AAS xx-xxx

THE DEVELOPMENT OF THE MSL GUIDANCE, NAVIGATION, AND CONTROL SYSTEM FOR ENTRY, DESCENT, AND LANDING

A. Miguel San Martin,[*] Steven W. Lee,[†] Edward C. Wong [‡]

On August 5, 2012, the Mars Science Laboratory (MSL) mission successfully delivered the Curiosity rover to its intended target. It was the most complex and ambitious landing in the history of the red planet. A key component of the landing system, the requirements for which were driven by the mission ambitious science goals, was the Guidance, Navigation, and Control (GN&C) system. This paper will describe the technical challenges of the MSL GN&C system, the resulting architecture and design needed to meet those challenges, and the development process used for its implementation and testing.

INTRODUCTION

The complexity of the Mars Science Laboratory (MSL) Entry, Descent, and Landing (EDL) was driven by a large improvement in the EDL performance required to meet the mission ambitious science goals. Amongst these improvements in performance was the reduction of the landing ellipse by almost an order of magnitude, and the increase in rover mass by more than a factor of four. The Guidance, Navigation, and Control (GN&C) system was a key component in the MSL EDL design, and played major role in achieving the required improvement in EDL performance.

To meet the challenge of reducing the landing ellipse size, MSL GN&C incorporated Entry Guidance for the first time in a Mars mission. In this technique, the center of mass of the entry capsule is displaced from its axis of symmetry in order for the capsule to trim aerodynamically with a non-zero angle-of-attack, thus generating a Lift vector. Then an Entry Guidance algorithm adjusts the direction of this Lift vector by commanding changes to the capsule attitude through RCS thruster firings, based on navigation data from an Inertial Measurement Unit (IMU).

---

[*] Mars Science Laboratory, Guidance, Navigation, and Control Chief Engineer, and Deputy EDL Phase Lead, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109.
[†] Mars Science Laboratory, Guidance, Navigation, and Control Project Element Manager, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109
[‡] Mars Science Laboratory, Guidance, Navigation, and Control Analysis Task Manager, Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109

To meet the challenge of landing a one ton rover, the design of EDL used the SkyCrane, a new landing technique that lowers the rover gently to the ground suspended with bridles from a propulsive stage. This new landing architecture imposed unprecedented functional and performance requirements on the design and implementation of the GN&C system. Such a system not only had to be stable during the operation of the SkyCrane but also had to achieve very low touchdown velocities to preclude damaging the rover wheels upon touchdown.

Designing and implementing a GN&C system that met all of these new unprecedented capabilities and stringent requirements required a major effort including the development of a new landing radar, the Terminal Descent Sensor (TDS), the development of throttleable terminal descent engines based on a Viking design, and the development of the GN&C algorithms and Flight Software that controlled the vehicle during the landing phase. In addition, the effort also required development and validation of detailed and complex models for simulations and analysis for end-to-end system performance evaluation.

This paper will describe the requirements and technical challenges of the MSL GN&C system, the resulting architecture and design needed to meet those challenges, and the development process used for its implementation and testing. It will also describe the testing campaign required to test and validate the models and the EDL/GN&C algorithms and flight software.

EDL/GN&C DRIVING REQUIREMENTS

In order to expand the number of places of scientifically interesting Martian sites accessible by MSL/Curiosity, the project early on adopted the use of Entry Guidance to reduce the size of the landing ellipse. The final size of the landing ellipse size was 20 x 7km (major axis, 99% confidence), which allowed the scientists to make a dramatic use of this capability by selecting Gale Crater for the landing site, and placing the ellipse between its rim and Mount Sharp at its center (Figure 1). In Figure 2 we can see how MSL landing ellipse compares with the ones from previous missions.

The large size and weight of the delivered payload was the second major requirement that had a profound effect on the design of EDL and its GN&C. MSL EDL had to place on Mars a rover with the size and weight, 900kg, of a small car. This also represented a large jump in capability when compared to previous missions (Figure 3). This technical challenge would have to be solved with the invention of an innovative landing architecture, the SkyCrane, which presented GN&C with several design challenges. The required reduction of the touchdown speed was one such challenge, as can be seen in Figure 4, which compares the MSL touchdown velocity requirement with previous missions.

