immediately after MOI but before the initiation of main phase aerobraking, a nominal mission Monte Carlo analysis was performed to estimate the duration of the aerobraking phase and final LTST. The simulation was initiated at the projected post-walk-in apoapsis altitude, which corresponded to an initial orbit period of 18 h. The atmospheric models for longitude-dependent density waves and orbit-to-orbit variability that were utilized were constructed from MGS data and were not necessarily representative of the Odyssey experience. Nevertheless, they represented the best estimate of atmospheric variability that was available at the time.

The results from this nominal mission Monte Carlo analysis were the first post-MOI predictions of conditions at the end of aerobraking and are summarized in Table 3. Nominal aerobraking termination (when h_a was reduced to 400 km) was estimated to occur between 2 January and 23 January, and the one-percentile low value of the final LTST was 1443 hrs, which was still 25 min above the 1418 hrs power degradation limit. The ability to calculate the 1- and 99-percentile values of these quantities was useful for demonstrating the potential for large spreads in the key aerobraking metrics.

A mean number of 25 ABMs was predicted to occur between the start of main phase and aerobraking termination, with 80% of them being up maneuvers (to raise the periapsis altitude). Down maneuvers only occurred during the first part of main phase, when the natural altitude drift tended to increase the periapsis altitude and move the peak heat rate on each orbit toward the lower corridor bound. The remainder of aerobraking, which occurred after passage over the pole, required up maneuvers to counter this effect. This part of the mission (second-half of main phase, endgame, and walk-out) required more maneuvers because it encompassed more drag passes due to the reduced orbit period. In fact, of the 370 total aerobraking orbits that were predicted, 115 occurred during the walk-out phase of the mission when the orbit lifetime constraint was dominant. The predicted number of walk-out drag passes was used by the navigation team to understand how many walk-out orbits could be eliminated by propulsively terminating aerobraking before the 400-km apoapsis altitude goal was reached. By the trade of the amount of fuel required for various aerobraking termination maneuvers against the number of drag passes that would be averted, it was possible to determine the point at which the overall mission risk was reduced by ending aerobraking early. During flight, approximately 30 m/s of fuel was eventually used to terminate aerobraking 2.5 days early, thus eliminating roughly 30 walk-out orbits.⁴

Another important result is evident from comparison of the initial conditions in Table 1 with the results in Table 3, where the inclination can be seen to decrease by approximately 0.2 deg over the course of aerobraking. This decrease occurred because the drag was aligned with the relative velocity vector, and the inclination was based on the inertial velocity, so that the angle between the two velocity vectors over time caused a shift in the inclination. The reduction was

Table 2 Input uncertainties implemented in Monte Carlo analysis

Quantity	Nominal value	Distribution	1- σ or min/max
C_A multiplier	1.0	Gaussian	$\pm 5\%$
C_N multiplier	1.0	Gaussian	$\pm 5\%$
C_Y multiplier	1.0	Gaussian	$\pm 5\%$
Density scale factor	1.0	Gaussian	$\pm 21\%$
Dust opacity τ	1.0	Uniform	0.3/1.6

Table 3 Nominal mission runout Monte Carlo statistics

Parameter	Mean	1-percentile	99-percentile	1-Sigma
Aerobraking duration	69.7	60.6	81.2	4.1
Aerobraking end date	12 Jan. 2002	02 Jan. 2002	23 Jan. 2002	—
No. of ABMs	25	19	33	3
No. of up maneuvers	20	15	27	3
No. of down maneuvers	5	3	7	1
ABM ΔV usage, m/s	17.9	15.2	21.2	1.3
Final inclination, deg	93.20	92.92	93.49	0.13
Final longitude of ascending node, deg	40.19	38.23	42.32	0.90
Final argument of periapsis, deg	6.54	-8.59	20.67	6.19
Final LTST	1506 hrs	1443 hrs	1522 hrs	7 min
Total no. of aerobraking orbits	370	321	428	22
No. of wall-out orbits	115	98	132	8

accounted for by biasing the inclination targeted by the MOI maneuver by 0.25 deg relative to the value desired for the final orbit. This bias was based on predictions from preflight aerobraking simulations. The results of these simulations also demonstrated that biasing the inclination slightly upward improved the final LTST because the higher inclination caused the line of nodes to precess more, thus slowing the reduction in LTST due to the motion of Mars about the sun.

