AAS 13-425

APPROACH AND ENTRY, DESCENT, AND LANDING OPERATIONS FOR THE MARS SCIENCE LABORATORY MISSION

Allen Chen,[*] Martin Greco,[†] Tomas Martin-Mur,[‡] Brian Portock,[§] and Adam Steltzner[**]

On August 5th, 2012, at 10:31 PM PDT, the Mars Science Laboratory (MSL) rover Curiosity landed safely within Gale Crater. Her successful landing depended not only upon the flawless execution of the numerous critical activities during the seven minute entry, descent, and landing (EDL), but also upon the operational preparations and decisions made by the flight team during approach, the final weeks, days, and hours prior to landing. During this period, decisions made by the flight team balanced operational risk to the spacecraft in flight with any resulting risks incurred during EDL as a result of those decisions. This paper summarizes the operations plans made in preparation for Approach and EDL and the as flown decisions and actions executed that balanced the operational and EDL risks and prepared the vehicle for a successful landing.

INTRODUCTION

During Approach, the last month of the seven and a half month cruise to Mars, preparations for entry, descent, and landing became the focus for the flight team. These preparations centered upon three primary objectives: 1) get the spacecraft ready for EDL, 2) get the spacecraft to the right place to start EDL at the right time, and 3) tell the spacecraft what it needs to know to execute EDL successfully.

To get the spacecraft ready for a successful EDL, the flight team developed and executed plans to put the flight vehicle in as close as to the ground tested configuration as possible. This included an operationally intensive multi-day "cold boot" activity to match the flight software and avionics state as closely as possible to ground testbeds, in an attempt to "fly as you test." Additionally, the team performed hardware checkouts and monitored autonomous preparations

[*] Mars Science Laboratory Entry, Descent, and Landing Operations Lead, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.
[†] Mars Science Laboratory, Entry, Descent, and Landing Activity Lead, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.
[‡] Mars Science Laboratory Entry, Descent ,and Landing Navigation Team Chief, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.
[§] Mars Science Laboratory Deputy Mission Manager, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.
[**] Mars Science Laboratory Entry, Descent, and Landing Phase Lead, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.

© 2013 California Institute of Technology. Government sponsorship acknowledged.

for EDL, including thermal conditioning, software mode transitions, and others to ensure the spacecraft was in the required state for EDL.

Trajectory correction maneuvers (TCMs) in approach represent opportunities to literally put the spacecraft in a position to succeed during EDL. The EDL system is designed to accommodate a range of delivery errors, which are the differences in position and velocity of the spacecraft at the start of EDL from the targeted position and velocity. Getting the spacecraft close to the target for the start of EDL reduces EDL risk. TCM opportunities were included and utilized in the approach operations plans. Each TCM decision requires the flight team to balance operational flight risk and EDL risk: doing a TCM increases spacecraft risk versus taking no action; allowing a delivery error to occur because a TCM was not executed pushes risk into the EDL system performance. TCM decision points, criteria, and guidelines were established to balance these risks; the criteria were successfully applied in flight to all approach TCM decisions.

Unlike its ballistic entry predecessors, Curiosity employed a guided entry system to increase delivered mass capability, increase landed elevation capability, and, perhaps most significantly, reduce miss distance from a given landing site target. To achieve these objectives, the spacecraft needs knowledge of its navigation state; its ability to "fly out" errors can never be better than its knowledge of its true state. During approach to Mars, navigation data is being collected nearly continuously with navigation solutions being produced several times a day. These solutions vary in quality with amount of data collected, time since the last maneuver, and distance to Mars among other factors.[1] A series of EDL parameter update opportunities, including navigated state updates, were designed into approach operations plans to guarantee that the EDL system had adequate knowledge at the beginning of EDL. As with any operational activity, parameter updates can introduce system risk through human error, spacecraft error, or infrastructure issues. The flight team developed parameter update guidelines and an analysis process; these were successfully utilized in flight decisions.

This paper details the operations plans, decision-making criteria, and the rationale behind both developed by the MSL team. It also summarizes their successful application during Curiosity's approach and the impact on the resulting EDL performance.

GETTING THE SPACECRAFT READY

Approach Activities

As Curiosity drew closer to Mars at the end its 8.5-month cruise, activities in preparation for EDL ramped up significantly and were added to the normal cruise operations activities. These included a "cold boot" series of activities, software and parameter loads, and checkouts. All of these activities were executed according to a schedule laid out well in advance and practiced by the operations team in several tests and team exercises.

The first major set of activities in preparation for EDL was a "cold boot," which was scheduled for and successfully executed at three weeks prior to landing. This multi-day set of activities involved clearing the flash memory and performing cold processor resets on both spacecraft computers. The intent of the cold boot was to put the vehicle in a condition closer to how it is tested in ground testbeds for the EDL scenario and remove the potential for problems like file system management corruption. Rather than attempt to accumulate months of cruise data and history in ground tests, EDL tests typically start assuming a clean avionics and software state. The in-flight cold boot is to ensure that we "fly as you test" and create the same tested state in flight.

Following the cold boot activities, the Throttle Valve Assemblies (TVAs) on the Descent Stage's Mars Lander Engines (MLEs) were exercised over their range of motion to ensure than any stiction over the 8.5 month cruise was at acceptable levels. Following the successful completion of the TVA exercise, the Second Chance software was loaded. Though never utilized in flight, the intent of the Second Chance software was to attempt to successfully complete EDL should the primary computer fail or reset during EDL. At 16 days before landing, the software was loaded for later installation.

A final checkout of the Terminal Descent Sensor (TDS) radar was performed 13 days prior to landing. As expected, the tests indicated that the TDS was healthy and ready for EDL. Following this checkout, the rover batteries were charged to 100% in preparation for EDL.

Final Approach Activities

At 8 days prior to entry, the vehicle entered a period dubbed Final Approach. This period included many EDL preparation activities including trajectory correction maneuver (TCM) opportunities, EDL parameter update (EPU) opportunities, final checkouts, and the monitoring of onboard autonomous EDL preparatory actions. The TCM and EPU schedule and as flown/executed performance are discussed in later sections of this paper. This section discusses the other activities performed during this crucial period.

In preparation for the EDL autonomous FSW behavior starting at E-6 days, the flight team successfully executed checkouts to critical avionics responsible for gathering inertial measurement data and controlling the MLEs. These were followed by configuring the fault protection, communication system, and pyrotechnic system for use during EDL.

Two days from Entry, the autonomous EDL FSW powered on the MLE's catalyst bed heaters in preparation for use during EDL. Shortly thereafter, the Second Chance software was enabled on the backup computer.

In the last 24 hours, several autonomous activities were executed and monitored by the flight team. These included configuring the attitude estimation system in preparation for entry, starting EDL communication data flow, powering on EDL instrumentation, and warming up entry reaction control system thruster catalyst beds.

At approximately one hour from entry interface, the last team spacecraft health and status poll was performed, with the main goal of ensuring that the system was ready to execute EDL in its current configuration and allow the DSN to turn off the uplink transmitters.

GETTING TO THE RIGHT PLACE AT THE RIGHT TIME

Trajectory correction maneuvers are designed and implemented by the operations team to put the spacecraft back on target by firing thrusters. Being on target at the start of EDL reduces EDL risk: the guidance system has less delivery error to correct and the spacecraft is closer to the majority of the simulation-analyzed conditions. But TCMs are not risk free: there is operational risk in designing, implementing, testing, and sending of the TCM and spacecraft risk in opening propulsion latch valves and firing thrusters. This section details the risk balancing factors influencing the TCM decisions, the criteria used to make the decisions in flight, and the actual decisions made during Curiosity's approach to Mars.