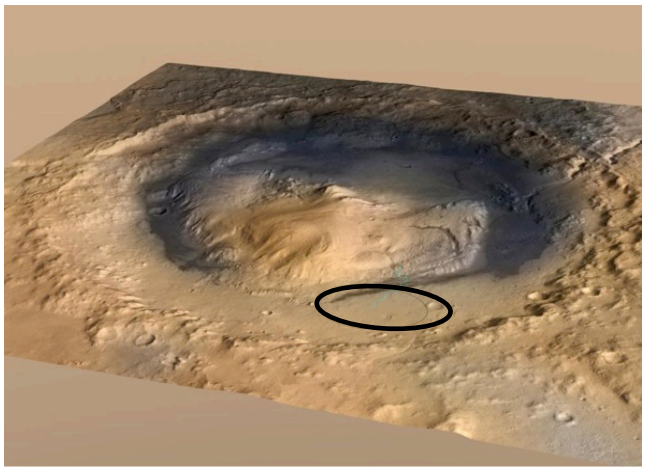

Figure 1. Gale Crater and Curiosity's Landing Ellipse

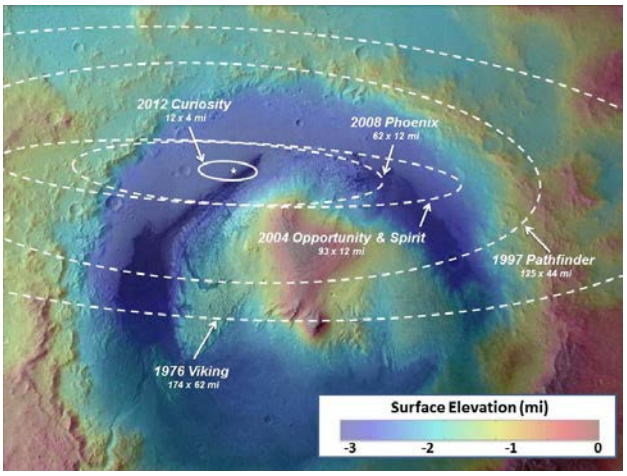

Figure 2. Gale Crater and Landing Ellipses for Several Missions

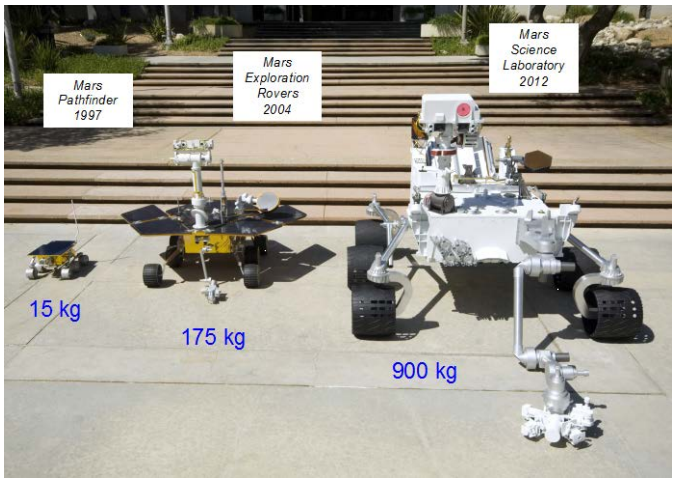

Figure 3. Comparison of MSL/Curiosity Rover with Previous Ones

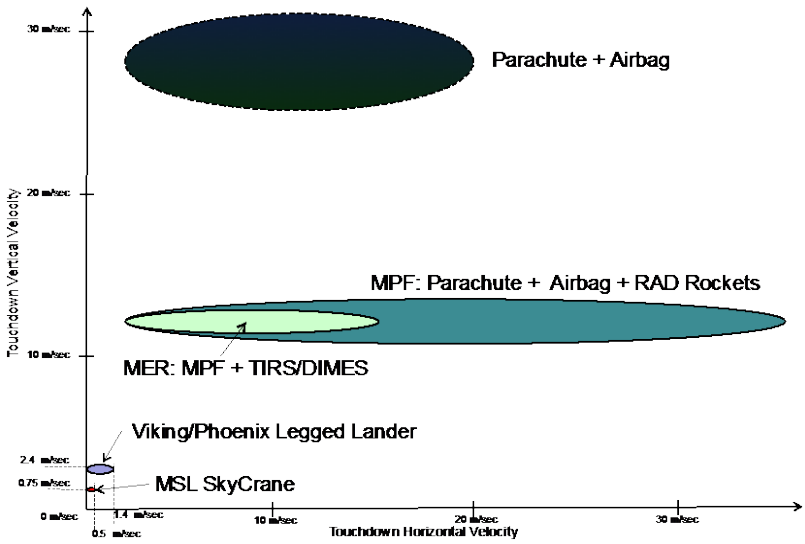

Figure 4. Comparison of MSL/Curiosity Touchdown Speed Requirements with Previous Missions

SYSTEM DESIGN DESCRIPTION BY PHASES

The problem of landing on Mars can be dissected into three phases, Entry, Descent, and Landing, as shown in Figure 5, for ease of design, management, and communication. The Entry phase starts with the spacecraft entering the Martian atmosphere at a point semi-arbitrarily chosen to be 125km altitude, with a velocity of 6.1 km/sec and ending as soon as the conditions to open the parachute safely are achieved, which for Curiosity occurs at Mach 1.7. During this phase 99% of the starting kinetic energy is dissipated through aerodynamic braking, which subjects the entry capsule to the largest decelerations and heating environments that must be handled by careful design of the Thermal Protection System (TPS). This phase is also the most important in determining the size of the Landing Ellipse. The Entry phase is followed by the Parachute Descent phase, which starts with the opening of the parachute, during which about 1% of the initial kinetic energy is dissipated. Given that the Martian atmosphere is only 1% of the Earth's atmosphere, the

terminal descent velocities during this phase are still extremely high, between 60 to 100m/sec, necessitating a third braking phase, the Landing Phase, in which the propulsion system takes the vehicle all the way until it is safely on the ground with zero ground speed and altitude.

For the rest of this section, we will describe these three phases in more detail, in particular their implication on GN&C design and architecture.

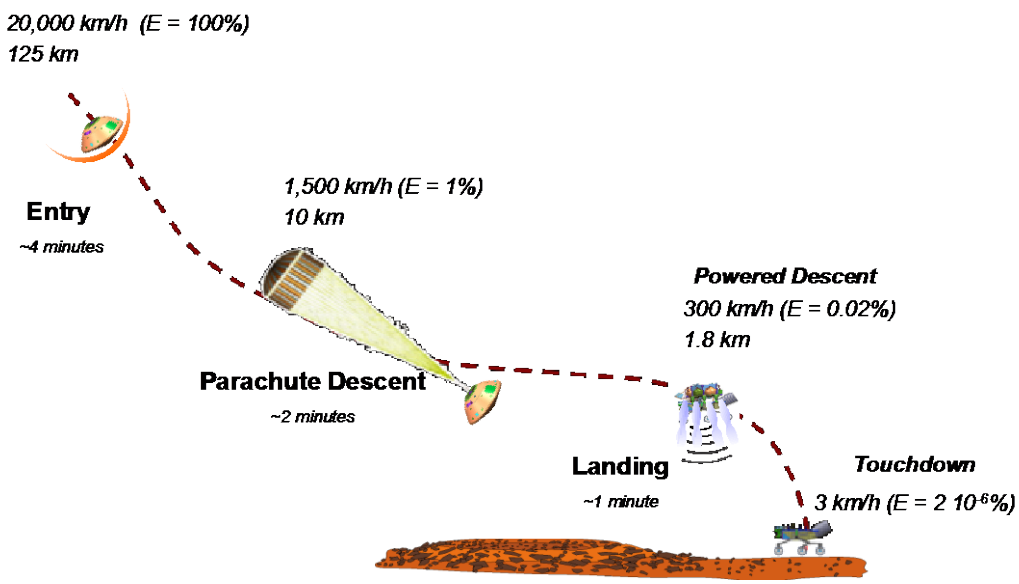

Figure 5. MSL/Curiosity Entry Descent and Landing Phases

Entry Phase

There are two types of Entry architectures that an EDL engineer has to select from for the mission: Ballistic vs. Lifting Entry (see Figure 6).

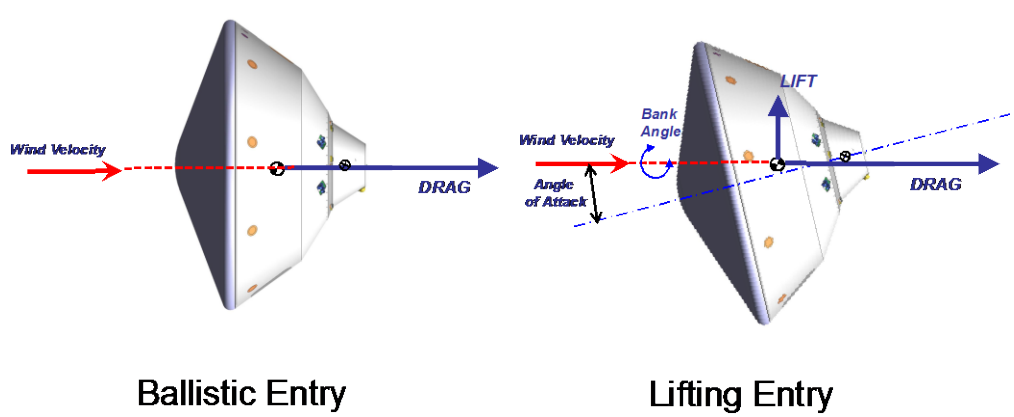

Figure 6. Ballistic vs. Lifting Entry

In a Ballistic Entry, the Center of Gravity (CG) of the capsule is aligned with its axis of symmetry, which results in a zero angle of attack and thus a zero lift force. On the other hand, in a Lifting Entry the CG is displaced away from the capsule axis of symmetry, which causes it to trim with an angle of attack and thus generate a Lift force. In a Lifting Entry architecture, a Reaction Control System (RCS) is required to control the direction of the Lift vector.

In a Ballistic Entry, the absence of Lift and of a mechanism to modulate Drag means that there is no control variable to guide the trajectory and thus reduce the size of the landing ellipse, which is hence dictated by navigation trajectory delivery errors, atmospheric dispersions (wind, density, etc.) and capsule aerodynamic dispersions. But in a Lifting Entry configuration, the capsule trajectory can be adjusted to reduce landing errors, by changing the orientation of the Lift vector through RCS actuated roll maneuvers, akin to the flight of an airplane but with much smaller Lift to Drag ratios (0.24 for MSL/Curiosity). We call this use of the Lift vector for this purpose "Guided Entry". In addition to reducing the size of the landing ellipse, Lift can also be used to improve access to high altitude landing sites, with vertical Lift used to fight gravity in order to slow the fall into lower altitudes while the capsule sheds its kinetic energy awaiting for the conditions needed for parachute deployment.