A representative time history of each aerobraking orbit's peak heat rate is shown in Fig. 5, along with the nominal corridor that was in place at the beginning of the mission. As aerobraking commenced, it was understood that the corridor would likely change periodically as trends in atmospheric behavior were observed and the level of success that the project team had in forecasting its behavior became apparent. This idea was reflected in the design of the initial corridor, which had a width of 0.18 W/cm^2 and an upper bound that provided a 100% thermal margin for the first week after walk-in (10–15 orbits) when the atmosphere models were being calibrated and an 80% margin for the rest of main phase and end game. During flight, the 100% margin limit was maintained much longer than a week because of higher than expected atmospheric variability.

The results shown in Fig. 5 are from one of the 1000 perturbed cases run in the Monte Carlo analysis, which is why there are a significant number of upper and lower bound violations. Recall that the POST simulation paused at the first slot in each 24-h planning interval so that a procedure could be carried out to predict the peak heat rates over the next day, thus enabling maneuvers to be scheduled that would prevent future corridor violations. These predictions were made using a nominal, unperturbed MG2K atmosphere that included the longitude-dependent wave model but not dispersions in the density scale factor. Because maneuver decisions were based on assumed nominal atmospheric behavior, when the simulation was propagated forward using the simulated atmosphere that included the random density scale factor perturbations, it was possible for the corridor limits to be violated. This process was designed to mimic the operational procedure in which ABM decisions made during flight would be based on predictions from a simulation that was run using a fixed atmospheric model.

Before flight, numerous trade studies were performed to determine the bounds of the heating control corridor. The corridor design was important because it affected how far the spacecraft could dip

into the atmosphere, thus establishing the pace of aerobraking as well as the likelihood of thermal damage from unanticipated orbit-to-orbit fluctuations in atmospheric density. The influence of the corridor on the aerobraking trajectory can be seen by comparing three different Monte Carlo analyses, each employing a different corridor. The baseline case utilized the nominal corridor that was planned for use at the beginning of main phase aerobraking (100% margin for the first 10 orbits, 80% margin for the rest of aerobraking, 0.18 W/cm^2 corridor width). Two additional cases were included that demonstrated the effect of a different upper bound (130% margin) and a different corridor width (0.12 W/cm^2). Figure 6 shows the frequency of violating various levels of thermal margin for each of the three corridors. For instance, in the baseline case, roughly 4% of the drag passes in all 1000 perturbed simulations had a peak heat rate that exceeded the upper corridor limit, that is, that had a less than 80% margin. For the case with the narrower corridor, the number of upper-limit violations increased because random fluctuations in density were more likely to push the peak heat rate over the upper corridor bound. Conversely, there were significantly less corridor violations for the 130% corridor case because the peak density during each drag pass was lower and, therefore, the density perturbations, which were applied as a multiplier on the nominal MG2K density, were smaller.

The effect of the corridor design on other trajectory parameters is shown in Table 4. Decreasing the corridor width reduced the mean aerobraking duration by more than a week and provided an additional 15 min of LTST margin. However, the more constraining corridor also required 14 more ABMs than the baseline case, although only 1 m/s of additional ΔV was required for the extra maneuvers. When the upper corridor limit was reduced to the 130% margin level, nearly 100 additional drag passes were required over the baseline case, and the aerobraking duration increased by more than 40%, causing the final LTST to drop below the 1418 hrs limit. Ultimately, the baseline corridor was selected for use during the beginning of aerobraking because there was a low probability of exceeding the flight allowable thermal limit, and the ABM frequency was within the project goal of 20–30 maneuvers.

A final aspect of the corridor design that was performed with the POST simulation was the development of the walk-out corridor, which was designed to maintain the project-mandated 24-h orbit




Fig. 5 Heat rate time history for perturbed Monte Carlo case.




Fig. 6 Effect of corridor design on frequency of various thermal margin violations.