Trajectory Correction Maneuvers and Risk Balancing

Trajectory correction maneuvers are fairly simple for a spinning spacecraft with MSL's design, when compared to maneuvers for other spacecraft. The cruise spacecraft configuration has balanced thrusters, making unintended delta-v's unlikely, and a simple proven maneuver execu-

tion sequence already demonstrated on the Mars Exploration Rovers and Mars Pathfinder missions. This does not, however, mean that TCMs can be considered risk free.

Human errors represent perhaps the largest risk in the TCM implementation and execution process. People are present throughout the TCM design and implementation process. Navigators determine the direction, magnitude, and recommended maneuver type. Attitude control engineers design the recommended maneuver and specify the thruster firings and durations in the TCM sequence. These steps are vulnerable to many different types of errors, including phasing. While any errors are likely to be caught in testing prior to execution on the spacecraft, there is definitely non-zero risk. As the TCM sequence moves from testing through the commanding processes other potential opportunities for errors exist. While these types of errors are very unlikely, the aggregate risk in the design, implementation, and commanding chain should be considered in TCM decisions.

TCM execution on the spacecraft itself also introduces non-zero risk. While earlier successfully executed TCMs during cruise mitigate some of this risk, each new TCM adds risk. Propulson latch valves and thrusters need to open and be closed again at the end of the maneuver, attitude control needs to function and reacquire celestial knowledge after the maneuver, and so on. Each maneuver is an additional opportunity for something to go wrong, e.g. a stuck open thruster valve or leaky latch valve. Typically, TCM execution risks are not a strong function of the TCM details: the risk is incurred by choosing to do a TCM and does not increase much with the size or design.

Even with tools like probabilistic risk assessment, it is difficult to numerically quantify the incremental risk of a TCM. Human error risks are soft risks to quantify; furthermore, not enough TCMs are executed to really get meaningful error statistics.

Choosing to forgo TCMs may result in accepting known interplanetary delivery errors. Accepting a different position and velocity than the targeted state means that the EDL system must accommodate the error. This drives risk into the EDL system for several reasons depending on the nature and size of the delivery error: the error may be too large for the entry guidance system to safely accommodate, the risk of recontact with the cruise stage may be significantly increased, EDL margins such as timeline and landing accuracy may be compromised, and/or the vehicle state may be outside the verified performance envelope.

Delivery errors can be broken down into three categories: 1) flight path angle errors, which are steep or shallow deviations from the targeted entry flight path angle; 2) crosstrack errors, which are errors perpendicular to the flight path angle/downtrack plane; and 3) time of flight errors, which result in the vehicle arriving at the entry point at a different time than targeted. These errors impact the EDL system in different ways.

Entry guidance sensitivities to delivery error are described in Mendeck and McGrew,[2] and summarized here with other EDL system sensitivities. Entering steeper than the targeted flight path angle causes the entry guidance system to have to fly a longer distance to the target than planned. This results in increased guidance prediction error and consumption of some guidance control authority. Larger guidance prediction error and overconsumption of guidance control authority leads to misbehavior of the guidance system, including landing accuracy degradation and altitude margin loss.

Entering shallower than the targeted flight path angle causes guidance to have to fly a shorter than planned distance to the target. This also results in increased guidance prediction error and consumption of some guidance control authority.

The targeted entry point is selected with a bias in a crosstrack direction to help the entry vehicle fly away from the probable path of the jettisoned cruise stage during entry. If delivery error results in a reduced crosstrack bias, the risk of recontact with the cruise stage or its components increases. Entering with a larger than targeted crosstrack bias does not have a strong negative impact on the EDL system unless the bias is extremely large.

While the EDL team has looked at performance away from the targeted entry conditions, the bulk of the analysis and verification and validation activities have dispersed entry conditions centered upon the entry target. This introduces "fly as you test" (or, perhaps more appropriately, "fly as you simulate") concerns if the actual spacecraft state is significantly different.

The vast majority of EDL trajectory Monte Carlo simulations were conducted with delivery error distributions centered on the entry target using Navigation "baseline" TCM-5 delivery errors - these errors essentially assume execution of the fifth trajectory correction maneuver, TCM-5, regardless of the size of errors pre-TCM-5. Choosing not to execute TCM-5 introduces an unplanned bias in the delivery error distributions. The EDL impacts resulting from uncorrected delivery errors are summarized in Table 1, below.

Table 1: Summary of EDL Sensitivities to Delivery Error

Error	EDL System Impact
Steep Flight Path Angle	• Reduced entry guidance control authority • Increased guidance prediction error • Risk of guidance saturation • Loss of altitude/timeline margin • Reduced testing experience
Shallow Flight Path Angle	• Reduced entry guidance control authority • Increased guidance prediction error • Reduced testing experience
Reduced Crosstrack Bias	• Increased cruise stage recontact risk • Reduced testing experience
Increased Crosstrack Bias	• Reduced testing experience • Slightly reduced entry guidance control authority

Trajectory Correction Maneuver decisions seek to balance the risks of doing or not doing a TCM and trade risk in cruise activities versus future risk during EDL. TCM risk is nearly invariant to the size of the delivery error to be corrected; by contrast, EDL risk from not doing a TCM does increase with delivery error. Qualitatively, once the delivery error is sufficiently large, the EDL risk outweighs the TCM execution risk. This is shown in notionally in Figure 1, below.

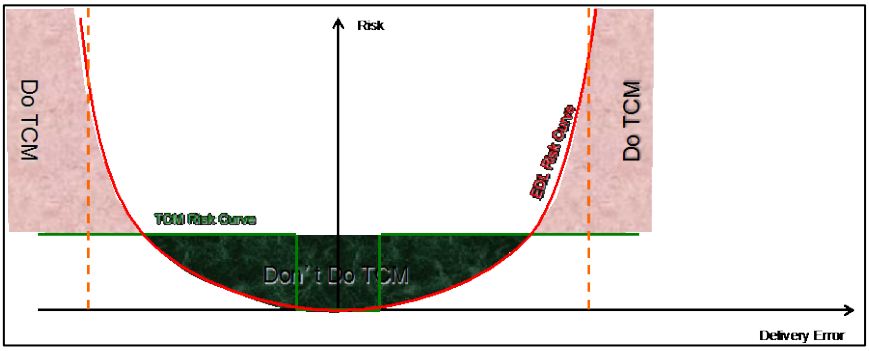

Figure 1: Qualitative Description of TCM Decision Risk Balancing

Trajectory Correction Maneuver Decision Criteria

While evaluating the TCM execution risk versus EDL risk necessarily involves some subjectivity, the MSL operations team constructed TCM decision criteria to guide the process. This allowed the cruise, EDL, and project management teams to pre-game the risk balancing considerations prior to the pressure-filled, time-critical period of Final Approach. Given that EDL risk is a function of delivery error, it makes sense to establish TCM criteria as a function of delivery error.

Like other interplanetary trajectories, MSL's trajectory and expected delivery state can be mapped on the B-Plane for ease of visualization. This allows quick assessment of the expected state and the type of error, be it flight path angle or crosstrack. Mapping TCM decision criteria on to the B-Plane enables rapid comparison of navigation estimate updates to the established criteria. Figure 2, below, shows the B-Plane and TCM-5 criteria for the MSL Gale Crater landing site and indicates the directions of each type of delivery error.

