AIAA 2002-4537
The Development and Evaluation of an Operational Aerobraking Strategy for the Mars 2001 Odyssey Orbiter

Paul V. Tartabini
Michelle M. Munk
Richard W. Powell

NASA Langley Research Center
Hampton VA 23681-2199

AIAA Atmospheric Flight Mechanics Conference & Exhibit
5-8 August 2002 / Monterey, CA

For permission to copy or to republish, contact the copyright owner named on the first page.
For AIAA-held copyright, write to AIAA Permissions Department,
1801 Alexander Bell Drive, Suite 500, Reston, VA, 20191-4344.

THE DEVELOPMENT AND EVALUATION OF AN OPERATIONAL AEROBRAKING STRATEGY FOR THE MARS 2001 ODYSSEY ORBITER

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

Abstract

The Mars 2001 Odyssey Orbiter successfully completed the aerobraking phase of its mission on January 11, 2002. This paper discusses the support provided by NASA's Langley Research Center to the navigation team at the Jet Propulsion Laboratory in the planning and operational support of Mars Odyssey Aerobraking. Specifically, the development of a three-degree-of-freedom aerobraking trajectory simulation and its application to pre-flight planning activities as well as 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

α	Angle of attack, deg
β	Sideslip angle, deg
CA	Axial Force Coefficient
CN	Normal Force Coefficient
Δh	Altitude change, km
Δh_P	Periapsis Altitude Change, km
ΔV	Velocity change, m/s
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
ρ	density, kg/km^3
ρ_0	reference density, kg/km^3
σ	Standard deviation of density scale factor
τ	Dust opacity factor
V_{ATM}	Velocity relative to atmosphere, m/s

Introduction

On May 28, 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 upon the Martian surface as well as 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

* Aerospace Engineer, Vehicle Analysis Branch, Aerospace System Concepts and Analysis Competency, member AIAA.
† Senior Research Engineer, Vehicle Analysis Branch, Aerospace System Concepts and Analysis Competency, member AIAA.

Copyright © 2002 by the American Institute of Aeronautics and Astronautics, Inc. No copyright is asserted in the United States under Title 17, U.S. Code. The U.S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental Purposes. All other rights are reserved by the copyright

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) was to conduct an independent verification of Mars aerobraking, one of the highest risk segments of the mission. To perform this independent analysis, NASA's Langley Research Center (LaRC) developed a three-degree-of-freedom (3DOF) flight simulation of the aerobraking phase of the mission. In addition to independently verifying analysis that had already been conducted, the simulation was used by NASA LaRC to perform numerous trade studies throughout the pre-flight planning stage of the mission and 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 October 24, 2001 following a seven-month interplanetary cruise and was initially propulsively inserted into an 18.6-hour 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 non-propulsively reduce its orbit period to the roughly 1.8-hour 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-a-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 removing energy from the spacecraft's trajectory for a brief segment of each orbit and incrementally shrinking its orbit to the proper size required for its mapping mission. Without aerobraking additional fuel would have been needed to provide a 1.08 km/s ΔV to propulsively slow the spacecraft down. The risk enters into aerobraking because the friction from the atmosphere that depletes orbital energy also causes surface temperatures to rise on the spacecraft. For Odyssey, this problem was compounded by the difficulty in accurately predicting the atmospheric conditions on Mars. This difficulty was due to the uncertainty in the atmospheric dynamics that produce orbit-to-orbit variability that is not predicted by current atmospheric models. Therefore, Odyssey's trajectory had to be closely monitored on each orbit and occasionally adjusted to control how deep Odyssey penetrated into the Martian atmosphere. Too deep and temperatures on the spacecraft could exceed the thermal limits of the solar arrays, too shallow, and 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 for establishing 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, demonstrate how it was used to provide data to the project prior to launch to aid in decision making, discuss its function as a testbed for screening operational strategies during pre-flight planning, and show how it was adapted to assist the project team in making daily operational decisions while the spacecraft was aerobraking. Since 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 truly understand how changes in strategy would affect the various trajectory parameters. Also, in the pre-flight planning stage of the mission it was discovered that small perturbations due to modeling uncertainties prevented the accurate prediction of many of these parameters when only a single aerobraking simulation was run. Consequently, it was necessary to run the simulation in a Monte Carlo fashion 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 prior to 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 pre-flight planning stage of the 2001 Mars Odyssey Mission, a three-degree-of-freedom (3DOF) simulation of the entire aerobraking phase of the mission was developed 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 non-standard 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 3DOF 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 the table 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. The initial orbit had a range of initial apoapsis altitudes that were considered in the aerobraking trade studies prior to 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 exact amount of oxidizer available for MOI was not known in advance.[4]

A typical simulation began at the apoapsis state prior to 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 effect of gravity (from both Mars and the Sun), atmospheric drag and aerobraking trim maneuvers were modeled. The equations of motion were integrated using a fixed step 4th order Runge-Kutta (RK4) integrator inside the atmosphere with a 1-second 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 (e.g., atmosphere, aerodynamics, etc.).