Table 4 Effect of corridor design on mean values of aerobraking trajectory results

Parameter	80% corridor, 0.18 W/cm^2 wide	80% corridor, 0.12 W/cm^2 wide	130% corridor, 0.18 W/cm^2 wide
Aerobraking duration	69.7	62.1	97.9
Aerobraking end date	12 Jan. 2002	04 Jan. 2002	09 Feb. 2002
No. of ABMs	25	39	22
ABM ΔV usage, m/s	17.9	18.9	18.8
Final LTST	1506 hrs	1521 hrs	1413 hrs
Total no. of aerobraking orbits	370	326	467
No. of walk-out orbits	115	114	120




Fig. 7 Relationship between periapsis heat rate and orbit lifetime.

lifetime constraint. Initial studies indicated that the lifetime constraint became dominant when the apoapsis altitude dropped below 1500 km. A technique was developed using Monte Carlo analysis to provide a statistical relation between heat rate and orbit lifetime, so that the same corridor control scheme used during main phase and endgame could also be employed during walk-out. The procedure was carried out at a discrete set of initial apoapsis altitudes between 400 and 1500 km. For each apoapsis altitude, 1000 perturbed simulations were performed using a range of input uncertainties that included aerodynamic coefficients, density perturbations, dust opacity, and initial longitude. The initial longitude was varied uniformly between 0 and 360 deg and had to be included as an uncertainty because it dictated which part of the longitude-dependent wave model would be sampled during the simulation. Perturbed cases that had more drag passes occurring during wave peaks tended to have lower orbit lifetimes.

Each simulation was started at periapsis with the altitude adjusted so that the initial heat rate for all 1000 cases was the same. The trajectory was then propagated forward, without maneuvers, until the apoapsis altitude fell below 300 km. The orbit lifetime was computed as the time between the first and last apoapses, and the 50-percentile value, derived from the results of all of the cases, was defined as the orbit lifetime at that particular apoapsis altitude and initial periapsis heat rate. Figure 7 shows a representative set of results where the variation of the 50-percentile lifetime is plotted for three different initial periapsis heat rates at a range of apoapsis altitudes. When it was noted where these results intersected the 24-h line (or any other line, for that matter), an upper walk-out corridor was constructed. During walk-out the corridor control strategy was modified to keep the periapsis heat rate of each drag pass beneath this upper limit to ensure that the orbit lifetime remained above 24 h. Although the lower limit was not important because the heat rate trended upward during walk-out, it was set at a small, positive value to prevent aerobraking from stalling out.

In addition to defining the heating control corridor, the POST simulation was also used to understand the effect that various atmospheric modeling parameters had on the aerobraking trajectory. This capability proved to be particularly useful when it became necessary to evaluate a new version of the Mars-GRAM atmosphere model midway through the aerobraking design process. Initial aerobraking design studies had been performed by using the Mars-GRAM 3.7 (MG37) model, which was similar to the Mars-GRAM model used to support MGS aerobraking. When the MG2K version became available and was incorporated into the Odyssey aerobraking simulation, the results differed significantly from those obtained previously with MG37. In particular, for every comparison case that was run, the aerobraking duration was 30–50% longer in the simulations that employed the MG2K model.

Further inspection of the results from both atmosphere models showed a large disparity in density scale height predictions over the range of altitudes and latitudes associated with Mars Odyssey aerobraking. The scale height (or lapse rate) is a measure of the rate at which density drops off with increasing altitude. This dropoff




Fig. 8 Example of scale height difference between MG37 and MG2K.

rate has a significant effect on the duration of aerobraking because it directly impacts the amount of energy that is removed during each drag pass.

This effect is illustrated in Fig. 8, where density profiles computed by both models are plotted for a typical drag pass. The two dashed lines show the density predicted by each model when the same trajectory was followed. Because the density predicted by MG2K was lower for the given altitude than that predicted by MG37, the corridor logic within the simulation would generally force the MG2K simulation to perform a deeper drag pass, thus enabling the peak density to match the MG37 value. When the altitude used to generate the density profiles in Fig. 8 was adjusted so that the peak density predicted by MG2K matched the peak value predicted by MG37, the curves still differed significantly because the lapse rates were different. Specifically, the MG2K case had a lower scale height, which caused the density to decrease at a faster rate than in the MG37 case. Consequently, less energy was removed from the drag pass because the integral of drag over the pass was lower. When this effect was repeated over many drag passes, the rate of aerobraking slowed significantly, causing the aerobraking duration to increase.