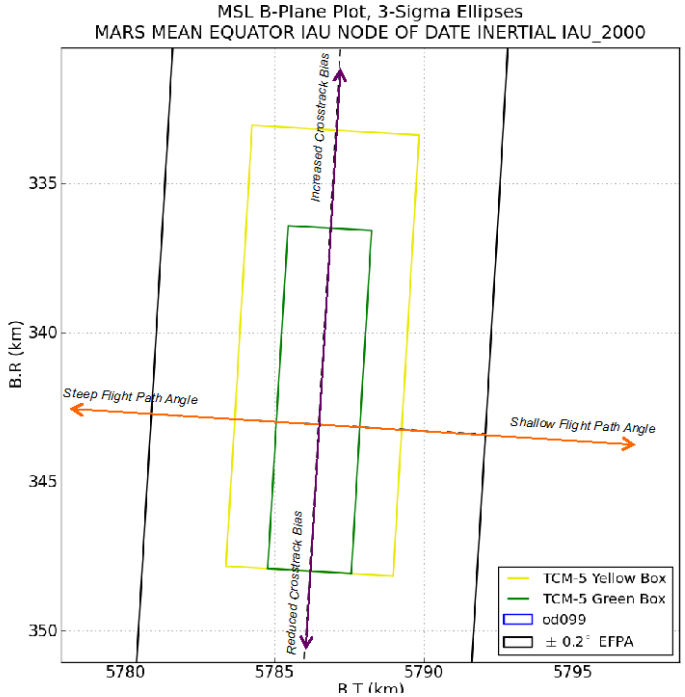

Figure 2: TCM Criteria Mapped on to B-Plane

The project established two TCM decision criteria, as shown by the yellow box and the green box in the above figure, with the target point shown at the intersection between the orange and purple lines. Acknowledging that the cruise operational risk of conducting a TCM depends on the conditions at the time of TCM, the operations team planned to use the green and yellow TCM criteria depending on the conditions at the TCM decision point. The green box, discussed below first, was intended for the nominal scenario, where the spacecraft is healthy and there are few, if any, other known risks to TCM execution.

In addition to the EDL system sensitivities identified and the desire to "fly as you test," the first time use of guided entry at Mars biased the EDL team towards TCM decision criteria that were very tight for the green box. Thus, the criteria are biased towards execution of planned TCMs (TCM-4 and TCM-5). A TCM threshold of 0.05° in flight path angle results in an expected delivery error distribution very similar to the navigation "baseline" TCM-5 delivery errors discussed previously and used for the bulk of EDL analyses.[3] This is shown in Figure 3, below.

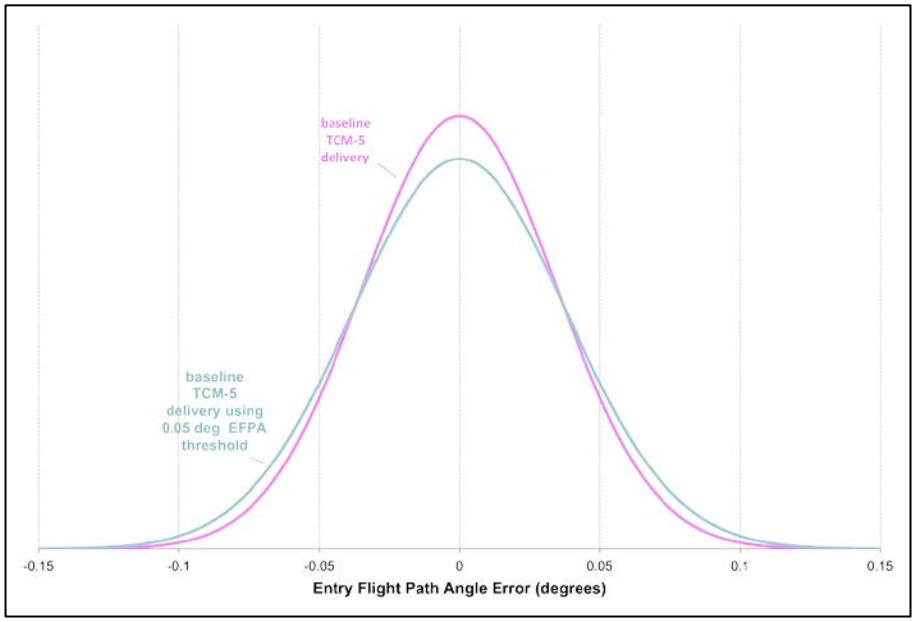

Figure 3: Comparison of Baseline Delivery Distribution vs. Distribution Assuming Green TCM Decision Criteria

As a result, a +/- 0.05° was used for the planned TCMs (TCM-4 and TCM-5/5X). If navigation solution centers were within this flight path angle error, no TCM was to be recommended. Within 0.05° of the flight path angle target, the delivery error was deemed to be sufficiently close to the entry target. Via the entry guidance sensitivities analyses performed, EDL system performance was known to be excellent and relatively unchanging within the 0.05° threshold.

As noted earlier, the EDL system has some sensitivity to crosstrack delivery errors, especially if the delivery error is in the direction that reduces the targeted crosstrack bias. Reducing the bias can significantly increase the risk of recontact with the cruise stage or its components. Since pre-flight navigation estimates of crosstrack delivery uncertainty were approximately +/- 4 km, the lower criteria for TCM execution was set at 5 km (vs. the nominal biased target of 8 km). Thus, in the worst case, there would likely still be at least 1 km of crosstrack bias in the correct direction.

The upper limit criteria for TCM execution was set at 12 km. The system is not very sensitive to larger than targeted crosstrack errors; 12 km was selected primarily by recognizing that the vast majority of trajectory simulation runs were within +/- 4 km (the expected cross track delivery error) of the targeted crosstrack bias of 8 km. Thus, cases with less than 12 km crosstrack bias would better meet the "fly as you test" intent. The crosstrack TCM criteria results in a "no-TCM" decision for crosstrack biases between 5 km and 12 km (-3 km to +4 km from the target point).

If the center of navigation solutions is within the green box, no TCM is recommended. The vehicle is close to the desired entry target and EDL risk is low. This is summarized in Figure 4, below.

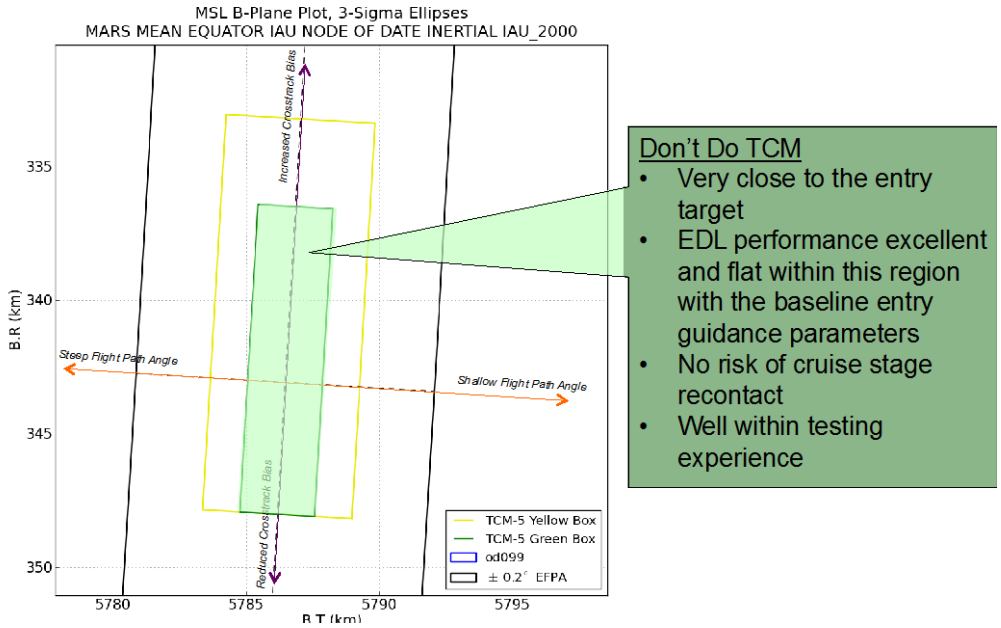

Figure 4: Green TCM Decision Criteria

The yellow TCM decision criteria for TCM-4 and TCM-5 were established to consider cases where the TCM build, test, and/or execution risk was different than expected. The operations team loosely termed these potential issues "spacecraft comfort" criteria and grouped them into subcategories of problems. These categories are summarized in Table 2, below.

