

Thermal Model Correlation for Mars Reconnaissance Orbiter

Ruth M. Amundsen, John A. Dec, Joseph F. Gasbarre
NASA Langley Research Center

ABSTRACT

The Mars Reconnaissance Orbiter (MRO) launched on August 12, 2005 and began aerobraking at Mars in March 2006. In order to save propellant, MRO used aerobraking to modify the initial orbit at Mars. The spacecraft passed through the atmosphere briefly on each orbit; during each pass the spacecraft was slowed by atmospheric drag, thus lowering the orbit apoapsis. The largest area on the spacecraft, most affected by aeroheating, was the solar arrays. A thermal analysis of the solar arrays was conducted at NASA Langley Research Center to simulate their performance throughout the entire roughly 6-month period of aerobraking. A companion paper describes the development of this thermal model. This model has been correlated against many sets of flight data. Several maneuvers were performed during the cruise to Mars, such as thruster calibrations, which involve large abrupt changes in the spacecraft orientation relative to the sun. The data obtained from these maneuvers allowed the model to be well-correlated with regard to thermal mass, conductive connections, and solar response well before arrival at the planet. Correlation against flight data for both in-cruise maneuvers and drag passes was performed. Adjustments made to the model included orientation during the drag pass, solar flux, Martian surface temperature, through-array resistance, aeroheating gradient due to angle of attack, and aeroheating accommodation coefficient. Methods of correlation included comparing the model to flight temperatures, slopes, temperature deltas between sensors, and solar and planet direction vectors. Correlation and model accuracy over 400 aeroheating drag passes were determined, with overall model accuracy better than 5°C.

INTRODUCTION

Many of the modeling methods used for the Mars Reconnaissance Orbiter (MRO) solar arrays, as well as previous similar aeroheating analyses, are described in previous papers^{1,2,3,4}. As a quick overview, the Thermal Desktop⁵ software was used to model the entire spacecraft and solar arrays, to calculate radiation exchange on the array, radiation to space, and absorbed solar and planetary fluxes. This information was mapped over to a thermal model of a single wing of the

solar array, which was implemented in Patran Thermal⁶. Predicted aeroheating from the drag pass was then applied to this model. Temperature predictions from the thermal sensor locations in the model were compared to the flight data from the thermal sensors. Two separate models were used only because this modeling method leveraged earlier work that made this the simplest method to accomplish the goal.

The position of the thermal sensors on the +X solar array is shown in Figure 1, on the Patran Thermal model of the array. The thickness of the array is exaggerated for viewing. On the non-cell (aeroheated) side, T-109 is the inboard panel sensor, T-110 is midway out the outboard panel, and T-309 is at the outermost corner. T-310 is opposite T-309, on the solar cell side. The overall size of each array (+X and -X) was about 2.5 m x 5.5 m x 2 cm, and each array was about 48 kg in mass. The thermal limit for the array was 175°C. Only the +X array was modeled, since behavior was expected to be symmetric.

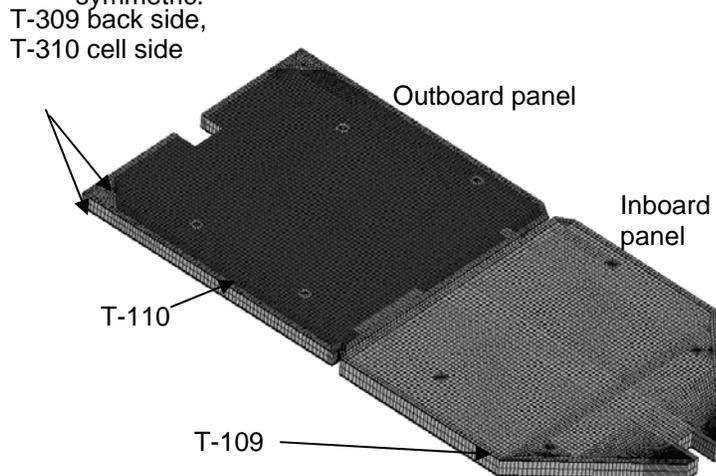


Figure 1. Thermal sensor locations on +X solar array, viewed from bare (non-cell) side.

CRUISE EVENTS MODEL CORRELATION

Several spacecraft maneuvers performed while en-route to Mars were very beneficial to the effort of thermal model correlation. In general, these events are

performed to check on thruster engine performance, antenna tracking, etc. The maneuver might involve pointing the spacecraft in a specific direction for a given period of time, then shifting the spacecraft to a new orientation. The benefit for the thermal model is to have a good set of flight temperatures, under a known set of orbital attitudes and distances from the Earth and Sun. A useful factor in this correlation was that the spacecraft had sensors to detect solar direction vectors, and the flight data included both solar and earth position vectors. These vectors were used as input to the thermal radiation model, so the orbit orientation of the radiation model was known to be correct. The distance to the sun for each event was computed from the trajectory. The radiation model was used to generate radiation conductors and heat fluxes, based on the known positions of the Earth and Sun. By using flight data for the solar orientation, most dependency on the accuracy of the radiation model was taken out. The temperature comparison was thus a check only of the radiation model optical properties and thermal model masses and connections. This was an extremely valuable effort, not only to give an initial check of model validity, but also to exercise the process of bringing in flight data and comparing it to model predictions. It also allowed initial correlation of the model to occur before the stress of aerobraking flight operations, and well before the solar array would encounter any significant aeroheating.

The first cruise event with flight data was Trajectory Correction Maneuver 1 (TCM 1). This event occurred on August 27, 2005, 15 days after launch. The solar direction flight data from this maneuver is shown in Figure 2. It can be seen that this maneuver involves just one major change of attitude. Using this data, and distances to the Earth and Sun from trajectory predictions, the thermal model was run "blind". In other words, these orientations were put into the Thermal Desktop radiation model, and heat fluxes were obtained. These heat fluxes were used in the run of the thermal model in Patran Thermal, producing temperature predictions at the thermal sensor locations. All of this was done without looking at the flight data from the sensors. The predictions were then compared to the flight data without adjusting the thermal model.

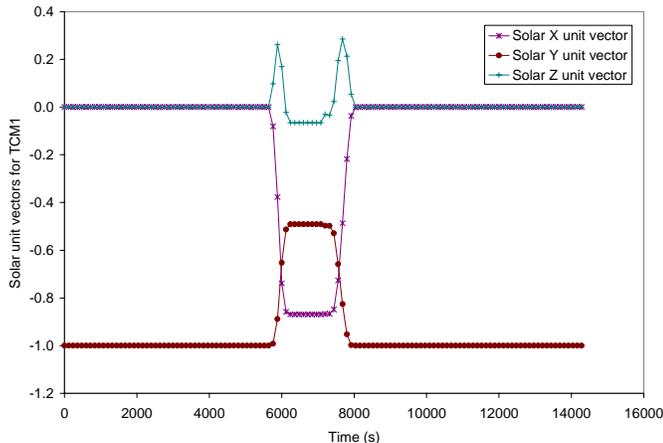


Figure 2. Flight data: solar attitude relative to spacecraft for TCM 1.

The results from the first "blind" run of the model, with no correlation, are shown in Figure 3. It can be seen that the initiation of the events occurs at the right time, and the prediction slopes are similar to the flight data slopes. The errors in the model can be seen more clearly by plotting the data in two different ways. Figure 4 is a plot of the difference of prediction versus flight. These results show that the prediction for T-310, the only sensor on the cell side of the array, was consistently too hot, and the other sensors' predictions were consistently too cold. This is shown in a different way in Figure 5, which plots the temperature difference between sensors T-310 and T-309 (located at the same point on the array, directly opposite each other). This figure demonstrates that the temperature delta through the thickness of the array was consistently too high. This result indicated that the thermal resistance in the model was higher than it should be, such that a higher-than-desired delta was maintained through the thickness of the array.

