JOURNAL OF SPACECRAFT AND ROCKETS
Vol. 42, No. 3, May–June 2005

Aeroheating Thermal Model Correlation for Mars Global Surveyor Solar Array

Ruth M. Amundsen* and John A. Dec†
NASA Langley Research Center, Hampton, Virginia 23681-2199
and
Benjamin E. George‡
U.S. Air Force, Kirtland Air Force Base, New Mexico 87111

The Mars Global Surveyor (MGS) spacecraft made use of aerobraking to gradually reduce its orbit period from a highly elliptical insertion orbit to its final science orbit. Aerobraking produces a high heat load on the solar arrays, which have a large surface area exposed to the airflow and relatively low mass. To accurately model the complex behavior during aerobraking, the thermal analysis needed to be tightly coupled to the spatially varying, time-dependent aerodynamic heating. Also, the thermal model itself needed to capture accurately the behavior of the solar array and its response to changing heat load conditions. The correlation of the thermal model to flight data allowed a validation of the modeling process, as well as information on what processes dominate the thermal behavior. Correlation in this case primarily involved detailing the thermal sensor nodes, using as-built mass to modify material property estimates, refining solar-cell assembly properties, and adding detail to radiation and heat-flux boundary conditions. This paper describes the methods used to develop finite element thermal models of the MGS solar array and the correlation of the thermal model to flight data from the spacecraft drag passes. Correlation was made to data from four flight thermal sensors over three of the early drag passes. Good correlation of the model was achieved, with a maximum difference between the predicted model maximum and the observed flight maximum temperature of less than 5%. Lessons learned in the correlation of this model assisted in validating a similar model and method used for the Mars Odyssey solar-array aeroheating analysis, which were used during on-orbit operations.

Nomenclature

a	=	acceleration (caused by atmospheric drag)
C_H	=	heating coefficient
C_p	=	specific heat
c_d	=	coefficient of drag
h_c	=	contact conductance coefficient
k	=	thermal conductivity
m	=	mass
q	=	aerothermal heating
T_{wall}	=	wall temperature
V	=	velocity (relative to the atmosphere)
V_f	=	view factor
α	=	solar absorptance
ρ	=	density

Introduction

THE Mars Global Surveyor (MGS) spacecraft launched in November of 1996 and arrived at Mars in September 1997. The initial orbit was highly elliptical, with an apoapsis of 54,026 km and a 45-h period. Aerobraking was used to slow the spacecraft and decrease the orbit altitude to allow insertion into a final science orbit. Aerobraking is the process of skimming the spacecraft through the planetary atmosphere at periapsis, decreasing its speed, and thus lowering apoapsis. This method saves substantially on the amount of fuel that must be used to achieve the science orbit, thus allowing for a greater scientific payload mass. The aerobraking effort was originally planned to take four months to bring the orbit apoapsis to 450 km. Early in the aerobraking, aerodynamic pressure caused one of Surveyor's two solar panels to bend backward slightly. This panel had been damaged shortly after launch. This event led to a pause in aerobraking and eventually a more conservative approach toward aerobraking, to put less stress on the damaged solar panel. Drag passes were designed to remain higher in the atmosphere than originally planned, thus decreasing the pressure load on the solar arrays but substantially increasing aerobraking duration. Aerobraking in this reduced fashion continued until March 1999, when the spacecraft was placed in the final science orbit.

The solar panels constitute the main drag on the spacecraft, and as low-mass "wings" they bear the brunt of aeroheating. The duration of aerobraking is often critical, because of high mission costs during that period and because the duration of aerobraking affects the final science orbit achieved. The Mars Odyssey mission was to use a similar mission plan as MGS, with solar arrays bearing the brunt of aerobraking. Thus, it was important to fully understand the MGS mission performance, to have the optimum plan for Mars Odyssey. In particular, it was important to determine the thermal behavior of the solar arrays because they were the limiting factor in aerobraking. If more aggressive aerobraking could be used (i.e., deeper cuts into the atmosphere on each pass), it would bring down the total time and cost for the Odyssey mission, as well as improving the ability to control the final science orbit. MGS flight data were used to correlate a thermal model of the solar arrays to ensure that the behavior during aerobraking was fully quantified. This information was used in the thermal analysis of Odyssey's aerobraking, which ultimately affected mission planning.

Because of the discovery of the damaged solar panel that led to the midmission change in MGS aerobraking operations, only the first 15 passes were done in the original flight corridor at relatively

Presented as Paper 2003-3765 at the AIAA 36th Thermophysics Conference, Orlando, FL, 23–26 June 2003; received 1 July 2003; revision received 29 January 2004; accepted for publication 27 February 2004.

*Aerospace Engineer, Mail Stop 431, Structural and Thermal Analysis Branch. Member AIAA.
†Aerospace Engineer, Mail Stop 431, Structural and Thermal Analysis Branch.
‡Lt., Spaceflight Test Engineer, Test and Operations Division, SMC Detachment 12; currently Flight Test Engineer, 1 S. Rosamond Boulevard, Edwards Air Force Base, CA 93524. Member AIAA.

high atmospheric densities with a comparatively high aeroheating environment. At the start of the aerobraking phase of the mission, the mission plan was to progressively lower the aerobraking altitude as a result of uncertainties in the density of the Mars atmosphere. Because correlation of the thermal model to high aeroheating environments was desired, the first 14 passes, which experienced fairly benign environments, were not used, and pass 15 was the first one considered for correlation.

Thermal Model Development

The thermal methods have been described in detail in an accompanying paper,[1] but enough will be described here to understand the driving factors in correlating the model. A thermal model of the spacecraft assembly was needed in order to capture the radiative environment of the solar arrays. Shading and view factors from the spacecraft helped determine the initial steady-state temperature of the solar array, as well as defining the solar and planetary flux environment that affect temperatures during a drag pass. This model captured only geometry and external surface properties for both the spacecraft and solar array. The orientation of the spacecraft with relation to the planet and sun was included in the telemetry data received from MGS.

For the detailed thermal solution, a full thermal model was developed for the solar array only. The solar arrays were mounted on gimbals with a low thermal conductance to the rest of the spacecraft, which minimized the conductive effect of the spacecraft on the array. The view factors to the spacecraft as well as solar and planetary fluxes from the radiation model were applied to the solar-array model. This method allowed much more detail to be captured in the solar-array thermal model without sacrificing solution time by including spacecraft elements in that model. Because the extent of damage to the solar array on the −Y side of the spacecraft was unknown, the data from those sensors were considered suspect, and the orientation with respect to the airflow was not reliably known. Also, because of the type of damage, that array was flown with the solar-cell side facing into the wind, which was not a correlation that would be used in the planning of any future missions. For those reasons, only the +Y array was modeled, with the analytically determined temperatures from that array being correlated to the +Y array flight thermal sensor data received in the telemetry.

Radiation Model

The radiation model of the spacecraft and the +Y solar array was developed in Thermal Desktop,[2] as shown in Fig. 1 (yellow arrow points to sun; blue arrow points to Mars). This model was used to calculate view factors from the solar array to deep space (Fig. 2) and also to calculate solar and planetary fluxes.