MSL incorporated Entry Guidance which required a Lifting Entry configuration, and the use of a guidance algorithm based closely on the algorithm developed in the 60's by Apollo [5]. This guidance algorithm consists of two parts: Range Control and Heading Alignment.

During Range Control, which starts when the accelerometers sense 0.05g's and ends when the navigated velocity is less than 1100 m/sec, a predictor-corrector algorithm commands bank angles to change the Lift force direction in order to control range while keeping cross-range errors within a given corridor by a performing, nominally three, bank reversals (Figure 7). If the algorithm predicts that the spacecraft is going to be short of the target, then it commands more lift-up to reduce drag (by going to the higher altitudes with lower atmospheric density) and fly longer. If the algorithm predicts that the spacecraft is going to overshoot the target, it then commands lower lift-up to increase drag (by going to the lower altitudes with higher atmospheric density) and fly shorter. In order to achieve this, the algorithm uses an on-board reference trajectory table and computes deviations from it based on the Inertial Navigator state data (Figure 8). From those deltas, the algorithm makes a prediction of the range to go until the conditions for parachute deployment are achieved. The difference between this predicted range and the current range-to-target is the range error that must be compensated for. Commanded lift-up is then computed based on the on-board reference trajectory value plus a correction proportional to the range error. Finally, the commanded lift-up (or L/D to be more specific) is used to compute the commanded bank angle.

During Heading Alignment, which starts after Range Control and ends with the start of the events leading to parachute deploy, Entry Guidance commands bank angle only to reduce the cross-track position errors that were left from the Range Control phase, and those that are currently being introduced. Range error is left uncontrolled during this phase because Range Control efficacy is diminished so close to the target. The dynamics of cross range control in Heading Alignment are analogous to the way an airplane controls its heading.

The Inertial Propagator, which uses the attitude rate and acceleration measurements from the DIMU (Descent Stage IMU), provides the estimated spacecraft state to Entry Guidance [10]. The attitude state estimate is initialized autonomously from Cruise Attitude Control System (ACS)

data prior to Cruise Stage separation. The position and velocity state is initialized from ground navigation solutions computed at the Entry minus nine minutes epoch, and uploaded to the spacecraft several days prior to landing, with opportunity to update them with improved ground navigation until a few hours before EDL.

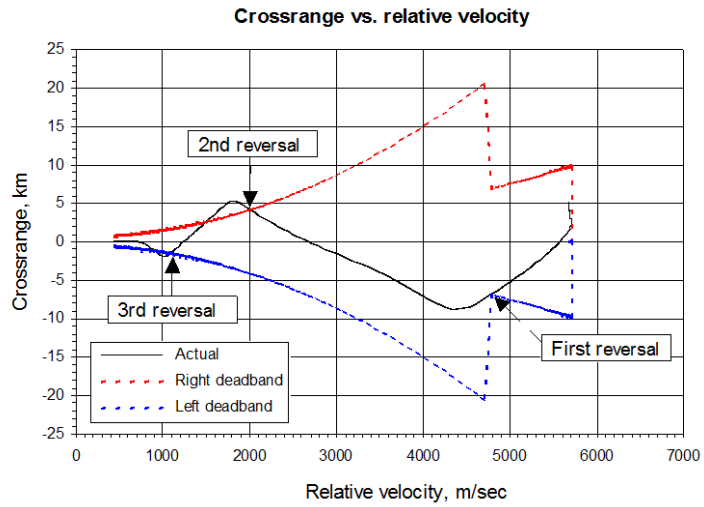

Figure 7. Crossrange Management During the Guidance Range Control Phase

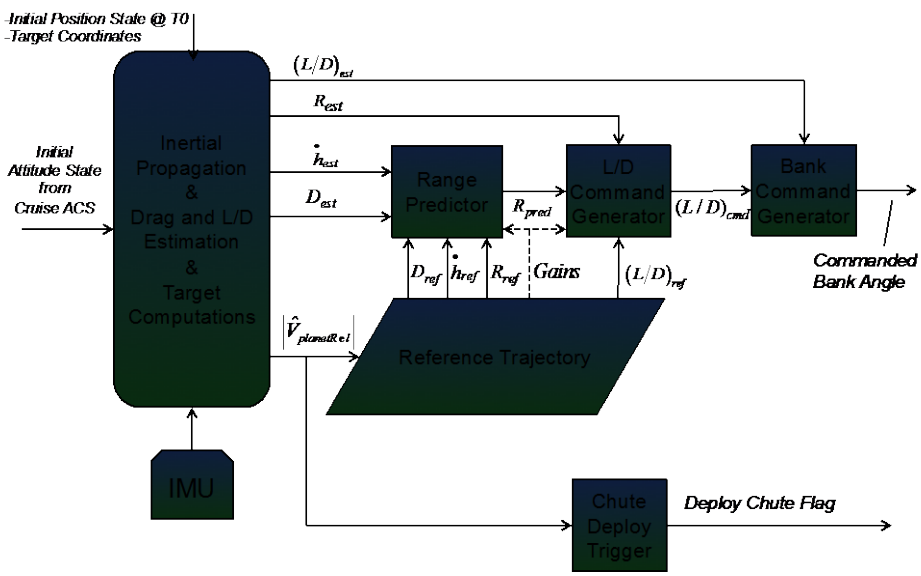

Figure 8. Entry Guidance Algorithm: Range Control

During this phase GN&C provides RCS based 3-axis Attitude Control functionality to stabilize the plant and control the direction of the Lift vector during atmospheric flight. The same controller is used during exo-atmosheric flight, prior to Entry Interface, to spin-down the capsule

from its cruise 2 RPM rate, and to turn it to the correct attitude for entry (i.e. angle-of-attack consistent with the atmosphere relative velocity).

The plant stabilization mentioned above is required because the capsule aerodynamics experiences a static instability at hypersonic speeds and a dynamic instability (unstable C_{mq}) at supersonic speeds. The latter is problematic in particular because it can result in growing capsule oscillations starting just before parachute deployment, which can exceed the maximum capsule angle of attack limit required for successful deployment and inflation of the chute.

The Entry Controller was implemented as a hybrid PD/dead-bands controller with feed-forward and RCS pulse-width-modulation (Figure 9) [6]. All commanded turns (i.e. turn-to-entry, bank reversals, turn-to-heading-alignment) were profiled and fed-forward, and they were implemented as coordinated turns (i.e. about the capsule trim axis).