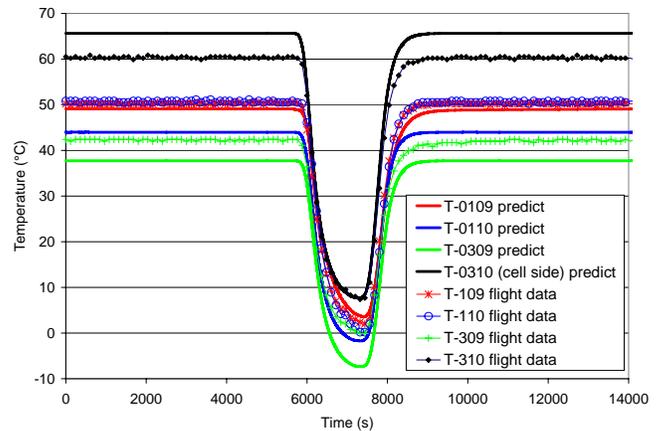


Figure 3. First run of TCM 1 thermal model versus flight data.

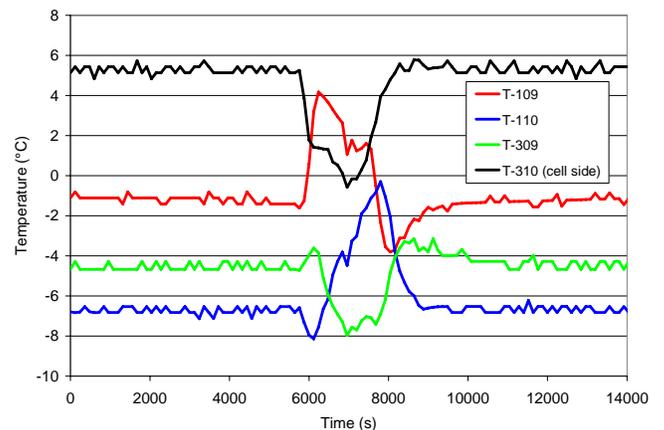


Figure 4. Temperature difference, predict versus flight for TCM 1 first run.

For each event, RMS error was used as a simple measure of the quality of the model correlation. The RMS error was calculated by taking the difference

between each sensor's flight data and prediction at each time point, and summing them via a root-mean-square (RMS) method, so that positive values did not cancel out negative ones. The RMS values give a measure of the model accuracy for that sensor over the entire timeline. The RMS values for the four sensors were averaged to arrive at a single quality measure. For this TCM 1 run, the overall average RMS error for all four sensors was 4.4°C. T-110 had the highest RMS error, at 6.4°C.

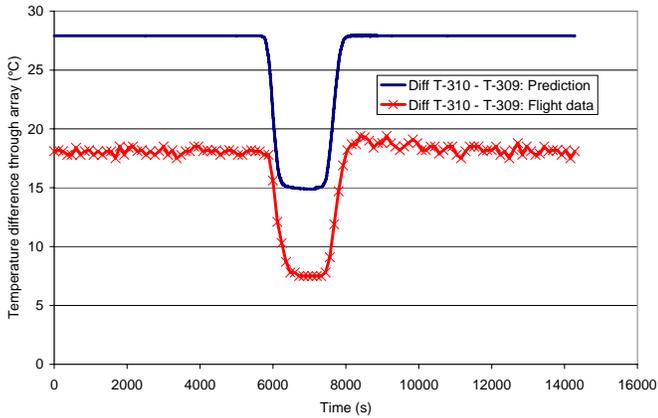


Figure 5. Temperature delta through the array (prediction versus flight) for TCM 1 first run.

Blind runs were made on five different in-cruise events before any correlation changes were made to the model to ensure that any changes would consider data from disparate events, and not just the model response to a single event. The second in-cruise event was Thruster Calibration 1, which was performed on September 15, 2005. This maneuver, shown in Figure 6, runs through three separate spacecraft orientations before returning to the base orientation, and thus offers more robust checking of the thermal model. Thermal model response, shown in Figure 7, shows a similar pattern as in TCM 1; T-310 is predicted too high, while T-309 and T-110 are too low. The predicted through-array delta is 6-12°C too high, as in TCM 1.

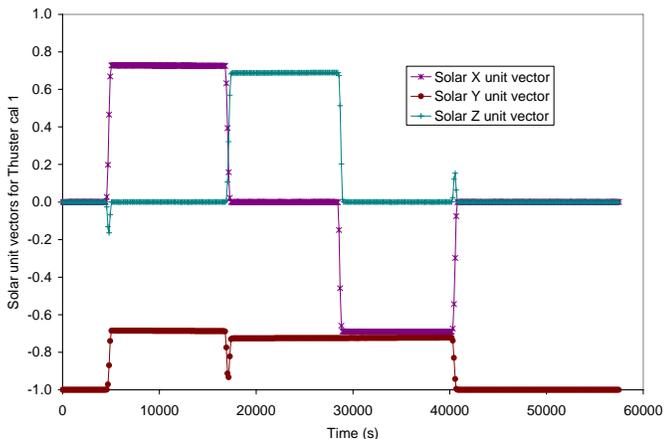


Figure 6. Solar orientation vectors for thruster calibration maneuver 1.

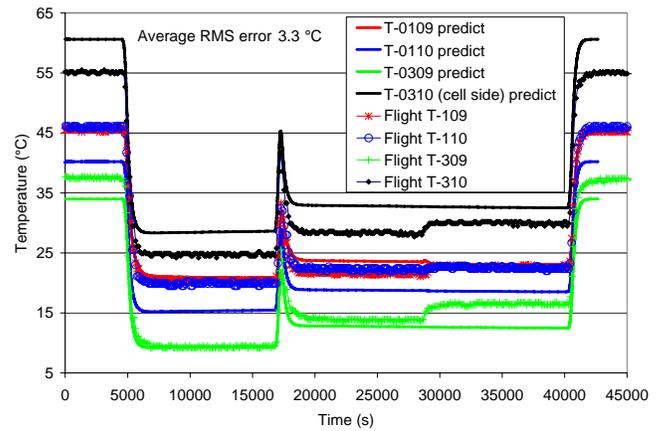


Figure 7. First run of thruster calibration 1 thermal model versus flight data.

The third in-cruise event was the “Stanford” calibration, performed on September 22, 2005. As shown in Figure 8, the orientation and timings for this event were quite different, with much more frequent and severe changes in orientation. The first run of the model is shown in Figure 9; as shown, the model predicts even these relatively rapid changes well and provides good confidence that the general thermal masses of the model are accurate relative to the flight hardware.

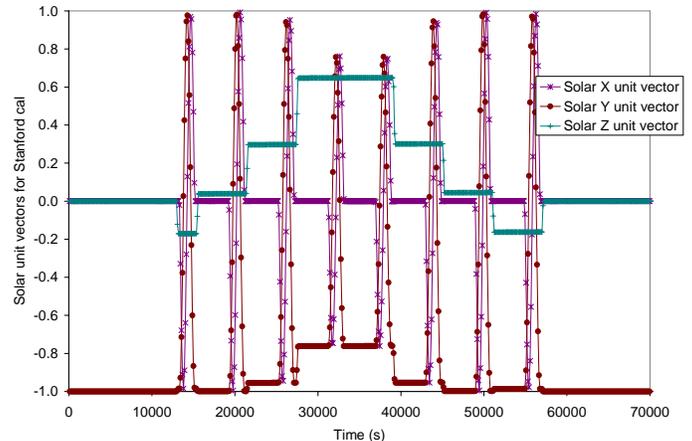


Figure 8. Solar orientation vectors for the Stanford calibration.

The fourth in-cruise event was thruster calibration 2, performed on November 2, 2005. The behavior is shown in Figure 10 and Figure 11, and exhibits a similar pattern as the previous events.

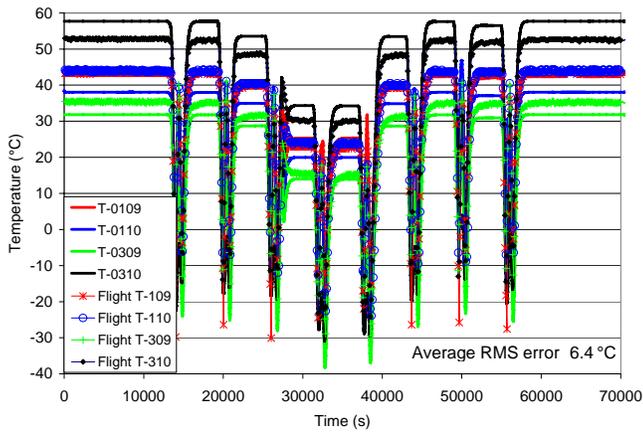


Figure 9. First run, Stanford calibration model versus flight data.