Thermal Model

The full three-dimensional thermal model was created in MSC/PATRAN.[3] This was done for two reasons. One was that existing FORTRAN code would allow simple inclusion of the aeroheating fluxes in PATRAN. The other was that the structural analysis

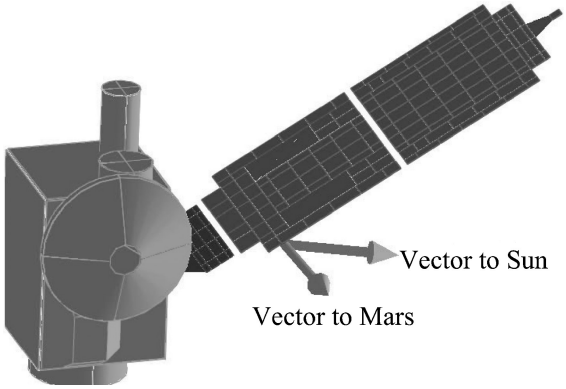

Fig. 1 Radiation model of spacecraft/solar array (just before aerobraking; cell side up, hot side down).

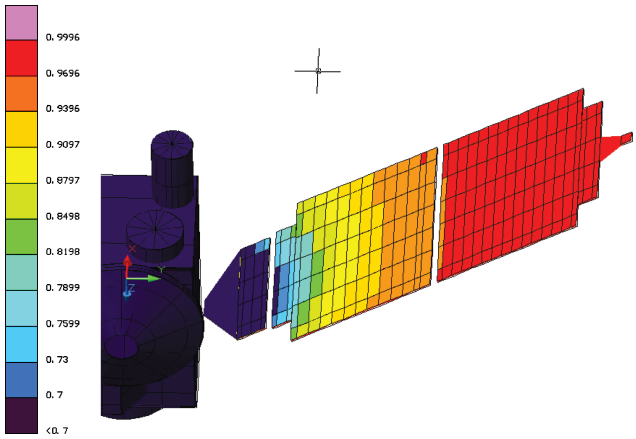

Fig. 2 View factors to space from solar array.

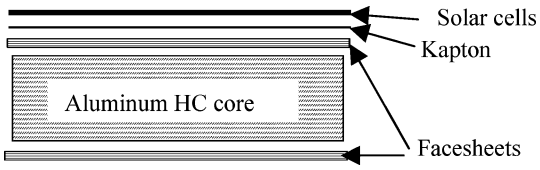

Fig. 3 Layup of solar array.

of thermal stresses would be easily accomplished by using existing methods. Normally, it is not more efficient to develop two models. However, in this case the time constraint of finishing the correlation before it was needed for Mars Odyssey on-orbit operations drove the use of existing, easily adapted methods. The radiation model in Thermal Desktop was necessary because orbital capabilities did not exist within MSC/PATRAN.

The solar array was mounted to the spacecraft at a gimbal and had a magnetometer mounted at the outside tip. The array was a sandwich construction that used graphite polycyanate facesheets with an aluminum honeycomb core. In places where structural connections were made, the facesheets were reinforced with doublers, and the aluminum core density was increased. The solar cells were mounted over more than 90% of the surface on one side. Because the thermal sensors that were used to correlate were in locations where coverage by solar cells was complete, it was decided to approximate the solar-cell coverage as 100% to simplify modeling. During the drag passes, the solar-cell side was oriented away from the direction of flight, so that the bare graphite side, often referred to as the "hot" side, received the aerodynamic heating.

The solar array was modeled in PATRAN as five distinct layers, as shown in Fig. 3. Each layer was modeled with plate elements, except for the aluminum core, which was modeled with solid elements. The layers were spaced apart, so that they could be connected via a contact conductance. This value could be varied to account for the adhesives used between the layers. Also, it accounted for the reduced contact area between the facesheets and honeycomb core. The solar-cell layer was made up of a combination of materials, as described next. The Kapton sheet between the solar-cell layer and graphite facesheet was 0.002 in. thick (0.051 mm). The facesheets were M55J/RS-3, 0.0075 in. thick (0.19 mm). The aluminum honeycomb core was 1.0 in. thick (25.4 mm). The film adhesive used between the core and facesheets was M1025A. The finite element mesh on the facesheets was customized to take the thermal sensor locations into account.

The solar-cell layer consisted of the following, from the outside in: a 0.005-in. glass coversheet, a layer of DC93500 adhesive (0.003 in. thick), the solar cells themselves, a metal backing to the solar cells, and a layer of CV2568 adhesive that held the cells down to the Kapton sheet. The solar cells were silicon on the outboard section of the array, GaAs on the inboard section, 0.008 in. thick. The overall thickness of this layer was 0.022 in. (0.5 mm). The properties of this layer were calculated as a weighted average of those five materials. The roughly 1 kg of wiring for the solar cells on each section was

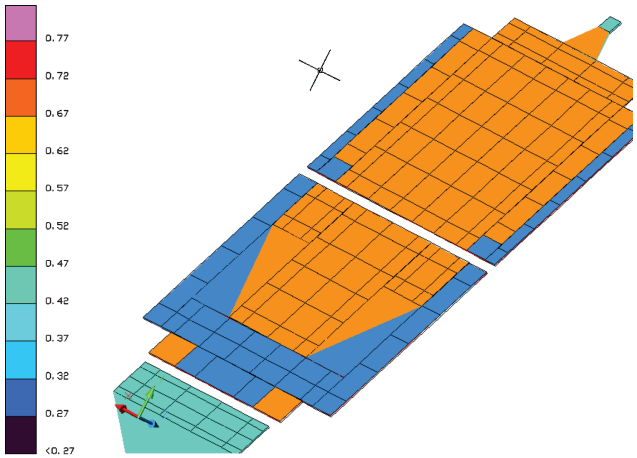

Fig. 4 White paint on hot side facesheet (α).

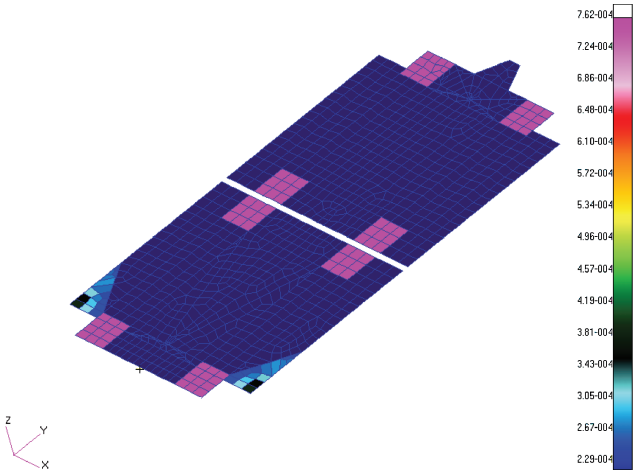

Fig. 5 Thicknesses on hot side facesheet (m).

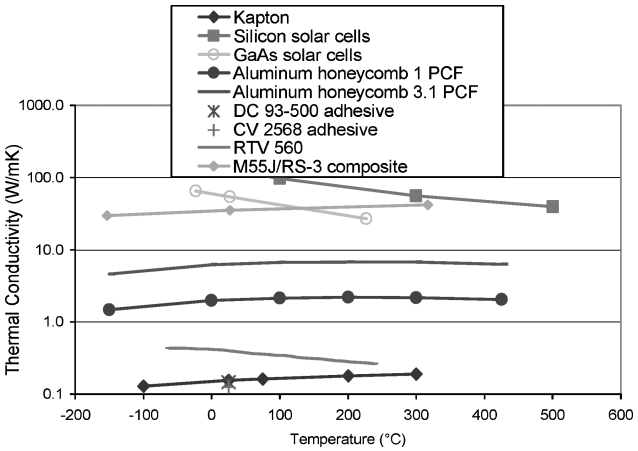

Fig. 6 Material thermal conductivities.

