Mars Science Laboratory
Entry, Descent, and Landing System Overview

Adam D. Steltzner , P. Dan Burkhart, Allen Chen, Keith A. Comeaux, Carl S. Guernsey, Devin M. Kipp,
Leila V. Lorenzoni, Gavin F. Mendeck*, Richard W. Powell**, Tommaso P. Rivellini,
A. Miguel San Martin, Steven W. Sell, Ravi Prakash David W. Way**

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-393-6708
Ravi.Prakash@jpl.nasa.gov

*NASA Johnson Space Flight Center, Houston, TX
**NASA Langley Research Center, Hampton, VA

Abstract—In 2012, the Mars Science Laboratory (MSL) mission will pioneer the next generation of robotic Entry, Descent, and Landing (EDL) systems by delivering the largest and most capable rover to date to the surface of Mars. In addition to landing more mass than prior missions to Mars, MSL will offer access to regions of Mars that have been previously unreachable. The MSL EDL sequence is a result of a more stringent requirement set than any of its predecessors. Notable among these requirements is landing a 900 kg rover in a landing ellipse much smaller than that of any previous Mars lander. In meeting these requirements, MSL is extending the limits of the EDL technologies qualified by the Mars Viking, Mars Pathfinder, and Mars Exploration Rover missions. Thus, there are many design challenges that must be solved for the mission to be successful. Several pieces of the EDL design are technological firsts, such as guided entry and precision landing on another planet, as well as the entire Sky Crane maneuver. This paper discusses the MSL EDL architecture and discusses some of the challenges faced in delivering an unprecedented rover payload to the surface of Mars.[1,2]

TABLE OF CONTENTS

1

[2] Final Version Updated 05/14/2010

1. INTRODUCTION

The Mars Science Laboratory (MSL) mission will continue the search for evidence of life on Mars through numerous scientific instruments aboard a 900 kg rover [1]. The MSL rover is much larger than and over five times as massive as the Mars Exploration Rovers (MER) launched in 2003. Landing constraints include landing the rover within an error ellipse of about 12.5 km (assuming no wind) up to an altitude of 1 km as defined by the Mars Orbiting Laser Altimeter (MOLA) program. As a point of comparison, MER delivered a 173 kg rover to an altitude of -1.44 km MOLA within an error ellipse of approximately 60 km. Landing at a higher altitude means that the spacecraft will need some lift to decelerate higher in the thin Martian atmosphere, and a tighter landing ellipse lends the need to use some of that lift for a guided entry. Thus, placing the MSL rover on the surface of Mars requires a novel method of Entry, Descent, and Landing (EDL) as a result of these more stringent landing requirements and since previous EDL methods do not scale well for such a large payload.

This paper will describe the MSL EDL sequence and some of the challenges that it faces. This paper also serves as an update for previous publications describing the MSL EDL process [2, 3].

2. VEHICLE CONFIGURATION

The spacecraft is composed of a cruise stage, entry aeroshell, descent stage (DS), and rover. The 4.5 m aeroshell is composed of a heatshield and backshell, and contains the DS and rover. Until the initial Orion mission launches and returns to the Earth, the MSL aeroshell will be the largest to enter any planet. These major spacecraft components are shown in Figure 1.

During cruise, the aeroshell is attached to the spin-stabilized cruise stage (Figure 2). Two cruise balance masses (CBMs) are attached to the outside of the aeroshell to maintain a zero center of gravity (c.g.) offset while the spacecraft is spinning during cruise. Shortly after the cruise stage is separated from the spacecraft, the CBMs are jettisoned prior to entry to enable a guided lifting entry by offsetting the vehicle center of mass from the aeroshell axis of symmetry.

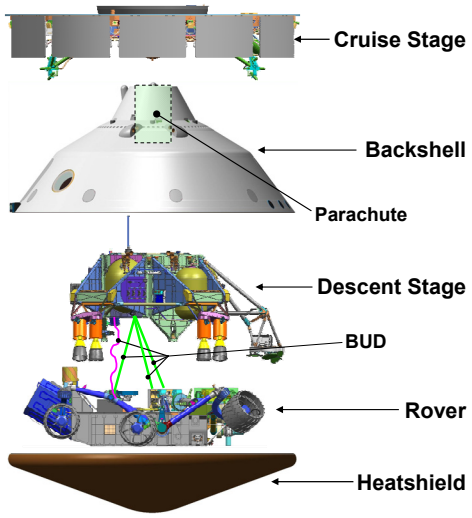

Figure 1 – Major Spacecraft Components

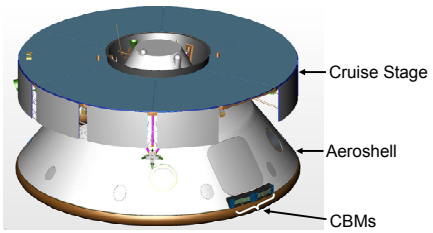

Figure 2 – Cruise Configuration

The entry configuration is shown in Figure 3a, along with approximate directions of lift, drag, gravity, and velocity vectors. The spacecraft retains this configuration until the parachute descent phase, shown in Figure 3b.

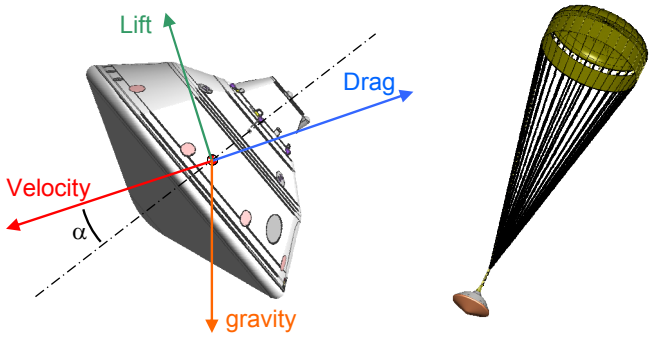

(a) (b)

Figure 3 – a) Entry and b) Parachute Descent Configuration

Following the supersonic parachute deployment, the heatshield is jettisoned to expose the both the DS and rover which comprise the Powered Descent Vehicle (PDV). Included on the descent stage are the terminal descent sensor (TDS), Mars Lander Engines (MLEs), and a bridle umbilical device (BUD). Once the spacecraft has reached an appropriate altitude and velocity, the backshell separates from the PDV. As the PDV descends, the MLEs decelerate the system even further (Figure 4).

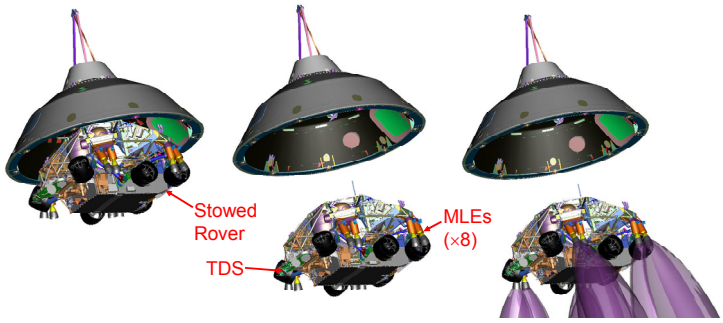

Figure 4 – Backshell Separation and Freefall Sequence

During the PDV descent, the rover separates from the DS and is lowered on the BUD which consists of three load bearing bridles and a single non-load bearing electrical umbilical. This configuration is referred to as Sky Crane. As the rover is being lowered on the bridle, the rover mobility system is deployed. Once the rover touches down on the ground, the bridle system is severed and the DS flies away to avoid contacting the rover. The Sky Crane and Flyaway configurations are shown in Figure 5a and Figure 5b, respectively.

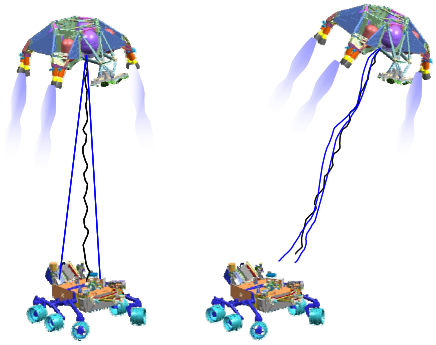

(a) (b)

Figure 5 – a) Sky Crane and b) Flyaway Configurations

The EDL sequence of events is shown in Figure 6, Figure 7, and Figure 8. The sequence of events is composed of various segments: Final Approach, EDL Start, Exo-Atmospheric, Entry, Parachute Descent, Powered Descent, Sky Crane, and Flyaway. Details of each segment are given in the following sections.

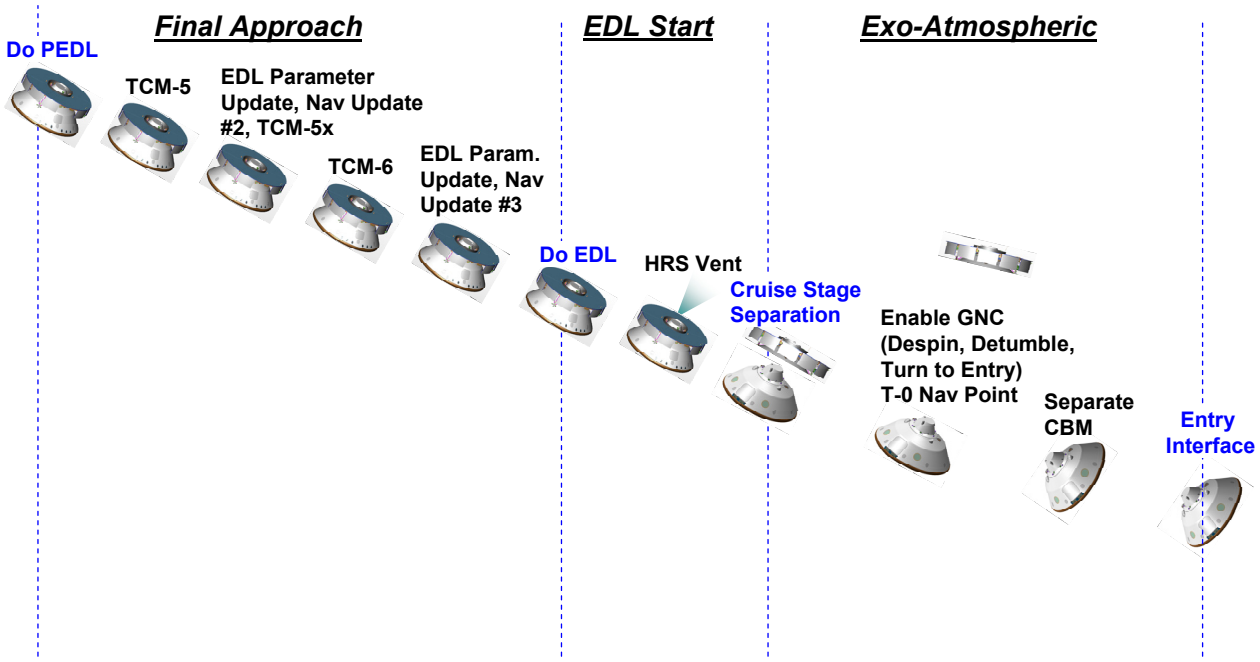

Figure 6 – MSL Entry, Descent and Landing Sequence of Events (1 of 3)