Figures 9a and 9b show a representative aerobraking trajectory plotted on top of density scale height contours that were generated with MG37 and MG2K, respectively. For this region of the atmosphere, it is clear that the structure of each atmospheric model is very different and that the scale heights predicted by MG37 are generally much larger than those predicted by MG2K. These modeling differences resulted in large discrepancies in many of the trajectory parameters that were predicted by the two simulations, a point that is demonstrated in Fig. 10, which shows a comparison of the final LTST envelopes computed from Monte Carlo analyses for a range of different initial orbit periods using both atmosphere models. The results shown are representative and were derived from earlier studies that utilized different initial conditions, preliminary simulation models, and a higher thermal margin corridor. Nevertheless, Fig. 10 shows that for the 99-percentile low final LTST to meet the 1400 hrs limit (the goal at the time of the study), the post-walk-in orbit period had to be less than 22 h with the MG37 model and less than 16 h with MG2K.

Odyssey's Atmospheric Advisory Group ultimately decided to select MG2K as the official atmosphere model for the project. This decision was made primarily because the scale heights predicted by MG2K were believed to be more accurate in the range of altitudes where aerobraking would be taking place based on data obtained during MGS aerobraking. MG37 was an appropriate choice for MGS aerobraking, which occurred at higher altitudes than Odyssey.

Flight Operations

During flight operations, NASA LaRC supported the flight navigation team in the daily aerobraking planning process by employing the POST simulation to evaluate maneuver options and to compare the effect of different atmospheric models on the predicted trajectory. The capability to perform Monte Carlo trajectory analyses was used extensively to provide a statistical basis for the options that




Fig. 9 Predicted density scale height contours: a) MG 37 and b) MG2K.




Fig. 10 Effect of atmospheric model on final LTST prediction.

were considered. The NASA LaRC results were used strictly to support the decision analysis, and all candidate maneuver choices had to be verified in the JPL flight-qualified simulation before an ABM decision was finalized.

NASA LaRC's participation in the daily operational process was an added risk-reduction measure taken by the project to augment the JPL navigation team. Because of the favorable 3-h time difference between the two centers, NASA LaRC was often able to screen many of the available maneuver options and deliver a package of initial results to JPL by the start of the JPL day. Also, having a completely independent simulation provided an effective method of validating the JPL simulation throughout aerobraking. This proved to be a valuable capability when a discrepancy that was detected

between the results of each simulation led to the discovery of an incorrect atmospheric modeling input. The 3-h time difference also aided the NASA LaRC effort because the JPL team would provide analyses they had done after the NASA LaRC day.

The operational procedure followed by NASA LaRC began each morning when a set of deterministic trajectory simulations, utilizing only nominal input data, was performed for each of the available atmospheric models. One of these models (called the navigation atmosphere model) was maintained by the navigation team and consisted of a multiplier that was applied to the baseline density predicted by MG2K. Additional models that included the effects of longitude-dependent density waves were continuously updated and refined throughout aerobraking as drag passes were completed. In general, the LDW models that were developed were not able to predict density accurately over a number of orbits and were, thus, used primarily to provide qualitative information on current atmospheric trends.¹⁵ Consequently, almost all of the eventual maneuver decisions were based on simulations that employed the navigation team's atmosphere model.

These initial deterministic simulations, which did not include maneuvers, were initiated from orbital elements derived from the most recent telemetry data and propagated forward over the next three 24-h planning intervals. During main phase and end game, the results from the various wave model cases were compared to a simulation that utilized a simplified gravity model (J_2 and J_3 terms only) to determine whether swings in heat rate were caused by predicted density waves or gravity-induced altitude perturbations. If there were no corridor violations caused by density waves for the orbits covered by the current planning interval, the LDW models were usually not considered any further in that day's analysis.

Monte Carlo analyses based on the deterministic simulations were performed for the cases that could affect the maneuver decision. In all of the Monte Carlo analyses that were conducted, 1000 perturbed simulations were run that included all of the uncertainties listed in Table 2 except for dust opacity, which was set at its nominal value of 1.0. The mean value of the density scale factor (the same multiplier used in the navigation atmosphere model) was the value that provided the best fit between MG2K predictions and data from recent drag passes. Likewise, its standard of deviation was based on the latest level of orbit-to-orbit density variability. Both statistical parameters were continuously tracked and updated any time they changed significantly. Midway through flight, the uncertainty in the aerodynamic force coefficients was lowered from ± 15 to $\pm 6\%$ for C_A and $\pm 10\%$ for C_N and C_Y , based on comparisons that were made between the aerodynamic database and accelerometer measurements.

