AAS 13-123

MARS SCIENCE LABORATORY ENTRY, DESCNET AND LANDING SYSTEM DEVELOPMENT CHALLENGES AND PRELIMINARY FLIGHT PERFORMANCE

Adam D. Steltzner,[1] A. Miguel San Martin,[2] and Tommaso P. Rivellini,[3]

The Mars Science Laboratory project recently landed the Curiosity rover on the surface of Mars. With the success of the landing system, the performance envelope of entry, descent, and landing capabilities has been extended over the previous state of the art. This paper will present an overview of the MSL entry, descent, and landing system, a discussion of a subset of its development challenges, and include a discussion of preliminary results of the flight reconstruction effort.

INTRODCUTION

On August 5th at 10:31 PM PDT, signals were received on Earth confirming the successful landing of the Mars Science Laboratory's Curiosity rover on the planet Mars. With this signal came the first proof of success of the ambitious new entry, descent, and landing system that had taken Curiosity to the surface. Curiosity weighs 900 kg and required the development of a new landing system. The MSL entry, descent, and landing system is novel, using old technologies in a new application and pioneering new technologies as well. Development was paced originally for a 2009 launch to Mars, but development issues across the spacecraft resulted in a slip to the 2011 launch window. With this slip, the development lifecycle was a total of eight years, from the inception of the entry, descent and landing (EDL) system architecture in 2004 to its final test at Mars in 2012.

The EDL system for Curiosity possesses many technological advancements, mostly in the form of novel uses of existing technologies (ex: Ka-band pulse pair Doppler multi beam radar for landing) but also some notable direct innovations (ex: the sky crane maneuver). Through out the development of this EDL system, architecture and design choices were made in an attempt to yield a robust system design that would be resilient to the ever present threats of spacecraft development such as shortfalls in piece part or subsystem performance, errors and lapses the understanding of the design challenges and requirements, and error in understanding the interconnectivity of the individual system elements. Through the development process, the MSL EDL architecture survived intact that is no major architectural redesign occurred (ex: adopting a different touchdown technique other than

[1] EDL Phase Lead, Systems Engineering Division, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.
[2] EDL Deputy Phase Lead and GNC Chief Engineer, Autonomous Systems Division, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.
[3] Mechanical Systems Lead, Mechanical Systems Division, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, CA 91109.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

Sky crane, or a different number of parachutes, or solid rocket motors as opposed to throttlable mono-propellant, etc.) That the over riding architecture survived, may be the result of some of the architecting and design practices deployed by the team. That said, in many instances, there were development challenges and struggles that resulted in changes to system elements (ex: heat shield thermal protection material selection) and additions in EDL complexity (ex: "dot-product trigger").

The object of this paper is to explore some of those development challenges and to give the reader insight into what was entailed in the development process for MSL. There is a brief section that attempts to describe the system margin practice used in the development of the MSL EDL architecture and give an example of how that may have protected the architecture during development. By design this is a narrative, and because of space constraints, a selected discussion of a subset of the challenges faced by the MSL EDL developers.

EDL OVERVIEW

The entry, descent, and landing process at Mars, or any other planet with an atmosphere, is a process of managed energy dissipation. The central challenge is to remove the kinetic energy of planetary arrival and prepare for a managed impact with the surface. At Mars, where the atmosphere is roughly about 1/100th that of Earth's, maintaining altitude and sufficient timeline to prepare for the surface contact is paramount. The job of dissipating arrival energy and preparing for managed impact is spread throughout the different segments of EDL.

Final Approach

During the Final Approach segment, at the end of the eight and a half month cruise to Mars, preparations for entry, descent, and landing became the focus for the spacecraft and flight team. These preparations center 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. These objectives result in activities such as hardware checkouts and parameter loads, trajectory correction maneuvers to put the vehicle on course for the targeted entry conditions, and navigated state updates to seed the entry system with accurate knowledge for use in guided entry. These activities require the cruise and EDL teams to work in concert to balance operational risks to the spacecraft in cruise with EDL risks associated with putting the vehicle in the best possible physical and software state to execute EDL successfully[1].

Exoatmospheric Segment

The exoatmospheric segment of EDL spans the period from cruise stage separation through arrival at the entry interface point, 3522.2 km from the center of Mars. During this period, the vehicle completes its transition from cruise to EDL configuration with the separation of the cruise stage at Entry – 10 minutes. At Entry-9 minutes, full EDL control of the spacecraft becomes active and the vehicle begins using reaction control thrusters to stop the residual spin rate from cruise and turn to the desire entry attitude. Two externally mounted tungsten masses, used to balance out the center of gravity during the spin stabilized cruise portion of the mission, are jettisoned, thus leaving the spacecraft with an offset center of gravity to generate lift during entry.

Entry[2]

The entry segment of EDL spans approximately 4.5 minutes and begins at entry interface and ends at supersonic parachute deploy. At entry interface, the vehicle is traveling at approximately 5.8 km/s at a targeted flight path angle of -15.5°. At this point, the entry propulsion system transitions from blow-down pressurization to a higher regulated pressure to provide sufficient thrust for entry control and the later powered flight segment.[3] When the spacecraft senses drag acceleration of 0.2 Earth g's, entry guidance begins. Using the lift provided by the center of gravity offset, the guided entry system is able to correct for errors in initial delivery state, atmosphere conditions, and aerodynamic uncertainties while also significantly increasing altitude capability by modulating the direction of lift. If lower than anticipated atmospheric drag is encountered, either because of

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

atmospheric density deviation from the expected atmosphere, or aerodynamic drag deficit from expectation, the vehicle flies to a lower altitude than the reference trajectory thereby increasing drag. If higher than expected drag is encountered, the vehicle lifts itself in the atmosphere. Nominally, the reference trajectory has the lift vector being stored in a predominately sideway direction, resulting in three bank reversals (think s-turns) being performed to keep the vehicle within the cross track error boundaries. The altitude and precision capabilities enabled by guided entry allowed the selection of never before considered landing sites on the Martian surface, sites that were far too cramped or too high to be considered previously.