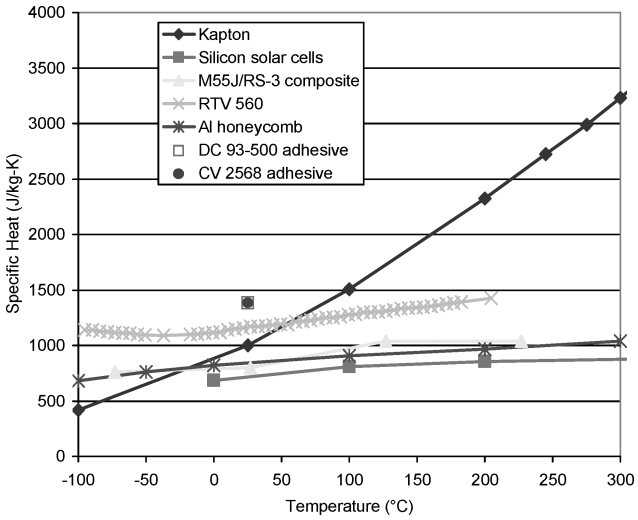

Fig. 7 Material-specific heats.

added to this layer as a smeared mass. The effective properties were a specific heat of 807 J/kg and a density of 3508 kg/m^3 for the GaAs panel and 2402 kg/m^3 for the silicon cell panel.

White paint was used over portions of the hot surface, and this was included in both the radiation model in Thermal Desktop and the radiation boundary conditions of the PATRAN model. These lower absorptivity regions on the edges and inboard corners of the array are shown in Fig. 4 on the radiation model. The purpose of the paint was to provide a lower absorptance for the solar flux and allow a lower temperature for the outer portions of the array. The pattern was chosen to correspond with the areas with the highest aeroheating flux, to help keep them within their thermal limits. Another modification included in the model was local thickening of the Kapton sheet in the highest heating areas. At the inboard array corners, the Kapton was laid up in layers to augment the thickness and thermal mass available to absorb the heating. This was included in the model by using a spatial field for the facesheet thickness, as well as changing the sheet thickness for all areas thickened with doublers. This is illustrated in Fig. 5. On the solar-cell side, the different optical properties of the two solar-cell types were also included in both models.

All material properties (except density and emissivity) were included as functions of temperature, as shown in Figs. 6 and 7.

Other parts included in the model were the hinges and magnetometer. The hinges were included as simple flat plates that correctly captured the mass of the hinge as well as the contact area on the array. The magnetometer mass was lumped over its contact area at the outer-most tip of the array.

The mass of the overall model was verified by a comparison with the as-built mass from flight assembly records. The model did not include the drag flaps, and the as-built mass of only the modeled portion was roughly 26 kg. The mass of the total thermal model was 23 kg, for an error of 11%. This was not unreasonable considering the approximations within the model, and the fact that some components mounted on the solar arrays were not modeled if they were known to be far enough from any thermal sensor not to affect sensor results. Because the model was too light, it was expected to lead to conservatively high temperature predictions. Several of the densities were varied during correlation to ascertain their effects on the model performance. The final correlated model has a mass of 23 kg.

Temperature Limits

The flight allowable temperature for the solar panel, used to set the aerobraking corridor, was 175°C. The thermal limits on each component material of the array are shown in Table 1.

Boundary Conditions

Boundary conditions included contact conductance between layers, convection, radiation, aeroheating, and solar and planetary fluxes. The contact between layers was calculated as the effective conductance as a result of the adhesive thickness. The contact between any two layers included the thickness of both layers and their thermal conductivities. For layers with less than 100% contact area (such as the aluminum honeycomb), the contact was reduced by the fraction of effective contact area. The entire effective conductance through the array was calculated in order to compare with tested values. This calculation is automatically performed within the PATRAN Thermal solver, but was also done manually as a verification. These calculations and the overall total effective conductance

Table 1 Material temperature limits

Material	Maximum use temperature, °C
Dow Corning 93-500 adhesive (between cover glass and solar cells)	200 continuous 300 short term (less than 2 h)
NuSil CV 2568 adhesive (between solar cells and Kapton)	265 continuous 400+ short term
Kapton film (cocured on graphite facesheet on solar cell side)	316 continuous 400 short term
Graphite composite facesheets (M55J/RS-3)	177 continuous 218 short term 254 ultimate (Tg of polymer)
Al 5056 honeycomb core	200
EA 9394 adhesive in potting and hinge areas	204
Ti-6Al-4V (hinges)	1000+
GaAs solar cells	150 continuous 300 short term
Bryte EX 1543-U adhesive film (cyanate ester between facesheets and Al HC core)	218 continuous 280 short term ultimate
Bryte EX 1541 20 PCF foam (densification foam at hinges)	177 continuous 280 short term ultimate
SN95/5 solder (solar cells)	220°C melt
Qualification test (backface temperature)	190
Flight allowable temperature	175

Table 2 Effective conduction through array thickness

Layer	Thickness, in.	Material	Added layer	Eff. hc, W/m^2K
Facesheet to core (each side)	0.0075	M55J/RS-3	0.003-in. adhesive layer, contact over 1% of area (honeycomb)	52
Kapton	0.002	Kapton	——	2755
Solar cells	0.0075 (GaAs) 0.008 (Si)	Solar cells	0.006-in. CV 2568 bond	822
Core	1	Al honeycomb	——	79
Total conductance	——	——	——	19

through the thickness of the entire array are shown in Table 2. The total calculated value of 19 W/m^2K was somewhat higher than the value from correlation of a two-node model in early ground testing (0.01 $W/in.^2C$ or 15.5 W/m^2K) (Ref. 4). However, it is difficult to compare the two values for several reasons. In the ground test, local effects and sensor mass could affect the correlated value. Radiation through the core was treated as a separate radiation boundary condition in this modeling, whereas in the two-node model it was included in the effective conduction value. Also, this effective conduction is shown only for the majority of the array, not for local regions such as the areas with doublers. The contact conductance values were varied during correlation, especially the most uncertain value, that of the facesheet to the core. The final correlated model used the values in Table 2.

Radiation

The radiation boundary conditions included both internal and external radiation. Radiation through the aluminum core, from one facesheet to another, was applied. View factors to both space and the spacecraft were included. The view factors from the Thermal Desktop spacecraft model were applied as fields in the appropriate regions, by mapping the viewfactors into the PATRAN model, as shown in Fig. 8. This figure clearly shows the blockage of the

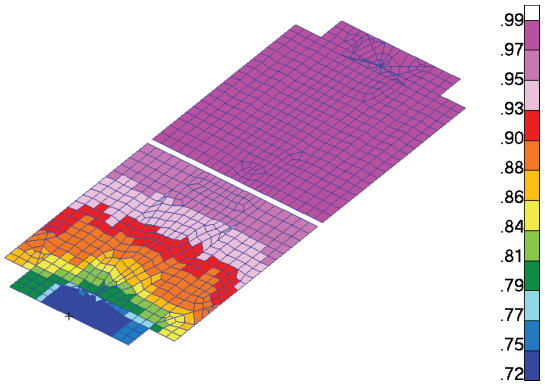

Fig. 8 Viewfactors to space from solar-cell side.

view to space on the solar cell side. Space temperature was fixed at 3 K (−270°C) during the exoatmospheric portion of the analysis, then was allowed to vary to simulate the Martian atmosphere/space combination during the drag pass. The radiation sink varied from −270°C (at the initial point of the transient) to −200°C (at the deepest atmospheric point).