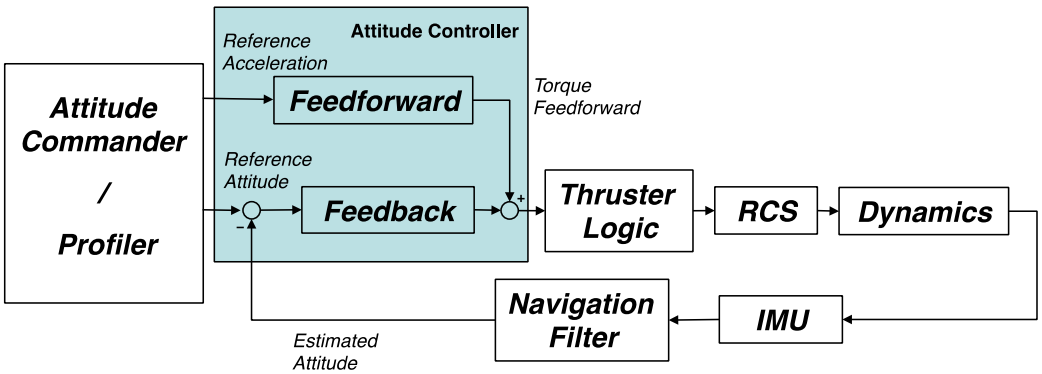

Figure 9. Entry Attitude Controller

During the Entry Phase GN&C provides triggers to indicate the start of the Guidance Heading Alignment phase and to start the activities leading to Parachute Deployment. All these triggers are atmospheric-relative-velocity-magnitude triggers that use inertially propagated velocity generated by the NavFilter [10].

Parachute Descent Phase

The big star during this phase is the supersonic parachute that has the function of slowing the vehicle down, first to achieve the right subsonic speed (~Mach 0.8) to separate the heatshield, thus allowing the use of the landing radar, and then to the right speed and altitude to start the next landing phase, which varies considerably with landing architecture.

The pitch and yaw (the transverse axes to the capsule axis of symmetry) capsule attitude dynamics during this phase is heavily influenced by parachute dynamics and winds. Pitch and yaw angles have an important influence on both landing-radar/Navigation-Filter performance and hardware separations (e.g. heatshield and backshell separations).

During this phase GN&C's main role is to determine the time to separate the heatshield (Heatshield Separation Trigger) and the time to separate from the backshell (Powered Descent Start Trigger) based on NavFilter outputs. The Heatshield Separation Trigger is also an atmospheric-relative-velocity-magnitude trigger like the ones used during the Entry Phase, but with a small twist to make it less sensitive to cruise attitude initialization errors. This trigger is called the Dot-Product trigger.

The Backshell-Separation trigger is an altitude vs. vertical-velocity trigger, that computes the right altitude to separate from the Backshell, in order to start Powered Descent, based on the vertical velocity. Both altitude and ground relative velocity are provided by the NavFilter using measurements from the landing radar and IMU [10]. The NavFilter first has to converge and produce an altitude and 3-axis ground relative velocity solution before the GN&C executive starts calling this trigger.

During this phase, in addition to the triggers, GN&C employs an attitude control system to damp the wrist mode attitude-rates and to control roll (the capsule axis of symmetry) attitude to point the radar antennas in a more favorable direction. The first was required to reduce the wrist mode oscillations to below the values required mechanically for Heatshield and Backshell separations. However, concerns about the potential damage of the parachute by RCS plumes resulted in opening the yaw/pitch deadbands to values that would result in RCS firings only with extreme attitude-rate. Roll control was disabled by choice of parameters. In both cases, Monte Carlo simulations showed that all requirements were being met even with no RCS control during this phase. Both capabilities were added as a robustness-enhancing feature in the first place, but when the issue of the potential parachute damage by RCS arose, the risk trade pointed to disabling it.

Landing Phase

The first part of Powered Descent, from parachute separation at 1.6 km to up to 20 meters altitude, is basically very similar to previous soft-landers such as Viking. It is a single body controlled in six-degrees-of-freedom by 8 throttleable engines and with altitude and 3-axis ground relative velocity provided by a Navigation Filter using measurements from a landing radar and the IMU.

MSL's Powered Descent Guidance is made up of four phases (Figure 10), each one with a different guidance objective to be achieved at the end of the given phase. For each phase, a polynomial function provides a reference trajectory (position, velocity, and attitude) to be followed, with its coefficients computed at the beginning of each phase by solving the two-point-boundary-value problem [8]. The Powered Approach phase guidance objective is to zero the horizontal velocity, reduce the descent velocity to 32 m/sec, align the lander attitude with the vertical, and perform a 300m divert maneuver to avoid the parachute/backshell from re-contacting with the lander. As far as altitude is concerned, the Power Approach phase aims at a point over the ground that is the desired altitude at the beginning of the Constant Deceleration phase, which is 142m, plus a Δh that corresponds to the altitude uncertainty (due to radar error and terrain relief features) at the beginning of the phase when the reference trajectory polynomial was computed. During Powered Approach the guidance law flies to the "inertial" point described above without introducing changes due to improvements in the altitude estimate knowledge, as the spacecraft gets closer to the terrain in which is going to land.

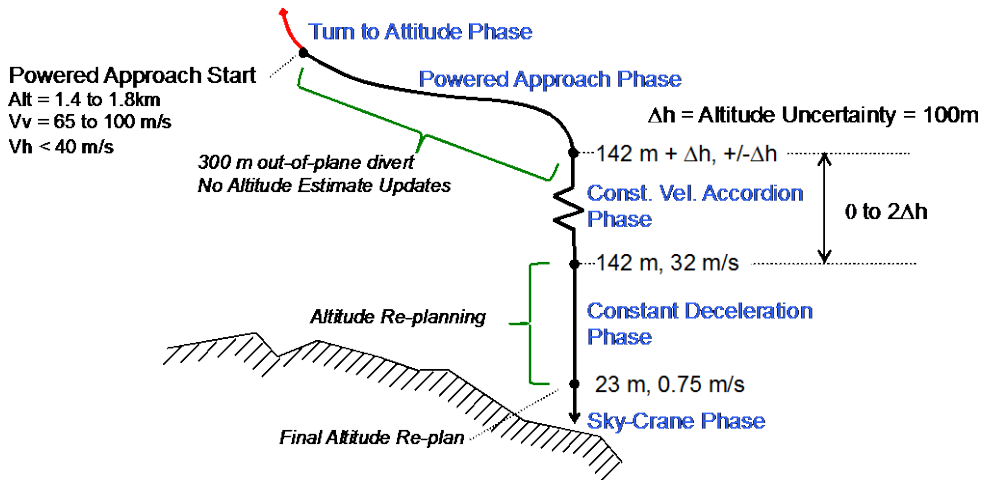

Figure 10. Powered Descent Phases

The Constant Velocity Accordion phase allocates altitude margin as an "accordion" to compensate for the altitude measurement error introduced at the beginning of Powered Approach. MSL flew with Δh = 100m to accommodate approximately +/-40m (3-sigma) of altitude error as determined by a Gale Crater EDL Monte Carlo simulation which included a high fidelity terrain model derived from Mars Reconnaissance Orbiter High Resolution Imaging Science Experiment (HiRISE) imagery [11].

During the Constant Deceleration phase the vertical velocity of the lander is brought down to the 0.75m/sec touchdown velocity, while continuing to follow a vertical profile (i.e. zero horizontal velocity). During this phase, there is also a reference trajectory re-planning capability used to adjust for minor changes in the altitude knowledge.