Table 2: "Spacecraft Comfort" Considerations for TCM Decisions*

Category	Criteria	Comments
Spacecraft Health	• Has a safing event occurred since TCM-4? • Are there any undispositioned alarms? • Is the battery below 100% SOC and/or do we have negative power margin? • Is telecom performance marginal in any way? • Has a SEFI occurred on either RCE since TCM-4? • Is spacecraft configuration different than configuration for TCM-4?	Avoid potential aggravation of any spacecraft health issues
Spacecraft Anomaly	• Is a spacecraft anomaly in progress or anomaly response team active?	Avoid dilution of anomaly investigation resources
Navigation	• Is the OD unstable? • Are there any issues with DDOR window execution or data collection?	OD stability also part of TCM decision criteria and navigation assessment
TCM Capability	• Was TCM-4 anomalous? • Are any cruise propulsion thermal conditions out of bounds? • Is ACS performance unstable?	
Solar Activity	• Is there any major solar activity?	SSA, SEU vulnerability
Testbed status	• Is the testbed unable to support TCM testing?	Baseline TCM-5 build and testing plan includes TCM testing on the MSTB

* Acronyms: SOC: State of Charge; SEFI: Single Event Functional Interrupt; RCE: Rover Computer Element; OD: Orbit Determination; DDOR: Delta Differential One-way Range; SSA: Sun Sensor Assembly; SEU: Single Event Upset; ACS: Attitude Control System; SSA: Sun Sensor Assembly; MSTB: Mission System Testbed

The categories focus on the health of the spacecraft to successfully execute a TCM, the space environment status, the confidence in the knowledge of the spacecraft's navigation state, and the ability to successfully test a TCM sequence. Violating the various considerations would prompt the operations team to fall back to the yellow TCM criteria. This acknowledges increased TCM operational risk by loosening the TCM decision criteria, making TCM's less likely.

The yellow criteria doubles the allowable flight path angle error to +/- 0.1° from +/- 0.5°. Give pre-flight estimates of the interplanetary delivery capability, relaxing the flight path angle criteria reduced the likelihood of TCM-5 from approximately 50% to approximately 20%.

To reduce the likelihood of a TCM further, the yellow criteria extend the allowable crosstrack error an extra 3 km, to 15 km of crosstrack bias, or 7 km from the nominal target. As noted previously, the EDL system is relatively insensitive to growth in crosstrack bias; increasing the allowable error reduces the ability to "fly as you simulate," but does not substantively damage EDL performance until the error grows very large.

If the "spacecraft comfort" criteria are triggered and the center of navigation solutions is within the yellow box, no TCM is recommended. The vehicle is close to the desired entry target and EDL risk does not outweigh the TCM risk when considering the "spacecraft comfort" issues. This is summarized in Figure 5, below.

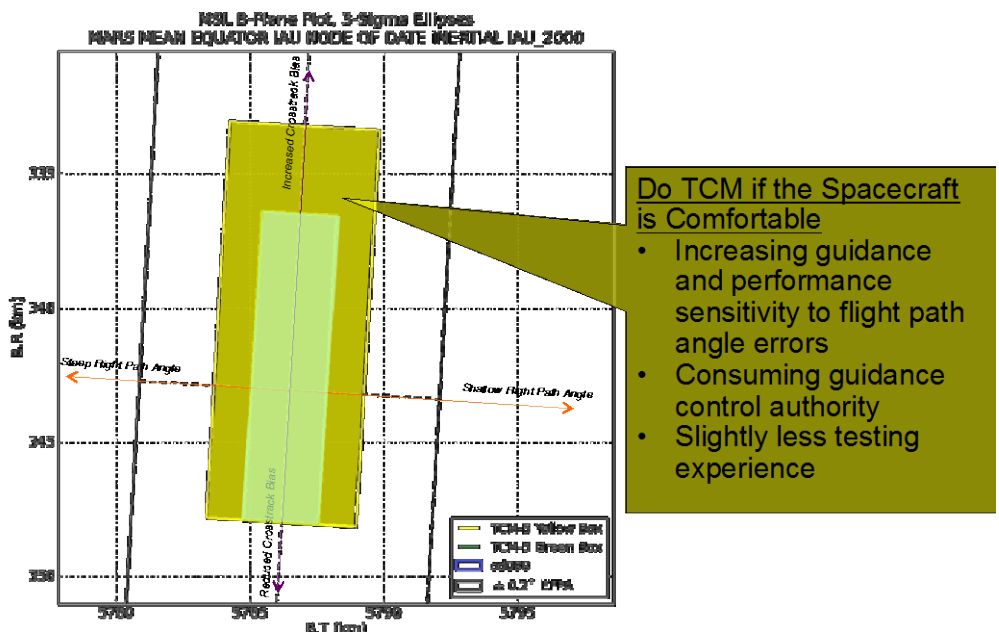

Figure 5: Yellow TCM Criteria

Contingency Trajectory Correction Maneuver Decision Criteria

In addition to planned nominal opportunities for executing a TCM, the project also carried a contingency opportunity for a maneuver, dubbed TCM-6. The maneuver opportunity was scheduled for execution nine hours before entry, if necessary. The intent of the maneuver was not to clean up residual delivery error; rather, it was planned as a last opportunity to "save the mission" should an error or issue occur resulting in the need for such a maneuver. The need might arise from a severe navigation error not detected until the final hours before entry or other anomaly that

might be pushing the vehicle off course. It would be executed to put the spacecraft back within the entry guidance's ability to correct any remaining delivery error and land safely and/or significantly reduce the risk of cruise stage recontact.

Because of the late execution of the maneuver and constrained timeline for building and testing the maneuver sequence, the contingency maneuver carries much higher operational risk when compared to a planned nominal TCM. As a result, the balance between TCM risk and EDL is significantly altered.

To reflect this new risk balance, the TCM decision criteria for the contingency TCM opportunity are greatly expanded, as shown in Figure 6, below, even when assuming the spacecraft is sufficiently healthy to perform a TCM. The expansion allows acceptance of additional EDL risk commensurate with the expected TCM risk. Accepting additional EDL risk entails trusting entry guidance to fly out the delivery error, while still correcting for other unknowns.