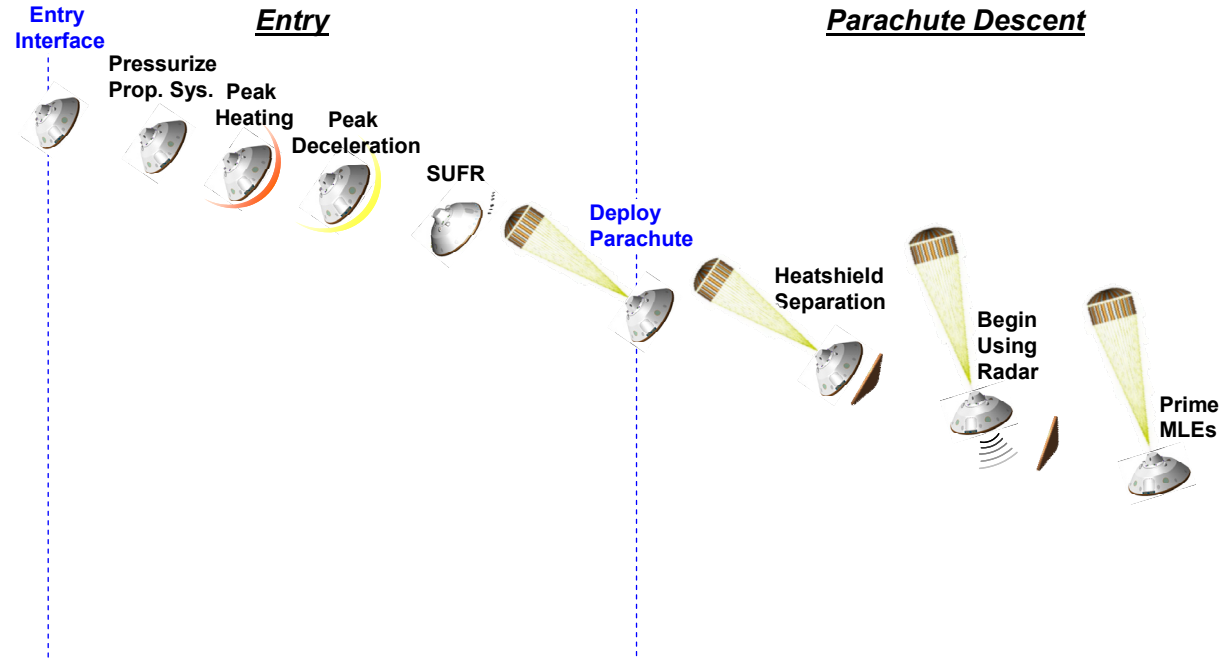

Figure 7 – MSL Entry, Descent and Landing Sequence of Events (2 of 3)

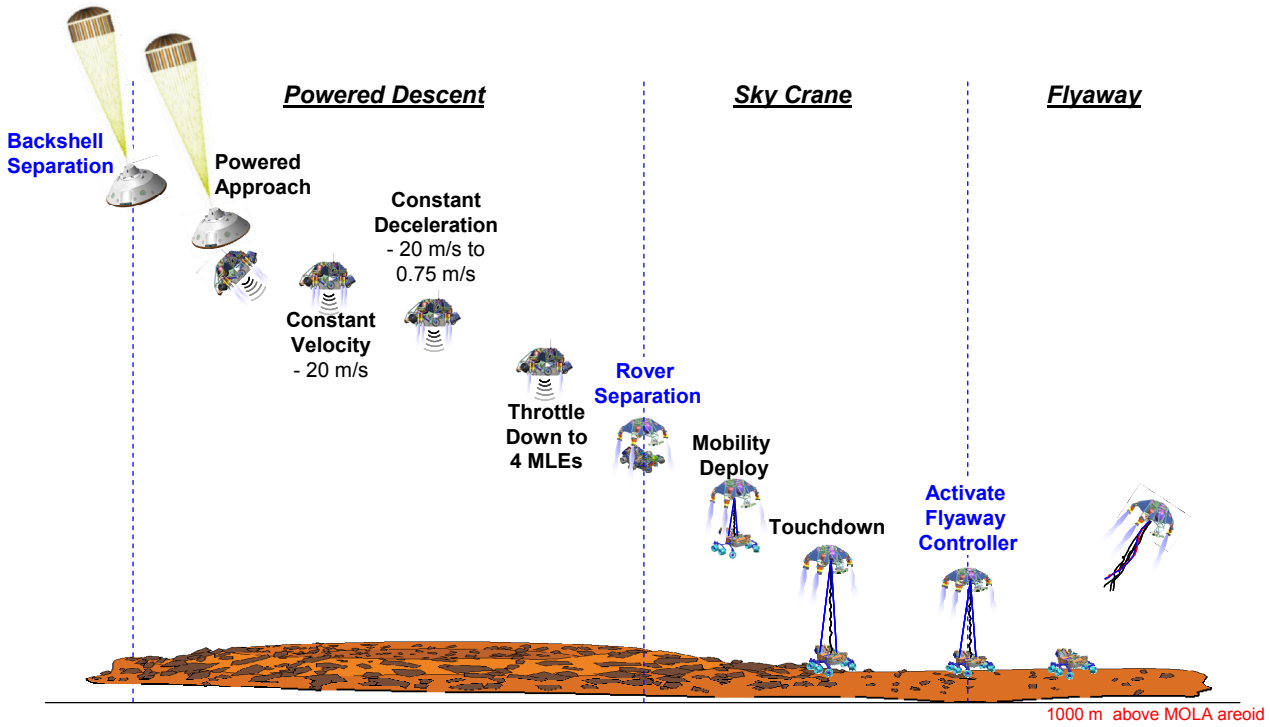

Figure 8 – MSL Entry, Descent and Landing Sequence of Events (3 of 3)

3. Approach and Exo-Atmospheric Flight

Approach Navigation

The fundamental objective of approach navigation for MSL is to ensure the spacecraft will arrive at the specified entry conditions at the correct time. The Earth-Mars transfer trajectory includes five planned trajectory correction maneuvers (TCMs) that are used as needed to adjust the entry target at Mars. The entry target is defined via optimization using both the Earth-Mars transfer trajectory and the EDL trajectory that ends with a safe landing at the desired surface target to minimize the overall TCM magnitude. Requirements are defined at atmospheric entry interface that capture the allowable contribution to landed position error from approach navigation. For MSL, these include a targeting uncertainty (delivery error) that is limited by control authority of the guided entry phase coupled with a position and velocity metric (knowledge error) that defines the initial knowledge uncertainties in these components. It is important to note that the only contributors to the surface position error from approach navigation are initial position and velocity uncertainty. By design, the requirement on targeting uncertainty is smaller than the available control authority of the entry vehicle, meaning that targeting errors can be flown out during EDL to the level of the onboard knowledge.

Covariance analysis is used to assess the impact of the various error sources (including spacecraft dynamics, Mars ephemeris errors, and signal path effects that degrade the Deep Space Network (DSN) radiometric data used for orbit determination) on the delivery and knowledge errors. Specifically, the influence of selected error sources and assumed error levels on the resulting target uncertainties is evaluated in order to define a robust navigation strategy that meets the imposed requirements. The data that are used for navigation analysis are two-way Doppler, 2-way range, and Delta Differential One-Way Range (DeltaDOR). The tracking schedule, parameters, and values used in developing the filter setup are based on past lander flight experience and the assumed MSL baseline design. Performance results for several selected cases, along with the delivery and knowledge requirements, are shown in Table 1.

Table 1- Delivery and Knowledge Requirements and Performance for Three Selected Cases

		Entry State Update	
Launch Day/Landing Latitude	**TCM-5 Delivery EFPA Error (deg, 3σ)**	**Knowledge RSS position error (km, to Entry)**	**Knowledge RSS velocity error (m/s, 3σ), Map to Entry**
Open/20N	0.08	1.5	1.1
Open/30S	0.07	2.0	1.5
Close/20N	0.09	1.9	1.3
REQUIREMENT	**0.20**	**2.8**	**2.0**

Transition to Entry

The transition from cruise to entry utilizes MER heritage where applicable, with incremental improvements. EDL execution will be performed by parameterized flight

software behaviors encompassing entry minus 5 days through the completion of DS flyaway. Additional flight software behaviors are running in parallel to EDL behaviors, and work is underway to minimize interactions between the parallel behaviors. Prior to cruise stage separation, EDL behavior is initialized, a final TCM is performed if necessary, the entry propellant and engine and thruster catalyst beds are preheated, vehicle attitude knowledge is initialized using an on-board star scanner, and the heat rejection system (HRS) is vented. Cruise stage separation occurs 10 minutes prior to entry, after which any maneuvering is performed using the entry RCS thrusters. The entry vehicle will then despin from its nominal cruise rate of 2 RPM to a zero spin, 3-axis stabilized state, then immediately turn to the entry attitude. Approximately 5 minutes before entry interface, two external cruise balance masses are jettisoned to create an offset center of gravity that provides a nominal lift-to-drag ratio of 0.24 at Mach 24. Atmospheric entry is defined at a nominal interface radius of 3522.2 km from the center of Mars.

4. ENTRY

Guided Entry

In contrast to the spin stabilized entries of MER and Mars Pathfinder (MPF), MSL utilizes an offset center of mass to create a nominal 18° angle of attack through peak heating and dynamic pressure, increasing to a 20° angle of attack just prior to parachute deployment. This angle of attack generates lift which is used to reduce the landing error ellipse size and increase the parachute deploy altitude. Entry guidance provides bank angle commands throughout entry that orient the vehicle lift vector to compensate for dispersions in initial delivery state, atmospheric conditions, and aerodynamic performance. This enables the vehicle to arrive at the supersonic parachute deployment velocity close to the desired downrange and cross-range position while maintaining a safe deployment altitude. Based on navigated attitude and the commanded bank angle, the entry controller generates roll, pitch, and yaw torque commands that are mapped into individual on/off commands for each of the 8 entry thrusters configured in pairs about the aeroshell.