At this point its is important to explain briefly the reasons for this choice of guidance law. In particular it is about the concept of paying for the altitude measurement errors, incurred at Backshell separation, in a single installment during the Constant Velocity Accordion phase, instead of paying it as you go through the Powered Approach phase, which is what Viking/Phoenix did as part of their gravity-turn guidance. First of all, the MSL EDL/GN&C engineers do not claim any fuel optimality in their approach. The emphasis instead was on ease of analysis and, hence, ease of Validation and Verification of the system performance in the presence of altitude estimate errors due to terrain relief. In addition, this way to account for and manage altitude knowledge errors was very synergistic with the process of levying landing site slope constraints and analyzing the probability of success at different proposed landing sites. For example, a safe landing site required terrain relief smaller than the Constant Velocity Accordion Δh, over the terrain length scale consistent with the distance from where the radar is illuminating the surface to where the spacecraft was landing, which was 1km approximately (Figure 11). With simple analysis we were able to bound the altitude knowledge error problem, thus providing a sanity check to complex Monte Carlo simulations.

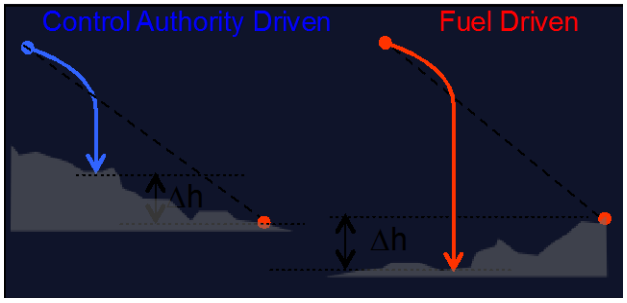

Figure 11. Altitude Knowledge Error Management

The SkyCrane phase starts at an altitude of 23m, with the lander already at the right ground relative velocity required for touchdown: 0.75m/sec vertical velocity and 0m/sec horizontal velocity. Since there is no need for large decelerations during this phase, four engines are throttled down to keep the remaining ones at a throttle level comfortably higher than the minimum allowed. The throttle-down to 4 engines occurs at the start of the SkyCrane phase and 2.5 seconds are allocated for the transients to subside before proceeding with Rover Separation (Figure 12).

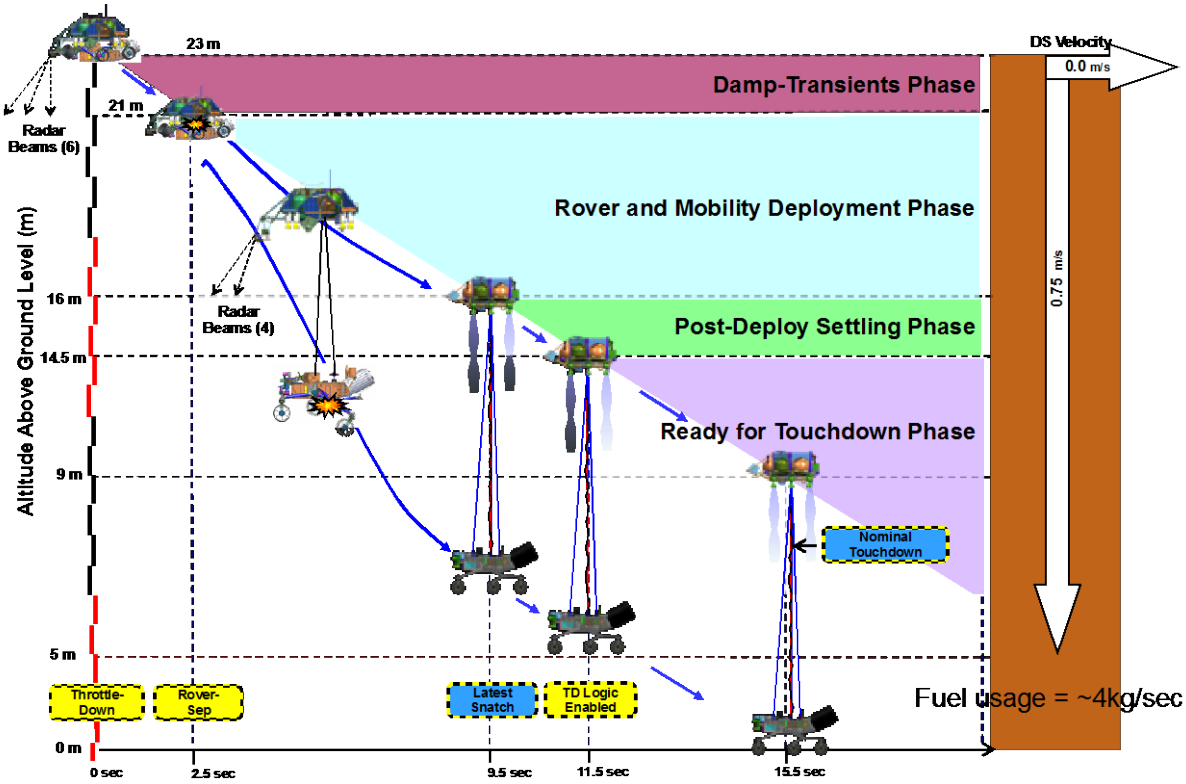

Figure 12. SkyCrane Phase

The final guidance trajectory planning also occurs at the beginning of the SkyCrane phase using the latest knowledge of altitude at that point. This planned trajectory allocates a 5m touchdown accordion to compensate for errors in the altitude knowledge at the start of the phase and the propagation errors incurred after that.

During this phase GN&C keeps the vehicle controlled and descending with the same velocity profile, even in the presence of the large transients caused by the Rover separation, Mobility deployments, and the touchdown event itself. All of these transients are handled by the GN&C feedback loops.

During the touchdown event, the constant velocity GN&C close-loop function continues until the full weight of the rover is supported by the ground, at which point the Touchdown Detection Trigger is activated. The Touchdown Detection Trigger works essentially by monitoring the commanded engine throttle required to support the system against gravity as it descends at constant velocity. When the engine throttle command is reduced to a value consistent with the weight of the Descent Stage only, then touchdown is declared. Because the weight of the Rover is roughly the same as the weight of the Descent Stage, the throttle command value is reduced by 50% through the touchdown event, thus providing a very strong and unambiguous signal.

The touchdown trigger starts the Fly-Away sequence in which GN&C first brings the vehicle to a stop with a constant vertical deceleration profile and then enters a zero velocity hold mode until Descent Stage control is transferred to the Fly-Away controller.

The Fly-Away controller is a simple GN&C algorithm designed to fly the Descent Stage to impact at a safe distance (~600m) from the Rover. This algorithm is hosted in the MLE Motor Controller computer (which also provides servo control of the MLE throttle valves) in the Descent Stage (the rest of GN&C was executed in the Rover computer) and uses IMU attitude-rate measurements to execute in close loop an attitude profile that turns the spacecraft about 45 degrees while it throttles up to fly the Descent Stage away.

Throughout all Powered Descent, six feedback control loops act together to follow the reference trajectory dictated by the guidance law, using position and attitude estimates from the Navigation Filter (Figure 13).