Aeroheating Boundary Condition

The aeroheating was calculated by using a two-step process. First, the atmospheric density experienced by the spacecraft was calculated using accelerometer data from the spacecraft. Next, by using the atmospheric density and the velocity relative to the atmosphere, the aeroheating on the solar array was calculated. More details on the Martian atmospheric models and the methods used for these predictions can be found in related publications.[5,6] The correlation of the atmospheric density to the accelerometers was made by using direct simulation Monte Carlo (DSMC), which modeled the interaction of atmospheric particles with the entire spacecraft.[7] DSMC calculates the average energy transferred to/from surfaces by individual molecules during the simulation through direct sampling. The relation between the density and the acceleration is captured in the equation

$$ma = \tfrac{1}{2} c_d \rho V^2 \tag{1}$$

The acceleration caused by drag was calculated at time intervals throughout the drag pass directly from the flight accelerometer data. Similarly, the velocity relative to the atmosphere was calculated at the same time intervals by using the postaerobraking pass reconstruction of the orbital trajectory. Next, for each time interval the aerothermal heating over the surface of the solar array was calculated for the given atmospheric density and relative velocity, once again by using DSMC techniques. This relationship is defined by the following equation:

$$q = \tfrac{1}{2} C_H \rho V^3 \tag{2}$$

The heating coefficient is a function of density, velocity, and orientation relative to the velocity. Chemistry effects on aeroheating are insignificant at both MGS and Odyssey aerobraking conditions.

All DSMC calculations were performed without chemistry being modeled; previous calculations have shown that chemical reactions have negligible effect on directentry stagnation heating at the densities of interest for MGS (<150 kg/km^3) (Ref. 8). Because of the large computational time needed to perform the DSMC calculations, some simplifying assumptions were made. The effects on C_H as a result of changes in velocity and orientation through the pass are negligible compared to the changes in density. Thus the calculations were made at a constant, average velocity for the nominal orientation. DSMC calculations were performed at prescribed intervals of density encapsulating the range of expected flight conditions, and interpolation was used to determine the heating coefficients over the surface of the array for the calculated density. An example for a single density is shown in Fig. 9. Next, by using Eq. (2), the aerothermal heating over the surface of the array was calculated at

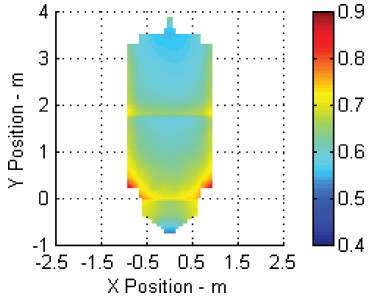

Fig. 9 C_H (at atmospheric density of 60 kg/km^3).

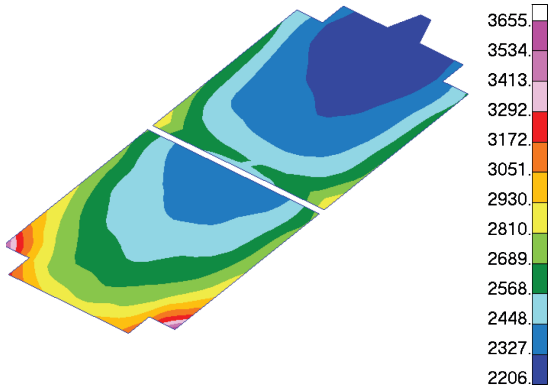

Fig. 10 Aeroheating near peak during pass 15 (W/m^2).

time intervals spanning an entire aerobraking pass. These aeroheating calculations were performed by using MATLAB® with an aeroheating flux array exported as the product for use in the PATRAN model.

The aeroheating flux array was a function of both physical position on the array and time within the drag pass. The aeroheating was interpolated in both time and space onto the PATRAN model on the entire exterior. User-developed FORTRAN was used within PATRAN to accomplish this interpolation. The mapping at one time point, on the hot side, is shown in Fig. 10. As shown, the inboard corners of the array received the highest heating, which was the reason for localized use of white paint and thickened Kapton.

Average heating on the edges of the array was taken to be 10% of the heating on the nearest point of the hot side. Earlier DSMC results[7] showed that the heating rates on the edges of the solar panels vary considerably. Because this model only considered the main front-side surface, some approximate method was needed to estimate the edge-heating values. The edge heating depends on the attitude of the vehicle with respect to the relative wind (spacecraft velocity). Edges facing away from the wind (leeward) had negligible heating, whereas edges that face into the wind (windward) can have heating rates that are roughly 20% of that on the nearest node on the front side of the panel. Because these edges contribute a small total heat load during an aerobraking pass, it was considered reasonable to assume that all edges were exposed to wind half of the time and thus use an average heating rate for the edges of 10% of the nearest node heating on the front side.

The aeroheating data included incident heating only, so that reflected heating needed to be calculated within PATRAN. Because the reflected energy is small relative to the incident energy for the expected range of wall temperatures, an approximate method could be used to approximate the reflected component using only the incident component normalized to the freestream kinetic energy. The following equation was developed empirically to approximate the reflected component, as in Ref. 7:

$$C_{H,\text{net}} = C_{H,i} - \lfloor 0.015 + 0.06 * (1 - C_{H,i}) \rfloor * (T_{\text{wall}}/300) \quad (3)$$

where the coefficients C_H are defined as

$$C_H = q \Big/ \left(\tfrac{1}{2}\rho V^3\right) \quad (4)$$

The empirical result was shown to be accurate within 2–3%. Because the reflected energy was typically 10% or less of the incident energy, the error in the total energy prediction was much less than 1%. This calculation was carried out within the existing FORTRAN code to subtract the reflected component from the incident heating.

Orbital Flux Boundary Condition

The orbital fluxes from the Thermal Desktop model were applied to the PATRAN model for both solar flux and planetary flux (albedo and IR). These fluxes were transient over the same period of time as the aerobraking analysis, starting at 150 s before periapsis, and ending 150 s after periapsis. The fluxes were applied by using mapping of the Thermal Desktop model results to the PATRAN mesh locations. A transient heat flux was generated for nodes on each side of the array. The viewfactors to sun and planet on the hot side were unobstructed by the spacecraft, as can be seen by the orientation in Fig. 1. Thus the flux on that side was uniform and could be applied across the entire hot side, modified only by optical property. The solar-cell side received an order of magnitude lower orbital heat flux than the hot side (8–20 vs 250–450 W/m^2, depending on optical properties). However, the inboard portion of the array was highly obstructed (as shown in Fig. 8) and thus had a nonuniform heat flux. The transient heat flux on this side was weighted by position; a factor was applied based on the heat-flux gradient across the array. Because the flux on the solar-cell side was so low, the influence of the solar-array efficiency (conversion of solar flux to electricity) was not taken into account in lowering the incoming flux. The planetary flux has a very minor effect on this assembly, and thus an assumption of constant Martian albedo (0.5) and emissive power (216 K planet average) was used.