The MSL entry guidance algorithm is divided into four phases. Entry interface marks the start of guided entry: guidance is initialized in the pre-bank phase and the controller commands bank attitude hold until the sensed acceleration exceeds 0.5 Earth g's. Once the sensed acceleration exceeds the specified trigger limit, the range control phase begins. During the range control phase, the bank angle is commanded to minimize predicted downrange error at parachute deployment. Throughout this phase, cross-range error is maintained with a manageable deadband limit by executing bank reversals as necessary. Peak heating and peak deceleration occur during this guidance phase. Once the navigated relative velocity drops below about 900 m/s, guidance transitions to a heading alignment phase to minimize residual cross-range error before parachute deployment. Just prior to parachute deployment, the vehicle angle of attack is adjusted to 0° by ejecting balance masses while the azimuth is aligned for better radar performance later during parachute descent. Parachute deployment is triggered at a navigated velocity of over 450 m/s.

Ellipse Size

The dispersed ellipse size at parachute deploy is driven by three factors. One is the navigated position knowledge error as the guidance cannot reduce the ellipse size any smaller. Another factor is the residual downrange error that results from a velocity-based parachute deploy trigger, although in MSL this contribution is secondary to the knowledge error magnitude. Third, the guidance accuracy is sensitive to the attitude initialization error prior to cruise stage separation. This error results in errors in the integrated altitude rate during entry, a quantity used by the guidance to predict the range flown. Greater attitude initialization errors result in greater range deploy errors as shown in Figure 9. Sufficiently large attitude initialization errors will dominate over other factors in the ellipse size.

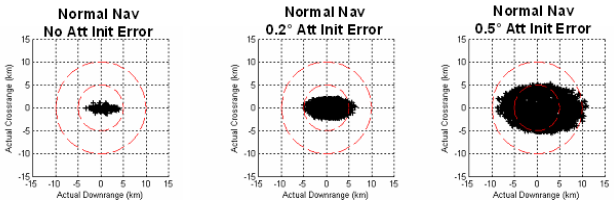

Figure 9 – Attitude Initialization Error Propagation

Aerodynamic

Given the use of the heritage 70° sphere-cone forebody geometry, aerodynamic analysis for MSL will utilize an aerodynamic database that draws upon databases generated and used by the successful Viking, MER, and Pathfinder missions. Through additional CFD analysis, wind tunnel testing, and ballistic range testing, the database will be augmented and refined to address MSL-specific risks. One such risk being addressed is the potential coupling of the entry RCS thrusters to the entry flow field aerodynamics, which may introduce unexpected control forces and augment backshell heating.

Aerothermal

An analysis of the entry aeroheating environment leads the team to expect smooth body transition to turbulence prior to peak heating, an occurrence which has not been predicted or observed in prior missions and will result in significantly higher heating rates. A combination of high ballistic coefficient, large aeroshell diameter, high atmosphere relative entry velocity, and a non-zero angle of attack

promotes this transition. Figure 10 shows the heat flux, pressure, and shear stress at their peak times on a nominal trajectory. Turbulence augments both heat flux and shear stress at the same leeside location and nearly at the same trajectory time. Pressure reaches its peak value in the stagnation region on the opposite side from the location of peak heat flux and shear stress. Studies and additional testing are underway to properly bound the heating environment and estimate associated uncertainties.

The augmented heating environment pushes the previously established qualification limits for high Mars heritage TPS materials like SLA-561V. As a result, the project has elected to baseline Phenolic Impregnated Carbon Ablator (PICA) as the forebody TPS while using SLA-561V on the backshell surfaces. PICA has been previously used on the Stardust reentry vehicle and is the baseline heatshield material for the Orion capsule.

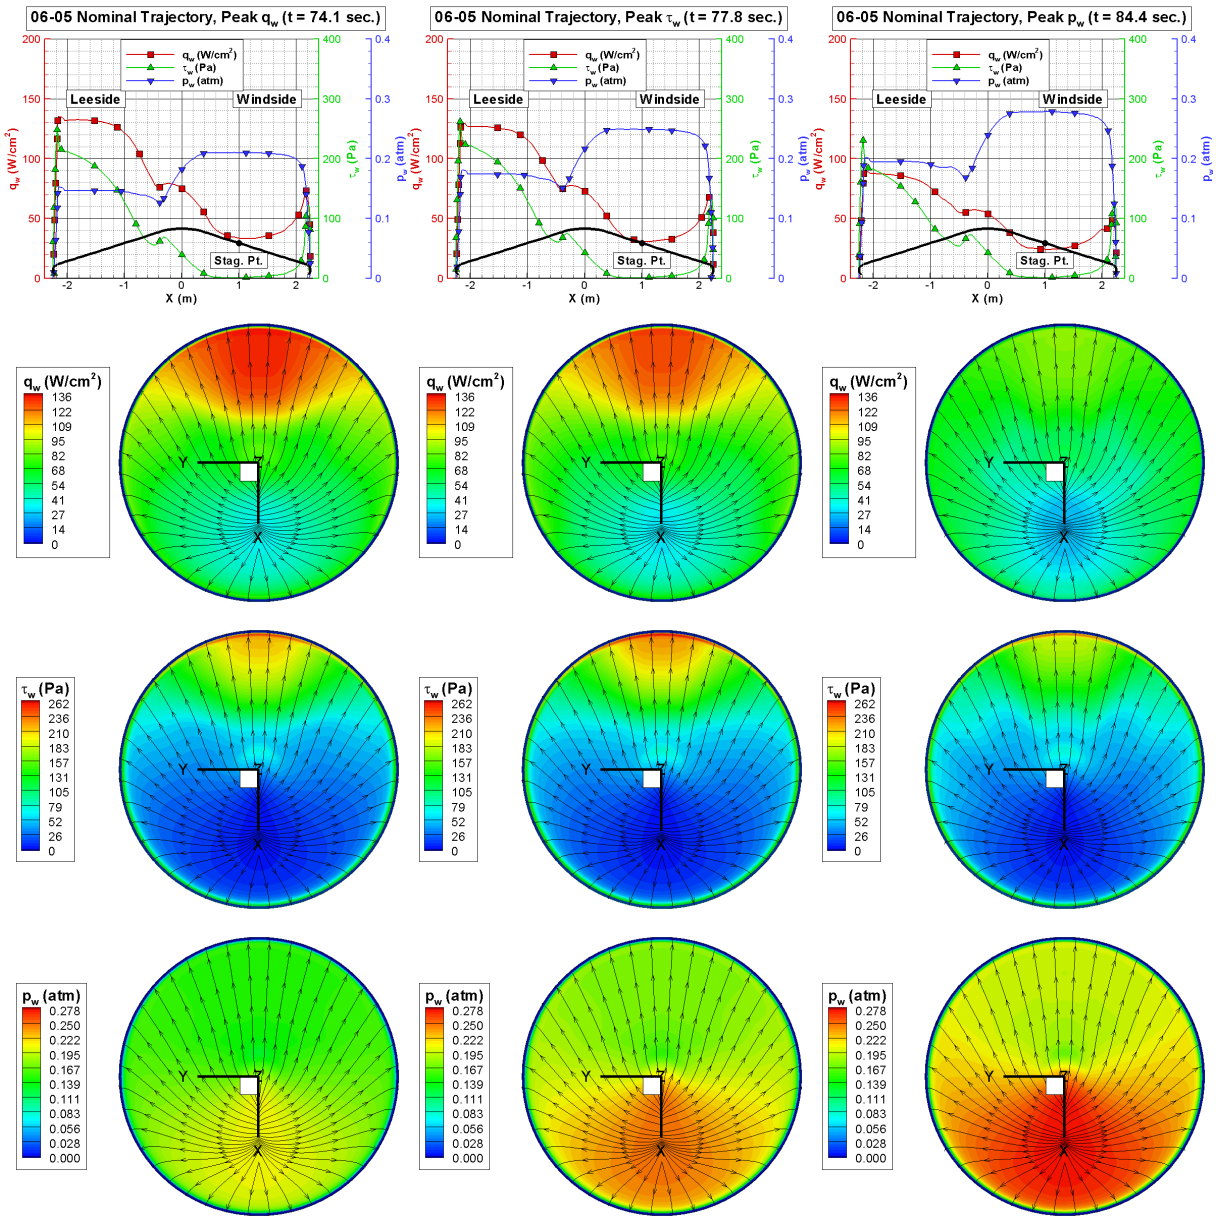

a. Peak Heat Flux **b. Peak Shear Stress** **c. Peak Pressure**

Figure 10 – LAURA 06-05 Heatshield Heat Flux, Shear Stress, and Pressure (No Uncertainties)

EDL Communications

A suite of X-Band and UHF antennas are utilized to maintain communications both directly to Earth and to Mars orbiting assets during the EDL mission phase. Direct to Earth (DTE) communications, the primary mode of communications throughout cruise and during the exo-atmospheric segment of EDL, will be through X-band low gain antennas. Due to signal strength constraints, DTE communications are limited to one-way semaphores from the spacecraft, as demonstrated by the MPF and MER missions. From the entry interface point through landing, UHF relay to the Mars Reconnaissance Orbiter is the primary communications path and has an expected bandwidth of 8 kbps. X-band and UHF antennas are mounted on the backshell, DS, and rover and will transmit data and semaphores sufficient for fault reconstruction.

Parachute Deployment

Parachute deployment is constrained to occur with an entry vehicle flight path angle of less than 5° in order to limit the off-axis loads the entry vehicle will experience during deployment. The MSL system reestablishes a nominal 0° angle of attack prior to parachute deploy by jettisoning six 25 kg ejectable mass balance devices (EBDMs) to re-center the center of mass to pre-entry conditions. The masses are ejected in two second intervals to minimize the attitude disturbance during the maneuver. This "straighten up and fly right" (SUFR) maneuver is triggered by a navigated velocity trigger commensurate with the parachute deploy trigger described below to allow 15 seconds for SUFR to complete. In addition to the SUFR maneuver, a 180° azimuth turn of the entry vehicle is conducted during this period to re-align the TDS radar beams in preparation for ground acquisition after heatshield deploy.