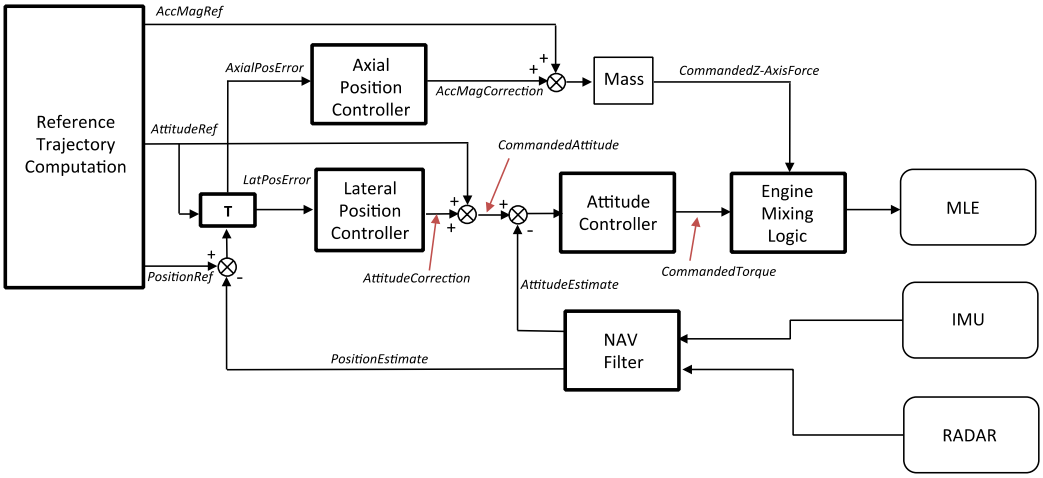

Figure 13. Powered Descent Controller

The position error is computed from the Navigation Filter estimated position and the Reference Trajectory desired position, and converted into the Reference Trajectory Descent Stage frame. In this frame, the z-axis of the position error is along the Descent Stage thrust direction (i.e. the vertical axis in its landing orientation) and is used by the Axial Controller to command the translational acceleration along that axis, which is directly controlled by the MLE throttle

valves. The x and y position errors go to the Lateral Controller that commands lateral acceleration changes by commanding pitch and yaw attitude deltas from the attitude reference. The desired resulting Descent Stage attitude is then commanded to the Attitude Controller, which request torques to the Engine Mixing logic, which are ultimately achieved by differential MLE throttle commands.

The single Axial and the three Attitude control loops have bandwidths in the 1 Hz range for good transient response and disturbance rejection. The two Lateral control loops have bandwidths in the 0.2 Hz range, low enough to maintain frequency separation with the inner 1Hz Attitude loops. Fortunately, the largest disturbances and transients occur in the high performance 1 Hz loops. All loops run at a 64 Hz sample rate in order to preserve phase margins.

No special control features were required during the two-body Sky Crane phase, beyond the capability to schedule different controller gains. The bridle axial and pendulum modes were phase stabilized.

GN&C RELEVANT HARDWARE

In order to perform the functions described in the previous section, GN&C required a set of state-of-the-art sensors and actuators, all mounted on the Descent Stage or DS (Figure 14).

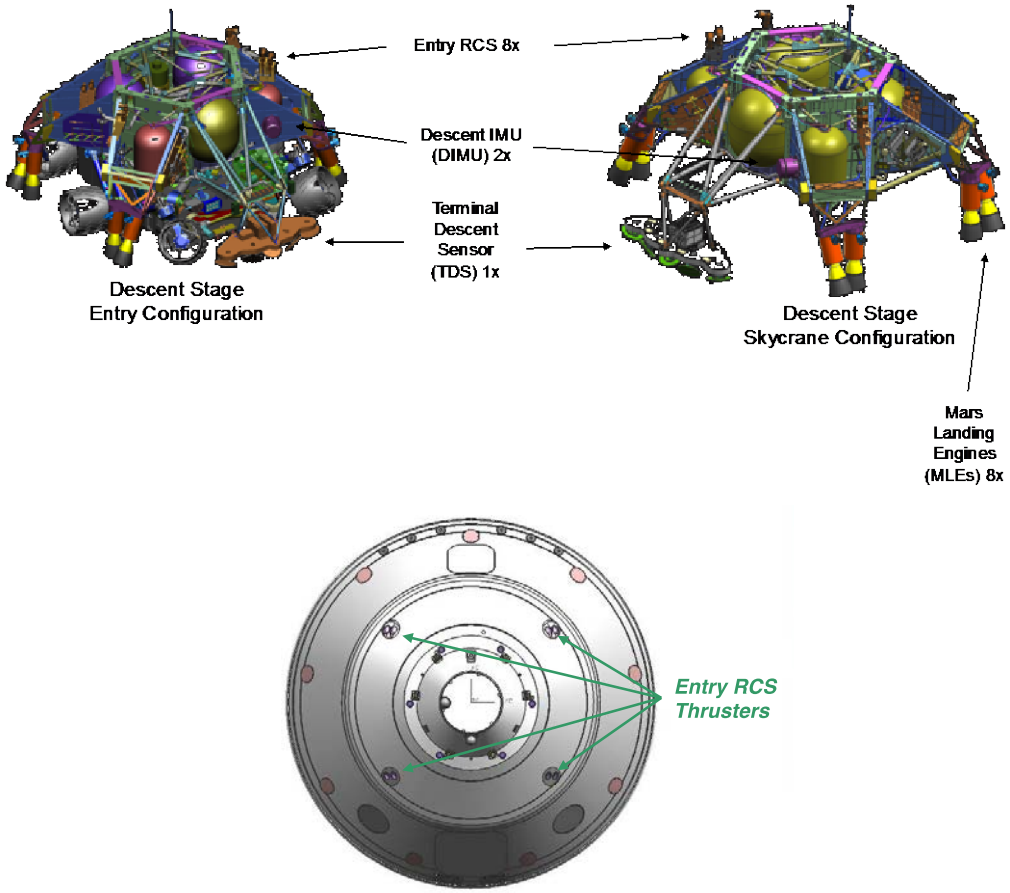

Figure 14. GN&C Relevant Hardware

As is typical in previous Mars soft-landers like Viking and Phoenix, MSL GN&C requires an IMU, to measure attitude rate and translational acceleration, and a landing radar, to measure 3-axis ground-relative velocity and altitude.

The IMU is the Honeywell MIMU and there are two located on the DS for redundancy, as can be seen in Figure 14 (DIMU for Descent stage IMU). The MIMU is a space-qualified unit that has flown successfully in space numerous times, including in the Phoenix mission that landed successfully on Mars in 2008 performing a similar function as with MSL.

MSL developed a new landing radar, the Terminal Descent Sensor (TDS), implementing a space-qualified design which employed the latest technology in order to achieve the accuracy and robustness required by the SkyCrane architecture and the reliability dictated by a flagship mission. The TDS is a Ka-band (35.75 GHz) pulse-pair Doppler radar design. It measures ground-relative velocity and slant-range along the boresight of each of its six narrow beam (~3.5deg width) antennas, at a rate of 20 measurements per second (where each measurement is a velocity/slant-range pair for a single antenna beam).

The antenna configuration consists of a nadir beam, for maximum altitude measurement accuracy during vertical flight prior to touchdown, three 20 degree beams, and two 50 degree beams which were added to guarantee that at least those two beams would not experience multipath interference from the rover during the SkyCrane phase (Figure 14 and 15). The beam sequence is programmable and is changed by commands from the EDL Executive, which selects the optimal beam sequence for the current phase of flight. For example, during the SkyCrane the beam sequence consists of only the two 50 degrees SkyCrane beams since the other beams were at risk of multipath interference with the rover.