Because not performing an ABM was always the first option considered, the initial Monte Carlo analysis that was performed each day was for the no-maneuver case. A typical heat rate corridor plot produced by this analysis is shown in Fig. 11. For each drag




Fig. 11 Predicted heat rate range from Monte Carlo analysis for no ABMs.

pass the mean, 1- and 99-percentile peak heat rates were plotted against the current heating control corridor and the flight allowable (FA) and immediate action (IA) limits. The IA limit was a further safeguard that was implemented to prevent the peak heat rate from exceeding the FA limit, and any violation of the IA line would force a periapsis-raise maneuver to be performed at the next apoapsis.

In general, the Monte Carlo analyses were used as a guide by the navigation team in making maneuver recommendations. However, the Monte Carlo results were often a major factor in ABM decisions, and in practice a periapsis-raise maneuver was always performed when a 99-percentile heat rate exceeded the FA limit and was usually performed when it exceeded the IA limit.

In Fig. 11, the 99-percentile heat rate obtained from using the density scale factor statistics from the current navigation wave model (mean = 0.8 and $\sigma = 0.20$) exceeded the FA limit. When the distance between the mean and 99-percentile heat rate values was linearly scaled, it was possible to estimate how the 99-percentile values would change for a different value of σ without running additional Monte Carlo cases. This technique was particularly useful in assessing the sensitivity of the heat rate statistics to short-term swings in orbit-to-orbit density variability. For the case shown in Fig. 11, a periapsis-raise maneuver was clearly required because there were 99-percentile violations of the FA limit predicted with both a nominal and a slightly lower estimate of σ .

When the no-maneuver Monte Carlo case had a 99-percentile heat rate that exceeded the IA line, a set of deterministic trajectory simulations was run that included an ABM at a slot before the violation. Each of these simulations utilized a different maneuver magnitude that was available to the spacecraft, corresponding to a predefined menu that was uploaded once per week. Figure 12 shows the predicted heat rate time history from these deterministic trajectories for the same case discussed in Fig. 11. Each of the three maneuvers that were considered lowered all of the predicted heat rates over the planning interval into the corridor. Monte Carlo analyses were then performed for the reasonable maneuver choices, to determine the effect of the ABM on the heat rate statistics. Figure 13 shows that, for the 0.85-m/s periapsis-raise maneuver, the 99-percentile heat rate was lowered significantly, and no violations of the IA line were predicted until orbit 201. Ultimately, this was the maneuver that was recommended by the navigation team because another maneuver was not likely to be needed for two days, and that was the desired ABM frequency at this time of the mission.

One phase of aerobraking in which Monte Carlo analysis was particularly useful was during the 3:1 resonance period when the orbit period was just over 8 h. During this time, Odyssey passed over the same three longitudes for several days, thus effectively amplifying the periapsis altitude variations caused by the Martian gravity field. The relationship between the gravity-induced altitude change Δh_p and longitude is shown in Fig. 14. Monte Carlo analyses predicted




Fig. 12 Effect of maneuver magnitude on predicted deterministic heat rate profile.




Fig. 13 Predicted heat rate range from Monte Carlo analysis for $\Delta V = 0.85$ m/s ABM.




Fig. 14 Altitude change as a function of longitude due to gravitational perturbations.




Fig. 15 Envelope of potential periapsis altitudes predicted during 3:1 resonance period.

that Odyssey would likely pass close to the 210°E longitude band (where a nearly 1.5-km drop in altitude was anticipated), thus potentially incurring this drop in altitude every three orbits. In Fig. 15, the 1- and 99-percentile Monte Carlo results were used to define the envelope of potential periapsis altitudes during the resonance period (orbits 70–90). Figure 15 indicates that the postresonance periapsis altitude had a total uncertainty of more than 14 km. By having advanced knowledge of this large uncertainty in periapsis altitude, the navigation team had more time to consider different maneuver strategies than if only a single deterministic simulation had been performed.




Fig. 16 Predicted range of orbit lifetime.