Parachute deployment is triggered [5] when the system reaches a specified navigated velocity, corresponding to about Mach 2.0, as determined by integration of the inertial measurement unit data in a navigation filter. The parachute deployment conditions are constrained by the heritage qualification of the 21.5 m geometrically scaled Viking parachute configuration to Mach less than 2.2. Vehicle loading constraints limit the dynamic pressure to less than about 700 Pa with a capsule angle of attack not more than 5°. Figure 11 illustrates how the dispersed MSL deployment conditions align with prior testing and mission experience.

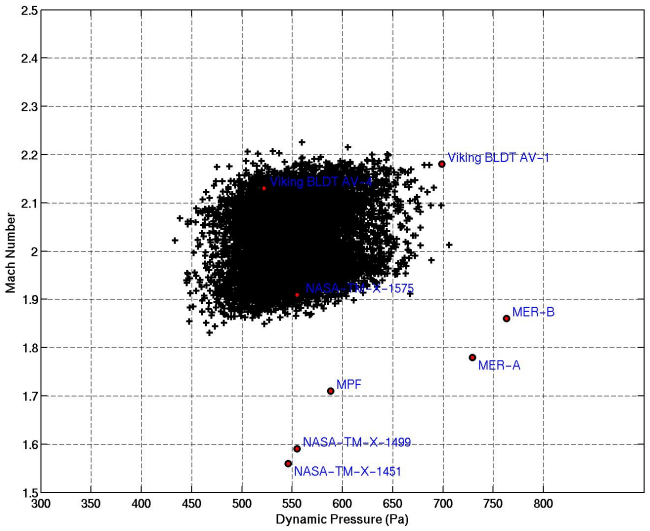

Figure 11 – Parachute Deployment Conditions

5. PARACHUTE DESCENT

During parachute descent, the spacecraft decelerates from over 450 m/s at parachute deployment down to approximately 100 m/s at backshell separation; thus the parachute system acts to burn over 95% of the remaining kinetic energy in just 50-90 seconds. In this short time descending under the parachute, the system undergoes a series of reconfigurations: jettisoning its heatshield, acquiring the Martian surface with the onboard TDS, and preparing the spacecraft to initiate powered descent.

Parachute Design

The high landed mass of MSL couples with the high landed altitude requirement to create demand for a very high performance parachute decelerator system. Design trades, including options with multiple parachutes, resulted in the baseline of a single 21.5 m diameter supersonic parachute. This parachute, which is scaled geometrically from the 16.15 m Viking disk-gap-band design, will be the largest parachute ever flown on Mars. The size of this parachute was chosen to minimize the departure from Viking heritage while providing sufficient performance to achieve a high landed altitude and reduce exposure to area oscillations.

Area Oscillations

Parachute area oscillations are a phenomena observed in historical flight test data, where the parachute's projected area oscillates notably during flight. This phenomenon is an issue of some concern because it subjects the parachute to repeated inflations at high Mach numbers, creating a dynamic environment involving high parachute structural loading and high aeroshell attitude rates. These oscillations, difficult to model computationally, have been observed to become more dramatic as inflation Mach number increases, but vanish at Mach numbers below 1.4. Time spent above

Mach 1.4 on the parachute should, therefore, be minimized. The larger parachute size of MSL reduces the on-chute ballistic coefficient and reduces the time the parachute is exposed to area oscillations.

Wrist Mode Damping

Estimating the oscillatory behavior of an entry capsule suspended underneath a parachute is an extremely dynamic and complex problem. Initial conditions at parachute deployment can pump large energies into the capsule wrist mode (rotation underneath the parachute about the capsule center of gravity) which will decay with time. Historical attempts to bound the wrist mode behavior and its time evolution following parachute deployment have failed to bound the behavior observed during flight (e.g. MER-B). While the MSL team believes we have attained a deeper understanding of wrist mode dynamics and the energies involved, the sensitivity of subsequent EDL events to high wrist mode energies led to the inclusion of an active wrist mode damping mechanism using the RCS thrusters. In cases where the capsule wrist mode frequency violates the "safe" flight envelope, RCS thrusters will fire to reduce the wrist mode frequency to acceptable levels.

Wrist mode damping is active throughout parachute descent and ensures a safe heatshield separation, good TDS surface acquisition, and a safe backshell separation.

Heatshield Separation

Following parachute deployment, the vehicle quickly decelerates to subsonic conditions. At this point, the spacecraft begins a series of critical reconfiguration events prior to initiating powered descent. The first of these critical events is separating the heatshield, which exposes the PDV to free-stream conditions and allows for the TDS to begin taking ground relative altitude and velocity measurements [5].

Heatshield separation must satisfy two requirements: positive separation from the flight system with no re-contact and satisfactory separation distance to ensure no more than one beam of the TDS is obscured after activation.

The first of these requirements is met by ensuring that (1) the push-off springs that create initial separation between the two bodies are sized sufficiently to avoid short term recontact between the rotating bodies and (2) sufficient ballistic coefficient difference exists between the heatshield and the entry vehicle to result in continuous positive separation. Because of the transonic drag characteristics of both the parachute and the heatshield, this second criterion is achieved by constraining heatshield jettison to occur at or below Mach 0.8. Since the determination of Mach number from navigated velocity is very sensitive to attitude errors, MSL has adopted a "dot product trigger" for initiating heatshield separation [5]. This trigger provides improved accuracy in deploy Mach number by accounting for an expected rotation in the navigated velocity vector caused by an initial attitude error at the start of EDL. The dot-product trigger velocity is tuned to provide for a 3-σ high separation at Mach 0.8.

The second requirement on heatshield separation exists to ensure that the TDS is taking measurements of the ground and is not getting significant return off the separating heatshield. Modeling of the antenna pattern shows that a minimum separation distance of 15 m is necessary to preclude the obscuration of multiple beams by the heatshield. The heatshield is expected to reach this separation distance within five seconds following heatshield jettison. Therefore, the MSL EDL timeline includes a five second hold following heatshield separation during which the TDS is assumed to be obscured and unable to provide any viable ground measurements. The TDS is in fact taking measurements during this five second hold and will likely provide some good data during this period. A robust filtering method is used that will flag and discard any direct measurements of the heatshield to prevent these false target measurements from impacting on-board altitude and velocity estimates.

Surface Acquisition

After the heatshield has been separated and achieved the necessary 15 m of separation, the TDS can acquire the ground and will begin taking direct measurements of spacecraft altitude and velocity relative to the Martian surface using a 3-axis Doppler velocimeter and a slant range altimeter. Prior to surface acquisition, the spacecraft state has been estimated using integrated IMU measurements from a known initial condition. This method provides state estimates sufficient for entry guidance, parachute deployment, and heatshield separation. However, large errors accumulated during atmospheric entry, combined with the precise navigation necessary for powered flight, create the need for accurate and robust altitude and velocity estimates as provided by the TDS. The TDS provides critical measurements that allow for the triggering and execution of the powered flight portion of MSL EDL.

The TDS consists of six independent radar beams: one beam aligned with the spacecraft axis of symmetry (e.g. aligned with the spacecraft velocity vector), three beams oriented in evenly distributed azimuth locations each at 20º elevation from the spacecraft axis of symmetry, and two "headlight" beams each at 50º elevation from the spacecraft axis of symmetry. Invoking azimuth control, whereby the spacecraft rolls to keep the "headlight" beams pointed downward, allows for maximal coverage of the areas on the surface closest to the eventual touchdown location; thus providing the best altitude measurements. Figure 12 shows the surfaces illuminated by the TDS beam pattern from backshell separation altitudes for steep (blue) and shallow (red) entries. The two beams on the far left are the 50º 'headlight' beams which are kept pointed downward via

azimuth control. Each black 'x' in the figure represents a discrete landing point from a 30,000 case Monte Carlo run. The figure illustrates that MSL may land as far as ~700 m from the location of the nearest surface measurement.

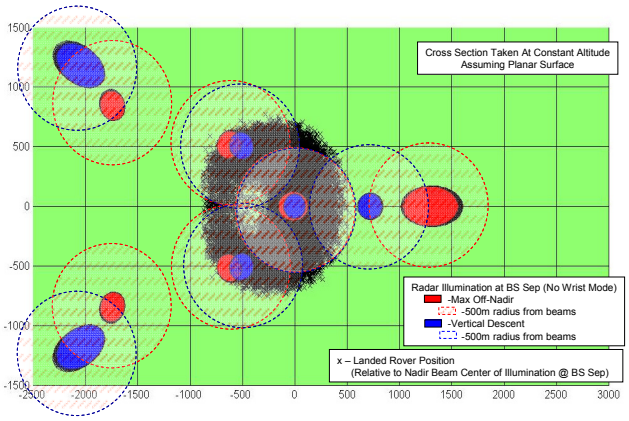

Figure 12 - Radar Illumination of the Surface from Backshell Separation Altitudes

In an ideal world, one would like the TDS to take altitude measurements relative to the exact location of touchdown. Due to variability in the approach trajectory and spacecraft orientation, this is not possible; we must make do with a system that provides altitude estimates relative to locations some distance from the final touchdown location. Hence, the system is vulnerable to large terrain variations across small length scales, limiting MSL to landing sites that are sufficiently flat. Proper selection of TDS pointing vectors and implementation of azimuth control during parachute descent significantly reduce this vulnerability.

Backshell Separation Trigger

Once the TDS is operating, the spacecraft continues to descend underneath the parachute with the PDV nested inside the backshell. The spacecraft is continually invoking both azimuth and wrist mode control using the RCS thrusters while approaching the backshell separation conditions.

The powered descent guidance algorithm triggers backshell separation, thus initiating powered descent, at an altitude between 1500 and 2000 m above ground level (AGL) and a velocity near 100 m/s as measured by the TDS. Just prior to triggering backshell separation, the propulsion system is reconfigured to enable powered descent by commanding each of the 8 MLE's to 1% throttle and firing 8 cross-strapped normally open pyro valves to begin the propellant flow to the MLE's.

6. POWERED DESCENT

Powered descent has two main goals: 1) deliver the spacecraft to the Sky Crane start condition (18.6m altitude, 0.75 m/s vertical descent velocity, 0 m/s horizontal) and 2) divert the spacecraft away from the initial trajectory to avoid potentially landing at the same location as the backshell and parachute, potentially damaging or tangling the rover.

During parachute descent, the TDS is used to determine the altitude and velocity of the spacecraft relative to the surface. As the spacecraft descends, the current velocity and altitude are used to compute the remaining altitude necessary to perform the Powered Descent and Sky Crane segments. Once this altitude is reached, Backshell Separation (BSS) is commanded, thus beginning the Powered Descent segment.