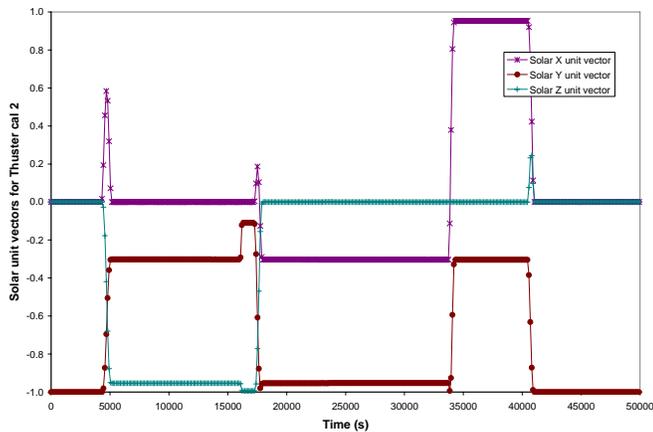


Figure 10. Solar orientation vectors for thruster calibration 2.

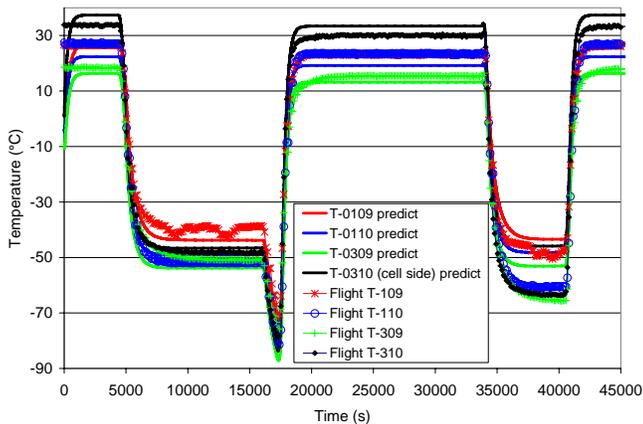


Figure 11. First run of thruster calibration 2 thermal model versus flight data.

The last in-cruise event used for model correlation was TCM 2, which was performed on November 18, 2005.

The solar orientations of this event are shown in Figure 12, and the initial model response is shown in Figure 13. Since most of the event used a constant orientation, the time scale in Figure 13 is expanded to only show the region of interest. One interesting feature of this plot is that, although the T-310 prediction is too high in the quasi-static portion of the timeline, it does not recover quickly enough during the transition, and thus the prediction is too low during the event. Since during the event, the bare side sensors are warmer than the cell side, this can also indicate that the connection between the sides is not strong enough in the model.

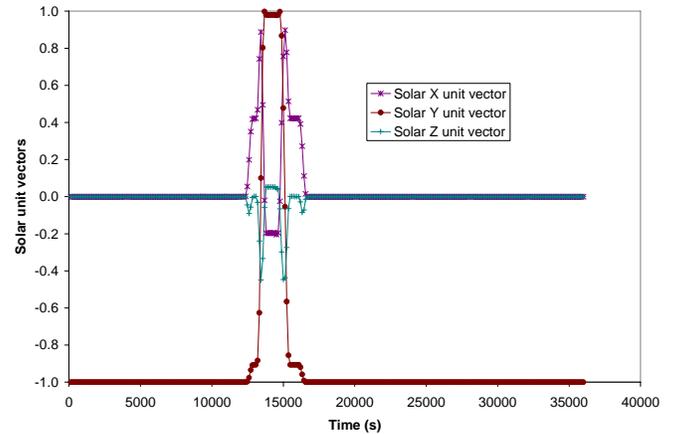


Figure 12. Solar orientation vectors for TCM 2.

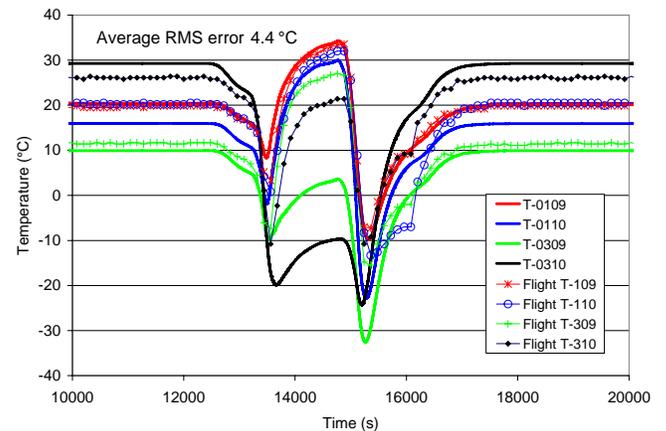


Figure 13. First run of TCM 2 thermal model versus flight data.

After the runs of the initial model on each of these maneuver events, the following observations could be made with regards to the model. First, overall, this is incredibly accurate performance for the first blind run of an uncorrelated thermal model. This performance was only possible because of previous experience with similar thermal models and their correlation, on Mars Global Surveyor (MGS) and Mars Odyssey. The MRO array model mass was verified before flight to match the

as-built mass to better than 2%. The material thermophysical and optical properties were considered well-correlated already from modeling and testing on MGS and Mars Odyssey. Optical properties (α/ϵ) used on the array were 0.81/0.85 for the solar cells, and 0.89/0.72 for the bare M55J composite side. The main changes considered for correlation were to boundary conditions and assumptions. The two types of error measured on this model both averaged about 5°C over all five events: both the error on the peak temperature prediction as well as the RMS error over the timeline of the event. In critique of the model, sensor T-310, the only one on the cell side, tended to be too high. Two possibilities were that the cell side was getting too much sun, or there was not enough conduction through the core. Flight data from sensors T-109 and T-110 tended to be similar when the array was in the sun, but they sometimes differed in the shade. The prediction of T-109 follows the flight data, but the T-110 prediction tended to be too low when the array was in the sun. This result also led to the conclusion of not getting enough conduction through the core. Since T-109 is inboard and T-110 is outboard, it could also have been due to the lack of connection between the two panels, but this is difficult to justify given the large distance lengthwise along the panels between the two. Sensor T-309 predictions are fairly accurate, although the prediction tended to be slightly too low in the sun. In each case, when the delta through the core is plotted, the prediction is consistently too high.

Based on these observations, the first change was to increase the conduction through the core of the array. This number is very difficult to calculate in advance, involving several thin layers of adhesive, conduction through the aluminum honeycomb cells, attachment area of the honeycomb to the facesheets, and radiation between the facesheets. The initial estimate of the convection coefficient, h_c , was 4.35×10^{-5} W/mm²-K for the conduction from each composite facesheet to the aluminum honeycomb core. This value was based on MGS and Mars Odyssey experience, but was purposely left slightly low in order to be conservative. Several higher values were used in trials to evaluate which yielded the best prediction. Values of 5×10^{-5} , 5.5×10^{-5} , 6×10^{-5} , 8×10^{-5} , 1×10^{-4} and 1.4×10^{-4} W/mm²-K were used. The value of 8×10^{-5} W/mm²-K gave the best overall fit to the data. This was determined by taking the best overall RMS value over three of the maneuver events (TCM 1 and thruster calibrations 1 & 2).

The difference predicted between T-109 and T-110 still was not as high as the difference as in the flight data. One reason for this is that a simplifying assumption had been made in the modeling. The view to space over the array is calculated in the Thermal Desktop software. However, it was assumed that whatever portion of the view was obstructed, would be viewing something roughly the same temperature as the array itself. Thus, the remaining Fij (view factor) that did not view space was neglected. An alteration was made to the model, to ensure that this "blocked" Fij was actually set to view

something with a temperature approximately equivalent to the spacecraft outer surface. This was helpful in correcting the difference between the inboard and outboard array (T-109 and T-110). Direct connections between the arrays, simulating the hinges, were also incorporated in the model, but had no effect on the temperature difference between T-109 and T-110.

With these two corrections to the model, the match between the final predictions and the flight data was much improved. Two measures were used to quantify the level of correlation of the model. One was the overall average RMS error for all four sensors. This value was always positive, and represented the overall quality of the correlation over the entire timeline. The second measure was to take the RMS of the difference between flight data and model prediction for the temperature delta between T-309 and T-310 over all time points: the error in predicted temperature delta through the thickness of the array.