Initial Thermal Results

The first runs of the thermal model were done with no knowledge of the flight data, to determine how close the model would be with no alterations based on postflight knowledge. This simulates to some extent a normal thermal progression, where the initial modeling is done before flight, and there is no knowledge of in-flight performance. The results are shown plotted on the three-dimensional model in Fig. 11. The through-thickness scale in this figure is exaggerated to allow the thickness to be visible. The cooling effects of the mass of the hinges, doubler areas, and magnetometer can be seen. Also, the temperature difference between the hot side and solar-cell side is substantial.

Thermal Correlation Process

The flight data consisted of thermal sensor measurements for four different locations on the solar array. There were two sensors on the solar-cell side and two on the hot side. The two on the hot side were applied to the underside of the graphite facesheet, with a hole through the core and solar-cell facesheet to allow the penetration. The sensors on the solar-cell side were applied directly to the outer facesheet. The two hot-side sensors were T211 and T213, and the cell-side sensors were T210 and T212, as shown in Fig. 12. In this figure only the facesheets are shown, and the through-thickness dimension is greatly exaggerated to allow visualization of the thermal sensor locations. The initial comparison to flight sensor data is shown in Fig. 13. The timescale starts before aerobraking, and a time of 0 s corresponds to the peak of the pass (periapsis). This first comparison was only for the hot-side sensors, because these defined the peak heating that can be sustained by the array and thus were most critical to correlate. Also because the point of this model was to correlate to changes that occurred during aerobraking, the initial condition was subtracted off, and only temperature changes for both flight and the model were compared. As is evident, the sensors in the model were responding too quickly to the heating and were getting too hot. The sensor predictions on the solar-cell side followed a similar pattern, although they were somewhat closer to the flight data. The hot-side sensors were over 60°C too hot in the model (much worse on the inboard side) and reached their peaks about 60 s before they should. The cell side sensors were about 30°C too hot in the

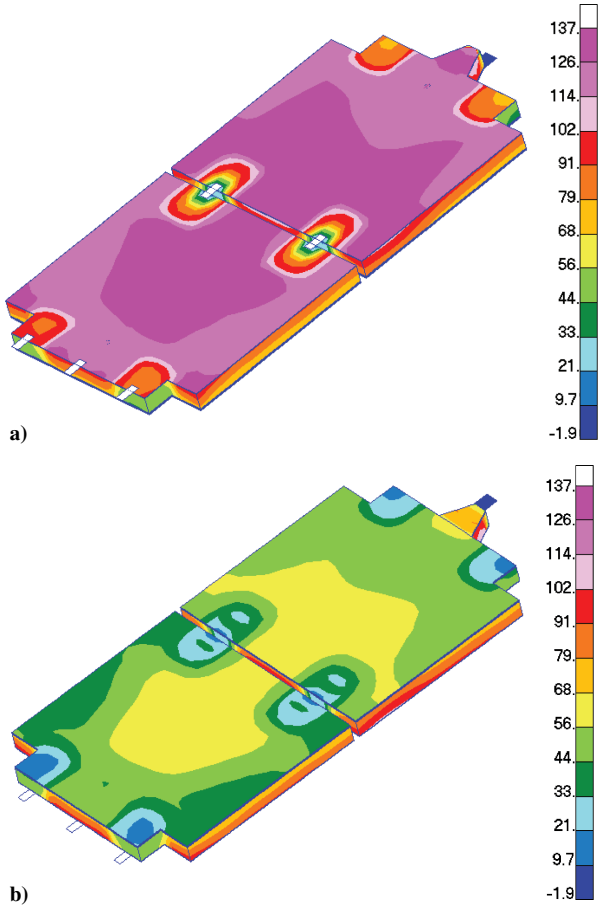

Fig. 11 Initial thermal predictions for MGS solar array near peak: a) hot side and b) cell side, °C.

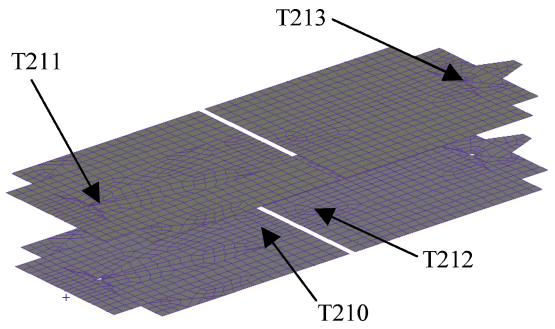

Fig. 12 Flight temperature sensor locations.

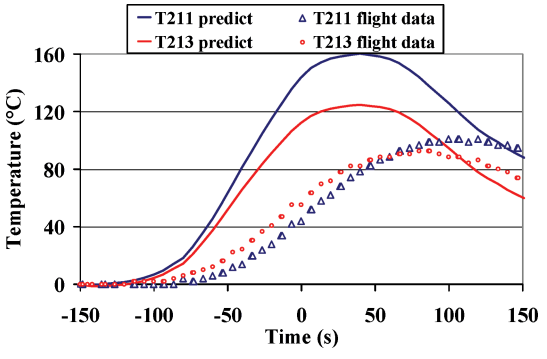

Fig. 13 Initial comparison of model to flight data (hot-side sensors).

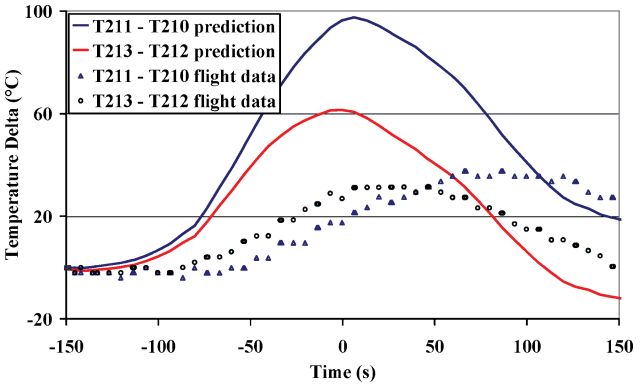

Fig. 14 Initial comparison of hot-to-cell-side differences.

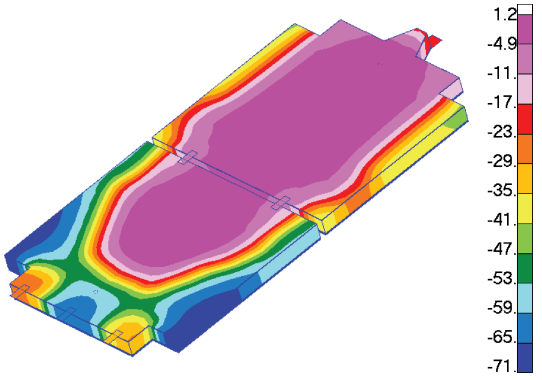

Fig. 15 Initial temperature condition (steady state in vacuum phase), °C.

model. The temperature difference between hot side and cold side was also too high and occurred too quickly.

Another interesting comparison was the difference between the hot-side and cell-side sensors, both on the outboard and inboard panels. This illustrates where the model was off on both through thickness and lateral conduction because the sensors were separated not only by the thickness of the panel, but also by most of the area of the panel as well. The initial comparison is shown in Fig. 14. It is obvious that the temperature difference between hot side and cell side was too high and was occurring too rapidly in the model.