Table 1. Orbital elements used to initialize POST aerobraking simulation

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

The aerobraking phase of the mission was divided into four segments: walkin, main phase, end game and walkout. Each of these segments had a different constraint that limited how far into the Martian atmosphere the spacecraft could penetrate. During walkin, the periapsis altitude of the spacecraft was gradually lowered in 7 predetermined steps over the course of 12 orbits from the initial capture orbit value of 292 km to 111 km, which was well within the sensible region of the upper Martian atmosphere. This careful transition to aerobraking altitudes was performed because of the known uncertainty in density predictions from the current Mars atmospheric models and the fact that no atmospheric data had ever been taken during winter in the northern latitudes, which was the season and location where Odyssey would be performing most of its aerobraking passes. It was difficult to model the walkin phase of the mission because the maneuver decisions were so subjective and depended upon the latest atmospheric observations. Consequently, simulations used for mission planning purposes were often initiated at Main Phase, immediately after walkin.

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:

$$\dot{Q}_{REF} = \frac{1}{2}\rho V_{ATM}^3 \qquad (1)$$

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 (e.g, 80% margin was a factor of 1.8 less than the flight allowable limit), which 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. That is, when a 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.

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 a number of reasons. First, it was desirable to complete aerobraking in the fewest possible number of orbits since 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. 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 shawdowed by Mars on each orbit grew longer. If the LTST of the descending node decreased below 2:00 PM, 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 2:18 PM in order to provide additional margin in the event of a major dust storm or other unplanned interruption of aerobraking.

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 since more ABMs were needed to adjust the trajectory to stay within the more constraining (narrower) corridor. 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, endgame and walkout, 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 (see Fig. 2). Walkout, as the name suggests, was the part of the mission where odyssey was gradually raised out of the atmosphere prior to aerobraking termination. During walkout 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-hr lifetime, which was defined as the time it took for the apoapsis altitude to decay to 300 km if contact was lost with the spacecraft and no ABMs could be performed. Near the end of walkout when the orbit period was less than 2 hours, an orbit with a 24-hr 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. The transition from endgame to walkout occurred when the apoapsis altitude was roughly 1500 km, which was the typical altitude where it was possible to meet the integrated heat load constraint that was dominant during endgame and not have a 24-hr lifetime. As in endgame, the lifetime constraint was parameterized as a heat rate limit, which was reflected in the heating control corridor used in the POST simulation.

Corridor Control

In the 3DOF POST simulation when a violation of the upper corridor limit was predicted, a small impulsive aerobraking trim maneuver 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 since more energy would be removed from each orbit. During flight operations actual maneuver decisions were more subjective and observed short-term changes in atmospheric variability were often taken into account. Consequently, the heating control corridor was used more as a guide 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 since the peak reference heat rate 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. Since 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.[7] For atmospheric modeling within the 3DOF 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 since 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 fact that the orbit period was decreasing throughout aerobraking and the regions of the atmosphere and gravity field through which the spacecraft passed were 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. In other words, 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 naturally decrease during the first part of aerobraking when the periapsis point moved northward towards the pole, and 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, and low in the corridor for the rest of the mission. This same strategy was employed during operations to keep the adjusted trajectory within the corridor for as long as possible in order to reduce the ABM frequency.

Simulation Models

A number of mission-specific models were incorporated into the 3DOF 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 3DOF trajectory, since effects due to solar radiation pressure were found to have a negligible effect on the 3DOF aerobraking trajectory and were therefore ignored. Mars' gravity was modeled using the JPL MGS75E gravity field.[9] Even though the 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 since it made a noticeable effect during the first 50-60 orbits of aerobraking when the orbit period was fairly large.

For altitude calculations the planet 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

below). 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 using an analytic ephemeris developed at NASA's Johnson Space Center that employs Van Flandern's low precision formulae.[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 3DOF simulation was created by applying free molecular and direct simulation Monte Carlo (DSMC) 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 3DOF 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 $\alpha = \beta = 0$ deg, 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 m^2 (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 using data obtained from previous missions including MGS, Viking and Pathfinder.[13,14]

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 atmospheric properties are computed) was defined as 170 km because accelerometer data from the MGS mission showed that sensitivity to the atmosphere was insignificant above 170 km.[15] Another significant input was the dust opacity, τ, which is a measure of the background dust level of the Martian 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 and had a nominal value of 1.0.

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 in 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.[16] Engineering judgment was used to extend the MGS data, which was 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 (3 wave peaks over 0 to 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 upon predictions of future spacecraft behavior. Also, only one ABM could be performed per day since the decision involved a significant

amount of analysis that was presented and discussed at daily program-wide 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 hours there was typically only one slot per planning period, but by the end of aerobraking there were as many as 5 slots in a 24-hr 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 in order 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-hr planning period. To determine how the trajectory would behave if an ABM was 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-hr 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, and at the 40% level when it was moving southward after passage over the pole. When the simulation reached walkout this procedure had to be modified since 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 3DOF 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. During this stage of mission planning, a simplified maneuver logic scheme was employed where maneuvers were performed if the peak heat rate of the previous drag pass had violated the heating control corridor. 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 implementation of the corridor-based ABM logic, in which only one crossing of a heating rate bound was needed to force the execution of a maneuver. As a result, small differences between the two simulations, from modeling uncertainties or even numerical errors, could lead to situations where a heating limit was just barely exceeded in one simulation, but not in the other. In such an instance, a maneuver would only be performed in the simulation that had the violation, and beyond that point, the trajectories would differ significantly and have very different maneuver histories.

