Guidance and Control Design for Powered Descent and Landing on Mars

Gurkirpal Singh, Alejandro M. SanMartin[3], Edward C. Wong[4]
Jet Propulsion Laboratory,
California Institute of Technology,
4800 Oak Grove Dr.,
Pasadena, CA 91109
(818) 354 3291
Gurkirpal.Singh@jpl.nasa.gov

Abstract—Future Mars landing missions must be capable of autonomously delivering highly capable and mobile rovers safely and gently in an upright orientation. The airbag landing system used to deliver earlier rovers (Mars Pathfinder and the two Mars Exploration vehicles) is incapable of landing the Mars Science Laboratory (MSL)-class rover. The design of a novel Sky-Crane landing concept to land the proposed Mars Science Laboratory rover is presented here. The descent is guided and actively controlled in six degrees of freedom. Terminal guidance is robust to terrain variations-induced altimeter noise. A Terminal Descent Sensor provides surface relative velocity and altitude measurements, the Inertial Measurement Unit measurements help propagate the vehicle attitude and positions. Guidance and control system commands eight throttle-able Mars Lander Engines to actively control the vehicle attitude and translations. Computer simulations demonstrate the viability of this concept in the presence of various environmental, configuration, and hardware imperfections. [1,2]

TABLE OF CONTENTS

1———————————

2 IEEEAC paper #1548, Version 3, Updated January 5, 2007
3 Chief Engineer, MSL GN&C
4 Task Lead, MSL GN&C Algorithms Development

Acronyms

CD	Constant Deceleration
CV	Constant Velocity
DS	Descent Stage
EDL	Entry Descent and Landing
GNC	Guidance, Navigation, and Control
IMU	Inertial Measurement Unit
MLE	Mars Lander Engine
MSL	Mars Science Laboratory
PA	Powered Approach
PDV	Powered Descent Vehicle
SC	Sky-Crane
TDS	Terminal Descent Sensor

1. INTRODUCTION

The Mars Science Laboratory (MSL) mission would deliver a large rover to the surface of Mars using a novel landing concept. The Sky-Crane landing approach has never been used before and offers several advantages, not the least of which the ability to deliver the Rover in a fully deployed, traverse-ready configuration. It obviates the need for an egress system and the related cost and reliability issues. The highly successful airbag landing concept previously used to deliver rovers to Mars is not capable of landing the MSL rover. Several other landing concepts, including the traditional legged landing, were explored before settling on the Sky-Crane approach.

MSL is required to deliver the rover in an upright attitude with zero lateral and a prescribed vertical velocity component. Rover-less vehicle, referred to as the Descent Stage, contains the propulsion system. In the Sky-Crane approach the lander enters a vertical descent phase before touchdown. A few seconds before touchdown, the rover separates from the rest of the lander. When separated, the two bodies remain temporarily connected with a triple-bridle and a data umbilical (Figure 1). The translation motions of the two bodies are coupled in the separated, 2-body configuration by virtue of the fact that they are roughly equal in mass. This coupling allows some control of rover lateral and vertical velocities.

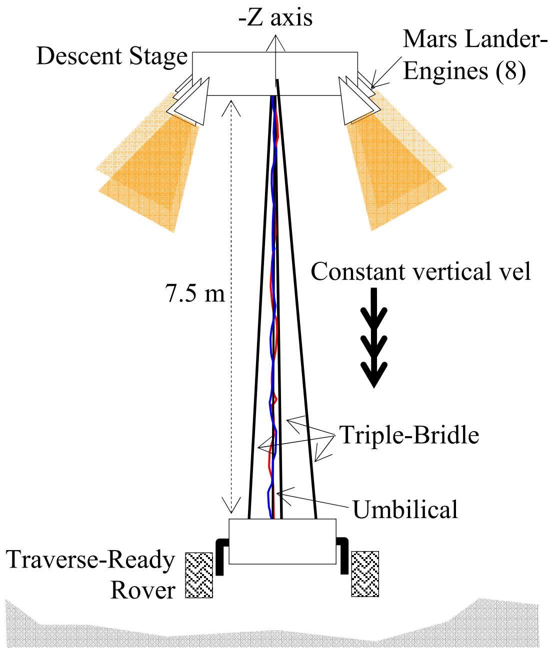

Figure 1. Sky-Crane Phase Configuration: Rover Separated from the Descent Stage (Top)