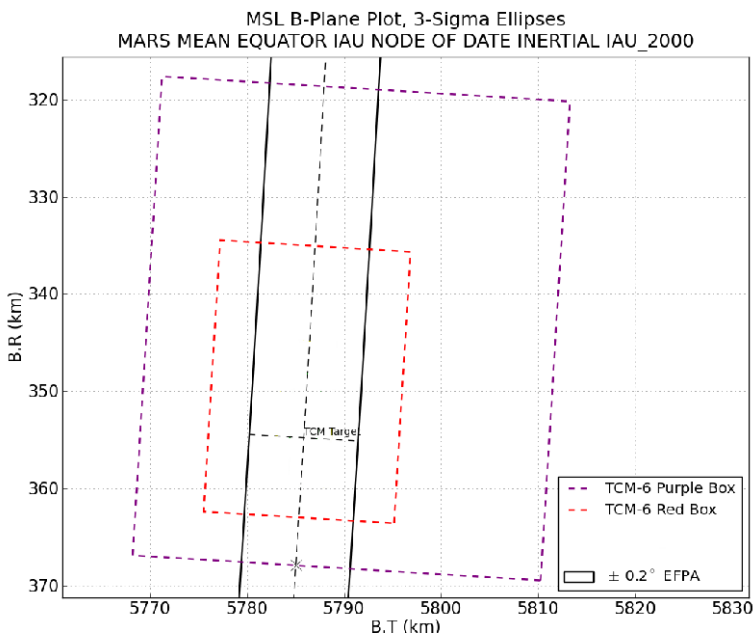

Figure 6: Contingency TCM Decision Criteria (for TCM-6)

Given the high perceived operational risk of performing a late contingency TCM, the flight path angle constraints were expanded to +/- 0.35°. As shown by entry guidance capability studies,[4] this causes some degradation of guidance performance and increased EDL failure rate. Additionally, crosstrack error thresholds are expanded to zero crosstrack bias (-8 km from the target) and 30 km crosstrack bias (+22 km from the target). This allows for increased risk of cruise stage recontact if near zero crosstrack bias or consumption of entry guidance control authority to correct the large crosstrack error.

The red box criteria assume a functioning spacecraft with no reservations with regards to executing a TCM. The operations team also developed a set of reference criteria to consider in cases where there may have been issues with executing a TCM. The purple box, dubbed the “10% death box” provided an estimate of the delivery errors that would result in an estimated EDL failure likelihood of 10%. This allowed the team to weigh the risk of executing a TCM under duress versus the potential impact to EDL.

Clearly, in the case of a contingency TCM, the project and operations team were ready to take on additional EDL risk to avoid executing a TCM. Fortunately, none of these contingency scenarios were necessary in flight.

As Flown Trajectory Correction Maneuver Decisions

The as flown interplanetary trajectory and accompanying navigation solutions, coupled with the TCM criteria developed, made for relatively simple TCM decisions. Very stable navigation solutions allowed for near certainty when evaluating expected delivery error versus the TCM criteria developed. Excellent TCM execution also simplified decisions, since there was little residual error to correct after each TCM.

TCM-4 was executed on July 28, 2012 at 10:00 PM Pacific Daylight Time (July 29, 2012, 5:00 AM UTC, Day of Year 211), as planned. TCM-4 corrected residual delivery error following the execution of TCM-3. The delivery error from TCM-3 resulted in an expected entry state well outside either then green or yellow TCM criteria boxes, as shown in Figure 7, below.

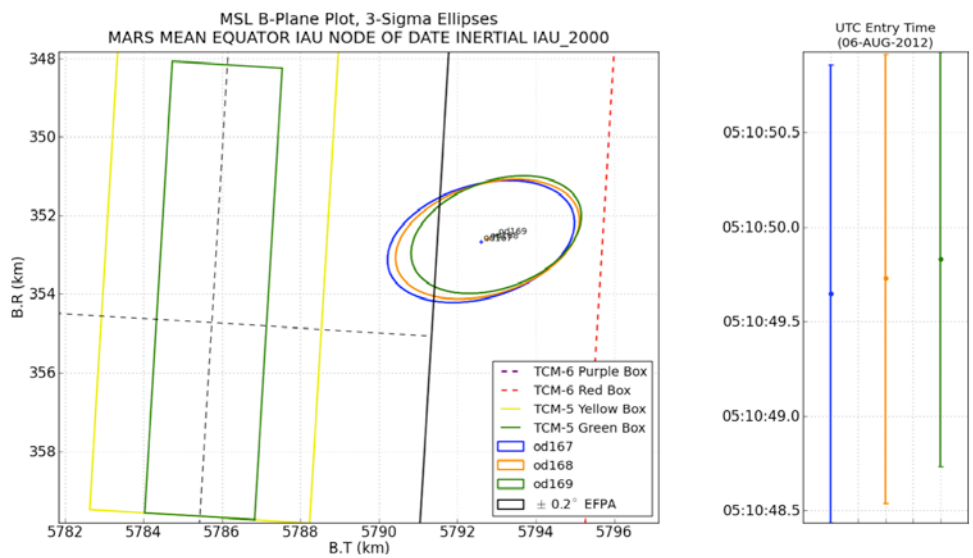

Figure 7: Navigated State at TCM-4 Decision Compared to TCM Criteria

With the spacecraft and ground systems healthy, a well rested operations team, stable navigation solutions, and favorable space weather conditions, the project team decided to build and execute TCM-4.

The next planned TCM, TCM-5, was scheduled for execution on August 3 at 10:00 PM PDT (August 4, 5:00 AM UTC, Day of Year 217); however, navigation solutions in the days following TCM-4 showed that the best estimates of the delivery error were exceptionally stable and well within the TCM green box criteria. This is shown in Figure 8, below.

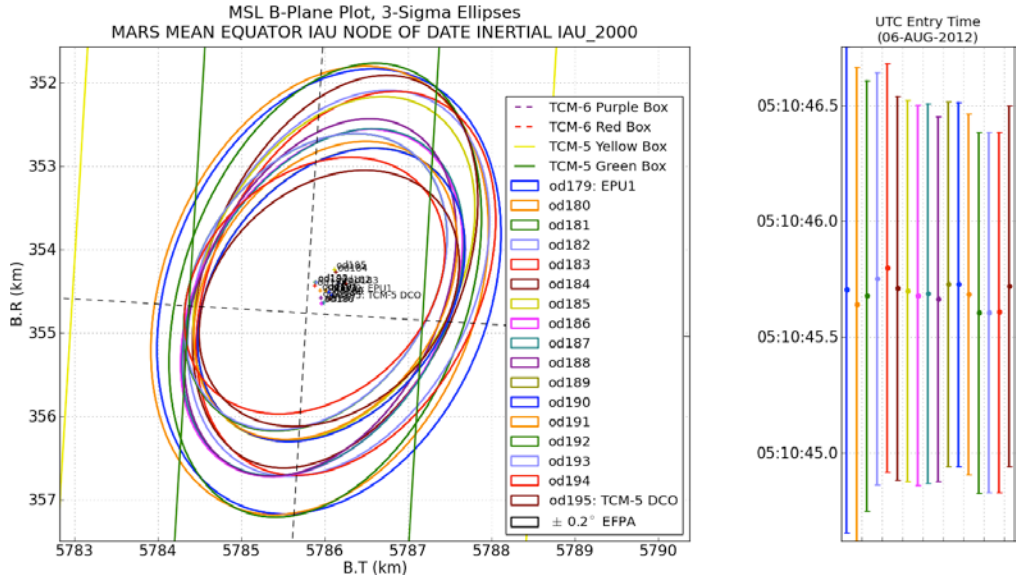

Figure 8: Navigated State Estimates at TCM-5 Decision Compared to TCM Criteria

While the spacecraft and team were prepared to build and test TCM-5, the closeness of the predicted delivery state to the target made it unnecessary and the TCM build and execution were cancelled.

The situation by the time of TCM-6, scheduled for execution nine hours before entry, was virtually unchanged, except that navigation uncertainties shrank as the spacecraft drew closer to Mars. As shown in Figure 9, below, the navigation solution remained stable and well within the green box. The shrinking uncertainties also showed the increased likelihood that the final delivery state would be well within the TCM criteria.