An example of this effect can be seen in Fig. 5, which shows a comparison of the peak heat rate from two slightly different simulations for the first 50 orbits. The only difference between the two cases is that in one, the atmospheric interface altitude was set at 170 km, and in the other the atmospheric model was always active. Even though the atmospheric density is negligible above 170 km, it was not zero, and therefore slight differences built up between the two cases until orbit 20, when a maneuver was triggered by a heating indicator violation in one simulation but not in the other. After this point the trajectories varied significantly, and ultimately there was a difference of 2.5 m/s in the predicted ABM ΔV usage, ~10% of the total.

The large disparity in the results of these two cases indicated that for long-term aerobraking simulations employing corridor-based ABM logic, small perturbations due to numerical modeling could cause significant changes in the predicted values of certain parameters

like 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 pre-flight planning, to evaluate a range of aerobraking design options, and throughout operations, to help make maneuver decisions and assess the progress made towards meeting the final LTST constraint.

For each Monte Carlo analysis, 1000 simulations were run in which dispersions were randomly applied to the aerodynamic force coefficients and the atmospheric density. Computational constraints prevented the use of more than 1000 samples, since a single aerobraking trajectory, which typically encompassed 300-500 drag passes, took approximately 45 minutes of CPU time on a 400 MHz Silicon Graphics Origin 2000 processor. Running only 1000 samples produced reasonable statistics since 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, since 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 to be 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 kept constant for the duration of the simulation. Short-scale dispersions were also applied to the atmospheric density 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 that acted as a scale factor, either amplifying or diminishing the density beyond the value derived from Mars-GRAM that included the effects of Mars-GRAM that included the effects of the LDW and dust opacity. The standard deviation of the density scale factor (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.

Table 2. Input uncertainties implemented in Monte-Carlo analysis

Quantity	Nominal Value	Distribution	One standard deviation or min/max
C_A multiplier	1.0	Gaussian	± 5 %
C_N multiplier	1.0	Gaussian	± 5 %
C_Y multiplier	1.0	Gaussian	± 5 %
density scale factor	1.0	Gaussian	± 21 %
τ, dust opacity	1.0	Uniform	0.3/1.6

Results and Discussion

Pre-flight Planning

In the months prior to arrival at Mars, the POST 3DOF 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 upon a number of trajectory parameters, including the final local true solar time, 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, and results from these studies facilitated the selection of Mars-GRAM 2000 as the official atmospheric model for the project. The capability to perform realistic mission simulations that accounted for various mission uncertainties Some highlights of these pre-flight analyses are presented here to demonstrate the capabilities of the simulation and the type of results that were generated in the months leading up to aerobraking.

Immediately after Mars Orbit Insertion but prior to the initiation of main phase aerobraking, a nominal

mission Monte Carlo analysis was performed to estimate the duration of the aerobraking phase and project the final local true solar time. The simulation was initiated at the projected post-walkin apoapsis altitude, which corresponded to an initial orbit period of 18-hr. 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 January 2 and January 23, and the 1-percentile low value of the final LTST was 2:43 PM, which was still 25 min. above the 2:18 PM power degradation limit. The ability to calculate the 1 and 99-percentile values of these quantities was especially important because the simulation utilized atmospheric models (for longitude-dependent density waves and orbit-to-orbit variability) that were constructed from MGS data, which was obtained during a different season and in a different location of the atmosphere. Nevertheless, the model represented the best estimate of atmospheric variability that was available at the time.

A mean number of 25 ABMs were 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 reduce the periapsis altitude and move the peak heat rate on each orbit towards the upper 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 walkout) 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 walkout phase of the mission when the orbit lifetime constraint was dominant. This result was used by the navigation team to understand how many walkout orbits could be eliminated by propulsively terminating aerobraking before the 400 km apoapsis altitude goal was reached. By trading 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 ~30 walkout 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 ~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 upon the inertial velocity, so the angle between the two velocity vectors caused a shift in the inclination. The reduction was accounted for by biasing the inclination targeted by the MOI maneuver by 0.25 deg, a value that was based upon predictions from pre-flight aerobraking simulations. The results of these simulations also demonstrated that biasing the inclination slightly upward improved the final LTST since the

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-02	02-Jan-02	23-Jan-02	-
No. of ABMs	25	19	33	3
No. of up maneuvers	20	15	27	3
No. of down maneuvers	5	3	7	1
ABM DV 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, HH:MM	3:06 PM	2:43 PM	3:22 PM	7 min
Total No. of Aerobraking Orbits	370	321	428	22
No. of Walkout Orbits	115	98	132	8

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. 6 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 100% thermal margin for the first week after walkin (10-15 orbits) when the atmosphere models were being calibrated, and 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. 6 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-hr 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. Since maneuver decisions were based upon assumed nominal atmospheric behavior, when the simulation was propagated forward using the actual 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 upon predictions from a simulation that was run using a fixed atmospheric model.