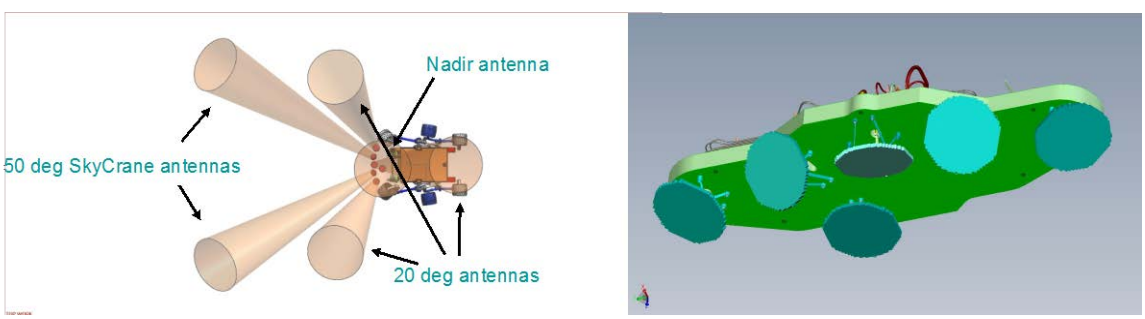

Figure 15. TDS Antenna Configuration

During exo-atmospheric flight and entry, eight RCS thrusters mounted on the Descent Stage and protruding through the backshell, were used for attitude control (Figure 14) [6]. Each provided 150 N of force during exo-atmospheric flight, and 250 N when the propulsion system was pressurized at Entry Interface.

During Powered Descent, eight throttleable Mars Landing Engines (MLE) were used for attitude and translation control (Figure 14 and 16). The MLE's are based on the Viking throttleable-landing-engines which, through a focused technology development program, were successfully revived with some important modifications. The MSL version does not have the 18 small nozzles used in Viking to minimize plume ground pressure since the SkyCrane architecture does not require it. It also has a new cavitating throttle valve that reduces the dependence of the flow rate on

downstream pressure drop. The engine can be throttled from a minimum thrust of 400 N to a maximum 3060 N.

The propulsion system is pressurized at Entry Interface thus maintaining a constant thrust level, in both the RCS thrusters and MLE's, despite fuel depletion.

Figure 16. MSL Mars Landing Engine (MLE)

GN&C ALGORITHM AND FLIGHT SOFTWARE DEVELOPMENT

MSL GN&C followed the same GN&C algorithm and flight software development methodology developed for Mars Pathfinder. That mission was one of the first developed under the Better, Cheaper, Faster philosophy being promoted by the NASA administration at that time. The Mars Pathfinder project decided then that a way to reduce costs while maintaining quality and reliability, was to have the GN&C analysts not only develop algorithms but to also to code them for flight and deliver an end-to-end integrated and tested GN&C flight software product, the GN&C flight software core; then the flight software team would integrate it with the rest of the software without any modifications. In conjunction with, and a complement to, this methodology, a subset of the GN&C analysts would stay with their product through the rest of its life cycle, including integration and testing in all the testbed environments, and ultimately participate in flight operations. This methodology contrasts dramatically with the more traditional compartmentalized approach in which algorithm development, flight software conversion, testing in multiple platforms, and operations, are all done by separate teams.

The advantages of the Mars Pathfinder approach, which in reality is the way things were done in the pioneering days of space exploration, is that eliminates extensive documentation, and empowers the engineers by providing them with a greater sense of ownership and accomplishment. By minimizing documentation the project not only saves money and schedule but also reduces errors typically introduced by the ambiguous language used in most documents and other person-to-person communications. Instead, the C programming language, which was used to code the flight software, becomes the unambiguous language of communication between all relevant personnel.

This methodology was also used successfully in the development of GN&C for the Mars Exploration Rovers. Given this successful track record, it was natural that it would also be adopted by MSL.

In order for this process to work, however, there is a set of rules and principles that must be followed strictly or the results could be disastrous.[§] Some of them are:

- Clear and natural interface (sensors, actuators, parameters, commands, telemetry, etc.) definitions between GN&C and the rest of flight software. The interface definitions should play to the typical strengths of the people in their respective disciplines. For GN&C devices, for example, the flight software team should manage its low level communications and data transport, power on/off sequencing and configuration, low level fault-protection. The GN&C team should manage device alignments and scale factors, fault protection that involve analysis of dynamics, or comparisons with other GN&C sensors. Most importantly, however, is that these interfaces allow the GN&C flight software core to fit naturally in all the different testing environments, from EDL Monte Carlo performance simulations to Hardware-in-the-loop ATLO tests.
- Clear and strict coding standards for the GN&C team. In general, GN&C analysts are not professional software writers by training and without proper direction can produce substandard code.
- GN&C analysts should write memory-less difference-equations subroutines, with separate argument structures for inputs, outputs, state (i.e. memory), telemetry, and parameters. GN&C team members are not allowed to create memory.
- An experience member of the flight software team should lead and provide support to the GN&C team on their software effort, and act as a liaison between both groups.
- This software expert should also write the executive and logic engine that calls the subroutines written by the GN&C team, and manage memory, parameters, and telemetry in a consistent and centralized way. This software module then should be integrated with the rest of GN&C in what we call the GN&C Flight Software Core (figure 17). The Core is then delivered back to the GN&C team, and becomes their responsibility to test it as unit before delivering it back to flight software and to the other testing environments.

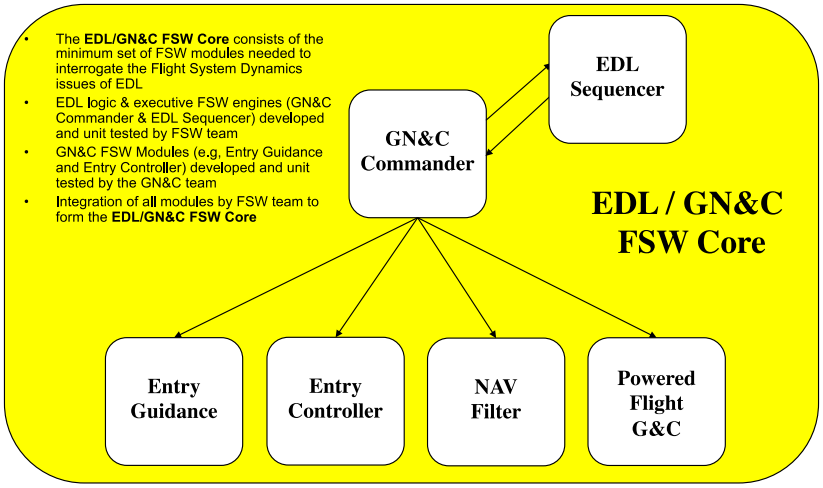

Figure 17. The EDL/GN&C Flight Software Core

[§] Perhaps the more traditional development approaches are not potentially as cost effective and efficient as Mars Pathfinder's, but they are probably more robust to improper implementation or sub-par personnel.

EDL/GN&C VALIDATION AND VERIFICATION (V&V)

EDL has the characteristic that is nearly impossible, and too costly, to do a high fidelity end-to-end test of it here on Earth and, therefore, the first time is tested in its entirety is on Mars and on landing day. The full sequence is automated and, unlike in other phases of the mission, there is little room for error, for there is no ground-initiated recovery possible.

It is for the reasons above that analysis and simulations using test-correlated hardware models become the main tool of the Validation & Verification campaign. A key factor for the success of this approach is the characterization of the key hardware with high-fidelity unit tests, and then the use of the test results to generate high fidelity models for analysis and simulations.

The main GN&C relevant hardware elements that received particular attention and emphasis were the TDS, the MLE's, the RCS, the DIMU, the structure flexibility, capsule aerodynamics, fuel slosh, and SkyCrane bridle flexibility. For each one of these elements, extensive hardware testing and characterization were conducted.