Several actions were taken to make the model more representative of flight. Most of these changes were additional detailing and refinements of the model. These are described in the following sections, beginning with the changes in the orbital initial condition.

Orbital Initial Condition

Initially, the model was intended to be run from a steady-state condition with a single global constant temperature. However, it became evident that the initial condition would substantially affect the transient solution. Thus, the model was run for a steady-state solution using the orbital orientation of the spacecraft just before initiation of aerobraking. This allowed an approximate gradient prediction across the array, for a more accurate response at the start of aerobraking. The main focus of the model was to correlate to aerodynamic heating, so that many of the preaerobraking spacecraft operations that occur in the vacuum phase were not included. Use of a steady-state temperature for the vacuum phase condition gave an accurate enough prediction of the thermal initial condition that no detailed transient in the vacuum phase was considered. The use of a steady-state result here rather than transient was felt to be justified for two reasons: there was no large change in orbital fluxes or shading at this point, and in the flight data it was seen that temperatures were fairly stable just before the beginning of aerobraking. The orientation of the spacecraft and +Y array with respect to the sun and planet is shown in Fig. 1. The temperature gradient produced by this initial condition, which was used as the starting temperature just

prior to aerobraking, is shown in Fig. 15. Use of this true prediction of initial condition meant that actual temperatures, rather than temperature changes, could be used in the correlation. One effect of this initial condition was that the precooling of the inboard end of the array (as a result of the white paint areas) can be clearly seen. This increases the temperature difference between the two hot-side thermal sensors (T211 and T213), making it much more similar to what was actually observed in flight.

Additional Model Refinements

Radiation within the thermal sensor holes was added, between the hot-side facesheet and the layer of Kapton over the hole in the solar-cell side. The solar-cell properties were modified to include the metal backing, wires, and other components that had not been broken out separately. This helped by about 5–10°C. The thickness of the facesheets was increased to include the effective thickness of the adhesive, which also helped bring temperatures down slightly. Radiation within the aluminum honeycomb was added, which helped by another 2°C, but still left the inboard hot side much worse than the outboard hot side.

The detail of the thermal sensors was increased drastically. In the initial model, the sensor temperatures were simply taken as the nodal temperature at the appropriate location on the facesheet. In subsequent runs, the lump mass for each sensor was added, with adhesive under them as their connection to the face sheet. The sensor wire mass was added, with radiative and convection transfer to the surrounding hole and facesheet, as appropriate. The optical properties for the GaAs solar cells, silicon cells, and white paint were refined and detailed on both the radiation and thermal models, when it was observed that the initial temperature played a large part in determining the final prediction. This helped improve the difference between the inboard and outboard panels. Changes like this were run in both the steady-state initial condition (vacuum phase) and the transient aerobraking model.

Myriad other minor changes were made, such as making sure all material properties were bounded over the thermal excursions experienced, adding radiation over the edges of the core, refining the hinges, and adding damper masses. Also, many properties were varied to experiment with their effects, such as aluminum honeycomb core density, facesheet conductivity, adhesive thickness under sensors, and honeycomb core contact area. One reason that honeycomb core density was varied is that the final as-built density was not necessarily the same as the virgin material; this material is so light that small amounts of adhesive and foam inclusion increase the density appreciably. Another small change was to increase the thickness of the facesheets in the areas where white paint was applied, which is a real effect that had not been considered in the original modeling. This did help by 2.5°C on the sensor directly under that painted area (T211).

In all, more than 60 runs were performed with changes to various parameters to improve the correlation. It is difficult to show all of the conditions of each run, but a short summary of some of the more significant runs is given in Tables 3–5. The number of each run is the chronological order in which they were done. The main change made in each run is given. The errors are shown between each sensor prediction and the flight data, in terms of the maximum temperature predicted. This is generally the most important parameter because it is used to set the aerobraking corridor limits. The average of these errors, as well as average time error and worst magnitude time error (of peak) are also given. The runs are grouped by those dealing with changes in mass, detailing of the thermal sensor areas, and optical property changes. Table 5 shows the runs after a correction was made to rectify an unfortunate typographical error in the equation of the reflected aeroheating, which had made the runs up to number 44 slightly conservative.

Final Thermal Results

The final corrected run (52c) compared to the pass 15 flight data is shown in Fig. 16. The hottest thermal sensor is predicted very well, and all others are conservative (i.e., prediction is somewhat too warm).

Because the intent of this model was mainly to predict the rise in temperature caused by aeroheating, and not focused on the preaerobraking vacuum orbit phase, the correlation was also evaluated only from the point of view of temperature rise during aerobraking. To

Table 3 Correlation runs with mass adjustments

Correlation run #	Initial	2	4	9	20	27	31
T210 error, °C	16.5	9.1	7.4	2.6	−2.7	8.9	5.3
T211 error, °C	59.3	49.8	47.5	40.3	19.2	21.9	13.7
T212 error, °C	14.2	4.7	3.0	−1.6	−6.8	5.1	1.6
T213 error, °C	32.1	23.0	21.3	16.5	4.1	3.8	−3.7
Avg temp error, °C	30.5	21.7	19.8	14.5	3.4	9.9	4.2
Avg time diff, s	−46.8	−43.2	−42.9	−40.0	−24.2	−30.0	−24.9
Worst magnitude time error, s	——	60.0	60.0	60.0	48.3	40.0	35.9
Run condition	——	Add wire/solder mass on solar, cell side	Facesheets up to 0.009 in. (0.2286 mm)	Gr-Ep at $C_p = 800$ instead of 690 J/kgK	Al HC dens up to 52 kg/m^3	Facesheets 0.0095 in. (0.241 mm), Al HC 30 kg/m^3	Al HC core up to 40 kg/m^3

Table 4 Correlation runs with detailing of sensors and optical properties

Correlation run #	22	32	35	38	42	44	18	39	23
T210 error, °C	−1.7	5.3	5.2	3.8	5.7	5.5	3.2	7.2	1.9
T211 error, °C	18.6	12.4	7.3	3.9	6.1	6.1	22.0	13.4	9.4
T212 error, °C	−5.8	1.6	1.4	−0.1	3.0	2.4	−2.8	3.2	−1.8
T213 error, °C	2.1	−4.9	−6.4	−8.0	−3.2	−3.2	6.2	0.3	−7.0
Avg temp error, °C	3.3	3.6	1.9	−0.1	2.9	2.7	7.1	6.0	0.6
Avg time diff, s	−20.0	−24.9	−20.8	−18.7	−5.8	−0.8	−23.1	−16.0	−26.8
Worst magnitude time error, s	47.5	35.9	29.0	24.9	21.0	14.6	40.0	22.6	40.0
Run condition	Radiation from portion of hole to sensor wire	Radiation from cell side facesheet to sensor wire	Add radiation from sensor wire to rest of hole	Add contact: sensor wire to hole	Add sensor section on wire, add sensor nodes on cell side	HC on sensors up to 10, W/m^{2}K	White paint α 0.92	Optical props: Gr-Ep 0.83/0.79, Ga/As α 89, white paint 0.18/0.9	Contact sheets to core incr from 50 to 100