Prior to 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 7 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 (i.e., had less than 80% margin). For the case with the narrower corridor, the number of upper-limit violations increased because the peak heat rate on each orbit was kept closer to the top of the corridor and large fluctuations in density was more likely to occur in the vicinity of the upper 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 percentage of the nominal density, were smaller.

The effect of the corridor design on other trajectory parameters is shown in Table 4. Decreasing the corridor

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-02	04-Jan-02	09-Feb-02
No. of ABMs	25	39	22
ABM DV usage, m/s	17.9	18.9	18.8
Final LTST, HH:MM	3:06 PM	3:21 PM	2:13 PM
Total No. of Aerobraking Orbits	370	326	467
No. of Walkout Orbits	115	114	120

width reduced the mean aerobraking duration by more than a week and provided an additional 15 minutes 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 2:18 PM 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 walkout corridor, which was designed to maintain the project-mandated 24-hour orbit 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 walkout. 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 8 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. By noting where these results intersected the 24-hr line (or any other line, for that matter), an upper walkout corridor was constructed. During walkout 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-hr. Although the lower limit was not important since the heat rate trended upwards during walkout, 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 using the Mars-GRAM 3.7 (MG37) model, which was similar to the Mars-GRAM model used to support MGS aerobraking. When the Mars-GRAM 2000 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, and has a significant effect on the duration of aerobraking because it directly impacts the amount of energy that is removed during each drag pass.

$$\rho = \rho_0 e^{\frac{-\Delta h}{h_s}} \qquad (2)$$

This effect is illustrated in Fig. 9, 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. Since 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 gen-

erate the density profiles in Fig. 9 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, the amount of energy removed from the drag pass was less because the integral of drag over the pass was lower. When this effect was repeated over many drag passes, the rate of aerobraking slowed down significantly, causing the aerobraking duration to increase.

Figures 10a and 10b 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. 11, which compares 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, the figure shows that in order for the 99-percentile low final LTST to meet the 2:00 PM limit (the goal at the time of the study), the post-walkin orbit period had to be less than 22 hr with the MG37 model and less than 16 hr 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 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 were considered. The 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.

LaRC's participation in the daily operational process was an added risk-reduction measure taken by the project to augment the JPL navigation team. Due to the favorable 3-hr time difference between the two centers, 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-hr time difference also aided the LaRC effort, as the JPL team would provide analyses they had done after the LaRC day.

The operational procedure followed by 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 accurately predict density over a number of orbits, and were thus used primarily to provide qualitative information on current atmospheric trends.[17] 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-hr 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 (J2 and J3 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 upon 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 ±6% for C_A and ± 10% for C_N and C_Y, based upon comparisons that were made between the aerodynamic database and accelerometer measurements.

Since 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. 12. For each drag 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 immediate action limit was a further safeguard that was implemented to prevent the peak heat rate from exceeding the flight allowable 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 usually performed when it exceeded the IA limit.

In Fig. 12, the 99-percentile heat rate obtained from using the density scale factor statistics from the current navigation wave model (mean = 0.8, σ = 0.20) exceeded the FA limit. By linearly scaling the distance between the mean and 99-percentile heat rate values 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. 12, a periapsis-raise maneuver was clearly required since 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 prior to 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. Fig. 13 shows the predicted heat rate time history from these deterministic trajectories for the same case discussed in Fig. 12. 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. Fig. 14 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 where Monte Carlo analysis was particularly useful was during the 3:1 resonance period when the orbit period was just over 8-hr. 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. 15. Monte Carlo analyses predicted that Odyssey would likely pass close to the 210 deg. E. longitude band, thus potentially incurring nearly a 1.5 km drop in altitude every three orbits. In Fig. 16, 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). The figure indicates that the post-resonance 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.

During the walkout phase of aerobraking, the criteria 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 that this occurred 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 the orbit lifetime at apoapses occurring after the maneuver had to be referenced to a new propagation carried out from the maneuver point to the lifetime condition. Results from a typical walkout Monte Carlo analysis are shown in Fig. 17. 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-hr in the daily planning interval. The 1-percentile low lifetime differed by about 6-hr from the mean lifetime, which was typical in these analyses. This result was important because it showed that the 24-hr mean lifetime requirement superseded another project mandate that required the 1-percentile lifetime to be greater than 8-hr.

In addition to supporting daily maneuver decisions, the LaRC simulation was the primary trajectory analysis tool that was used to predict the aerobraking termination date. This was a key parameter that 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 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 2:18 PM 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. 18, which shows the results of mission simulations beginning 13 orbits after walkin, 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 2:18 PM 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. Using this approach, fictitious trajectories were simulated that had maneuvers performed 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 in Fig. 19 and were used to gauge the effect of maintaining various positions within a given corridor on the final LTST.

Each week the total number of aerobraking orbits was predicted using a Monte Carlo analysis that employed the most recent corridor control strategy. This plot, shown in Fig. 20, 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 the spacecraft began to 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 January 11, 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 local true solar time of the descending node at the end of aerobraking was 3:04 PM, 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.