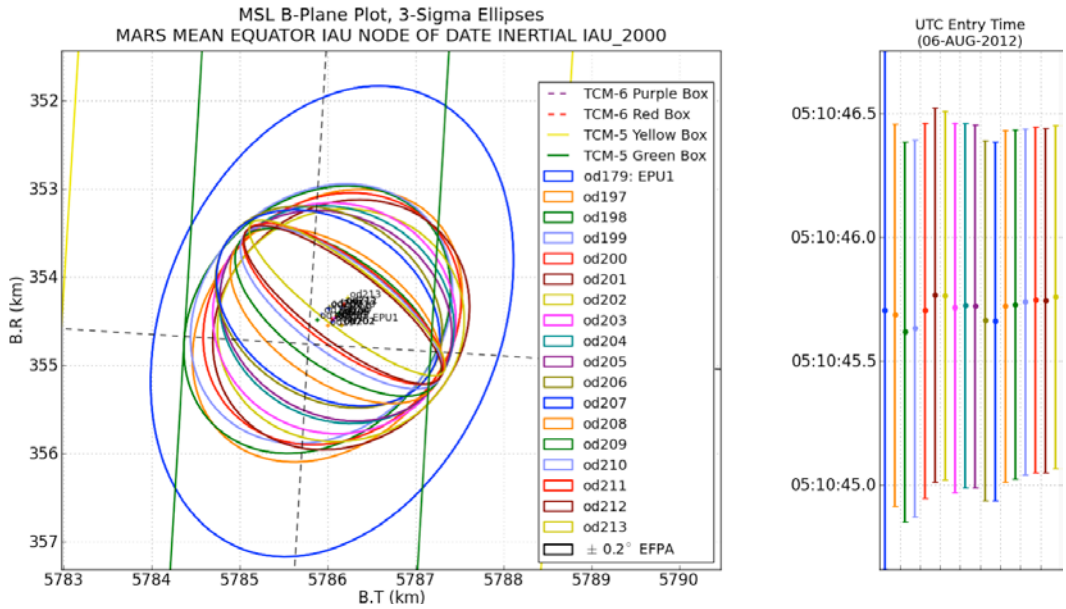

Figure 9: Navigated State Estimates at TCM-6 (Contingency) Decision

Given the lack of change in navigation solutions since TCM-4 and their shrinking uncertainty, the operations team had extremely high confidence that ultimate delivery state was fairly well known. The minimal error from the targeted entry point made the decision to cancel the contingency TCM opportunity nearly self-evident.

TELLING THE SPACECRAFT WHAT IT NEEDS TO KNOW

MSL's EDL system included the first time flight of an entry guidance system on Mars. Whereas previous expeditions to Mars were ballistic during entry and thus most sensitive to delivery errors, MSL is also quite sensitive to onboard knowledge. The guidance system can only perform as well as the quality of the onboard knowledge; if the guidance is misled by incorrect state knowledge, the spacecraft will end up in the wrong place. As a result, providing the spacecraft with EDL parameters capturing the best knowledge of the spacecraft state was a critical component of approach operations.

EDL Parameter Updates and Their Impact

Nominal EDL Parameter Updates involve changes to the time and navigated state (and associated parameters) at an event dubbed "Tzero," the instant that the EDL software assumes full control of the spacecraft. This occurs one minute after cruise stage separation and nine minutes before the spacecraft reaches the entry interface point.

Tzero time and navigated state updates directly impact entry guidance performance, as the vehicle will attempt to control its path based on its navigated state propagated from the Tzero time. The first estimate of the Tzero time and navigated state is provided to vehicle six days before entry.

Subsequent EDL parameter updates are needed when the best estimate of the Tzero time and/or navigated state differs significantly from the onboard Tzero time and/or navigated state. Because the entry guidance assumes the onboard spacecraft state knowledge is truth, any difference between the onboard knowledge and truth is reflected in system performance, especially in landing miss distance. This is shown in the example in Figure 10, below.

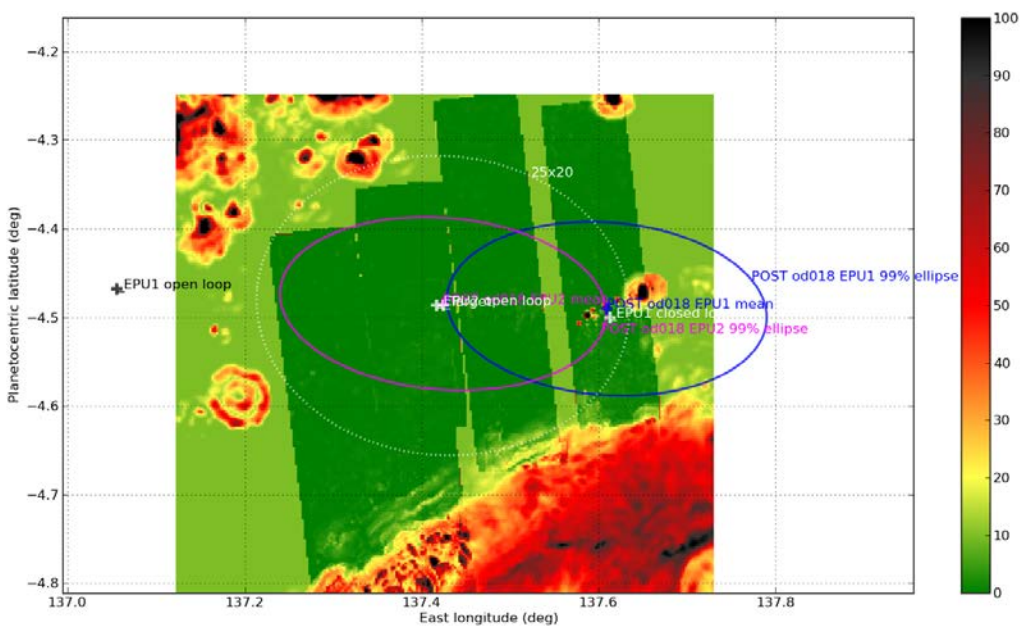

Figure 10: Example Describing Impact of EDL Parameter Updates

In this simulated non-flight example, the best estimate of the navigated state of the spacecraft changed between the time of EPU 1 and EPU 2. Since EPU 1 is onboard in this example, the simulated spacecraft has a knowledge error that it does not know about. As a result, it files down-range of the landing target and results in the blue "EPU1" ellipse, which is shifted and centered long of landing target. In this example, the onboard guidance thinks it has done its job and hit the target; it has not hit the actual desired target because the knowledge error. If EPU 2 is sent to the spacecraft instead in this example, the landing ellipse is re-centered back onto the landing target, as desired. Thus, updating the vehicle with the best knowledge restores nominal EDL performance.

Contingency EDL Parameter Updates

In off-nominal cases, EDL parameter updates may also involve updates to the Entry Guidance Parameter File (EGPF) that shapes the entry guidance performance by changing guidance gains. Contingency EGPFs were prepared for anomalously large delivery error situations, where the spacecraft would enter either steeper or shallower than the targeted entry flight path angle (EFPA) of -15.5°. Steeper or shallower flight path angles result in changes in the range between the entry point and the landing target. The contingency EGPFs help optimize the entry guidance for the different than expected range to target. Even though some control authority is consumed by accepting the flight path angle delivery error, the contingency EGPFs help ensure that the guidance behavior is orderly and that EDL performance is preserved, where possible.

The EDL team developed two contingency EGPFs, one for shallow flight path angle conditions and one for steep flight path angle conditions. As shown by the entry guidance design map sensitivity work, the baseline/default guidance parameters on board work well for a wide range of flight path angle errors, up to +/- 0.2°. The contingency EGPFs were developed for errors beyond +/- 0.2°. Even with the contingency EGPFs, performance degrades for increasingly large flight path angle errors, but the EGPFs help ensure that the degradation is gradual rather than steep.