This behavior used repurposed Apollo guidance algorithms to form a simple and robust guidance approach that allowed the reduction of landing miss distance from 70+ km down to under 10 km[4]. The guidance scheme is broken into two different regimes. First, range control from entry interface down to approximately 1100 m/s, where the down track drag acceleration is matched to attempt to be over the desired target when the parachute deployment conditions are met. During this portion of flight, the vehicle experiences peak heating and peak deceleration.

The second part of guided entry is heading alignment; in which cross track errors that may have accumulated due to the bank reversals and other disturbances during range control are corrected. The spacecraft flies ~100 km downrange during heading alignment at a near constant altitude between 10 and 15 km. At the end of heading alignment, the vehicle transitions into the "Straighten Up and Fly Right" (SUFR) maneuver, which rolls the spacecraft while separating balance masses to change the center of gravity of the vehicle in preparation for parachute deploy.

Figure 1. Cruise Balance Masses Being Ejected from the Entry Vehicle Just Prior to Atmospheric Entry (Left); Entry Balance Masses Being Ejected Just Prior to Parachute Deployment (Right).

The six tungsten masses used to create the cg offset needed for guided entry are ejected one at a time so as to reduce the spacecraft's angle of attack down to zero (Figure 1) degrees. This is done to avoid large swinging pendulum-like motion and subsequent high angular rates when the parachute opens.

Parachute Descent

Once the capsule has slowed its speed to approximately 400 m/s (Mach 1.7), a parachute is deployed. This parachute represents the largest extraterrestrial parachute ever flown. It is also the largest supersonic parachute ever used. The Viking design disk-gap-band parachute, developed by Pioneer Aerospace, slows the capsule, initially at 6 Earth g's. At a speed around 150 m/s (Mach 0.7) the heat shield of the capsule is released and allowed to fall away. Heat shield separation is very important, since altitude knowledge error can be as large as +/- 3km at this point; this is not close to accurate enough to allow landing without directly sensing the ground. Once this has occurred, a JPL built six-beam Ka-band pulse-pair Doppler radar is used to look for the surface and perform range and velocimetry measurements. These measurements are used to determine the next transition in vehicle flight configuration. At and altitude of 3000 m, the Mars Lander Engines (MLEs) on the Descent Stage are primed in preparation for backshell separation and Powered Flight.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

Figure 2. Separation from the Backshell and Initiation of Powered Descent

Powered Descent[5]

At a height of about 1600 m above the Martian surface while travelling around 85 m/s almost vertically downward, the spacecraft separates from the backshell and parachute and begins powered descent (Figure 2). The triggering of the backshell separation event is based on the vehicle's velocity and altitude and takes into account the amount of time required to complete the backhsell separation events. Using 8 throttle-able hydrazine engines, with a resurrected Viking catalyst bed design by Aerojet Corp, the spacecraft diverts 300 m perpendicular to the plane of the entry trajectory to avoid the parachute/backshell and flies a profile to cancel the horizontal velocity and bring the vertical velocity to a specified value nominally ~240 m above the surface.

Once the vehicle has been delivered to a steady vertical descent flight condition at ~240 m above the eventual touchdown site, the vehicle nominally flies through the "constant velocity accordion," descending vertically for approximately 100 m. During the constant velocity accordion, nominally 100 m in size, a speed of 32 m/s is maintained until about 140 m above the surface. The constant velocity accordion is to protect against altitude estimate errors by the vehicle during parachute descent where the radar beams may not be pointing at the final touchdown site and therefore not measuring the altitude to the final touchdown site. At ~140 m above the final touchdown site, the vehicle then decelerates at about 8 m/s^2 for about 115 m to a velocity condition of 0.75 m/s. The vehicle then throttles down four of the eight engines and prepares for the sky crane maneuver

Skycrane

At approximately 21 m above the landing site, the spacecraft begins the "sky crane" maneuver. The maneuver begins when the rover is separated from the descent stage, which contains the propulsion system and GNC sensors (IMU and radar), and is lowered below on a 7.5 m bridle system. During the deployment of the rover, the descent stage continues its 0.75 m/s trajectory vertically downwards toward the surface. During deployment, the mobility system is also released from its stowed condition. Once fully deployed at the end of the bridles, with the mobility released and ready to touchdown, the rover and descent stage continue downward at 0.75 m/s until contact with the Martian surface occurs. There is another portion of constant vertical flight at this stage, called the "ready for touchdown" accordion nominally of 5 m in length. This accordion exists, again, to compensate for possible range errors by the radar used in the planning of the sky crane timing. After

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

the rover has touched down, the descent stage continues downward until persistent reduced throttle settings have been confirmed using a 1.5 second rolling window. At that time, touchdown is declared and flyaway begins.

Flyaway

After touchdown is declared, descent stage motion is stopped and the bridles are cut free. Control is transferred to the Descent Stage. The Flyaway controller on the Descent Stage executes an attitude and throttle profiles to fly it on a preprogrammed trajectory to a distance of at least 150. This leaves the rover wheels down on the surface of Mars and ready to begin commissioning.