This paper discussed the contributions made by NASA LaRC in support of the JPL navigation team throughout the planning and operational phases of the aerobraking mission. In particular, the development of an independent aerobraking simulation, which began as a risk-reduction measure to validate early aerobraking studies, and evolved into a useful operational tool that provided capabilities never before employed in an aerobraking mission is described. The inclusion of predictive maneuver logic that mimicked operational decision-making was an important feature that enabled the simulation to act as a testbed for various aerobraking control strategies during pre-flight planning. In addition, the ability to perform Monte Carlo trajectory analysis allowed for the estimation of the effect that various mission uncertainties, especially atmospheric variability, had on the aerobraking trajectory. This Monte Carlo capability was used for the first time in an aerobraking mission to provide statistical information for maneuver decision-making. Ultimately, the Mars Odyssey aerobraking experience demonstrated the value of having a capable, independent simulation for the support and validation of the flight navigation team.

Acknowledgments

The authors would like to thank John Smith, Julia Bell and Robert Mase of the Jet Propulsion Laboratory 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 would also like to acknowledge the inputs made by the members of the NASA Langley 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, May 28, 2001

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

3 Braun, 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.

4 Smith, J. C. and Bell, J. L., "2001 Mars Odyssey Aerobraking," AIAA Paper 2002-4532, Aug., 2002.

5 F. T. Krogh: "Variable Order Integrators for the Numerical Solution of Ordinary Differential Equations." TU Doc. CP2308, NPO-11643, Jet Propulsion Laboratory, Pasadena, California, 1969.

6 Tolson, 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.

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

8 Lyons, 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.

9 "The Gravity Field of Mars: A 75th Degree and Order Model," Sjogren W.L., D.N. Yuan, A.S. Konopliv, and A.B. Kucinskas, Fall Meeting, AGU, San Francisco, CA, Dec 15-19, 2000.

10 Lyons, D. T., "Preliminary Mars Planetary Constants and Models for the Mars Sample Return," JPL IOM 312/99, DTL-1, Jan. 20, 1999.

11 Braden, E.M., "Analytic Ephemeris," JSC Internal memorandum ET4-91-07.

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

13 Justus, C. G. and James, B. F., "A Revised Thermosphere for the Mars Global Reference Atmospheric Model (Mars-GRAM Version3.4)," NASA TM-108513, July 1996.

14 Justus, C. G. and James, B. F., "A Revised Thermosphere for the Mars Global Reference Atmospheric Model (Mars-GRAM Version3.4)," NASA TM-108513, July 1996.

15 Bougher, S., Private Communication, Feb. 17, 2000.

16 Dwyer, A. M., Tolson, R. H., Munk, M. M. and Tartabini, P. V., "Development of a Monte Carlo Mars-GRAM Model for Mars 2001 Aerobraking Simula-

tions," Advances in the Astronautical Sciences, Vol. 109, 2001, pp. 1293-1308.

17 Tolson, R. H., et. al., "Application of Accelerometer Data to Mars Odyssey Aerobraking and Atmospheric Modeling," AIAA Paper 2002-4533, Aug. 2002.

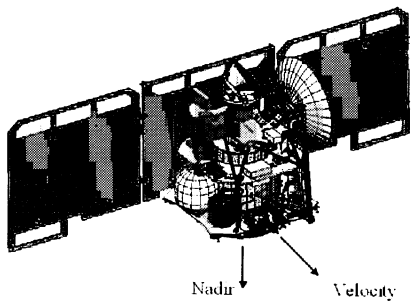

Figure 1. Mars 2001 Odyssey Orbiter in aerobraking orientation.

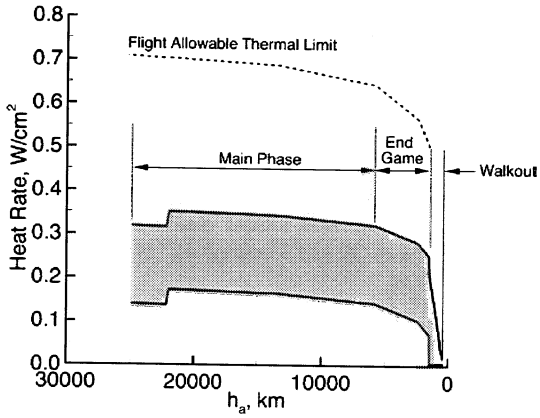

Figure 2. Representative heating control corridor used in the POST simulation.

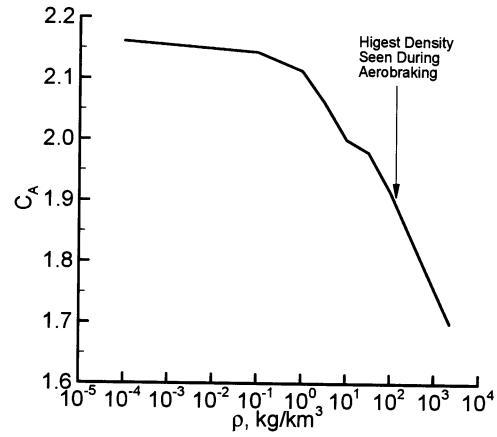

Figure 3. Variation of axial force coefficient with atmospheric density for $\alpha = \beta = 0$ deg.