The final predictions versus flight data for TCM 1 are shown in Figure 14. The average RMS error for all four sensors was 2.2°C, an improvement from the 4.4°C RMS error before correlation. These errors were very low, which was expected since the solar orientation data was from flight data, and no aeroheating was applied. The basic response of the model to orbital fluxes and radiation was being correlated here. The predicted versus flight deltas are shown in Figure 15. These results indicate that the flight deltas between each of the sensors were matched very well by the final predictions. In particular, the delta through the array, indicated by the difference between T-309 and T-310, had an RMS error of only 2.7°C after correlation. This result was a huge improvement over the 9.5°C RMS error seen in the through-array delta from the initial run of the model. The delta across the outboard (OB) panel, and the delta from inboard (IB) to outboard, tracked qualitatively well, but did not respond in exactly the same way to the maneuver. Part of the reason for this was that they both used T-110 in their calculation. Each of the other sensors kept within a 2°C band of error, whereas T-110 jumped to about 6° error during the maneuver. Detailed corrections to the T-110 connection were made later during aerobraking, as discussed below.

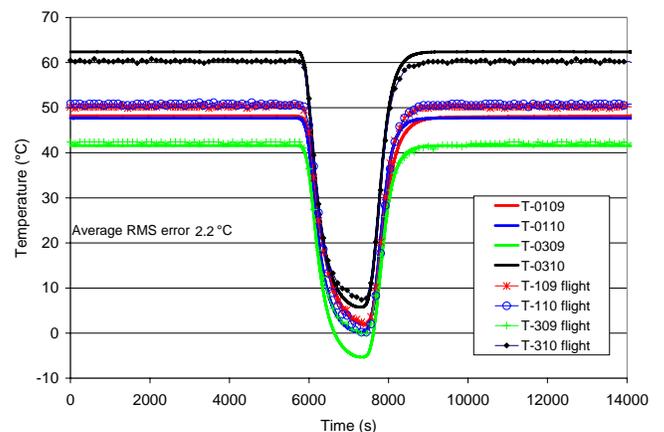


Figure 14. Final prediction versus flight data for TCM 1.

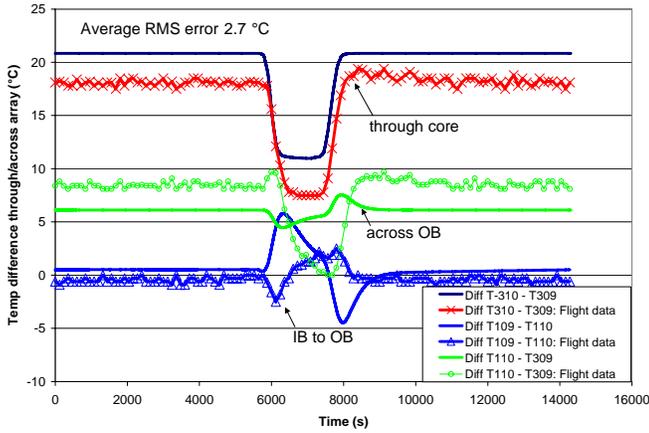


Figure 15. Array deltas (final prediction versus flight) for TCM 1.

The final model prediction for thruster calibration 1 is shown in Figure 16. This maneuver had excellent correlation, with an RMS error averaged over all four sensors of 1.3°C. The correlation of the deltas between sensors is shown in Figure 17; the delta through the array was a very exact prediction, with an RMS error from flight data of only 1.2°C. Although the IB to OB delta also matched very well, the delta across the OB panel was not matched as well. This maneuver involved orientations where the OB panel was partially shaded. Because the spacecraft radiation model was not exact, and specularities and surface shapes of MLI are difficult to match, the shadows on the OB panel are not matched exactly in the radiation model. This result means that the thermal gradients on the OB panel will not be predicted exactly. However, since in general the solar arrays would be undergoing aerobraking while in the shade, and so gradients of this type on the OB panel would not be occurring during the maximum heating of the array, this result was not of great concern.

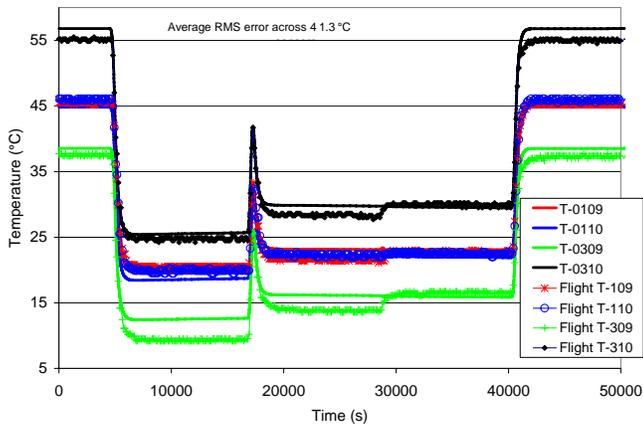


Figure 16. Final prediction versus flight for thruster calibration 1.

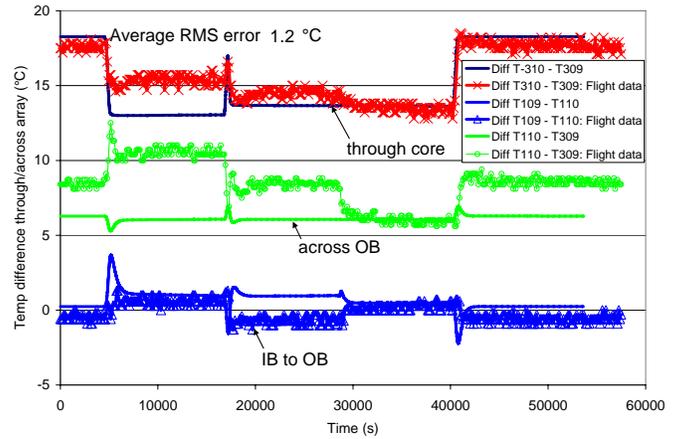


Figure 17. Deltas across array (final prediction versus flight) for thruster cal 1.

The correlated model prediction versus flight data for the Stanford calibration is shown in Figure 18. This maneuver was very useful for correlation. Looking at the slopes of each sensor, it can be seen that the timing of events was matched very well, which gave great confidence that the basic thermal mass and connection of the model was accurate. Since the events are so compressed, an expanded time plot of one cycle is shown in Figure 19. It can be seen that sensors T-309 and T-310 lagged behind slightly, so that during the fast reversals, they did not reach the same maximum as the flight data. However, in the cases where the array had more time to come to equilibrium, they reached the same temperature. Since the model was to be used mainly for aeroheating, which is a relatively slow process, this behavior was acceptable. The response of the delta through the array was also well-matched both in static behavior, and during changes, as shown in Figure 20.

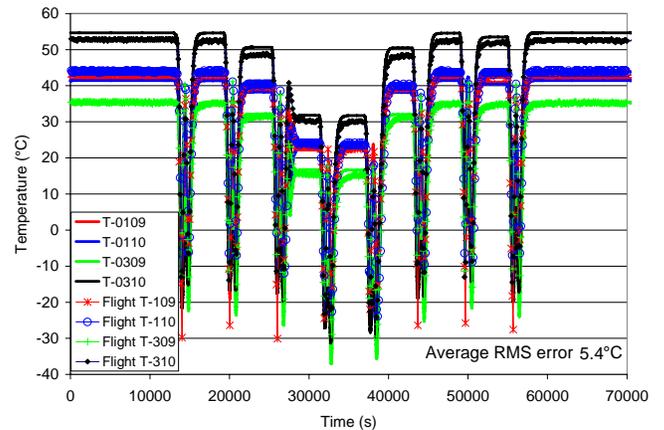


Figure 18. Stanford calibration final thermal prediction versus flight.

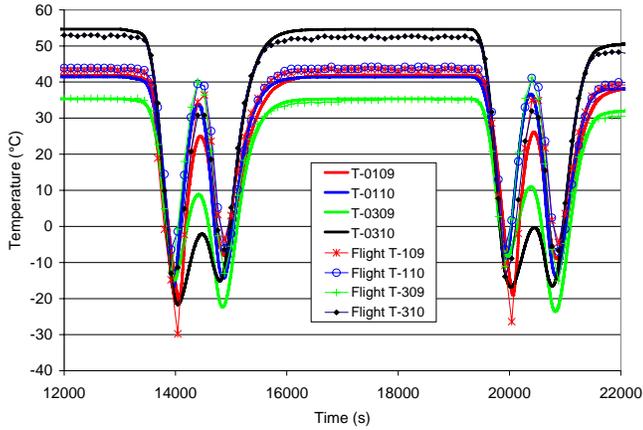


Figure 19. Expanded-time-scale plot of the Stanford calibration.