At initiation of BSS, separation nuts are fired to release the PDV from the backshell. For one second, the PDV freefalls out of the backshell to provide sufficient separation to avoid inadvertent recontact when maneuvering begins. Once this one-second freefall is complete, the eight MLEs are throttled up from their 1% near-shutdown condition and the PDV begins a 2.2 second period during which any residual attitude rates from the BSS event are removed and the PDV assumes a pre-defined attitude for the beginning of powered descent. This timeline is shown in Figure 13.

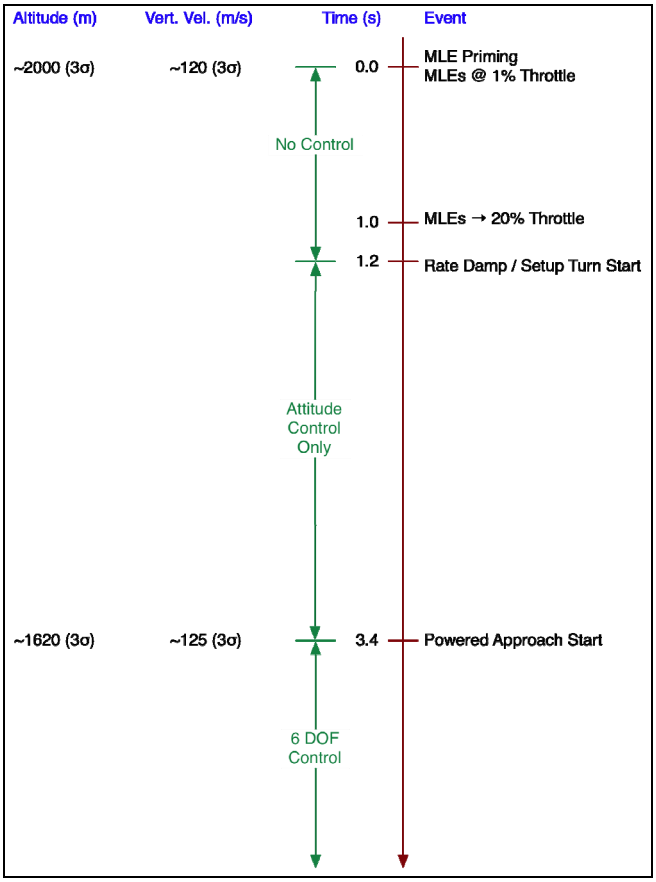

Figure 13 - Backshell Separation Timeline

The Powered Descent segment consists of four sub-segments:

1. Powered Approach
2. Constant Velocity Accordion
3. Constant Deceleration
4. Throttle Down

Powered Approach

During Powered Approach, the PDV follows a 3-D polynomial trajectory which was computed at BSS. As the PDV follows the polynomial, horizontal velocity is smoothly brought to zero while vertical velocity is simultaneously brought to 20 m/s. The end point of the trajectory is about 100 m above the surface and 300 m perpendicular to the plane of the entry trajectory. Since the PDV is actively slowing, the parachute and backshell will actually travel past the PDV and reach the surface ahead of the PDV. The 300 m divert distance is adequate to ensure the PDV does not land on the parachute or backshell. Once the endpoint of the Powered Approach trajectory is reached, the Constant Velocity Accordion begins.

Constant Velocity Accordion

When the altitude is computed for BSS, the spacecraft is still traveling horizontally and the TDS may not be illuminating the exact point on the surface where landing will occur. This, as well as inherent system errors, will contribute to an error of up to 50 m in knowledge of altitude at BSS. To accommodate this, a period of constant vertical velocity is used to fly out the altitude error. This is termed the Constant Velocity Accordion.

Since the next sub-segment (Constant Deceleration) begins at an altitude of 50 m, the target altitude for the beginning of the Constant Velocity sub-segment is set to 100 m. This will allow for the case where the surface is 50 m closer than initially calculated. In this case, the length of the Constant Velocity Accordion is zero. In addition, enough fuel must be allocated for the Constant Velocity phase for the case where the surface is 50 m further away than initially calculated, in which case 100 m of altitude will need to be traversed.

The Constant Velocity sub-segment ends when the 50 m Constant Deceleration altitude is achieved.

Constant Deceleration

Beginning at an altitude of approximately 50 m above the surface, the PDV begins the constant deceleration segment. During this sub-segment, the PDV is decelerated from 20 m/s to 0.75 m/s. This is done at a constant deceleration rate roughly equivalent to 90% throttle setting.

The Constant Deceleration sub-segment ends at an altitude of 21 m above the surface at which point the Throttle Down sub-segment begins.

Throttle Down

At this point in the landing sequence, more than half of the initial 400 kg of fuel has been consumed. In order to maintain thrust equal to weight, the MLEs would need to be throttled back to thrust levels on the order of 20-25%. Since the MLEs operate less efficiently at these throttle settings, four of the MLEs are throttled back to their near-shutdown condition of 1%. This allows the four remaining MLEs to function in the more efficient range of 50% throttle.

The transition from eight to four MLEs introduces disturbances to the system. Therefore, a 2.5 second period of time is allotted for the disturbances to settle allowing for predictable and stable conditions for the next major segment of the landing: Sky Crane.

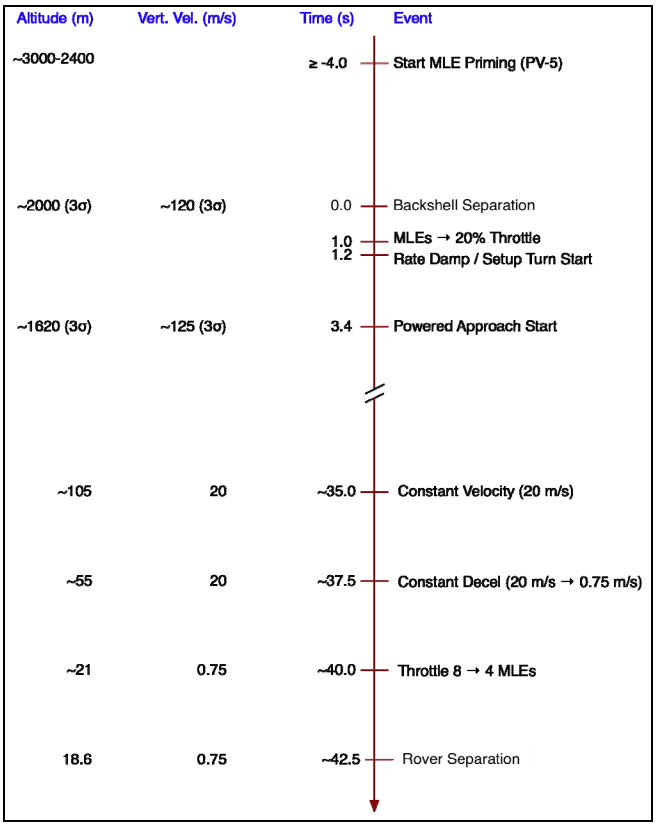

Figure 14 - Powered Descent Timeline

7. Sky Crane And Flyaway

The touchdown technique employed by the MSL design is the most innovative portion of the EDL architecture. The technique, referred to as the Sky Crane maneuver, involves lowering the lander on a triple bridle from the slowly descending DS until the bridles are fully extended to a length of 7.5 m. A 0.75 m/s constant velocity vertical descent is maintained until rover touchdown is detected via persistence of bridle offloading as inferred from DS throttle commands. Figure 5a shows the Sky Crane configuration.

Implementation of the Sky Crane architecture presents many advantages over historical touchdown methods, namely airbags and legged landers. The two body architecture keeps the engines and thrusters away from the surface, mitigating surface interactions like dust excavation and trenching, while enabling closed looped control throughout the touchdown event. The bridle decouples the touchdown event and associated disturbances from the DS controller. Additionally, rather than using a traditional touchdown sensor, touchdown is detected though a persistence of reduced throttle commands necessary to maintain the constant descent rate.

Due to the persistence of tethering during touchdown and low touchdown velocities, the system has greater touchdown stability and experiences lower impact loads than other landing systems. High stability and low loading, on par with rover driving loads, means that a separate touchdown system is not required and the egress phase can be eliminated. Rather, the rover's rocker-bogey suspension, which is specifically designed for surface interaction, is the touchdown system and it is properly positioned to begin operations immediately after touchdown.

Sky Crane Profile

The Throttle Down segment ends with the PDV descending at a rate of 0.75 m/s at an altitude of 18.6 m. At this point, separation pyros are fired to release the rover. Once the rover is released, the PDV is two separate vehicles: the DS and the Rover.

As the DS maintains a constant vertical velocity of 0.75 m/s, the rover is lowered on a triple bridle to 7.5 m below the DS through the use of an electromagnetic brake connected to a spool containing the three bridles. All of the bridles pass through a confluence point on the DS which is nearly collocated with the DS center of mass. In doing so, the Rover imparts minimal disturbance on the DS. Since all three bridles pass through a single point, it is impossible for differential loading of the bridle to produce moments on the DS. Figure 15 shows the BUD.

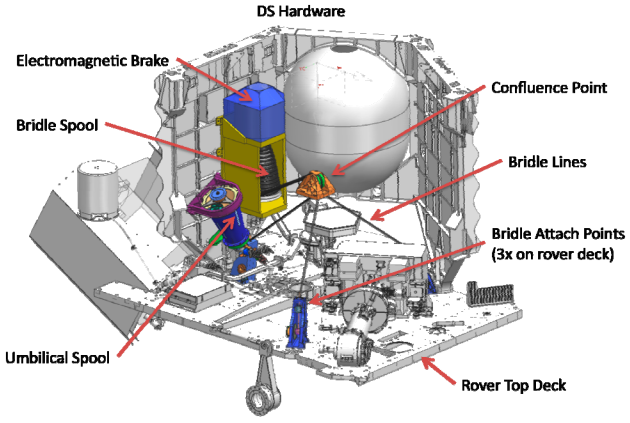

Figure 15 - Bridle and Umbilical Device (BUD)

As the Rover is being lowered by the BUD, the mobility system is simultaneously being deployed. Figure 16 shows the sequence of the BUD deployment and mobility deployment.

Seven seconds after Rover separation from the DS, the bridle reaches its fully deployed length of 7.5 m at which point the bridle is at the end of travel and motion stops – this event is called "snatch". Two seconds of post-snatch settling time is allotted to allow the DS to damp out any disturbances introduced by the snatch event. At this point, the system is ready for touchdown and the touchdown logic is enabled.