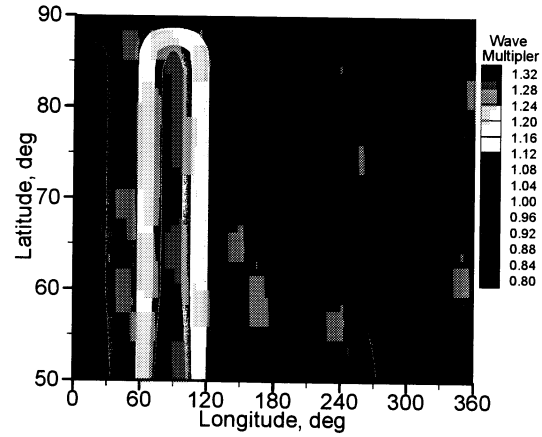

Figure 4. Variation of the MGS LDW-derived wave multiplier in the region of Mars Odyssey aerobraking.

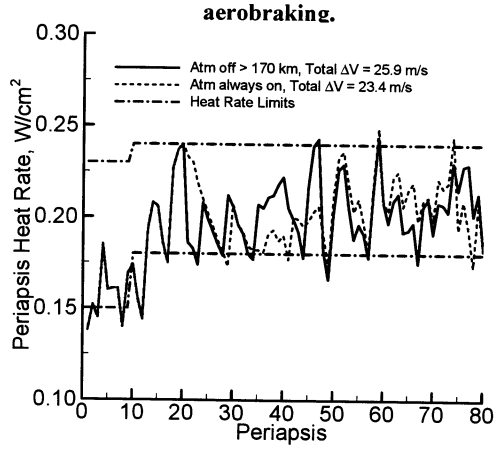

Figure 5. Effect of small perturbations on maneuver sequencing.

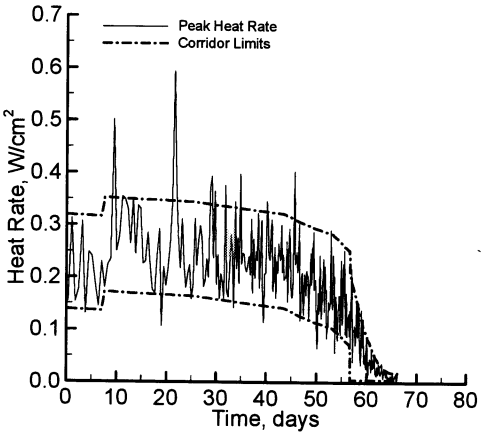

Figure 6. Heat rate time history for perturbed Monte Carlo Case.

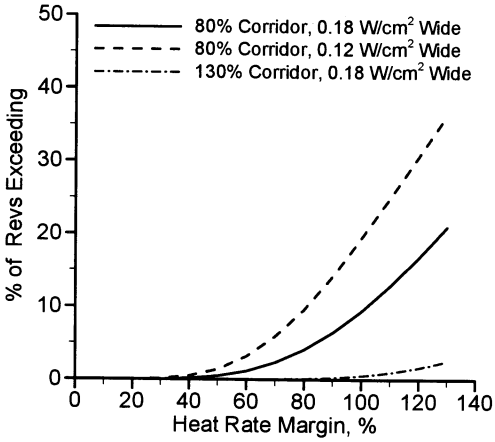

Figure 7. Effect of corridor design on the frequency of various thermal margin violations.

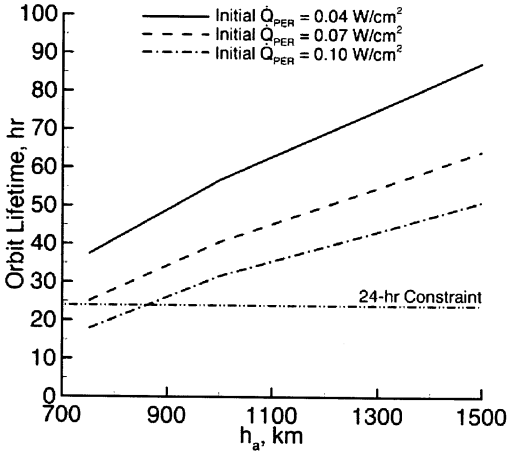

Figure 8. Relationship between periapsis heat rate and orbit lifetime

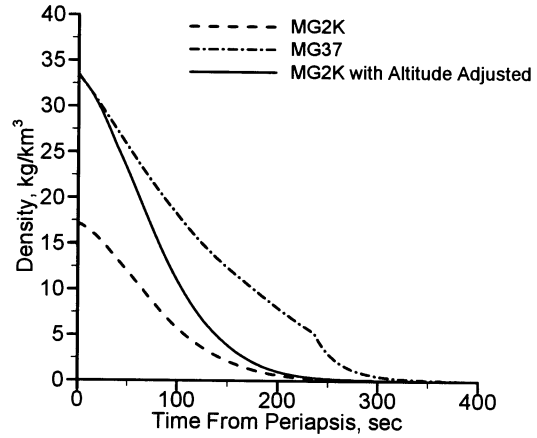

Figure 9. Example of scale height difference between MG37 and MG2K.