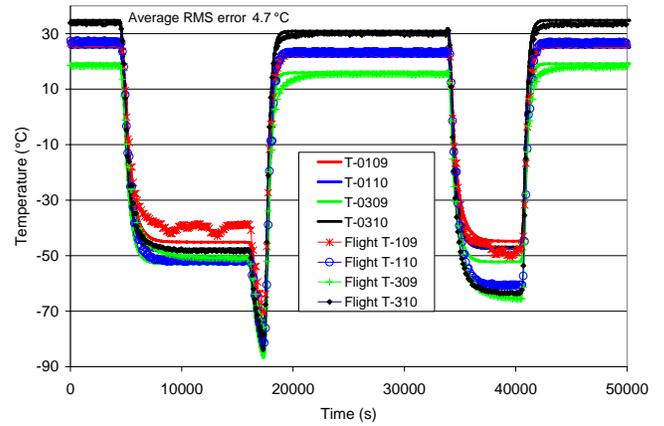


Figure 21. Final prediction versus flight for thruster calibration 2.

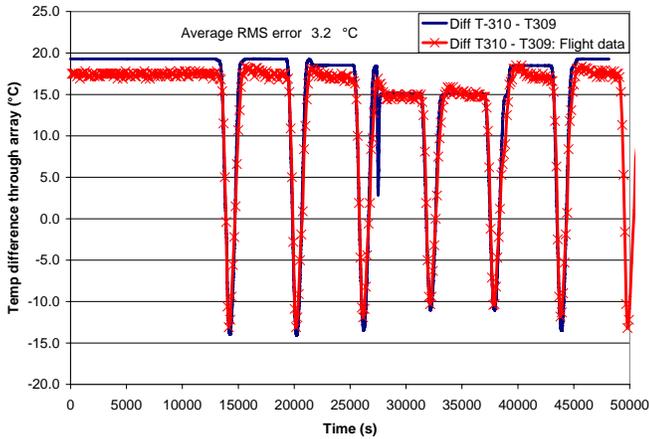


Figure 20. Temperature delta through array (final prediction versus flight) for the Stanford calibration.

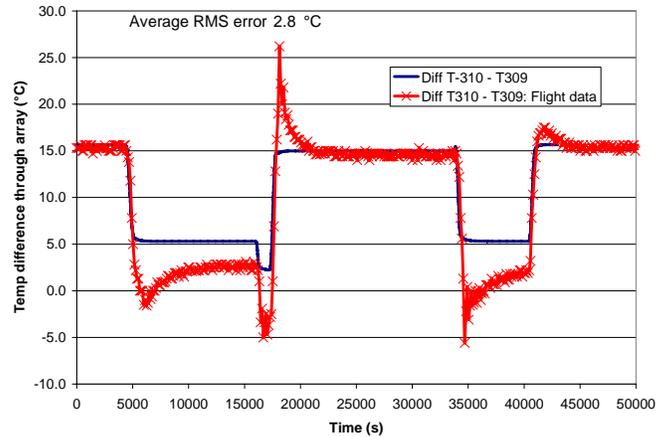


Figure 22. Delta through array (final prediction versus flight) for thruster cal 2.

The post-correlation model results versus flight for the thruster cal 2 event are shown in Figure 21 and Figure 22. The prediction of the temperature delta through the array improved, but the overall model RMS error over the four sensors was still 4.7°C. The error was especially high in the time region from 35000 to 40000 seconds, where the model prediction was too high. One possible reason for this high error is that there was a relatively large power draw from the solar cells at that time. This would decrease the solar energy that was actually producing heating of the array, since it would instead be used to produce electricity. This feature was not included in the model, again because this would not materially affect maximums seen during aerobraking (since the array would be in the shade during aerobraking and not undergoing power draw). A supporting piece of evidence that this was the correct reason, was that the +X and -X array data were different during this time period. The orientation of the solar angles during this time period should not cause that asymmetry.

Final results for TCM 2 are shown in Figure 23 and Figure 24. The overall RMS error decreased from 4.4°C to 4.0°C, and the RMS error for the delta through the array improved greatly, going from 7.6° to 1.0°C.

A summary of the overall model accuracy pre- and post-correlation is shown in Table 1. In general, both the average RMS error as well as the error in the delta through the array improved substantially. It was extremely useful to have the data on these cruise phase events for model correlation, prior to aerobraking, to make all possible improvements to the model before entry into the atmosphere. For each of these events, the +X flight data was compared to the -X data, and behavior was symmetric, as expected.

AEROBRAKING MODEL CORRELATION

Mars Orbit Insertion, MOI, occurred on March 10, 2006, to successfully capture the MRO spacecraft into the intended initial 35-hour-period orbit. From that time, aerobraking was used to shrink the orbit to the desired science altitude. The first several orbits were used for checkout of systems and science experiments. The first pass that actually dipped into the top of the Mars atmosphere was pass 14, at the end of March. Since the objective of the model was to predict the peak temperature experienced during aerobraking, a parameter to measure model correlation quality in addition to the RMS error was added. The parameter was simply the error in predicting the maximum temperature experienced by any sensor on the array. Since the maximum temperature observed by any sensor was the value used to determine the limit line that the spacecraft would fly into the Mars atmosphere, the error in this value was felt to be one of the most important measures of the model predictions.

On the first several passes after pass 14, the temperatures predicted by the model were too high (average RMS errors of about 30°C), as shown in Figure 25 for pass 14. These passes had negligible aeroheating, so there is no perceptible peak around the 814 second time of periapsis, and the peak error value is not meaningful. The model error was not due to an erroneous prediction of aeroheating, because the predictions were uniformly high, including the initial temperature prediction at the start of the pass. This meant that there were modeling inaccuracies in the orbital radiation model. Thus, steps were taken to improve that model. These initial passes with insignificant aeroheating were very useful as a grace period to improve the model with respect to planet and solar fluxes, orbit orientation, etc. Several errors were found and corrected that brought predicted temperatures down. These errors included a solar flux that was too high, as well as conservative margin factors on Martian albedo, aeroheating and solar loads that had been inadvertently left on. The solar flux was too high because the model had been run for the in-flight cruise events; it was not corrected after the last cruise event to the true lower value it had when the spacecraft was at Mars. This was a good lesson learned for the team, that in-flight events are useful to correlate to, but care must be taken to ensure that every aspect of the model is returned to the values it will have for flight operations. After these corrections, the correlation over most of the orbit was much better, as shown by the initial temperatures in Figure 26. The average error in the initial temperatures decreased from 27°C to 4°C for pass 14 when these corrections were made. However, the change in temperature as the spacecraft went into shade at about 200 seconds was still not correct, as shown both in Figure 26 for pass 26 and in Figure 27 for pass 28. The decrease in temperature was not sharp enough, which left the temperatures too high at the start of the aeroheating at 750 seconds, so that the maximum

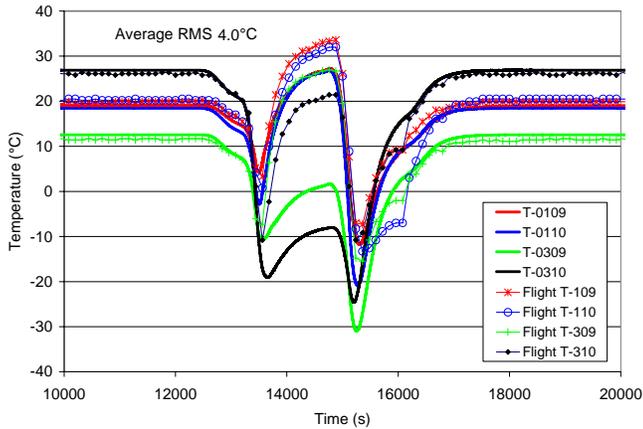


Figure 23. Final prediction versus flight for TCM 2.