Table 5 Final correlation runs, with correction for reflected aeroheating

Correlation run #	44c	45c	46c	48c	50c	51c	52c
T210 error, °C	0.2	5.2	−1.5	−1.5	3.4	−6.5	−0.7
T211 error, °C	0.5	8.1	−1.7	−4.0	3.2	−3.0	4.4
T212 error, °C	−2.5	2.6	−2.5	−2.5	2.8	−8.0	−1.6
T213 error, °C	−10.2	−4.6	−10.2	−10.2	−4.4	−7.6	−1.8
Avg time diff, s	−0.8	−5.6	0.8	0.8	−2.4	−8.5	−3.7
Worst magnitude time error, s	14.6	21.0	8.2	8.2	14.6	24.4	12.7
Run condition	All runs now corrected for accurate Qrefl	Al HC $\rho = 25$ kg/m^3	GaAs $C_p = 566$, $\rho = 3508$	White paint area (under T211) to 12.5 mil (0.32 mm) Al HC $\rho = 40$	Al HC $\rho = 25$	From 48 c, HC contact $h = 52$	Combine runs 50 and 51

Table 6 Correlation quality measures

Correlation measure	T210	T211	T212	T213
Difference in peak temperature change, °C	−0.7	4.4	−1.6	−1.8
Difference in peak timing, s	−12.7	−0.7	−0.7	−0.7
Difference in peak actual temperature, °C	7.5	8.0	8.0	−2.8
RMS difference in temperature change over time, °C	2.6	6.7	2.1	2.9
Largest difference by time point, °C	5.5	12.0	4.4	6.2
Difference in peak temperature change, %	−1%	4%	−2%	−2%

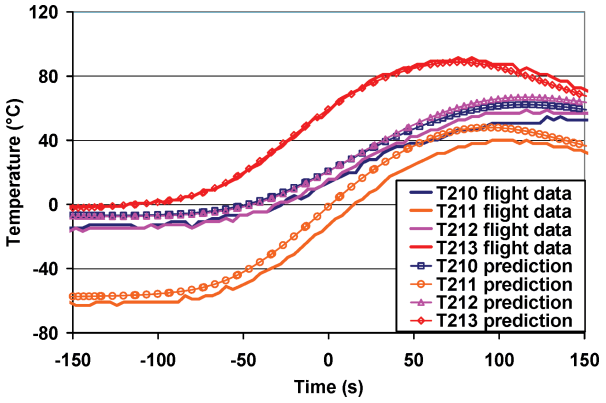

Fig. 16 Final correlation, actual temperatures.

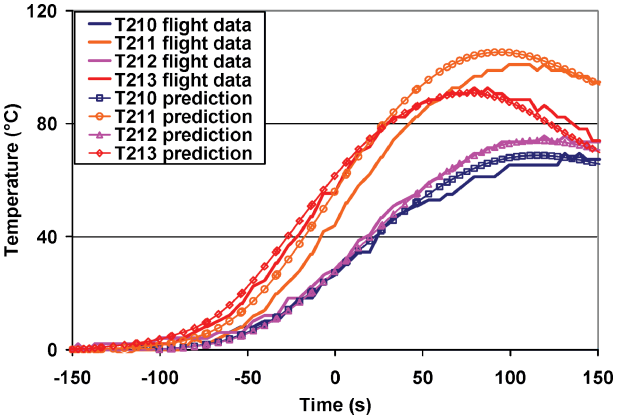

Fig. 17 Final correlation, as temperature change.

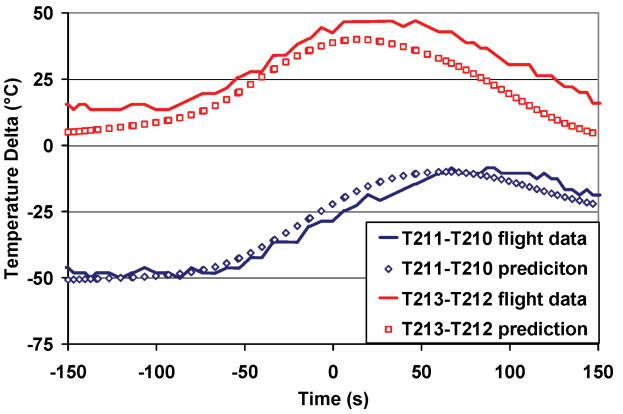

Fig. 18 Temperature differences, hot to cell side.

do this, the initial temperature of each sensor and prediction were removed from the data, leaving only the change from the start of aerobraking. The results of this are shown in Fig. 17. All sensor responses look very good, except that the prediction for sensor T213 is rising a little too rapidly. This might be caused by the low mass of the model, or by excessive conductance for the adhesive under the sensor. The hot-side-to-cell-side differences are shown in Fig. 18, and these have obviously improved greatly in the correlation.

Overall quality of the model was evaluated in several ways, which are listed in Table 6. The primary measure of correlation was the change in temperature caused by aerobraking. This difference in peak temperature change between model and flight, regardless of the timing, is shown in the row "difference in peak temperature change." The worst value on a sensor was 4.4°C. Another measure was the time difference between when peak temperature was reached between the model and flight, shown in the row "difference in peak timing." The difference in actual peak temperature is listed next in the table; differences of about 8°C reflect the inaccuracy of the starting temperatures. The rms difference in the aeroheating temperature rise was calculated, which gives a measure of the accuracy of the model at all times, not just at the peak. The worst rms difference was 6.7°C. The largest overall difference between prediction and flight at all time points was worst for T211 at 12°C. The final row gives the error of the model as a percentage, based on the first row values (peak temperature rise). If the rms of these is taken, it gives an overall correlation error of 3%. All of these measures point to a robust model that reflects in-flight performance accurately.

After this was completed, orbits 10 and 13 were selected at random to verify the correlation and were run with the same model. Correlation plots are shown in Figs. 19 and 20. The overall quality values (rms error for peak temperature prediction) for these were 5 and 6% for passes 10 and 13, respectively. In these lower aeroheating passes, the solar and planetary fluxes make up a larger proportion of the effect, so they are not expected to achieve as good a correlation.

The three-dimensional thermal maps of this correlated model are also useful to visualize the overall thermal condition of the array. In all of these figures, the thickness dimension is exaggerated to improve visualization. The initial predicted temperature gradient

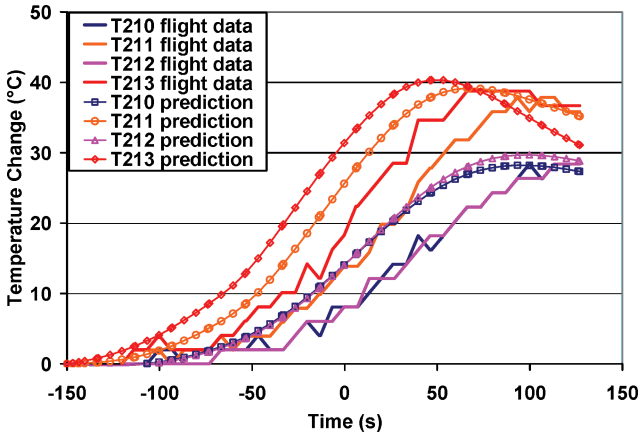

Fig. 19 Correlation for pass 10.