THERMAL PROTECTION SYSTEM MATERIAL

The initial plans for the MSL EDL system assumed the design would use SLA-561V for the thermal protection material on the forward facing heat shield. This material had been used on the Viking, Mars Pathfinder, Mars Exploration Rovers, and Phoenix missions with great success. The SLA material, like many lower density TPS materials, is a mixture of many component materials; cork, silicon, phenolic and a host of other materials make up the slurry that is cured on the heat shield and forms a protective coating with a slightly spongy texture and a density between that of cork and rubber.

MSL has a heat shield that is sufficiently large that the subsonic boundary layer, the source of a great fraction of the heating during entry, was thought to possibly be able to trip to turbulent flow prior to peak heating. This was a concern since turbulent boundary layers are more efficient at transferring heat to the TPS material. Turbulence would augment the previously anticipated heating rates. Further, turbulent heating augmentation is not well understood, for this application and there was substantial debate amongst the team as to what level of heating we should expect.

The previous missions were though to have experienced the following heating rates are shown in Table 1.

Mission	**Peak Heating (W/cm^2)**
Viking (1/2)	26 (predicted)
Mars Pathfinder	108(predicted)/100(as flown)
Mars Exploration Rovers	44(predicted)
Phoenix	47(predicted)
Mars Science Laboratory	225(design)/ 120(predicted)/105*(as-flown)

*Initial estimate

Table 1, Peaking Heating Rate Data for Mars Missions

As development proceeded, three phenomena converged that eventually overran the plans to use SLA. First, the assumed heating environment grew as more and more concern was voiced regarding the possible heating that we might experience due to the fully turbulent boundary layer flow that was expected. Second, uncertainty as to how much heating augmentation would occur when considerations of roughness considered also drove the design heating rate upwards. Third, testing of SLA at relatively high heating rates and pressures revealed that complex material response behaviors, including melt flow, vaporization and sublimation, made it very difficult to understand the margins present, because earth testing can not exactly duplicate the Mars flight conditions.

Once the SLA material was pushed up into a thermal region for which we could not model the material response with confidence, our knowledge of the margins we would hold in flight was undercut. No testing in arc jets facilities here on Earth can exactly duplicate the conditions of entry on Mars. It is confidence in the TPS material response to the range of possible at Mars conditions

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

that serves as margin for flight. Once the SLA had been pushed up into a region where its material response was very complicated, we could no longer have confidence in our margins assessment.

The last blow to the use of SLA came in a set of tests in which the material seemed to vaporize very rapidly. This response was not predicted and although the conditions of the tests had some uncertainty to them, they were likely within the possible flight envelope, or at least not far from it. Under these circumstances, the project elected to make a very late change in the heat shield TPS material, choosing to go forward with Phenolic Impregnated Carbon Ablator (PICA), a much more simple material system, that had been shown to handle much higher heating conditions.

This new material, PICA, was at the time the leading candidate for use on the Orion heat shield for return missions from the International Space Station. Further, this material had been used on the Stardust mission, which had experienced very high heating rates on Earth return.

The cost of PICA came in development delay, dollars, and mass. PICA is a great high heating rate ablator, but the density is higher than SLA and the thermal conductivity is also greater. In our application, PICA was a very mass inefficient choice. At our low heating rates, by PICA standards, our TPS would be sized by thermal conductivity considerations and the resulting heat shield would have over 100 kg of additional mass.

Thanks to the efforts and professionalism of the folks at Fiber Materials Incorporated (FMI), NASA Ames Research Center, Lockheed Martin Space Systems, the aeroshell provider and integrator of the TPS material, and the staff at JPL, the heat shield with new TPS tested and installed was ready for the 2009 launch.

In hindsight, the as flown heating rates and aero thermal environment as detailed in [Edquist paper] indicated that SLA may have been able to be used for the MSL flight. At the very least, a reduction in the required thickness of the PICA or SLA TPS materials may be achievable with the as flown conditions, if there were to be a re-flight of a similar design in the future.

BRIDLE, UMBILICAL AND DESCENT RATE LIMITER (BUD) DEVELOPMENT

At the heart of the sky crane maneuver is a winch-like system that is used to lower the rover below the decent stage. The BUD subsystem consists of an electromagnetic brake connected to a tapered spool, three nylon suspension bridles, a confluence point pulley, an umbilical deployment device and three rover mounted bridle exit guides.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

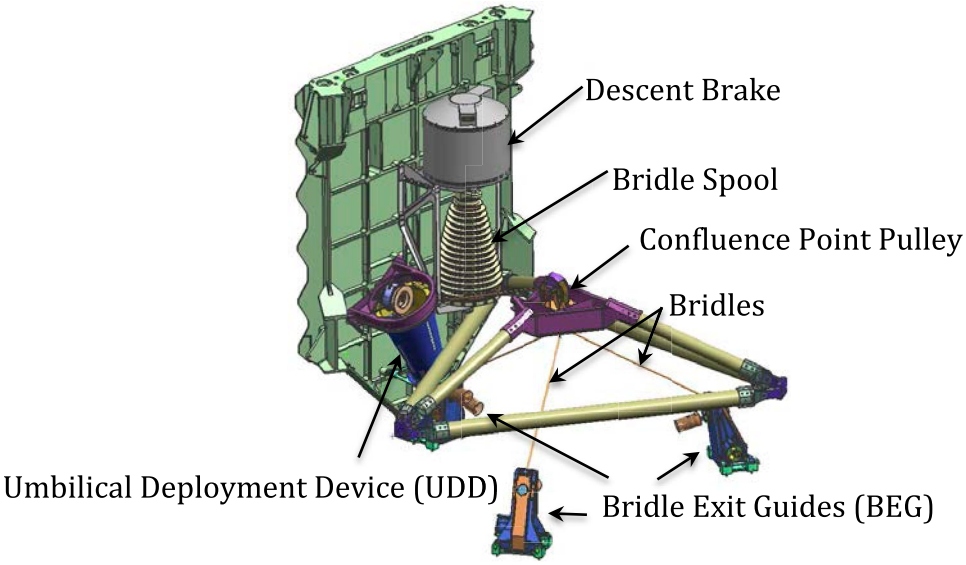

Figure 3: BUD assembly components shown inside the descent stage.