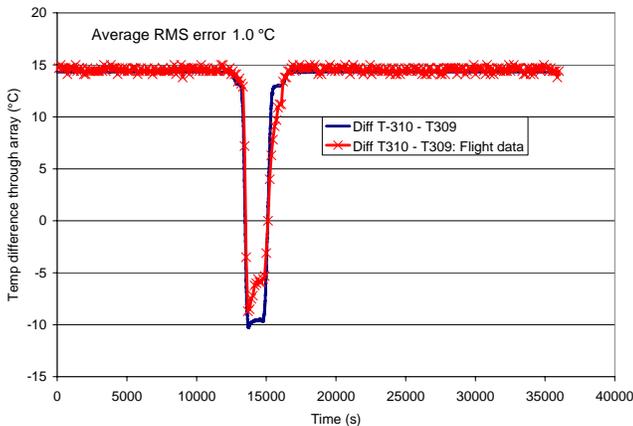


Figure 24. Delta through array (final prediction vs. flight) for TCM 2.

Table 1. Cruise Phase Model Correlation Summary

	Pre-correlation		Post-correlation	
	Average RMS error (°C)	Delta through array RMS error (°C)	Average RMS error (°C)	Delta through array RMS error (°C)
TCM 1	4.4	9.4	2.2	2.8
Thruster cal 1	3.3	7.0	1.3	1.2
Stanford calibration	6.4	9.3	5.4	3.2
Thruster cal 2	5.5	6.7	4.7	2.8
TCM 2	4.4	7.6	4.0	1.0

temperature observed in aeroheating was over-predicted.

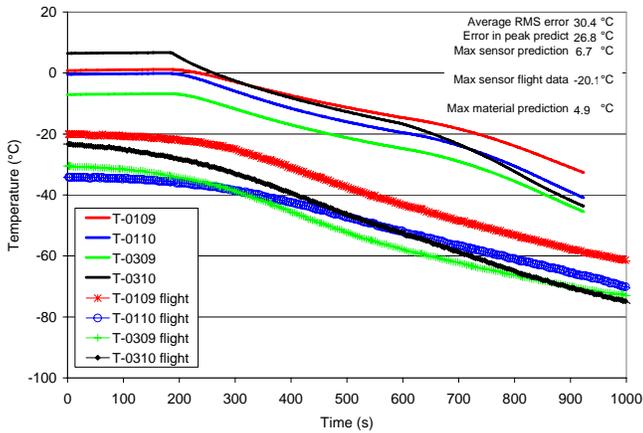


Figure 25. Initial model prediction vs. flight data for pass 014.

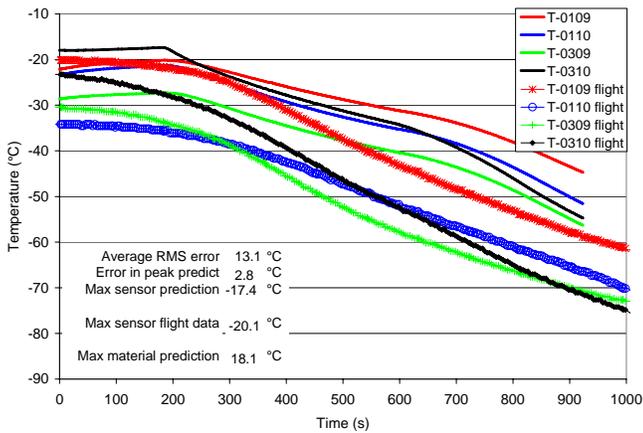


Figure 26. Prediction with correct solar flux vs. flight for pass 014.

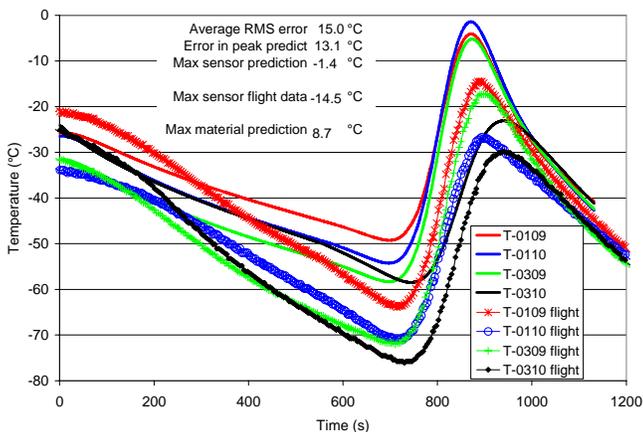


Figure 27. Original model versus flight data for pass 028.

The reason for this difference in slope on the cooldown was not immediately evident; however, it was discovered that the temperatures could be artificially corrected by a temporary expedient of copying over files. To allow a more accurate prediction while the physical reason for this was investigated, this “artificial” method was used for a short period while the real physical explanation was being investigated. This method was as follows. From the radiation model, there were 20 time steps, corresponding to 20 points around the orbit. The sixth time step was in the shade. By taking the heating file for the sixth time point, and copying it over the heating files for time points three through five, the model could be forced to produce a prediction that much more closely matched the flight data, as shown for pass 28 in Figure 28. As shown, the slope of the prediction matched the flight data much more closely; the RMS error was cut by nearly 50%, and the absolute error on the peak temperature was decreased from 13°C to 1°C.

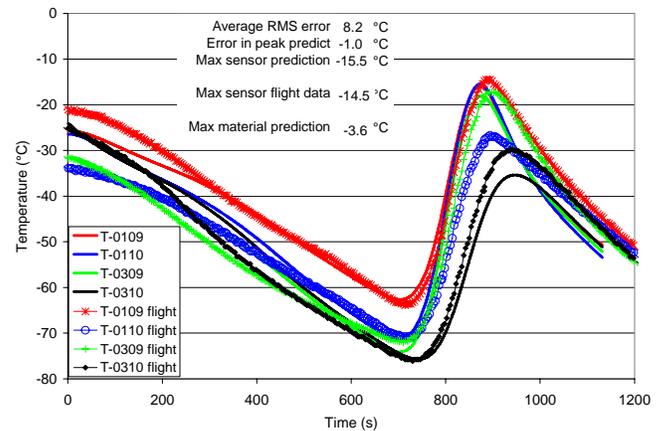


Figure 28. Artificially corrected model versus flight data for pass 028.

This artificial-fix model was used for several passes to make predictions while the reason for the orbital model error was evaluated. The error was specific to the rate of temperature decrease when the spacecraft was slewing to the aeropass orientation. Quasi-static temperatures during the vacuum phase of the orbit, near apoapsis, were predicted fairly well. Thus, suspicion centered on the model’s representation of vehicle orientation just before the aeropass. An aid in resolving this issue was the flight data available on the solar and planet orientation. Embedded in the flight data were x, y, z unit direction vectors to both the sun and Mars. These vectors could be compared to the output from Thermal Desktop on model orientation. The solar data matched fairly well, but not perfectly, and the data for the planet vector had major deviations. The vehicle orientation during aeropass had been defined to the thermal team as “nadir-pointing”, which normally has a very specific meaning in orbital terms, and indicates that the bottom face of the spacecraft points continuously in the nadir direction, i.e., directly at the planet. After renewed consultation with the navigation team, it was determined that the aeropass was not actually run nadir

pointing. It was run inertially-fixed; in other words, the orientation of the spacecraft stayed fixed inertially (but not with respect to the planet vector) over the course of the aeropass. The fixed orientation was chosen so that at perapsis, the spacecraft would be nadir-pointing. In flight, as it turned out, the spacecraft was flying at a 7° angle of attack (AOA). Thus, the fixed orientation for the entire drag pass was selected to give that 7° angle at the perapsis.

This true flight orientation was incorporated in the Thermal Desktop model by inserting trackers to maintain Earth point of the antenna over the majority of the orbit, but slewing to the fixed orientation for the period of the drag pass. An illustration of this analysis, for four of the 20 time points, is shown in Figure 29. After making this change, the solar and planetary direction vector predictions agreed quite well with the flight data. The thermal model was then re-run using these solar/planet heating files. The result was a much improved correlation without the false copying over of the early time point files. These results are shown in Figure 30 for pass 49.