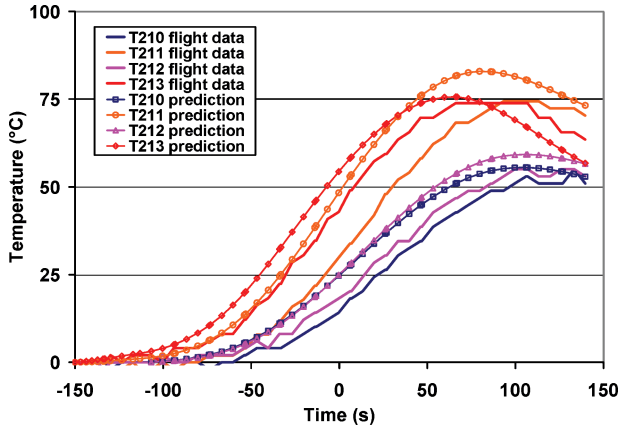

Fig. 20 Correlation for pass 13.

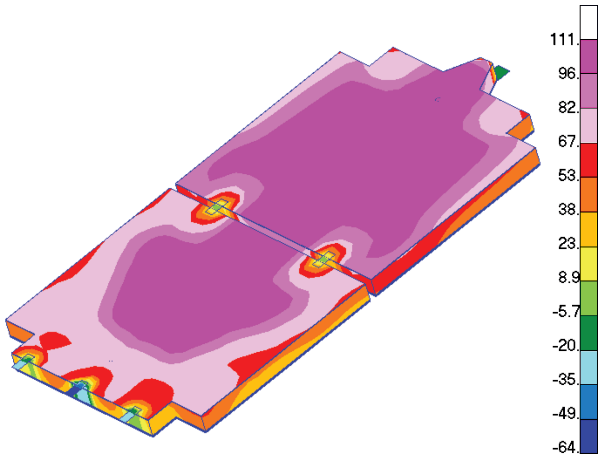

Fig. 21 Final thermal map at 70 s, hot side (°C).

before the start of aerobraking was already shown in Fig. 15. The white painted areas are shown to be effective in cooling the inboard corners of the array, which are expected to receive the highest aeroheating fluxes. The gradient prediction at the time of peak temperature is shown in Figs. 21 and 22 for the hot side and cell side, respectively. The inboard corners have been effectively precooled, so that they do not become the hottest parts of the array. The mass of the hinges can be shown to locally slow down temperature rise. The effects of the relatively high conductance along the facesheets, and lower conductance through the core, can be seen.

A benefit of this full three-dimensional model is that the maximum temperature prediction for each component or material can be calculated and compared with its maximum service temperature. Because MGS had already undergone aeroheating at the time of

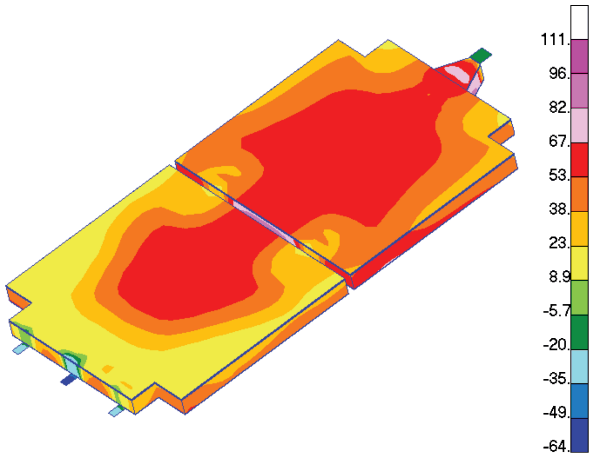

Fig. 22 Final thermal map at 70 s, cell side (°C).

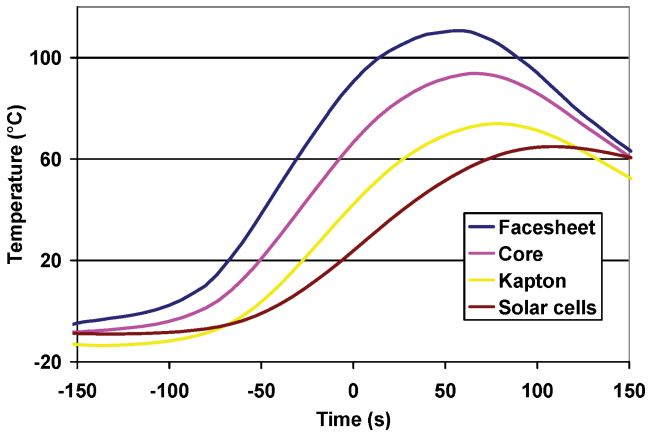

Fig. 23 Material maximum predictions.

this modeling, this capability was not of use to this program, but in general this information can be used in the design process as well as in the trajectory planning. An example is shown in Fig. 23. All materials remained well within their service limits, as expected because the aeroheating levels actually experienced were much less than the design target.

Conclusions

A full three-dimensional thermal model of the MGS spacecraft was developed, and predictions at sensor locations were compared to the flight data. Good correlation was obtained for the highest heating pass (pass 15), as well as for two other earlier passes. A measure of correlation quality was the rms value from four sensors of the flight-to-prediction difference in peak temperature change caused by aerobraking. With use of this measure, model accuracy was in the range 3 to 6%. Lessons learned in the model correlation were utilized in the Mars Odyssey project thermal modeling effort. Small details such as increased mass caused by paint thickness, sensor mass, and radiation through thermal sensor penetrations were found to have substantial impact in achieving a correlated model. The thermal state before the start of aerobraking also played a large part in determining the response.

Acknowledgments

The assistance of Michael Griffin at Lockheed Martin Astronautics in providing drawings, details of the assembly, and advice is gratefully acknowledged. Thanks are owed to the entire Mars Odyssey flight operations team at Langley, in particular Richard Wilmoth for help with the aeroheating calculations, Dick Powell for his support, and Bob Tolson for initiating this effort.

References

[1]Dec, J. A., and Amundsen, R. M., "A Thermal Analysis Approach for the Mars Odyssey Spacecraft's Solar Array," AIAA Paper 2003-3764, June 2003.

[2]*Thermal Desktop User Manual*, Ver. 3.2, Cullimore and Ring Technologies, Inc., Littleton, CO, Aug. 2000.

[3]*MSC/PATRAN User Manual*, Ver. 2000 (r2), MacNeal–Schwendler Corp., Santa Ana, CA, Aug. 2000.

[4]Griffin, M., to Bathias, S., "Solar Array Thermal Model," Internal Memo, Lockheed Martin, 27 May 1999.

[5]Hanna, J. L., Tolson, R., Dwyer, A., and Dec, J., "Autonomous Aerobraking at Mars," *5th International ESA Conference on Guidance Navigation and Control Systems and Actuator and Sensor Product Exhibition*, ESA, Frascati, Italy, Oct. 2002.

[6]Tartabini, P. V., Munk, M. M., and Powell, R. W., "Development and Evaluation of an Operational Aerobraking Strategy for Mars Odyssey," *Journal of Spacecraft and Rockets*, Vol. 42, No. 3, 2005, pp. 423–434.

[7]Wilmoth, R. G., Rault, D. F., Shane, R. W., and Tolson, R. H., "Rarefied Aerothermodynamic Predictions for Mars Global Surveyor," *21st International Symposium on Rarefied Gas Dynamics*, Universite de Provence, Marseille, France, July 1998.

[8]Moss, J. N., Wilmoth, R. G., and Price, J. M., "DSMC Simulations of Blunt Body Flows for Mars Entries," AIAA Paper 97-2508, June 1997.

T. Lin
Associate Editor

Color reproductions courtesy of NASA Langley Research Center.