During the walk-out phase of aerobraking, the criterion for making ABM decisions became orbit lifetime. The same basic simulation and Monte Carlo approach used during main phase and end game were also used here, with the only difference being that each trajectory was propagated forward until the orbit lifetime constraint ($h_a < 300$ km) was met. Once the time of the constraint violation was known, the lifetime at any previous apoapsis was simply the difference between the final time and the time of that apoapsis. Whenever an ABM was performed in the simulation, another propagation had to be conducted to determine the orbit lifetime at apoapses occurring after the maneuver. Results from a typical walk-out Monte Carlo analysis are shown in Fig. 16. For this case, an ABM had to be performed at apoapsis 271 to satisfy the project requirement that the mean orbit lifetime always be above 24 h in the daily planning interval. The 1-percentile low lifetime differed by about 6 h from the mean lifetime, which was typical in these analyses. This result was important because it showed that the 24-h mean lifetime requirement superseded another project mandate that required the 1-percentile lifetime to be greater than 8 h to accommodate potential outages of the deep space network.

In addition to supporting daily maneuver decisions, the NASA LaRC simulation was the primary trajectory analysis tool that was used to predict the aerobraking termination date. This key parameter was tracked each week by the project to gauge aerobraking progress. Early in the mission, higher-than-expected atmospheric variability forced the project to take a more conservative approach to aerobraking by maintaining an upper corridor limit that provided a 100% thermal margin for most of main phase. The original plans going into aerobraking had called for the upper bound to be raised to the 80% level after the first 10–15 aerobraking orbits. The increased conservatism hampered the progress of aerobraking and created concern that the final LTST would fall below the 1418 hrs limit. Mission simulations that modeled various corridor control strategies were performed each week and compared to the Monte Carlo results from the current corridor scheme. This process is illustrated in Fig. 17, which shows the results of mission simulations beginning 13 orbits after walk-in, when the decision was made to continue using the 100% corridor. An 18-min penalty in final LTST was incurred by changing from the original strategy to the 100% corridor. The additional case indicated that the corridor would have to be lowered by approximately 0.08 W/cm^2 before the mean value of final LTST fell below the 1418 hrs limit. These types of results were useful in providing a measure of the cost of conservatism, which was helpful when daily maneuver decisions were being made to protect against potential heat rate violations.

Another technique was devised to estimate the theoretical best- and worst-case final LTST values that could be attained at various points in the mission. In this approach, fictitious trajectories were simulated that performed maneuvers at every apoapsis to force the peak heat rate of each orbit to equal the upper (best-case) or lower (worst-case) corridor limit. Typical results of this analysis are shown




Fig. 17 Effect of different corridor control strategies on mean final LTST prediction.




Fig. 18 Predicted best- and worst-case LTST profiles.




Fig. 19 Weekly progression of the predicted total number of aerobraking orbits.

in Fig. 18 and were used to gauge the effect of maintaining various positions within a given corridor on the final LTST.

A final parameter that was tracked using Monte Carlo analyses was the total number of aerobraking orbits. Each week this parameter was predicted using a Monte Carlo analysis that employed the most recent corridor control strategy, and a running plot was maintained throughout aerobraking. This plot, shown in Fig. 19, provided a measure of aerobraking progress and showed the effectiveness of various changes that were made in the corridor control strategy. After four weeks of aerobraking, the mean number of aerobraking

orbits predicted by the Monte Carlo analysis had stabilized near 420. At this time, the spacecraft began aerobraking in the polar regions of the atmosphere, and the orbit-to-orbit density variability decreased significantly. Mission Monte Carlo simulation results were used by the project to aid in the decision to increase the upper corridor to the 80% margin level during passage through this more quiescent part of the atmosphere. This shift in strategy significantly improved the rate of aerobraking, and the total number of predicted aerobraking orbits decreased to roughly 360. After the periapsis point moved over the pole, the 100% corridor was reinstated.

Summary

On 11 January 2002, a propulsive maneuver was performed to raise Odyssey's periapsis altitude out of the atmosphere, thus marking the end of the successful aerobraking phase of the mission. The spacecraft performed 332 drag passes without exceeding the design limits of the spacecraft, despite higher than expected levels of atmospheric variability during much of main phase. Furthermore, the LTST of the descending node at the end of aerobraking was 1504 hrs, well above the limit that was chosen to prevent power collection problems. The satisfaction of these mission constraints was due in large part to the ability of the flight operations team, which was composed of numerous geographically dispersed organizations, to work together. The effective coordination of the many diverse roles played by the various participants helped reduce risk and enhance the probability of mission success.