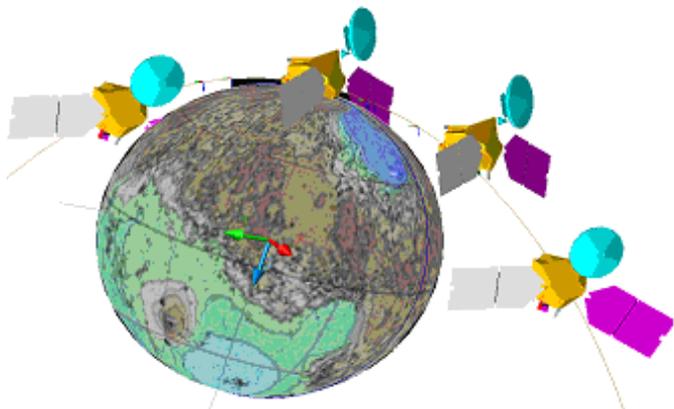


Figure 29. Thermal Desktop model: orientation change using trackers.

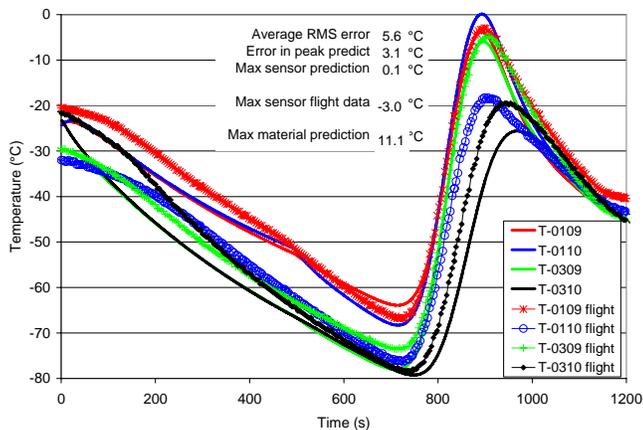


Figure 30. True corrected model versus flight data for pass 049.

The in-plane temperature gradient for pass 64, using this model, is shown in Figure 31; the influence of the 7° angle of attack in warming the upper edge of the model is clear. Also, it is obvious that connections between the inboard panel and outboard panel, such as due to hinges, have merely local effects. By this point in time, the aeroheating during the pass was becoming more of a factor, and the peak during aeroheating actually rose above the initial temperature. A huge advantage of the MRO orbit orientation was that the spacecraft entered the shade just before the drag pass, as shown by the temperature drop during the first 700 seconds of Figure 30. This gave about a 50°C cushion of cooling that could be used up by the aeroheating, without breaking above the quasi-static initial temperature caused by solar heating during the vacuum phase of the orbit.

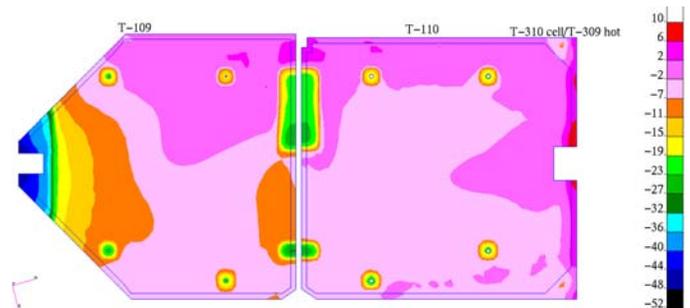


Figure 31. Thermal map of peak temperatures on pass 064 (°C).

There was an additional minor error observed in the thermal modeling. Overall, the sensor T-110 had a much higher error than the other sensors. This is the sensor located midway down the outboard panel, on the hot side. As can be seen in Figure 30 as an example, the flight data for sensor T-110 tended to track with the sensor on the cell side, T-310. However, the predictions for this sensor tended to track with the other hot side sensors, T-109 and T-309. The in-cruise events had shown problems with correlating T-110, but not exactly the same behavior as the aerobraking passes. Thus, it was assumed that this behavior was linked to sudden heating on the hot side. Since the tendency of T-110 to track the cell side rather than hot side sensor persisted over all orbits, regardless of aeroheating severity, it was hypothesized that the sensor might have some enhanced connection to the cell side, due to a variation in manufacturing at that location. Many sources could lead to a slight increase in the effective connection between facesheets at that sensor's location: a local increase in honeycomb density or state of crush, a higher influx of adhesive, lay-up of wires in that location, or embedded hardware elements not included in the model. It was felt that this assumption was more reasonable than assuming a local decrease in the aeroheating in that area, both because the behavior was so repeatable, and because there was no known physical entity present that would have caused a local decrease in heating. Also, it would have been too coincidental that the aeroheating would decrease just

the right amount to allow the sensor to track the cell side temperature. As was done during the in-cruise events, the flight data for the +X array versus the -X array was continually evaluated, to ensure that the behavior observed on the T-110 sensor was not simply a local anomaly. The two arrays were always found to have similar behavior; thus it was assumed that the cause was some hardware element present on both arrays. Another possibility was that the T-110 sensor (and its duplicate on the -X array) were actually installed on the cell side of the arrays, and mis-represented on drawings; this would also be characterized by a local increase in conduction.

Several values for a local increase in conduction under sensor T-110 were evaluated in the modeling. Results were graded by looking at the errors on all four sensors over several orbit passes. The best overall results came from adding a localized conduction of $1.6 \times 10^{-4} \text{ W/mm}^2\text{-K}$ over a 1500 mm^2 area under T-110. This correction was used for passes 113 and beyond. An example of the improved behavior is shown in Figure 32.

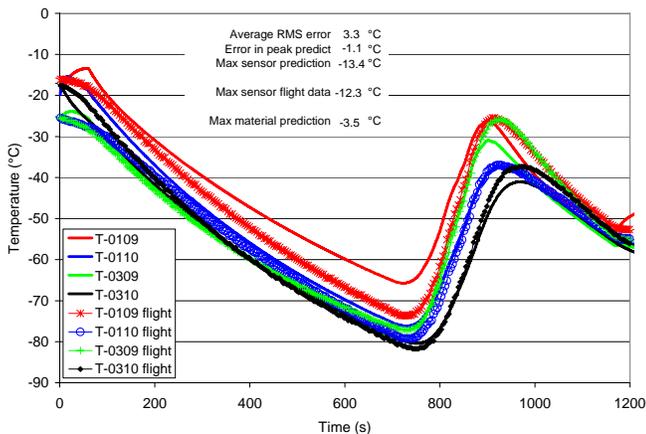


Figure 32. Flight data versus model prediction for pass 166.

Several times during aerobraking operations, it was noted that the flight thermal peak observed during aerobraking would undergo a slow shift over several passes where it would start to deviate slightly from the model prediction. This was corrected by changing the model's aeroheating accommodation coefficient, or the factor on the absorption of aeroheating. This factor was changed to values ranging from 0.85 to 1.0 at different times during the mission, in order to more accurately represent the aeroheating behavior. There was no apparent correlation of the value used and any mission parameter such as atmospheric density, velocity, or array initial temperature.

One factor that often affected the error observed in the thermal model was that initially, the density profile used was predicted before the pass. A0 was the multiplier on the density in the atmosphere model used to more closely represent densities recovered from

accelerometer data. In the initial predictions, A0 was assumed to be 1; the assumption was that the model would closely represent actual densities. As operations progressed, the error in the atmosphere model required density multipliers (or A0) as large as 5, or 500% what was originally modeled for those altitudes. This error in atmosphere modeling translated to error in the temperature predictions from pre-operations to the actual operations phase. Also, the maximum densities varied orbit-to-orbit by 40% 1-sigma, making orbit-to-orbit temperature predictions no better than what the atmosphere model would suggest. When the atmospheric density was under-predicted, the aeroheating was also under-predicted, and the thermal model then under-predicted temperatures. Partially in order to isolate this behavior, and to be able to quantify the error in the thermal model alone, the analysis process was changed in June 2006 (pass 100), about two months into aerobraking. Instead of using the pre-pass predicted density profile, the reconstructed density from flight data available after the pass was used for aeroheating to allow much better correlation of the thermal model. The density profile taken from flight data was usually available within 12 hours after the drag pass, which still allowed rapid correlation with the post-pass flight temperature data. An example of the model-to-flight comparison using reconstructed flight densities was shown above in Figure 32. The correlation using this method to the highest temperature pass during aerobraking, pass 262, is shown in Figure 33. Correlation of the model, even at this relatively high heating pass, was very good, with 4.5°C RMS error, and 0.8°C error on the peak prediction.