The electromagnetic brake and spool taper profile were tuned to allow the rover to descend 7.5 meters below the descent stage in approximately 5.5 seconds using gravity as the only motive force. The descent rate profile was tuned to manage loads imparted to the descent stage, the rover chassis and most importantly the rover mobility system.

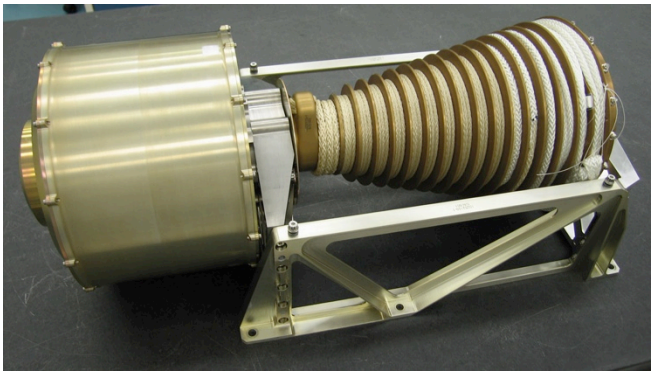

Figure 4: Descent brake assembly flight unit.

The descent brake gearbox had modest margins on the maximum allowable speed in some of the bearings in the gearbox as well as gear loading under momentary impulse. As described below, the mobility deployment events that take place during rover descent on the descent brake cause very high momentary loads into the descent brake. Since both the descent brake and the mobility deployment hardware had structural margins that were either slightly negative or very small, a highly detailed ADAMS simulation was used to refine the mobility deployment timing to finesse the descent brake margins and mobility margins. This was not the desired position to be in, but nonetheless the "pencil sharpening" effort bore fruit and resulted in a final selection of a mobility deployment sequence, which was acceptable. Figure 5 illustrates the final predicted descent brake torque-speed profiles generated from a 2000 case Monte-Carlo simulation. The curves excursion outside the red dashed boundary indicates the region of negative margin.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

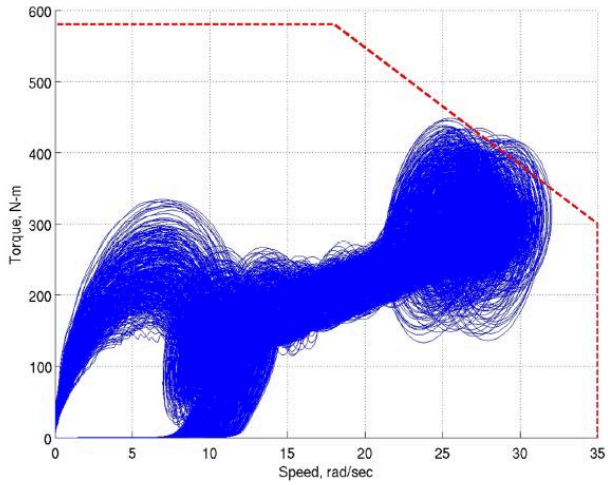

Figure 5. Torque-Speed plots of 2000 simulations. Each blue curve represents a rover descent on the descent brake. The non-linear "loop" in the curves reflects the transients caused by rover mobility deployment. The red dashed lines represent the performance limits of the mechanism.

An electrical umbilical is used to connect the electronics on the descent sage to the electronics on the rover. The deployment of the electrical umbilical was performed using an Umbilical Deployment Device (UDD). The UDD also contains a retraction spring device with a leash to the umbilical. The retraction device is used to maintain tension on the umbilical during the touchdown even to prevent slacking and recontact with the top deck. After the umbilical is severed, the retraction spring also pulls the umbilical up and away from the rover, again to avoid recontact with the rover.

Figure 6: Umbilical deployment device flight unit.

There was an extensive test program to develop and demonstrate the reliability of all of the BUD devices. Early testing focused on cable management during the deployment as well as loads characterization. As hardware matured the test programed evolved into characterization of descent speed, elastic modes and post cable cut cable dynamics. A highly detailed analytical model of the descent rate limiter and suspension cables was generated. This model was used for all of the subsystem and system level simulations of the sky crane maneuver.

Preflight statistical analysis predicted a mean descent time of 5.3 seconds. Telemetry from the mission indicates that the complete deployment took 5.1 seconds, indicating that the descent brake performed exceptionally well. The imagery of the top deck of the rover after landing shows nominal configuration of the BEGs indicating that they operated as expected.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

Figure 7: Rear bridle exit guide image taken from Curiosity imager. Image indicates nominal operation of the device.

MOBILITY DEPLOYMENT

Due to volume constraints inside the aero shell the rover's mobility system was required to be stowed during launch and EDL, but needed to be fully deployed prior to touchdown. Unlike any other rover mission the MSL sky crane landing system used the rover's mobility system as the landing gear. The design loads for the mobility system components were driven by a combination of touchdown loads as well as surface operations loads. In a few cases some components had deployment loads as their driving load case.