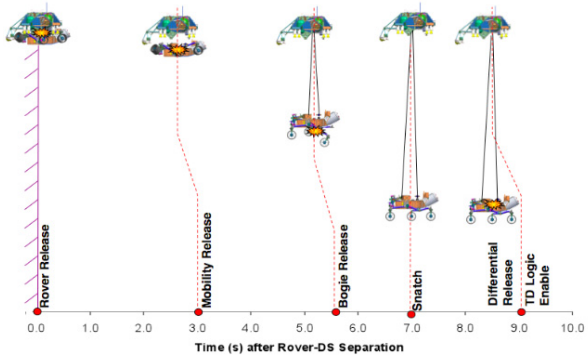

Figure 16 - BUD and Mobility Deploy Sequence

Touchdown Logic

While the DS is following a constant velocity reference trajectory, the commanded vertical thrust is equal to the weight of the system. After touchdown, the rover weight is supported by the surface, the bridle is offloaded, and the commanded vertical thrust is reduced to almost half of its previous value. This reduction in commanded thrust will persist after touchdown because the constant vertical descent reference trajectory ensures persistent offloading of the bride. The touchdown algorithm takes advantage of this inherent offloading by relying on the commanded vertical thrust to sense the touchdown event.

The touchdown logic is enabled 9 seconds after Sky Crane start (Rover Separation). Once enabled, a sliding 1 second window buffer of throttle commands is examined. At every point, the data is subjected to two tests to determine if touchdown has occurred. First, the data is examined to determine if there is a persistence of a constant state, i.e., the commanded throttle nearly constant over the window. If it is flat to within a settable tolerance, the average value is computed. If that average value is within an expected tolerance of the command necessary to support the DS mass only, touchdown is declared.

Flyaway

Once touchdown is declared, the DS halts vertical motion and the triple bridles are cut. The BUD has built-in retraction springs to retract the now free bridles away from the Rover top deck. At this point, control is transferred to

the Flyaway Controller on the DS and the command to cut the umbilical is issued.

Once the flyaway controller on the DS assumes control, it first holds the current altitude for 187 msec to allow sufficient time for the umbilical to be cut. After the requisite hold time, the MLEs throttle up and the DS ascends vertically for a predetermined amount of time. Then, the DS begins to execute a turn to approximately 45° pitch. The DS holds this attitude with the MLEs at 100% until the fuel depletes. The hold, ascent, and turn take place within 2 seconds, and the remaining time is variable depending on the amount of fuel remaining. The DS will then ballistically fall to the surface at a distance of at least 150 m from the Rover. The flyaway timeline is shown in Figure 17.

Entry Descent and Landing is considered complete when the kinetic energy of all hardware is zero relative to the Martian surface.

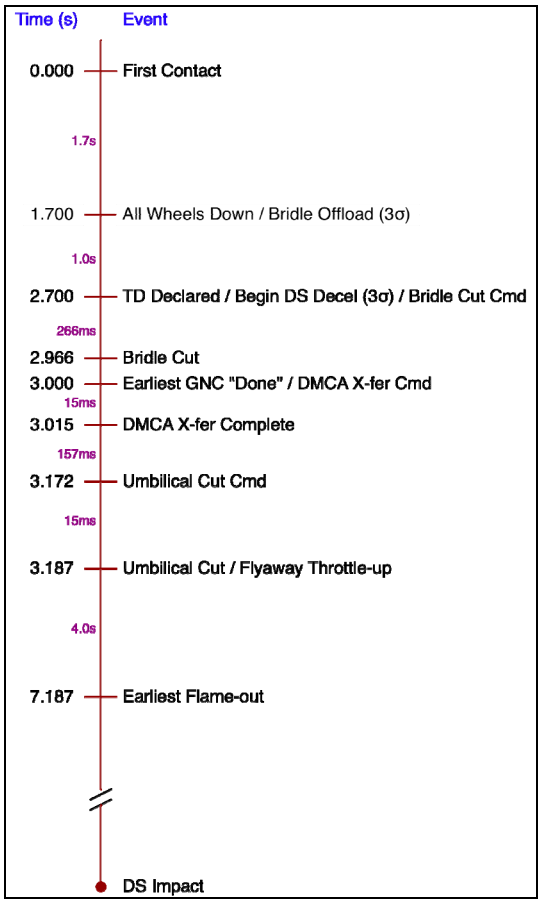

Figure 17 - Flyaway Timeline

Landing Stability and Loads

Understanding surface interactions is crucial to validation of Sky Crane performance. In particular, an understanding of rover stability and loading relative to landing site terrain is key. The terrain is characterized by surface slope on varying length scales and rock abundances for different sized rocks. Initial analysis demonstrates the ability to land on rover length scale slopes of up to 15° with vertical and horizontal velocities up to 0.85 m/s and 0.5 m/s, respectively, which are well within anticipated dispersions. Similarly, bounded analysis of worst case touchdown loads, including specific wheel impact cases, shows anticipated torques, loads, and accelerations are in family with the traverse design loads for the rocker-bogey system. At worst case, these touchdown analyses are a factor of 1.5 worse than the traverse design loads.

Plume Effects

The unique "Sky Crane" landing concept introduces concerns over rocket plume interactions not present in "Viking-Style" landers. These concerns include the potential for damage to the rover due to direct plume impingement or contamination from combustion products. CFD modeling of the plume as well as touchdown geometric and terrain analysis was performed to ensure there would be no direct plume impingement on the rover during the touchdown event. The worst-case scenario is shown in Figure 18 where the Rover lands on a 15° slope and a 0.55 m tall rock. The solid gold color is the plume boundary which could be damaging to the Rover.

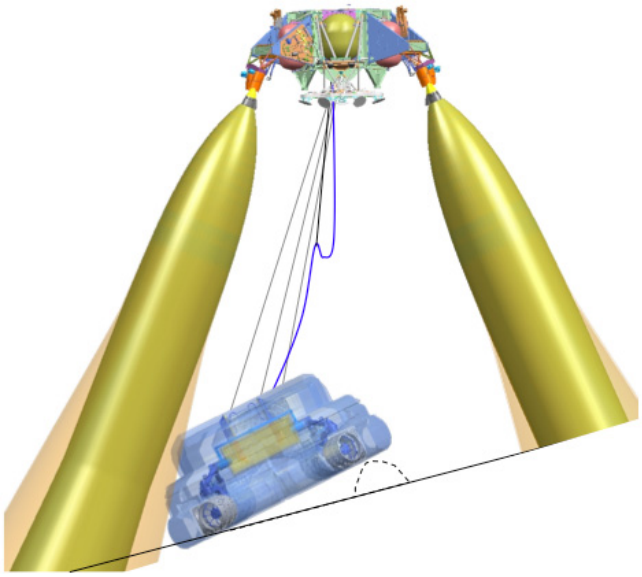

Figure 18 - Rover/Plume Clearance at TD

8. LANDING SITE SELECTION

The MSL landing site selection process consists of a series of workshops attended by members of the international Mars Science Community. The number and schedule of the Landing Site Workshops (LSW) is driven by MSL project milestones. The objective of each workshop is to systematically narrow the number of candidate landing sites so to provide a list of sites to be further scrutinized by the Mars Reconnaissance Orbiter (MRO).

The first workshop occurred in early Summer 2006 during which the MSL project presented a description of the MSL mission, its science objective, science instruments, and the landing site selection processes. During the workshop, specific landing sites and regions with specific high science priority recommended by the science community were identified and prioritized. The outcome was a list of approximately 30 sites, designated in nearly equal portions as "high priority," "moderate priority," or "best effort priority".

The MSL project then participated in the MRO Science Planning Process to coordinate the acquisition of data at the sites identified in the list so to have each site imaged as a single pass with the following approximate swath width and resolutions of 6 km, 0.3 m/pixel from the HIRISE instrument; 10 km, 20 m/pixel with 512 spectral bands between 0.4 and 4 microns from the CRISM instruments, and 30 km, 6 m/pixel from the CTX instruments.

While MRO was supporting data acquisition for science context and topographic characteristics, the Mars Atmosphere Science community supported the MSL project to preliminarily assess atmospheric features at the various Mars regions.

The second workshop occurred in October 2007 and had the objective to select the 6 landing sites with the most promising scientific return and in line with the safety guidelines. The final result of this down selection identified 3 landing sites in the Mars Northern hemisphere (Nili Fossae Trough, Mawrth Vallis, and N. Meridiani) and 3 in the Southern hemisphere (Miyamoto, Holden Crater, and Eberswalde). These sites will now be analyzed in greater details by MRO and by the EDL system team. MRO will provide additional insight to each site to investigate science characteristic and to analyze topographic features of interest to the landing site safety assessment.

The landing site safety assessment guidelines were made available by the MSL project to the science community in form of a landing site selection users' guide. Together with mobility and surface operations constraints, the EDL engineering constraints are summarized in Table 2 and Table 3.

As described in [6], the preliminary assessment of MSL EDL sensitivities to Martian environment showed that the system is sensitive to the atmosphere before and at heatshield separation when the system performance is based on inertial sensors measurements and at maximum winds when approaching the ground. Before heatshield separation, the EDL system is found to be sensitive to density variations from 8 to 30 km altitude, to horizontal wind uncertainties from 4 to 15 km MOLA, to maximum vertical winds from 1 to 5 km above ground level, to speed of sound uncertainties from 3 to 15 km MOLA. After heatshield separation, the EDL performance is dependent upon maximum vertical winds which, from 1 to 5 km AGL, are expected to be less than 20 m/s and upon max horizontal winds around 30 m/s in the first 10 m above the landing ellipse. After heatshield separation, the system is sensitive to terrain characteristics when the performance is based on the interaction with the ground.

Also after HS separation, the performance is based on the interaction with the ground and the EDL system is sensitive to maximum terrain relief, in particular the safety guidelines drive the site investigators to look for:

- Length scales of 2-5 m: maximum slope of 15º over all scales. At this length scale the system needs to ensure stability and trafficability of the rover in the touchdown condition.
- Length scales of 200-1000 m: maximum relief of 43 m over all scales (maximum slope varies because the maximum relief applies over all length scales). At this length scale, the system needs to ensure proper control authority and fuel consumption during powered descent.
- Length scales of 1 to 2 km: Maximum relief of 43 m at 1 km (~2.5º slope), linearly increasing to 720 m at 2 km (20º).
- Length scales of 2 to 10 km: Maximum slope of 20º over all scales. At this length scale, the system needs to watch for radar spoofing in preparation of powered descent.