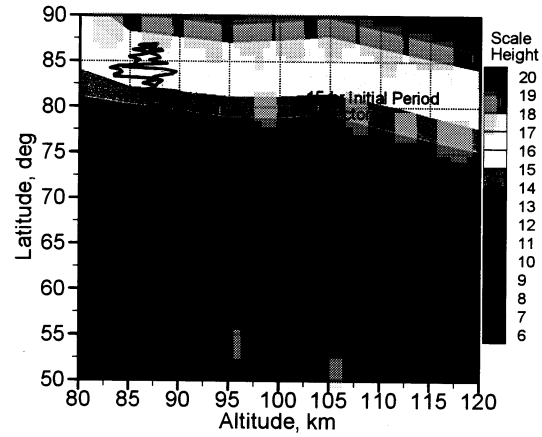

a) Mars-GRAM 3.7

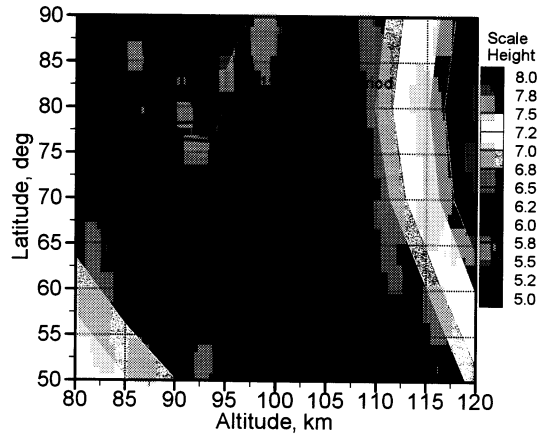

b) Mars-GRAM 2000

Figure 10. Predicted density scale height contours.

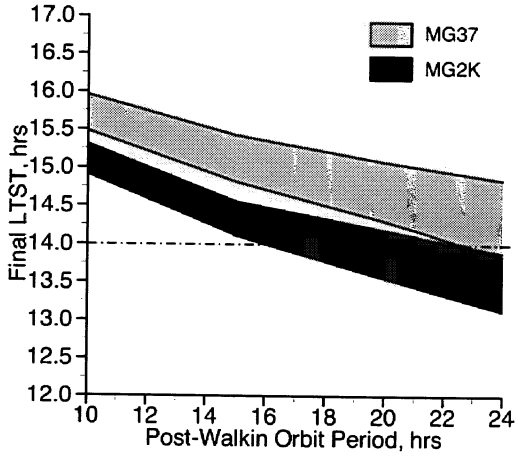

Figure 11. Effect of atmospheric model on final LTST prediction.

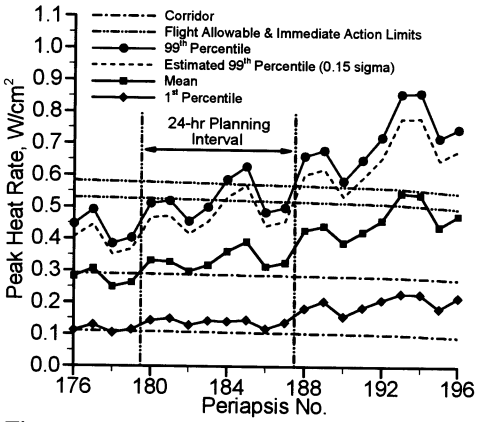

Figure 12. Predicted heat rate range from Monte Carlo analysis for no ABMs.

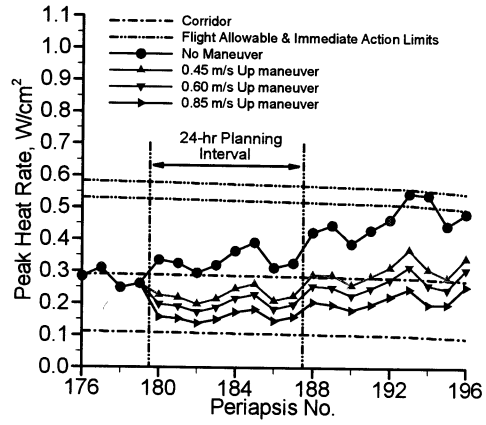

Figure 13. Effect of maneuver magnitude on predicted deterministic heat rate profile.

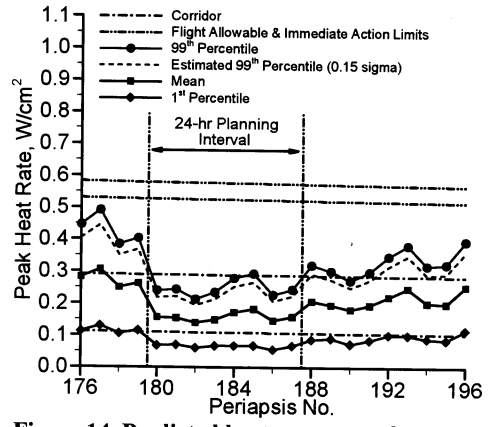

Figure 14. Predicted heat rate range from Monte Carlo analysis for ΔV=0.85 m/s ABM.