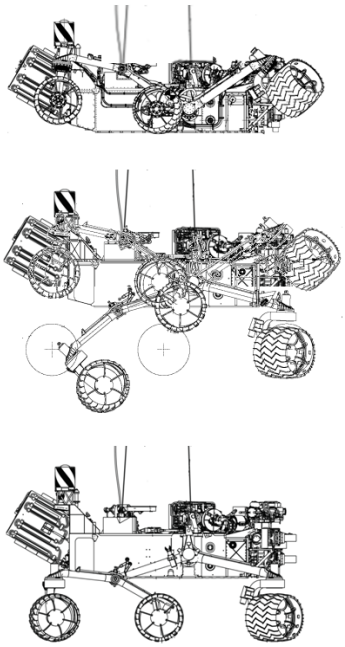

Figure 8: Rover mobility deployment sequence.

In order to maximize reliability and minimize complexity, the rover mobility is passively deployed and does not make use of any rate limiters or shock absorbers. This brute force approach results in a more challenging analysis and verification effort but pays off in reliability.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

The decision to deploy mobility during rover descent on the BUD was made early in the project when the timeline and fuel margins were exceptionally tight. As the project matured the statistical timeline margins relaxed a little, but never enough to warrant the added risk of moving mobility deployment to take place after BUD descent was complete. As it turns out, the as-flown mission profile had ample time to stage these activities sequentially but there was no way to take advantage of this in real time.

The mobility deployment V&V program was focused on verifying hardware functionality from the component level all the way up to the system level, and on verifying the math models used in all of the sky crane and touchdown simulations. Component tests generated performance data that was fed into the master ADAMS simulation. Multiple deployments of the mobility system were done under controlled ambient conditions as well as under simulated Mars ambient conditions. The highest level of fidelity test was called the Full Motion Drop Test (FMDT). The FMDT used a high fidelity descent stage and BUD subsystem along with a very high fidelity rover model. Both the rover and BUD hardware were essentially identical to the as flown hardware. The FMDT was a full-scale separation and mobility deployment of a rover from a descent stage using a BUD subsystem to control the descent. The test included every pyrotechnic event and deployment performed during the sky crane maneuver including all of the mobility deployments. To compensate for Earth gravity, the rover mass was reduced to approximate its Mars weight so that the BUD descent profile was as flight like as possible. The difference in gravity did affect some of the timing and accelerations in the test but did not jeopardize any of the hardware.

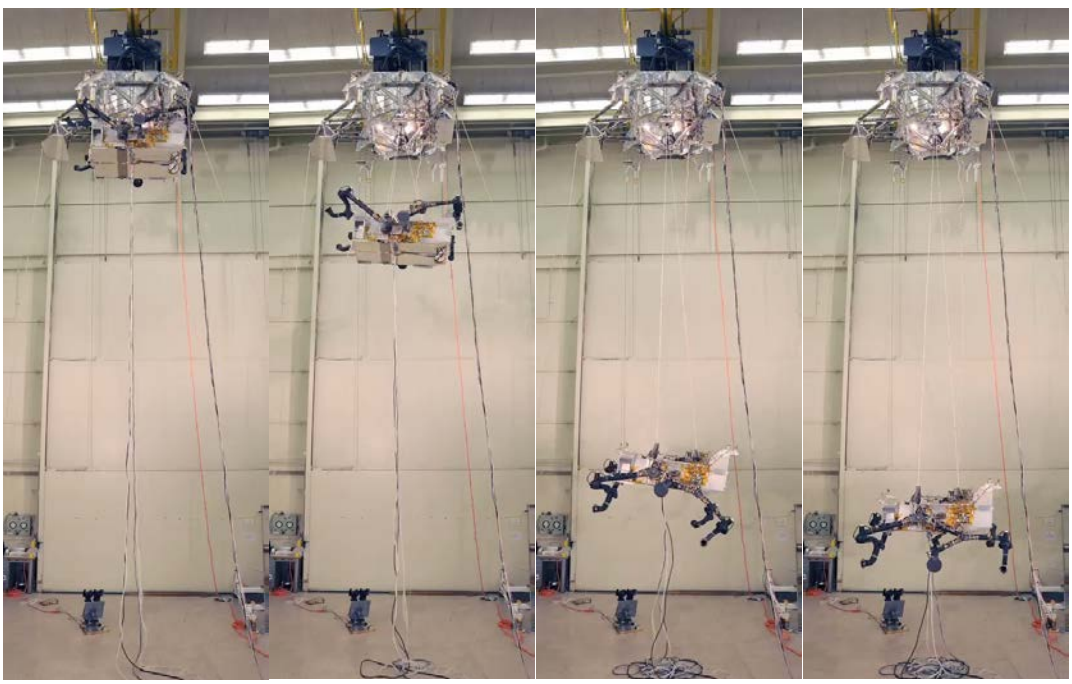

Figure 9: Rover and descent stage full motion drop test.

RCS THRUSTER CONFIGURATION

During Entry, the RCS thrusters are used to bank the vehicle (coordinated roll and yaw) to steer the lift vector and influence the vehicle's flight through the atmosphere. The RCS thrusters are attached to the descent stage structure and protrude through initially cowls and later, after a design, ports on

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

the backshell. The pointing directions of these thrusters had been set relatively early in the development of the EDL system, around the 2006 time frame.

The initial layout of the RCS thrusters was evaluated for control authority at the time of the design. After development work had commenced the aero thermal and CFD analysis of the configuration indicated possible concerns on two fronts. First, the thruster cowls, which are shown in Figure 10, appeared to receive unacceptable levels of aerodynamic heating.