Acknowledgments

The authors thank John Smith, Julia Bell, and Robert Mase of the Jet Propulsion Laboratory, California Institute of Technology, for the opportunity to work with them on Mars Odyssey aerobraking and for their many contributions to the development of the aerobraking simulation. The authors also acknowledge the inputs made by the members of the NASA Langley Research Center Odyssey Team. A special thanks to Charles Davis for the hard work and long hours spent keeping the computers operating throughout the intense period of flight operations.

References

¹"Odyssey Finds Ice in Abundance Under Mars' Surface," NASA Press Release, 28 May 2001.

²Brauer, G. L., Cornick, D. E., and Stevenson, T., "Capabilities and Applications of the Program to Optimize Simulated Trajectories (POST)," NASA CR-2770, Feb. 1977.

³Braun, R. D., Spencer, D. A., Kallemeyn, P. H., and Vaughan, R. M., "Mars Pathfinder Atmospheric Entry Navigation Operations," *Journal of Spacecraft and Rockets*, Vol. 36, No. 3, 1999, pp. 348–356.

⁴Smith, J. C., Jr., and Bell, J. L., "2001 Mars Odyssey Aerobraking," *Journal of Spacecraft and Rockets*, Vol. 42, No. 3, 2005, pp. 406–415.

⁵Krogh, F. T., "Variable Order Integrators for the Numerical Solution of Ordinary Differential Equations," Jet Propulsion Lab., California Inst. of Technology, TU Document CP2308, NPO-11643, Pasadena, CA, May 1969.

⁶Tolson, R. H., Keating, G. M., Cancro, G. J., Parker, J. S., Noll, S. N., and Wilkerson, B. L., "Application of Accelerometer Data to Mars Global Surveyor Aerobraking Operations," *Journal of Spacecraft and Rockets*, Vol. 36, No. 3, 1999, pp. 323–329.

⁷Alexander, M., "Mars Transportation Environment Document," NASA TM-2001-210935, March 2001.

⁸Lyons, D. T., Beerer, J. G., Esposito, P., Johnston, M. D., and Willcockson, W. H., "Mars Global Surveyor: Aerobraking Mission Overview," *Journal of Spacecraft and Rockets*, Vol. 36, No. 3, 1999, pp. 307–313.

⁹Yuan, D. N., Sjogren, W. L., Konopliv, A. S., and Kucinskis, A. B., "The Gravity Field of Mars: A 75th Degree and Order Model," *Journal of Geophysical Research*, Vol. 106, No. E10, 2001, pp. 23,377–23,401.

¹⁰Lyons, D. T., "Preliminary Mars Planetary Constants and Models for the Mars Sample Return," Jet Propulsion Lab., California Inst. of Technology, IOM 312/99, DTL-1, Pasadena, CA, Jan. 1999.

¹¹Van Flandern, T. C., and Pulkkinen, K. F., "Low-Precision Formulae for Planetary Positions," *Astrophysical Journal Supplement Series*, Vol. 41, 1979, pp. 391–411.

¹²Takashima, N., and Wilmoth, R. G., "Aerodynamics of Mars Odyssey," AIAA Paper 2002-4809, Aug. 2002.

¹³Justus, C. G., and James, B. F., "A Revised Thermosphere for the Mars Global Reference Atmospheric Model (Mars-GRAM Version 3.4)," NASA TM-108513, July 1996.

¹⁴Dwyer, A. M., Tolson, R. H., Munk, M. M., and Tartabini, P. V., "Development of a Monte Carlo Mars-GRAM Model for Mars Odyssey 2001 Aerobraking Simulations," *Journal of the Astronautical Sciences*, Vol. 50, No. 2, 2002, pp. 191–211.

¹⁵Tolson, R. H., Dwyer, A. M., Hanna, J. L., Keating, G. M., George, B. E., Escalera, P. E., and Werner, M. R., "Application of Accelerometer Data to Mars Odyssey Aerobraking and Atmospheric Modeling," *Journal of Spacecraft and Rockets*, Vol. 42, No. 3, 2005, pp. 435–443.

R. Mase
Guest Editor