Table 2 – EDL Engineering Constraints

Altitude	Density	Horizontal Wind	Vertical Wind	Speed of Sound
20 to 30 km MOLA	≤ 15% uncertainty			
6.5 to 20 km MOLA	≤ 10% uncertainty	≤ 25 m/s uncertainty		≤ 7% uncertainty (6.5-15km)
3 to 6.5 km MOLA		≤ 20 m/s uncertainty	≤ 20 m/s uncertainty	≤ 7% uncertainty (3-6.5km)
1 to 5 km AGL			Max ≤ 20 m/s	
0 to 10 m AGL		Max ≤ 30 m/s		

Table 3 – EDL and Surface Constraints for Landing Site Selection

Engineering Parameter		Requirement	Notes
Latitude		45°N to 45°S	Sites poleward of 30° have degraded EDL comm.
Elevation		≤ +1 km	MOLA-derived elevation.
Landing ellipse radius and azimuth		≤ 12.5 km downtrack; 10 km crosstrack	Allowing for wind-induced uncertainty during parachute descent.
Terrain Relief/ Slopes	2 to 10 km length scale	≤ 20°	Radar spoofing in preparation for powered descent.
	1 to 2 km length scale	≤ 43 m relief at 1 km, linearly increasing to 720 m at 2 km	Radar spoofing in preparation for powered descent.
	200 to 1000 m length scale	≤ 43 m relief	Control authority and fuel consumption during powered descent.
	2 to 5 m length scale	≤ 15°	Rover landing stability and trafficability in loose granular material.
Rock Height		≤ 0.55 m	Probability that rock higher than 0.55m occurs in random sampled area of $4m^2$ should be < 0.50%. Suggests low to moderate rock abundance.
Radar reflectivity		Ka band reflective	Adequate Ka band radar backscatter cross-section (> -20dB & < 15dB).
Load bearing surface		Not determined by dust	Albedo < 0.25; radar reflectivity > 0.01 for load bearing bulk density.
Surface winds for thermal environments		≤ 15 m/s (steady), ≤ 30 m/s (gusts)	Constraints apply over all seasons and time of day, at 1 m above the surface.

9. SYSTEM PERFORMANCE SIMULATION

POST2 Simulation

The POST2 [7] end-to-end MSL EDL performance simulation leverages the versatility and heritage of POST2 and adds specialized user routines specific to the MSL mission. This simulation begins one minute following Cruise Stage Separation, nine minutes prior to the nominal atmospheric entry interface, and ends with DS impact, following rover touchdown and execution of the flyaway maneuver. Major EDL events modeled in the simulation include turn to entry, propulsion system pressurization, hypersonic guided entry, supersonic parachute deploy and inflation, subsonic heatshield jettison, terminal descent sensor start, descent engine warm-up and powered descent initiation, powered approach, rover separation and BUD deployment, rover touchdown detection, and DS flyaway. EDL system performance is assessed using a Monte Carlo approach by simulating thousands of entries with random spacecraft and environmental perturbations, generating statistics, and evaluating trends and outlying cases. These random perturbations (Monte Carlo input variables) represent uncertainties in the atmospheric models, aerodynamic characteristics, mass properties, vehicle control authority, initial states, sensor performance, etc.

As the mission design life cycle progresses from conceptual design to operations, the fidelity of the various simulation models are increased to reflect the maturity of the overall system design. Early in the design process, performance Monte Carlo runs are typically conducted in Three Degrees of Freedom (3-DoF) with simplified models, allowing quick assessments of design trades. Later, as day of entry approaches, the performance simulation will be conducted in Six Degree of Freedom (6-DoF) or Multi-body with the most detailed models of the system to verify required performance and evaluate the control system performance, sensor interaction, parachute dynamics impact, detailed "Sky Crane" dynamics, etc. In addition, the simulation allows the degrees of freedom to be adjusted within a single run to tailor the fidelity to the specific problem being analyzed. For example, if the "Sky Crane" dynamics are the subject of the investigation, the portions of the simulation prior to terminal descent can be run in a 3-DoF mode to provide realistic starting conditions with minimal CPU usage and then the fidelity of the simulation can be increased to Multi-body and 6-DoF for the remainder of descent. The choice of the fidelity distribution within different events is a simple modification to the POST2 input, thus allowing the same master input deck to support multiple degrees of fidelity during the EDL simulation.

MSL performance Monte Carlo runs are given an alphanumeric designation (i.e. MSL 05-21d) and tracked via a configuration control process. The process provides source control, contains the pertinent simulation inputs and outputs, and parameter handling in an operations-like manner. The use of the configuration control process facilitates the passing of data among team members and ensures that everyone is utilizing consistent assumptions, while also preparing the team for operations. The use of such configuration control is especially important because the design team is spread over multiple NASA centers.

Initial Conditions

Initial State (EI -9 min)

The initial vehicle position and velocity states are supplied in the Mars-Centered Inertial (MCI) Cartesian coordinates at 9 minutes prior to the nominal entry interface time. The Mars Mean Equator plane and the Mars Prime Meridian vector of date define the particular choice of MCI coordinate frame in use. A detailed approach navigation analysis is performed to assess the expected delivery and knowledge covariances, based on the navigation strategy employed. This analysis is similar to that performed for the Mars Odyssey, MER, and MRO missions.

Flight Path Angle

Because MSL is utilizing a guided hypersonic entry, the choice of entry flight path angle is an important design parameter that affects the overall system performance. The nominal flight path angle is chosen when constructing the

entry guidance reference trajectory to maximize the parachute deploy altitude while reserving sufficient performance margin to remove the expected delivery errors and respecting maximum heat rate, heat load, deceleration, and lofting limits [8]. Once the nominal flight path angle has been chosen, the nominal initial state is adjusted to target the desired landing site.

Initial Attitude Error

With sufficient hypersonic Lift-to-Drag ratio (L/D) to fly-out the expected delivery dispersions, the parachute deploy footprint is dominated by the residual knowledge error. The initial attitude error (IMU misalignment at the last navigation upload prior to entry) contributes to guidance range error through integrating errors in the estimated altitude rate. Because of the EDL system sensitivity to this parameter, a conservative assumption is modeled in the simulation. The attitude error is assumed to be at a fixed, "worst-case" magnitude in a nearly uniform random direction. Sweeps of Monte Carlo runs are performed at different levels of attitude error to quantify the EDL system sensitivity to this parameter. As the landing site selection proceeds, specific arrival geometry for candidate sites can be used to refine the size and shape of the attitude initialization error.

Atmosphere

The MSL project is currently applying a variety of atmosphere models and techniques to capture the range of potential conditions the vehicle might experience on landing day. While the landing season and time of day is relatively invariant for the 2009 launch opportunity, the range of atmospheric conditions may vary widely due to landing site latitude, regional terrain, and local topography. Dispersed atmospheric profiles are generated from site specific atmosphere model "base profiles" using tools included in the Mars Global Reference Atmospheric Model (MarsGRAM) 2005 version. Site specific modeling is in progress for candidate landing sites in support of the site selection process. The resultant dispersed profiles are integrated in the end to end performance simulation to provide site specific performance for margin assessment.

Monte Carlo Products

Outputs

Any parameter contained in the list of POST output variables can be captured at critical events during a Monte Carlo run. Currently in the MSL simulation, over 600 output variables are captured at over 20 events for each of the Monte Carlo cases. These outputs are used to generate statistics and plots that facilitate the evaluation of the vehicle design and EDL system performance. Statistics for particular parameters of interest are reported in the configuration control spreadsheet. Some examples of typical Monte Carlo products are discussed next.

Statistics

Typically, Monte Carlo runs are conducted with 8,000 randomly perturbed cases in addition to the nominal. Experience has shown this number to be a reasonable balance between statistical accuracy and computer run time. However, runs with as many as 100,000 cases are periodically performed to ensure that the statistics generated with 8,000 cases are representative of the larger data sets.

Figure 19 is an example of a statistical quad-chart reported on a Monte Carlo output variable (e.g. altitude at supersonic parachute deploy). A histogram of the data is plotted in the upper-left quadrant. In the upper-right quadrant, a cumulative distribution function is shown. The quantile-quantile plot (or q-q plot) in the lower-right quadrant is used to assess whether or not the given sample is drawn from a Gaussian distribution. If the distribution is Gaussian, the points plotted in the q-q plot will fall along a linear line with an intercept equal to the mean and a slope equal to the standard deviation. The text in the lower-left quadrant reveals the minimum, maximum, mean, and various percentiles of the data. By convention, most performance metrics are tracked by a 0.13 percentile or 99.87 percentile statistics. This convention is used to provide requirements with a high probability (3-sigma under the assumption of a Gaussian distribution) of success that is insensitive to the underlying distribution or the number of cases run.

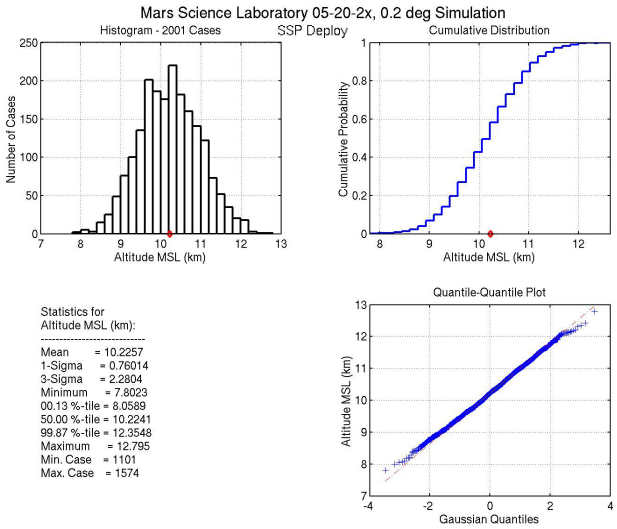

Figure 19 - Statistical Quad Chart

Scatter Plots

Parachute Deploy Conditions

Figure 11 is an example of a scatter plot reported for a Monte Carlo performance run. This figure shows the supersonic parachute deploy conditions in Mach and dynamic pressure. Figures such as this are used by the design team to track system requirements, such as those for safe deployment of the parachute. This figure superimposes the conditions of high-altitude Earth tests and previous

Mars flights to compare the current design with test and flight heritage. Another utility of scatter plots such as this is the ease of identifying outlying cases. Any outliers are subjected to more detailed forensic analysis, and if a fundamental design flaw is uncovered, the design is modified.