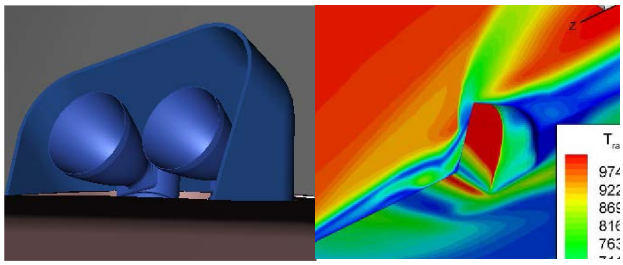

Figure 10, Initial RCS Thruster Configuration and Aero-heating

The second concern was about possible loss of control authority when using the RCS jets, because of interaction with the wake bubble behind the vehicle. Behind a blunt body, there is a recirculation zone of near stagnant atmosphere. In this bubble of stagnant gas, the mass flow of the RCS jet exhaust can pressurize regions of this bubble and effect the pressure distribution on the backshell. Although these pressure changes are relatively small, they can act over a wide area and possibly sum to an appreciable force, possibly reducing, or even reversing the thruster's net torque upon the vehicle. The concern is at least that reduction of thruster control authority may result in a vehicle with a slower and more sluggish response time to bank commands. In the worst case such sluggishness or the reversal of thruster control authority would result in an unstable control system, an uncontrollable spacecraft and a loss of the mission. Figure 11 shows Mach contours at Mach 18 and illustrates the subsonic bubble behind the vehicle. Figure 12 illustrates pressure distributions on the backshell that can result in adverse torques, that is torques in opposition to the applied RCS torque.

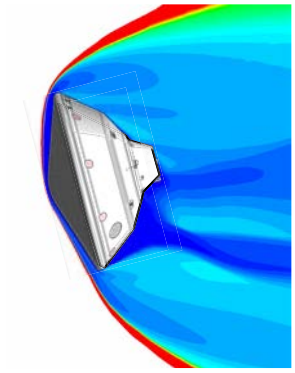

Figure 11, Mach Contours at Mach18 Flight Condition

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

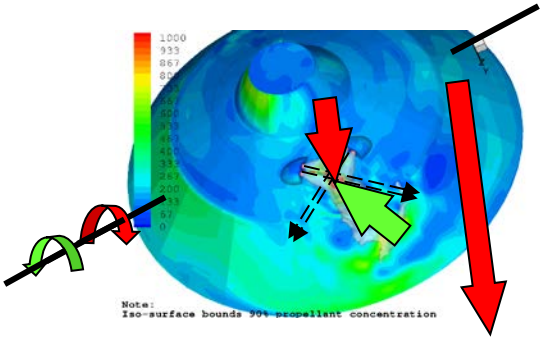

Figure 12, Adverse Aero-RCS Interaction Torque

The design response to this was to reassess our RCS thruster configuration and pointing directions and ultimately make a change to the design. This was relatively painful since the resulting design change required re-routing of propulsion fuel lines and changes to the primary structure to support the updated configuration. The design fix was agreed upon by the project because of the importance of clean high margins in domains for which testing is very difficult or impossible. The reorientation of the RCS thrusters resulted in improved torque margins against the possibility of torque reversals. This analysis technique was shared with the Phoenix (PHX) project, and analysis of the PHX aero-RSC interaction showed that thrust reversals about the yaw axis were possible. The PHX project chose to not use the RCS system during entry, as had been the plan. During flight no aero-RCS interaction were detected, suggesting that the interactions, if present, were very small.

PARACHUTE TESTING

Parachute testing focused on strength verification and mortar deployment verification. A series of full-scale deployment tests were performed at the National Full Scale Aerodynamics Complex (NFAC). Ultimately, 10 load cycles starting with a 25% overload of the peak expected load and gradually dropping to approximately 10% overload were performed. The load cycles tests were performed to demonstrate canopy robustness to load oscillations seen while flying at Mach numbers greater than 1.5. The mortar deployment tests were aimed at demonstrating the adequacy of the parachute pack and deployment bag, and to verify the ballistics of the mortar itself.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

Figure 13: Parachute structural verification testing performed at the National Full Scale Aerodynamics Complex.

Telemetry from the mission indicates that the parachute performed as expected. It deployed at Mach 1.7 and produced descent rates and system dynamics within 10 % of preflight predictions.

SYSTEM DESIGN MARGINS

Surviving development challenges and disruptions requires having the margin to spend in recovering from those problems. This margin is commonly discussed when it comes to cost and schedule, but less frequently treated within a design context. Building the system design with technical margins within it is what allows the design to survive the errors in design, oversights in design requirements and failure in implementation. A measure of how effectively a system's design was margined is the extent to which the fundamental design architecture survives the development process unperturbed.

In the system design for the MSL EDL system, a JPL heritage, "overlapping requirements" approach was used. That is, requirements were partitioned at a system level, such that subsystems had overlap in the performance requirements. This is a method of transparently holding system technical margin in such a way that the underlying technical risks can be understood. Within each subsystem, the risks of meeting requirements can frequently be known, and the consequence of not receiving the interface requirement similarly understood. Naturally, for subsystem-to-subsystem technical margins, there is no general way of trading them against each other; they must be treated on a case-by-case, risk-by-risk basis.

For example, consider the powered descent vehicle separation from the parachute and backshell (BS Sep). Successful separation requires that the angular rates at the time just prior to the separation are within some given limit and that within that limit, the mechanical separation of the hardware will be successful.

During the flight of the MER Spirit rover, the team had been surprised by the amount of angular motion of the rover system under the parachute. As the early development of MSL began, the team wanted to protect itself against the possibility that we could be surprised by angular rates (we call them wrist mode rates) in flight again. To combat that threat, and keep the MSL radar viewing the ground at acceptable angular rates and off-nadir viewing angles, the descent stage reaction control system (RCS) thrusters had been configured to allow use of the RCS thrusters while the vehicle was under the parachute. Preliminary computational fluid dynamics simulation had shown that the thrusters plume would not be a concern to the parachute.