EDL Parameter Update Decision Guidelines

The EDL parameter update process carries inherent human factors and spacecraft risks. Parameter update decisions weigh the process and operational risks versus the EDL risks assumed by not making an update.

EDL Parameter Update criteria were established by the team to set levels below which EDL parameter updates should not be considered. Below these thresholds, the candidate parameter updates would have negligible impact on EDL performance, including landing ellipse location; consequently, there is no reason to take on the operational risks associated with executing an update. It is also possible that below these thresholds, the changes in navigated state or Tzero time are just a result of stochastic motion of the navigation solution. Thus, the thresholds help the team not needlessly chase estimation noise. The minimum update threshold was set at 300 m in position, 1 m/s in velocity, and 50 ms in time.

Being above the minimum threshold for a parameter update does not, however, mean that a parameter update should be executed. Because of the landing ellipse shifting effect of onboard state updates, the motion of the expected landing location, both direction and magnitude, on the surface must be considered. Landing location influences two primary overall mission factors: 1) landing terrain hazards that may be encountered and 2) drive distance to the desired science targets. The direction and magnitude of the landing ellipse shift depends on the difference in navigated state when compared to the onboard state. Because there is not a single science target, not all ellipse shifts are created equal. As a result, the operations team prepared to evaluate all shifts, providing landing hazard values, maps of the accompanying landing uncertainty, and distances to

science targets for project management to make educated decisions. So, while there were guidelines for minimum thresholds for onboard state updates, there were no explicit "must do" parameter update triggers.

In off-nominal situations, the flight team was prepared to make more invasive parameter updates, including the uplink of pre-prepared Entry Guidance Parameter Files (EGPFs). As described earlier, EGPFs alter the entry guidance performance. EGPFs were prepared for anomalously large delivery error situations, where the spacecraft would enter either steeper or shallower than the targeted entry flight path angle (EFPA) of -15.5°.

Per the entry guidance sensitivity work performed, the steep and shallow contingency EGPFs designed were applicable for flight path angle errors >0.2°. Thus, the criteria were clear: if flight path angle error was larger than 0.2°, the appropriate steep or shallow EGPF was warranted, as shown in Figure 11, below.

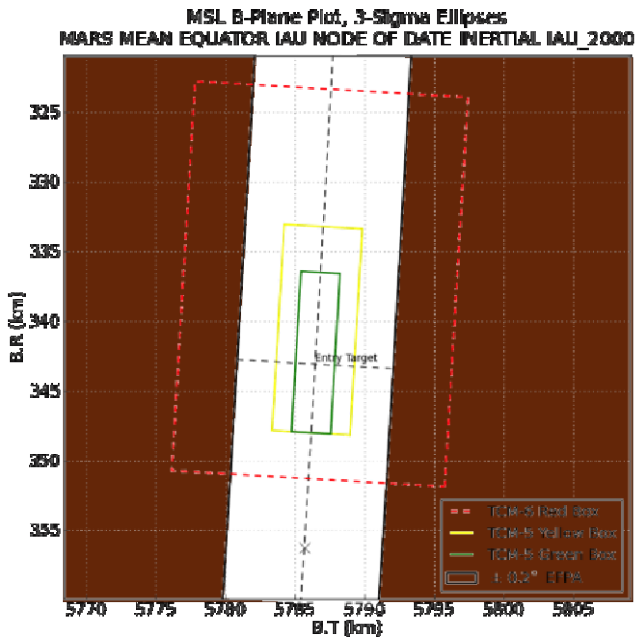

Figure 11: Entry Guidance Parameter File Update Criteria

As Flown EDL Parameter Updates

The excellent performance and stability of navigation solutions during Curiosity's approach to Mars simplified nearly all of the EDL parameter update decisions. While the team was prepared to turn around rapid parameter updates until two hours before landing, the lack of motion in the navigation solutions made this unnecessary. In fact, Curiosity successfully landed with an onboard state estimate sent to the spacecraft almost one week earlier. This section walks through the parameter update decision opportunities and the decision made.

The EDL timeline engine behavior requires an initial state, contained within an EDL parameter update to begin execution. The t-zero time specified by the EPU is used to set the execution times of all timepoints. Per the Approach, EDL unified timeline, EPU 1 was built and tested and ultimately sent to the spacecraft at 9:00 PM PDT on July 30 (4:00 AM July 31, DOY 213), as planned. At the time of the EPU 1 build, test, and execution cycle, there were no spacecraft, ground system, flight team, or space weather threats to the parameter update. Post TCM-4 orbit

determination indicated that TCM-4 execution was successful with a slight, ~4.6% overburn. All post TCM-4 navigation solutions were stable and consistent, as shown in Figure 12, below.

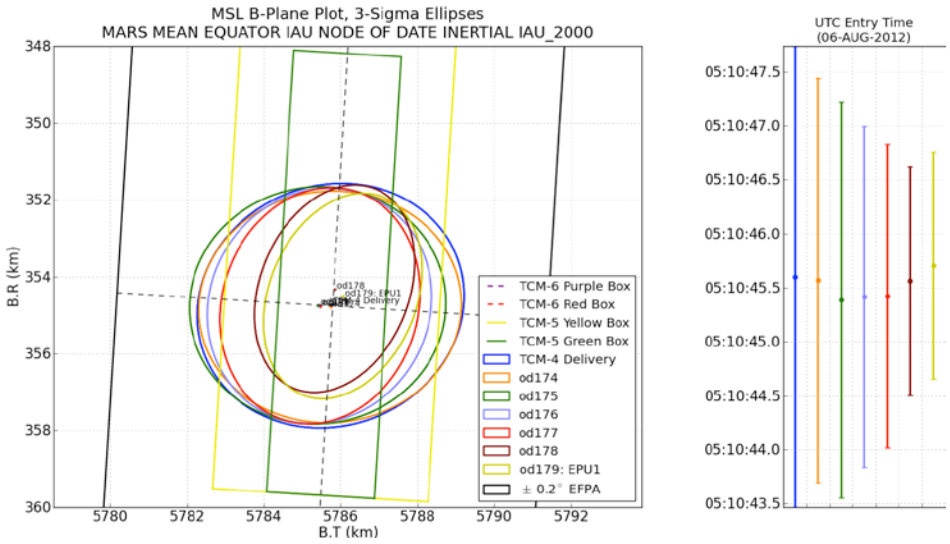

Figure 12: Navigation Solutions for EPU 1

Given the need to send an initial onboard state to the vehicle, EPU 1 was build, tested, and executed as planned.

Per the Approach/EDL unified timeline, EPU 2, planned for one day before entry, was the second opportunity to make EDL parameter updates. Given that the navigated state estimates were extremely stable, as shown in Figure 13, below, very similar to the onboard state from EPU 1, and TCM-5 was cancelled, the project elected to cancel EPU 2. This also allowed the operations team to focus on another issue discovered that day.