Radar Constraints

Figure 20 is another example of a scatter plot, showing the vehicle's altitude and off-vertical angle at the moment when radar surface acquisition efforts are initiated (5 seconds after heatshield jettison). In this figure, pink lines illustrate the radar performance envelope for both 5-beam and 3-beam acquisition of the surface. Scatter plot points appearing above and to the right of these lines represent cases where the time-on-radar is maximized and thus the best radar solution can be obtained. Three additional curves indicate the mean, 0.13-percentile, and 99.87-percentile off-vertical angle statistics at different altitudes. These additional curves represent the envelope of potential trajectories during parachute descent and help the design team understand how the parameters that effect radar performance change with proposed design modifications.

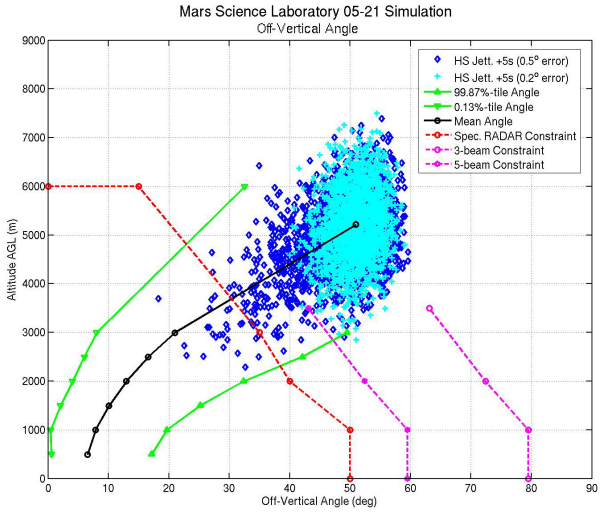

Figure 20-Heatshield Jettison Conditions

Profile Plots

For most Monte Carlo runs, the stored data is limited to a relatively small (several hundred) number of variables at specific critical events. However, it is often of interest to see how particular parameters vary between the discrete events. To facilitate this, individual cases from a Monte Carlo run can be rerun and saved. As many as 300 variables can be stored in a profile at each simulation time step. This capability allows that individual runs can be subjected to detailed forensic analysis, but also groups can be systematically investigated to better understand sensitivities. For example, Figure 21 shows 25 profiles of bank angle as a function of velocity during the guided entry.

Plots such as this are used to evaluate guidance and control performance (e.g. is the guidance saturating, are the bank reversals consistent, is the control system able to provide the guidance commands, etc.).

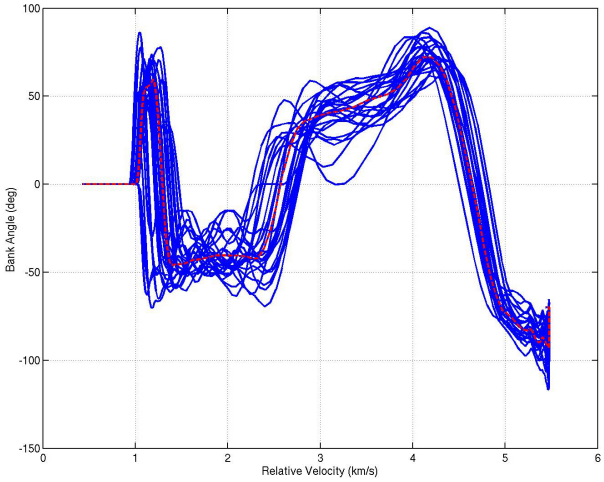

Figure 21 - Bank Angle Profile

10. SUMMARY

The EDL approach described here to land a rover on the surface of Mars is a unique method that has never been employed before. The MSL EDL sequence is a result of a more stringent requirement set than any of its predecessors. Most notable among these requirements is landing a 900 kg rover in a smaller landing ellipse than any previous Mars lander. Thus, there are many design challenges that must be solved for the mission to be successful. Several pieces of the EDL design are technological firsts, such as guided entry and precision landing on another planet, as well as the entire Sky Crane maneuver.

The development of the MSL EDL system will continue over the next two years. The MSL EDL system described herein extends current delivery capabilities in terms of mass delivered, altitude attained and landing accuracy. This system will enable a notable extension in the advancement of Mars surface science by delivering more science capability than ever before to the surface of Mars.

ACKNOWLEDGMENTS

The authors would like to thank the many members of the MSL team across various NASA centers who have contributed to this research. The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

[1] *Mars Science Laboratory mission website*, http://marsprogram.jpl.nasa.gov/msl/, Jet Propulsion Laboratory, visited Dec 2007.

[2] Way, D.W., Powell, R.W., Chen, A., Steltzner, A.D., San Martin, A.M., Burkhart, P.D., Mendeck, G.F., *Mars Science Laboratory: Entry, Descent, and Landing System Performance*, IEEE Aerospace Conference Paper No. 2007-1467, Big Sky, MT, March 3-10, 2007.

[3] Steltzner, A. D., Kipp, D. M., Chen, A., Burkhart, P. D., Guernsey, C. S., Mendeck, G. F., Mitcheltree, R. A., Powell, R. W., Rivellini, T. P., San Martin, A. M., Way, D. W., *Mars Science Laboratory Entry, Descent, and Landing System*, IEEE Aerospace Conference Paper No. 2006-1497, Big Sky, MT, Mar. 2006.

[4] Kipp, D. M., San Martin, A. M., Steltzner, A. D., and Essmiller, J. C., *Mars Science Laboratory Entry, Descent, and Landing Triggers*, IEEE Aerospace Conference Paper No. 2007-1445, Big Sky, MT, Mar.

[5] B. Pollard, A. Berkun, M. Tope, C. Andricos, J. Okonek, and Y. Lou, *Ka-band Radar Terminal Descent Sensor*, NASA Tech Brief NPO-44462.

[6] Lorenzoni L, Steltzner A, Sanmartin M, Chan Allen. "Preliminary Assessment of Mars Science Laboratory Entry, Descent and Landing System Sensitivity to Martian Environment". IEEEAC # 1087 IEEE Aerospace 2007.

[7] Program to Optimize Simulated Trajectories: Volume II, Utilization Manual, prepared by: R.W. Powell, S.A. Striepe, P.N. Desai, P.V. Tartabini, E.M. Queen; NASA Langley Research Center, and by: G.L. Brauer, D.E. Cornick, D.W. Olson, F.M. Petersen, R. Stevenson, M.C. Engel, S.M. Marsh; Lockheed Martin Corporation. Version 1.1.1.G, May 2000.

[8] G. Carman, D. Ives, and D. Geller, Apollo-Derived Mars Precision Lander Guidance, AIAA 98-4570, AIAA Atmospheric Flight Mechanics Conference, August 10-12, 1998, Boston, MA.

BIOGRAPHY

Ravi Prakash *joined JPL in 2005 as a Systems Engineer in the Entry, Descent, and Landing Systems and Advanced Technologies group. Prior to joining the MSL team, he led a terminal descent sensor trade study for the Orion spacecraft and was involved in the design of various aspects of the Orion vehicle including the heatshield separation system and retro-rocket landing system. He has a B.S. in Aerospace Engineering from The University of Texas at Austin and a M.S. in Aerospace Engineering from the Georgia Institute of Technology.*

Dan Burkhart *is a Senior Staff Engineer in the Guidance Navigation and Control section at JPL. Since joining JPL in 1995, he has worked on Mars Global Surveyor, Mars Polar Lander, Mars Odyssey, Mars Express, and various navigation-related roles on MSL before selection as the Mission Design and Navigation Team EDL Simulation Lead. He holds a B.S. in Aerospace Engineering from West Virginia University, and an M.S. and Ph.D. from The University of Texas at Austin.*

Allen Chen *is an Engineer in the Systems Engineering section JPL and has supported the Mars Science Laboratory project as a flight systems engineer since 2002. He holds a B.S. and M.S. in Aeronautics and Astronautics from the Massachusetts Institute of Technology.*

Keith Comeaux *currently holds the position of the Mars Science Laboratory Entry Verification and Validation lead. Before joining JPL, while at Boeing Satellite Systems in El Segundo, CA, he held the positions of lead flight systems engineer on Boeing's GOES-R development team and integration and test manager on the Anik F2 702 communications satellite. He received his Ph.D. in Aeronautics and Astronautics from Stanford University and his B.S. in Mechanical Engineering from Louisiana State University.*

Devin Kipp *is a Systems Engineer at JPL working in the Entry, Descent, and Landing Systems and Advanced Technologies group. Devin has been a member of the MSL EDL team since 2003. He has a B.S. in Aerospace Engineering from the University of Washington and a M.S. in Aerospace Engineering from the Georgia Institute of Technology.*

Leila Lorenzoni *received the Doctor Degree (cum Laude) in Electronic Engineering from the University of Rome, Tor Vergata, Rome, Italy in 1996. Her thesis in robotic systems supported the implementation of the first IORT (Intra Operative Radio Therapy) system. She spent a few years in aircraft maintenance and overhaul. Leila then joined the Italian Space Agency as a technical liaison for a mission developed in cooperation with NASA and she was a resident at JPL where she mainly worked on the development and integration of the payload SHARAD on board of the Mars Reconnaissance Orbiter.*

Gavin Mendeck *is an engineer in the Descent Analysis Group (DM42) of the Flight Design and Dynamics Division at the NASA Johnson Space Flight Center. Gavin is currently the entry guidance lead for MSL..*

Steven Sell *joined JPL in 2006 in the Entry, Descent, and Landing Systems and Advanced Technologies Group and is currently the Powered Flight EDL Systems Engineer for MSL. Prior to joining JPL, Steve served as the Special Projects Group lead for Payload Systems, Inc. where he was responsible for several university payloads for ISS and Shuttle. Steve holds a B.S. from Florida Institute of Technology and an M.S. from the University of Maryland.*

Adam Steltzner *joined JPL in 1991 and has worked on flight projects including Galileo, Cassini, Mars Pathfinder, and the Mars Exploration Rovers. Adam holds a B.S. from University of California, Davis, M.S. from the California Institute of Technology, and a Ph.D. from the University of Wisconsin, Madison. Adam was the mechanical systems lead for EDL on MER and is currently the Chief Engineer for EDL on MSL.*

David Way *is an Aerospace Engineer in the Exploration Systems Engineering Branch at the NASA Langley Research Center. His areas of expertise are flight mechanics, trajectory simulation, and entry system performance. He holds a B.S. in Aerospace Engineering from the United States Naval Academy, as well as an M.S. and Ph.D. in Aerospace Engineering from the Georgia Institute of Technology. David is currently the LaRC EDL Flight Mechanics Lead for MSL.*