The guidance navigation of control subsystem (GNC) was given a requirement to damp the on-chute angular rates to less than 30 deg/sec about the transverse axes and less than 10 deg/sec about the axis of symmetry. This was well within the radar performance envelope of 90 deg/sec and would ensure best possible radar performance. If the mechanical subsystem (Mech) had been given the same requirement, there would have been no wasted effort and no requirements overlap. There would also be no robustness to errors made in design, our understanding of the flight risks, or implementation.

In the EDL system design, the requirements placed on the mechanical subsystem were to successfully separate at up 50 deg/sec about any transverse axis and 20 deg/sec about the symmetry axis. This required that mechanical build stronger guide rails to control that separation than would have been required otherwise. This could be seen as a mass hit, but in the development of a new system, caution and conservatism, are, most frequently, your friends.

During development, mechanical requested a reduction of their requirements to 50deg/sec about any axis. This request was made because the accommodation of the original requirement was

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

leading to wide spread backshell structural design accommodations rather that local design reinforcements. In the face of a wide spread mass consumptive design response to what was truly an engineering judgment based design overlap requirement, the request for requirements relaxation was accepted at the system level.

Four years later in 2010, further CFD solutions of the RCS thruster plume jets had brought up the concern that hot gases, possibly hot enough to decompose Technora, might impinge on the Technora parachute suspension lines. Unfortunately, there were no tests that could be performed to help uncover whether this concern was truly warranted or not. With such an unanswered concern, the team began to consider if the on-chute RCS damping of the angular rates was the correct risk balance.

Further, in the intervening years, work at understanding the wrist mode angular motion on the parachute had progressed and the final models suggested that the maximum angular rates would be just about 50deg/sec, 3-sigma. Although the final modeled wrist mode rate was acceptable from a mechanical separation standpoint and the mechanical components all had healthy margins at this level without the use of the RCS system, the memory of being surprised during MER-A dynamics was still in the team's mind and there was skepticism as to how far the parachute system dynamics model could be trusted.

Balancing the risk across all the various considerations, the final selection of GNC rate dead-bands was set at 50 deg/sec, to inhibit RCS use unless the parachute dynamic modeling was very wrong and we were presented with another dynamics surprise during flight.

PRELIMINARY FLIGHT PERFORMANCE

The performance of the EDL system in flight was largely within the expectation of preflight predictions. Table 2 shows the as flown values for many of the quantities of interest.

Table 2. Flight Predictions vs. As Flown Performance[6]

Description	Average Prediction	Range of Prediction	As Flown Value
Peak Heating, Best Estimate/No Uncertainties (W/cm^2)	113	96 – 130	109
Peak Pressure (atm)	0.315	0.301 – 0.331	0.305
Peak Deceleration (Earth g's)	12.71	12.26 – 13.25	12.61
Number of Banks Reversals	3	2 -4	3
Parachute Deployment Mach	1.70	1.55 – 1.87	1.75
Parachute Deployment Dynamic Pressure (Pa)	494.9	431.7 – 564.8	493.6
Peak Parachute Deceleration (Earth g's)	6.69	5.34 – 8.37	6.07
Radar Lock-up Altitude (km)	6.82	5.71 – 7.64	8.35
Navigation Filter Velocity Error at Radar Lock-up (m/s)	N/A	N/A	0.76
Navigation Filter Altitude Error at Radar Lock-up (m/s)	N/A	N/A	113
Propulsion System Priming Time (s)	17.9	13.9 – 23.8	17.3
Backshell Separation Altitude (m)	1662	1597 – 1726	1674
Velocity at Backshell Separation (m/s)	77.4	59.2 – 95.3	78.6

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

Powered Descent Duration (s)	53.8	50.2-58.1	55.6
Fuel Consumed during Powered Flight (kg)	287.7	279 – 297	270
Constant Velocity Accordion Start Altitude (m)	242	220 – 264	248
Constant Deceleration Start Altitude (m)	142.5	140.6 – 144.5	142.7
Rover Separation Altitude (m)	21.6	20.5 – 22.8	21.5
Velocity at Rover Separation (m/s)	0.75	0.65 – 0.85	0.77
Rover Deployment Duration (s)	5.4	5.0 – 5.8	5.3
Touchdown Velocity (Descent Stage Vertical) (m/s)	0.75	0.67 – 0.82	0.60
Touchdown Velocity (Descent Stage Horizontal)(m/s)	0.04	0.00 – 0.09	0.12
Landing Miss Distance (Range) (km)	0	0 – 8.11	2.385

There were two notable anomalies during the landing night sequence of events. First the x-band radio, which was telemetering semaphore tones directly to Earth, turned itself to stand-by mode approximately seven seconds prior to the expected loss of signal due to Earth setting. During landing night, this was assumed to be a misinterpretation of the timing of flight telemetry. Upon closer inspection of the data, it is clear that the radio set itself to stand-by mode, for an as of yet undetermined reason. This is a subject of active investigation.

The second anomaly was the unexpectedly low touchdown velocity. This arose, in the lead author's opinion, from a combination of errors in judgment during the design phase of the mission. Because of issues with the near field stay out regions of the Ka-band Radar, the six-beam antennae lay out was forced to change. During this change, it was decided that only two beams dedicated to operation during the sky crane maneuver would be sufficient. This use of only two beams meant that third direction of the velocity vector would have to be estimated using the IMU and the assumption of a well-known local surface gravity. Preflight trajectory simulation analysis used the J3 gravity model for Mars and the onboard navigation filter utilized the J2 gravity model.[7] These are reduced from the 85 x 85 spherical harmonic gravity model available. In fact, there is a notable difference between the gravity models due to the topography at the Gale Crater landing site (not surprising considering that some large impact formed the crater).