The testing of the TDS can be given as an example of the large commitment by the project to test and characterize the elements of EDL and GN&C that form the building blocks of the EDL/GN&C simulations. To test this important GN&C sensor, an extensive, and expensive, Field Test campaign was implemented to test the radar under as much of the flight dynamic envelop and conditions (altitude, velocity, angular rates, view angles, terrain RF characteristics, etc.) as possible. This required multiple testing venues and platforms to cover this wide range of envelopes, such as an F18 fighter airplane, to test the high altitude and high velocity part of the envelop, a helicopter to test the middle and lower altitudes and velocity ranges, and a tower to test the last 100meters of flight, including the SkyCrane (Figure 18) [11].

Figure 18. TDS Testing Venues for Field Test Campaign

It was also important to test EDL/GN&C under all the environmental conditions encountered on Mars (atmospheric density, winds, terrain, etc.) and with the highest fidelity models of the spacecraft hardware and of the EDL/GN&C software. For this purpose, a high fidelity Monte Carlo simulation was developed by the Langley Research Center using the POST (Program to Optimize Simulated Trajectories) simulation engine. This Monte Carlo simulation incorporated the EDL/GN&C Flight Software Core described above, and thus fulfilled the double duty of testing the performance and functionality of the algorithms and their flight software implementation. In other words, the POST Monte Carlo simulation became an important EDL/GN&C flight software debugging tool. This resulted in fewer software bugs creeping into other testbeds with more and higher fidelity spacecraft avionics, where bugs are more difficult and costly to find, analyze, and fix.

Finally, to make sure that the EDL/GN&C Core worked with the rest of the flight software, avionics, and spacecraft hardware, a series of testbeds were implemented with varying degrees of increasing fidelity. The highest fidelity one incorporated an unprecedented degree of hardware-in-the-loop capability: TDS Radio Frequency (RF) stimulation, RCS driver signal detection, MLE valve position sensing, and pyro and power driver sensing. The DIMU signal was stimulated through a special port and did not require the de-matting of the unit. This simulation setup was also utilized during flight spacecraft Assembly, Test, Integration, and Launch Operations (ATLO).

MSL LANDING RESULTS

MSL landed successfully in Gale Crater on August 5, 2012, only 2.2 km east from its target, thus successfully demonstrating the first use of Entry Guidance on Mars (Figure 19). The SkyCrane, the other new element in this EDL, also performed flawlessly [9]. Detailed analysis of EDL telemetry do not indicate any significant problems or shortcomings with the system design, development, and testing.

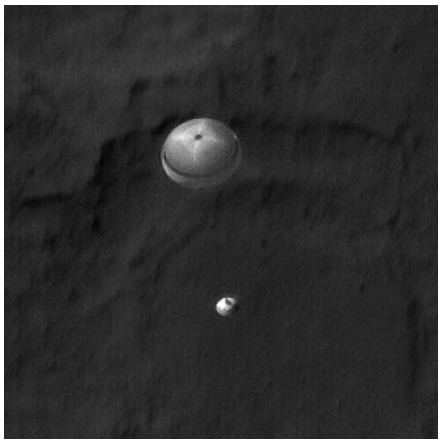

Figure 19. Mars Science Laboratory during Successful EDL August 5, 2012

HiRISE Image Credit: NASA/JPL/University of Arizona

CONCLUSIONS

This paper described the EDL/GN&C requirements and technical challenges involved in landing an almost one-ton rover in a small landing area inside Gale Crater. It described the GN&C architecture and high-level design highlighting the first-ever Guided Entry on Mars and the revo-

lutionary SkyCrane landing technique. It also gave a brief description of the algorithm and flight software development strategy and V&V approach.

The successful landing of MSL/Curiosity on August 5, 2012, validates the EDL/GN&C design, the development, and testing methodology employed. Analysis of EDL telemetry indicates that the system performed with ample margins thus paving the way for the application of this new landing architecture in future exciting missions.

ACKNOWLEDGEMENTS

The work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

© 2013 California Institute of Technology. Government sponsorship acknowledged.

REFERENCES

[1] N.A. Holmberg, R. P. Faust, H.M. Holt, "Viking '75 Design and Test Summary, Volume 1, Lander Design," NASA Reference Publication 1027, November 1980

[2] A.M. San Martin, E.S. Bailey, "The Mars Exploration Rover Transverse Impulse Rocket System: An Active System to Compensate for Martian High Winds During Landing," 27th Annual AAS Guidance and Control Conference, Breckenridge, CO, February 2005.

[3] A. Johnson, R. Willson, J. Goguen, J. Alexander, D. Meller , "Field Testing of the Mars Exploration Rovers Descent Image Motion Estimation System," *Proc. IEEE Int'l Conf. Robotics & Automation (ICRA'05)*, April 2005.

[4] M.R. Grover III, B.D. Cichy, P.N, Desai, "Overview of the Phoenix Entry, Descent and Landing System Architecture." AIAA Paper 2008-7218, Proceedings of AIAA Guidance, Navigation, and Control Conference, Honolulu, HI, August 2008.

[5] Mendeck G. F. and Carman G. L.,"Guidance Design for Mars Smart Landers Using The Entry Terminal Point Controller," AIAA 2002-4502, August 2002.

[6] Brugarolas P.B., San Martin A.M., Wong E.C., "The Entry Controller for the Mars Science Laboratory," AAS/AIAA Space Flight Mechanics Meeting, Kauai, HI, Feb. 2013

[7] A.A. Dyakonov, M. Schoenenberger, W.I. Scallion, J.W. Van Norman, L.A. Novak, C.Y. Tang, "Aerodynamic Interference Due to MSL Reaction Control System." *41st AIAA Thermosphysics Conference,* June 2009, San Antonio, Texas

[8] G. Singh, A.M. San Martin, E.C. Wong, "Guidance and Control Design for Powered Descent and Landing on Mars," IEEEAC paper #1548, Version 2, Updated December 11, 2006

[9] S. Sell, A. Chen, J. Davis, M. San Martin, F. Serricchio, G. Singh, "Powered Flight Design and Reconstructed Performance Summary for the Mars Science Laboratory Mission," AAS/AIAA Space Flight Mechanics Meeting, Kauai, HI, Feb. 2013

[10] F. Serricchio, A.M. San Martin, E.C. Wong, "Mars Science Laboratory Navigation Filter," AAS/AIAA Space Flight Mechanics Meeting, Kauai, HI, Feb. 2013

[11] James F. Montgomery, James H. Bodie, Joseph D. Brown, Allen Chen, Curtis W. Chen, John C. Essmiller, Charles D. Fisher, Hannah R. Goldberg, Steven W. Lee and Scott J. Shaffer, "Implementing the Mars Science Laboratory Terminal Descent Sensor Field Test Campaign," Proceedings of the 35th Annual AAS Rocky Mountain Guidance and Control Conference, AAS 12-091, pp. 647-664, Breckenridge, CO, February 2012.

[12] A. McEwen, E. Eliason, J. Bergstrom, N. Bridges, Candice J. Hansen, W. A. Delamere, J. Grant, V. Gulick, K. Herkenhoff, L. Keszthelyi, R. Kirk, M. Mellon, S. Squyres, N. Thomas, C. Weitz, "MRO's High Resolution Imaging Science Experiment (HiRISE)," Journal of Geophysical Research – Planets, 2006