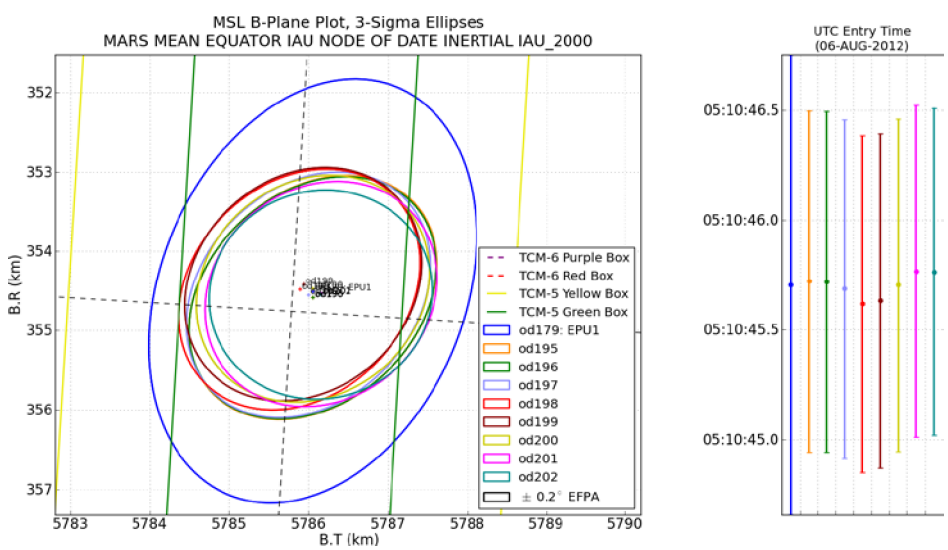

Figure 13: Navigation Solutions for EPU 2

The situation continued to be the same for EPU 3, if not improved by the shrinking navigation uncertainty as the spacecraft approached Mars, as shown in Figure 14, below. Compared to the estimate used to create EPU 1, the onboard state estimate, the EPU 3 navigation solution had changed by about 600 m on the ground (300 m at the B-plane, 500 m at entry). This difference was deemed to be more than just estimate noise, but the resulting landing ellipse bias was not very significant either in magnitude or in terrain hazard probability and did not substantively affect distances to science targets. As a result, while the operations team was ready to support the EPU 3 build, it was cancelled.

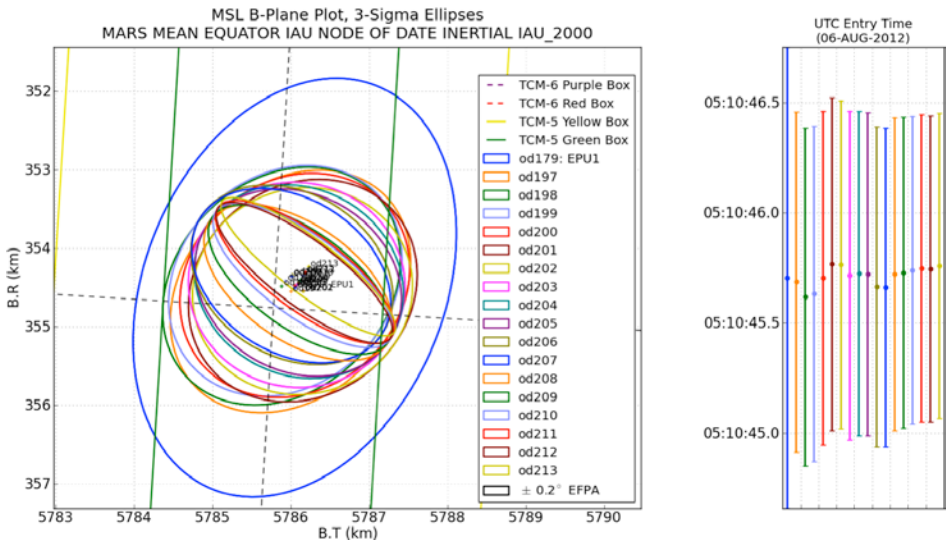

Figure 14: Navigation Solution for EPU 3

By the time of EPU 4, scheduled for uplink at Entry – 2 hours, the situation remained the same with even further reduction in navigation uncertainty. The best estimate solution remained stable and close to the onboard EPU 1 state, as shown in Figure 15, below. Distances from the onboard state were extremely similar to those at the EPU 3 decision point.

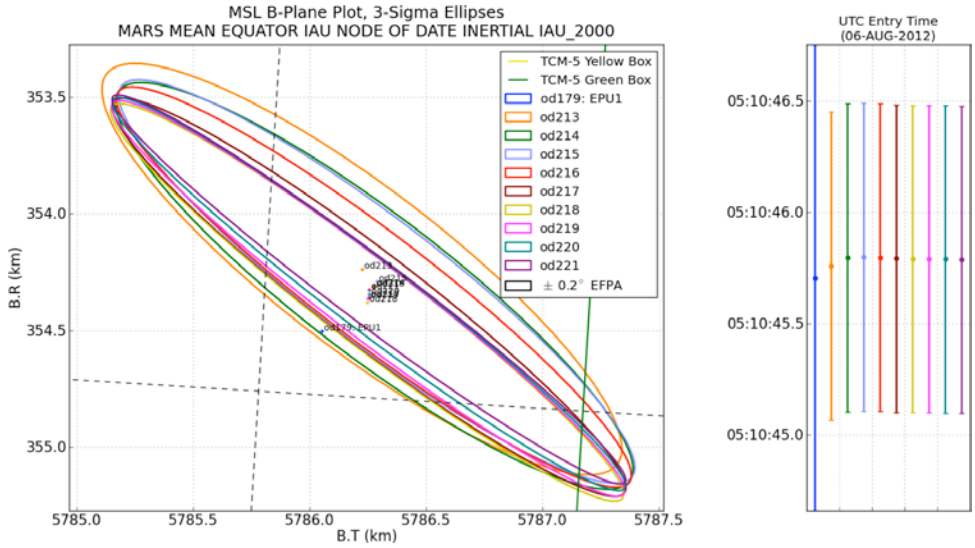

Figure 15: Navigation Solution for EPU 4

As a result of the minimal motion of the navigation solution and minimal impact on EDL performance, the EPU 4 opportunity was also cancelled. Ultimately, the navigation performance was so stable and accurate that Curiosity entered the atmosphere and landed successfully with a six-day-old estimate of the delivery state. Per the post-flight reconstruction performed by the navigation team,[5] that six-day-old estimate was only a few hundred meters off from the reconstructed delivery state.

CONCLUSION

In the weeks and days leading up to MSL successful landing on August 5, 2012, the MSL flight team successfully executed and monitored a number of critical activities that prepared the spacecraft for EDL. Planned activities included software and parameter updates, hardware checkouts, and opportunities for trajectory correction maneuvers. Success in these activities was predicated by preparation and the establishment of decision criteria for TCMs and parameter updates. With excellent TCM execution and navigation solution performance and stability, decisions made by the flight team were relatively simple and placed the spacecraft in nearly the optimal physical and software state, thus contributing to a near flawless landing.

ACKNOWLEDGMENTS

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1] Tomas Martin-Mur, Gerhard Kruizinga, and Mau Wong, "Mars Science Laboratory Interplanetary Navigation Performance," AIAA/AAS Space Flight Mechanics Meeting, AAS 13-232, February 2013.

[2] Gavin Mendeck and Lynn McGrew, "Post-Flight EDL Entry Guidance Performance of the 2011 Mars Science Laboratory Mission," AIAA/AAS Space Flight Mechanics Meeting, AAS 13-419, February 2013.

[3] Tomas Martin-Mur, Gerhard Kruizinga, and Mau Wong, "Mars Science Laboratory Interplanetary Navigation Performance," AIAA/AAS Space Flight Mechanics Meeting, AAS 13-232, February 2013.

[4] Gavin Mendeck and Lynn McGrew, "Post-Flight EDL Entry Guidance Performance of the 2011 Mars Science Laboratory Mission," AIAA/AAS Space Flight Mechanics Meeting, AAS 13-419, February 2013.

[5] Fernando Abilleira, "2011 Mars Science Laboratory Trajectory Reconstruction and Performance from Launch Through Landing," AIAA/AAS Space Flight Mechanics Meeting, AAS 13-113, February 2013.