On landing night, the error in the gravity model that had been used to construct the navigation filter z-axis velocity propagation resulted in an acceleration upward and along the spacecraft x-axis of a magnitude of approximately 450 μg's.[8] This resulted in the vertical motion of the descent stage and rover slowing as they approached the surface, while simultaneously picking up horizontal velocity. At touchdown, the vertical velocity of 0.6 m/s was well outside any prediction made during the development and analysis process (0.75 m/s +/- 0.1 m/s, 3σ). The lower than planned vertical velocity resulted in reduced touchdown loads on the spacecraft; the increased horizontal velocity was still within the anticipated range.

The root cause of this anomaly can be traced back to a design error in laying out the beams of the radar. By devoting only two beams to the landing measurement, the system had to rely on a priori knowledge of the environment (in this case, local gravity). This reliance on detailed prior knowledge of the environment is antithetical to the design principals used in the formulation of the MSL EDL architecture. This was acknowledged at the time this design choice was made but it was believed

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

that we understood the local gravity well enough to rely on an Earth based model and estimation, rather than direct measurement of the critical quantity of interest, in this case, the third component of the descent stage's ground relative velocity.

There is some poetry in the fate of the landing. That this was really the prime lapse, acknowledged at the time, in the team's design practice and that it would come to pass that the resulting system vulnerability exposed us to risk that was realized during flight, is a great lesson for future expeditions into the unknown of spaceflight engineering. The resulting touchdown velocity error would have been very likely survivable even if the sense of the error had been reversed (that is higher touchdown velocity of approximately 0.9 m/s). Such a landing, however, would have been survived on the margins carried within the mechanical design of the rover; that was certainly not the design intent.

SUMMARY AND CONCLUSION

The object of this paper was to explore some of those development challenges and to give the reader insight into what was entailed in the development process for MSL. The discussion must be understood to be incomplete. Space limitations require that we omit many important development struggles such as, the challenges of developing a descent radar, the concerns over blowing sand and dust and the possibility for spoofing the radar measurements, ground-plume interaction physics and the challenge of correctly understanding the phenomena, the challenge of testing of the radar using helicopters, towers and fighter jets, struggles with the thermal characterization of the MLE plumes and there effect n the rover, causing MLE nozzle redesign, etc. The development of the MSL EDL system represents a rich opportunity to understand what steps can be taken to set a system design up to survive the development challenges, and the kind of challenges that may present themselves.

Though the development of the MSL EDL system resulted in an EDL system that performed a very successful landing, it must be appreciated that luck played no small part in the success of the development and flight performance. As the rover grew in mass through the development process, from its initial estimate of 500kg to its final value of 900kg, the EDL system survived. It survived however, because the landing site altitude requirement could be softened. Near the beginning of the development process for the EDL system as designed today, the rover was 500kg and we were required to land at up to 2.0 km MOLA and by the time that the rover was at 900kg, our performance was closure to -1.0 km MOLA. It was the fact that science was interested in water, and water flows downhill, so to speak, that meant the sites of interest were low enough that the reduced landing altitude performance did not jeopardized the EDL design.

Good engineering judgment and a healthy appreciation for the challenges that await a spaceflight system development are always needed. There are steps to be taken to maximize the probability that a system design or architecture will survive the development process, but, in the end, a large measure of luck will likely be needed.

Acknowledgement

The authors would like to thank the members of the entry, descent and landing team from across the country that made this landing system possible, members include staff at NASA LaRC, NASA JSC, NASA ARC, as well as industry and academic partners. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.

[1] Allen Chen, Martin Greco, Tomas Martin-Mur, Brian Portock, and Adam Steltzner, "Approach and Entry, Descent, and Landing Operations for the Mars Science Laboratory Mission," AAS 13-425, AIAA/AAS Spaceflight Mechanics Meeting, Lihue, HI, February 2013.
[2] Allen Chen, Robin Beck, Paul Brugarolas, et al., "Entry System Design and Performance Summary for the Mars Science Laboratory Mission," AAS 13-422, AIAA/AAS Space Flight Mechanics Meeting, Lihue, HI, February 2013.
[3] Raymond Baker and Art Casillas, "MSL Descent Stage Integrated Propulsion Subsystem: Development and Flight Performance," AAS 13-462, AIAA/AAS Space Flight Mechanics Meeting, Lihue, HI, February 2013.
[4] Gavin Mendeck and Lynn Craig McGrew, "Post-Flight EDL Entry Guidance Performance of the 2011 Mars Science Laboratory Mission," AAS 13-419, AIAA/AAS Spaceflight Mechanics Meeting, Lihue, HI, February 2013.
[5] Steven Sell, Allen Chen, Jody Davis, et al., "Powered Flight Design and Reconstructed Performance Summary for the Mars Science Laboratory Mission," AIAA/AAS Space Flight Mechanics Meeting, AAS 13-424, Lihue, HI, February 2013.
[6] David Way, Jody Davis, and Jeremy Shidner, "Assessment of the Mars Science Laboratory Entry, Descent, and Landing Simulation," AAS 13-420, AIAA/AAS Space Flight Mechanics Meeting, Lihue, HI, February 2013.
[7] Fred Serricchio, A. Miguel San Martin, and Edward Wong, "The MSL Navigation Filter," AAS 13-418, AIAA/AAS Space Flight Mechanics Meeting, Lihue, HI, February 2013.
[8] David Way, "Preliminary Assessment of the Mars Science Laboratory Entry, Descent, and Landing Simulation," IEEE-2013-2755, IEEE Aerospace Conference, Big Sky, MT, March 2013.

Copyright, 2013 by California Institute of Technology. Government sponsorship acknowledged.