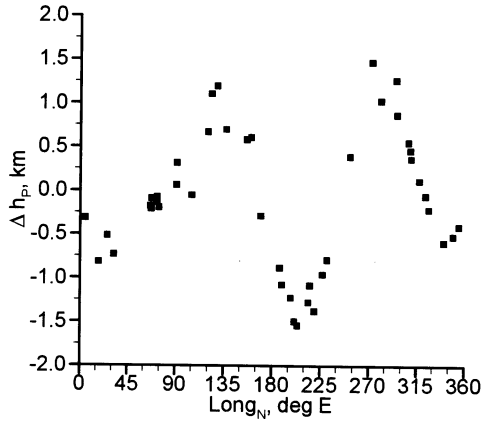

Figure 15. Altitude change as a function of longitude due to gravitational perturbations.

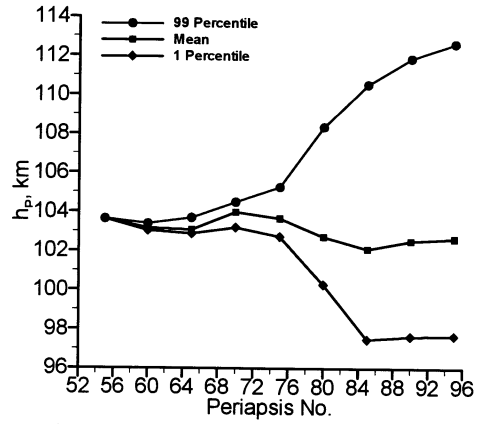

Figure 16. Envelope of potential periapsis altitudes predicted during 3:1 resonance period.

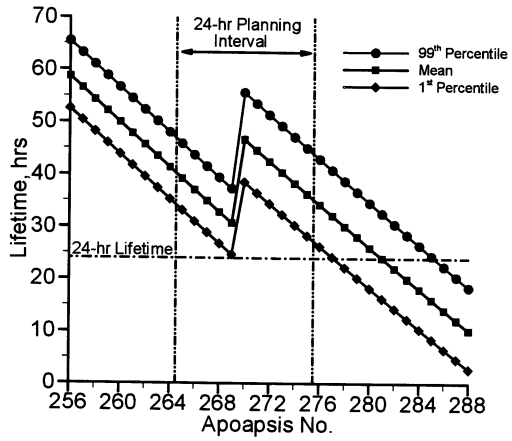

Figure 17. Predicted range of orbit lifetime.

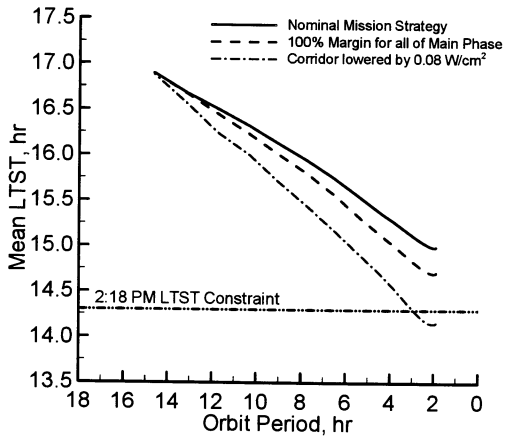

Figure 18. Effect of different corridor control strategies on mean final LTST prediction.

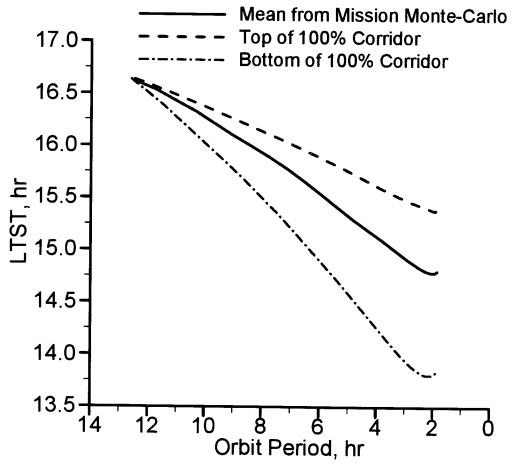

Figure 19. Predicted best and worst-case LTST profiles.

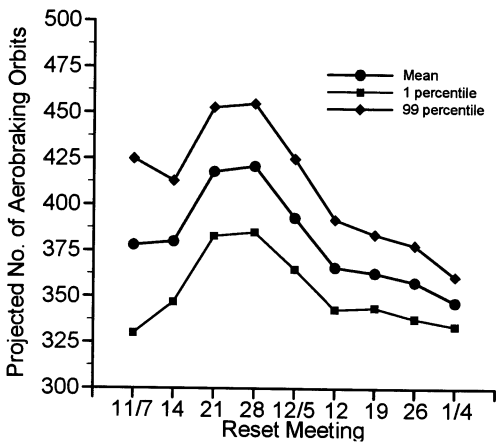

Figure 20. Weekly progression of Monte Carlo results.