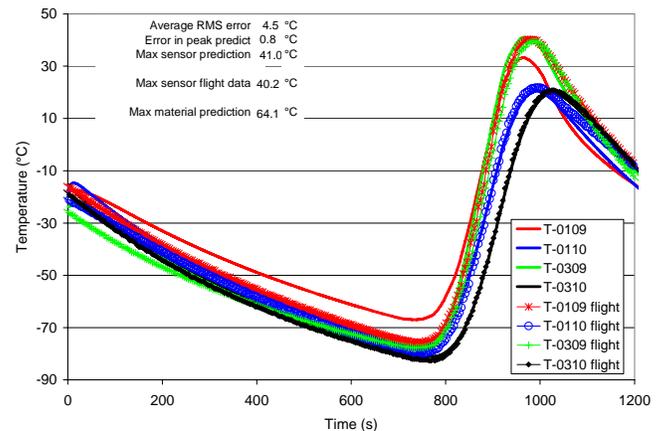


Figure 33. Flight data versus model prediction for pass 262.

There were two additional reasons that the thermal model was no longer run before each drag pass. First was that the analysts had a response surface model available⁷ that allowed for quick predictions of future passes with good accuracy, without running the full thermal model. The second reason was that the aeroheating had been extremely benign, with the solar panel temperature rarely rising above 0°C, and there

was no plan to make the aerobraking more aggressive. In fact, the solar arrays experienced higher temperatures during the in-cruise events TCM 1 and thruster calibration 1 (55 to 60°C) than ever observed during aerobraking (41°C on pass 262). Since the solar panel temperature was usually around 0°C, about 170°C below its thermal limit, there was no urgency in having a detailed thermal model run before each pass. Data over the entire mission of peaks on the flight sensors, as well as predicted peaks both on the sensors and overall maximum on the array, are shown in Figure 34.

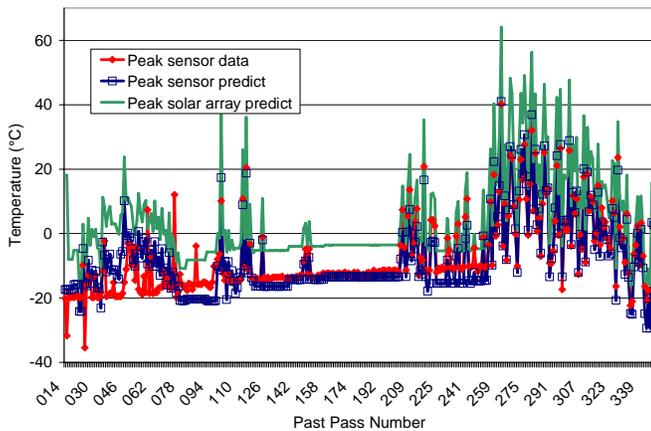


Figure 34. Mission thermal data: flight and model predictions.

After making this change in process to use reconstructed densities, the average error observed in the thermal model decreased dramatically. While using the pre-pass predicted A0, the error in predicting the peak temperature had been off as much as 25°C for some passes, both high and low, due to errors in predicting A0. After drag pass 100, when the change was made to running the thermal model with an actual flight density profile, the peak error was never more than 10°C, and only rarely went over 5°C. The RMS average of peak model error from that point to the end of the mission, over 250 drag passes, was 3°C, as shown graphically in Figure 35. This average error is done as an RMS, because taking a straight average of the peak errors on each pass would allow the positive and negative errors to cancel each other out, yielding an artificially low total error. The average of the RMS error for each pass (the total error observed over the timeline of each pass, averaged as an RMS value) was about 5°C.

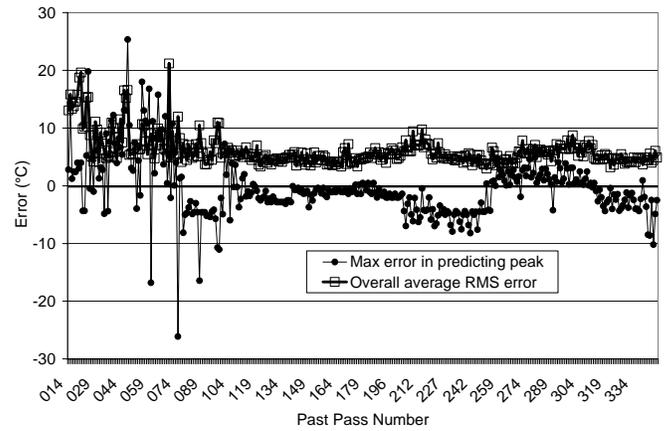


Figure 35. Overall RMS and peak errors over MRO aerobraking operations.

CONCLUSION

The thermal modeling of the solar arrays for Mars Reconnaissance Orbiter in general correlated very well with flight performance. An overall average error as low as 5°C over hundreds of passes was possible because of the knowledge gained from earlier thermal modeling of Mars Global Surveyor and Mars Odyssey aerobraking. Use of flight data from the cruise events such as thruster calibrations was very helpful in verifying the accuracy of the model, as well as allowing corrections to occur before aerobraking operations. Use of other flight data in addition to temperatures, such as solar and planet vectors, was critical in facilitating correlation. Much credit goes to the spacecraft development team for including measurement and rapid downlink of these parameters. The most significant changes to the model included increasing the through-array conductivity, changing the spacecraft orientation through the drag pass, and correcting the solar flux level at the planet. For the locations of these sensors, connections between the panels such as hinges were not found to be important factors. One important though obvious lesson learned was that, after the cruise events have been used to correct the model, model parameters must be returned to the values expected during flight operations at the planet. The overall aeroheating during flight operations was extremely benign, which led to the thermal predictions not being a critical part of navigation team decisions. However, the lessons learned for correlation and modeling methods will enable more efficient and ever more accurate thermal modeling on future aerobraking missions. The performance of the fully-correlated model in producing an average peak model error of 3°C and time-averaged RMS error of 5°C, over hundreds of orbit passes, is exceptionally good for a model and physical scenario of this complexity.

ACKNOWLEDGMENTS

The support of the entire Mars Reconnaissance Orbiter Navigation team is gratefully acknowledged; in particular the support of Jill Prince, Scott Striepe, Jeremy Shidner, Richard Wilmoth and Derek Liechty has been invaluable. Also, the support of Christian May and Neil Tice at Lockheed Martin is gratefully acknowledged.

ACRONYMS

AOA	Angle of Attack
IB	Inboard
MGS	Mars Global Surveyor
MLI	Multi-Layer Insulation
MOI	Mars Orbit Insertion
MRO	Mars Reconnaissance Orbiter
OB	Outboard
RMS	root-mean-square
TCM	Trajectory Correction Maneuver

REFERENCES

1. Dec, John A., Gasbarre, Joseph F., and Amundsen, Ruth M., "Thermal Modeling of the Mars Reconnaissance Orbiter's Solar Panel and Instruments During Aerobraking," 07ICES-64, 37th International Conference On Environmental Systems, Chicago, Illinois, July 2007.
2. Amundsen, Ruth M., Gasbarre, Joseph F., Dec, John A., "Thermal Analysis Methods for Aerobraking Heating," 16th Thermal & Fluid Analysis Workshop (TFAWS 05), Orlando, Florida, August 8 - August 12, 2005.
3. Dec, John A., and Amundsen, Ruth M., "A Thermal Analysis Approach for the Mars Odyssey Spacecraft's Solar Array," AIAA 2003 Thermophysics Conference, Orlando, Florida, June 2003.
4. Amundsen, Ruth M., Dec, John A., and George, Ben E.; "Aeroheating Thermal Model Correlation for Mars Global Surveyor (MGS) Solar Array," *AIAA Journal of Spacecraft and Rockets*, Volume 42, Number 3, May-June 2005, pp 464-473.
5. Thermal Desktop User Manual, Cullimore and Ring Technologies, Inc., Version 4.7, October 2004.
6. MSC/PATRAN User Manual, MacNeal-Schwendler Corporation, Version 2005, December 2004.
7. Dec, John A., "Probabilistic Thermal Analysis During Mars Reconnaissance Orbiter Aerobraking," AIAA-2007-1214, 45th AIAA Aerospace Sciences Conference, Reno, Nevada, January 2007.