Terminal guidance is the key in all controlled landing applications. Several, previously used approaches [1], [2] were explored as candidate methodologies for MSL. For instance, the gravity turn guidance employed by Viking [2] does a poor job of controlling the rover motions in the two-body Sky-Crane phase. Forcing the Descent Stage to remain on a fixed vertical path (path following) yields far better results. Better damping of the rover pendulous motion is realized in this fashion, which helps minimize the lateral touchdown velocity component.

2. MISSION DESCRIPTION

The proposed mission would place a large rover on the surface of Mars. According to the current plans, following a guided entry segment , a supersonic parachute is inflated to slow the vehicle down to speeds at which the terminal descent segment may begin. The vehicle at the start of terminal descent is made up of the Descent Stage (DS) and the Rover, rigidly attached together, and referred to as the Powered Descent Vehicle (PDV). The propulsive elements, eight throttle-able Mars Lander Engines (MLEs), are located on the DS. DS –Z axis points up. The eight MLEs are arranged in clusters of two in the four XY quadrants. All MLEs have an outward cant angle to avoid plume-impingement on any part of the Rover. This cant angle also enables control of rotations about the vehicle Z axis

In the current plans, the rover processor houses all flight algorithms. Both the Terminal Descent Sensor (TDS) and the Inertial Measurement Unit (IMU) are located on the DS. Several events transpire while the PDV is still attached to the parachute. Heat shield is jettisoned so that the downward looking Terminal Descent Sensor (TDS) may be exposed to the Martian terrain. The TDS provides surface-relative altitude and velocity measurements. State estimates derived from these measurements are used to trigger a pyrotechnic event which separates the PDV from the backshell and the attached parachute. Prior to this separation event, all 8 MLEs are primed and warmed up in anticipation of start of the terminal descent segment. The MLEs are warmed-up at about 2.5 km above the Martian surface. Terminal descent start altitude is a function of surface relative velocities. This phase of the mission begins when the lander is approximately 2 km above the Martian surface. The lander is commanded in six degrees of freedom to follow a fixed, profiled descent trajectory. Horizontal velocity is removed first, following which the lander enters a vertical descent phase at a designated altitude. During this segment a controlled separation of a fully deployed, traverse-ready rover from the rest of the PDV (the Descent Stage) takes place. A bridle, measuring 7.5 m when fully deployed, connects the two bodies. The bridle is fully deployed before the rover actually makes contact with the surface at a (small) nominal vertical and 0 horizontal velocity component. The PDV with the fully deployed Rover is shown in Figure 1. After depositing the rover on the ground and touchdown confirmation, a bridle cut is performed and the descent stage performs a controlled fly-away to land at least 200 meters away from the Rover.

3. GUIDANCE, NAVIGATION AND CONTROL ARCHITECTURE

Figure 2 depicts the terminal descent Guidance, Navigation, and Control architecture. The primary algorithms are: State Estimator, Guidance (*aka* Trajectory Commander and Control), Attitude Commander and Control, and Thrust Allocation Logic. All functions are executed regularly in a single rate group (in the 60 – 70 hz range).

The State Estimator function provides surface-relative position and velocity estimates, inertial-relative attitude and angular rate estimates, and the coordinate transformation from the surface frame to the inertial frame to the rest of GNC entities. The State Estimator starts propagating attitude and angular rate estimates before entering the Martian atmosphere. Attitude propagation is subsequently driven only by the IMU data. The position and velocity state propagation is initially in the inertial frame. The initial conditions for this propagation are provided by a ground uplink prior to atmospheric entry. Once the TDS data becomes available, the position and velocity state propagation then switches to a surface relative frame defined once. As opposed to the rest of the functions depicted in Figure 2, the State Estimator function is executing before atmospheric entry. The rest of the

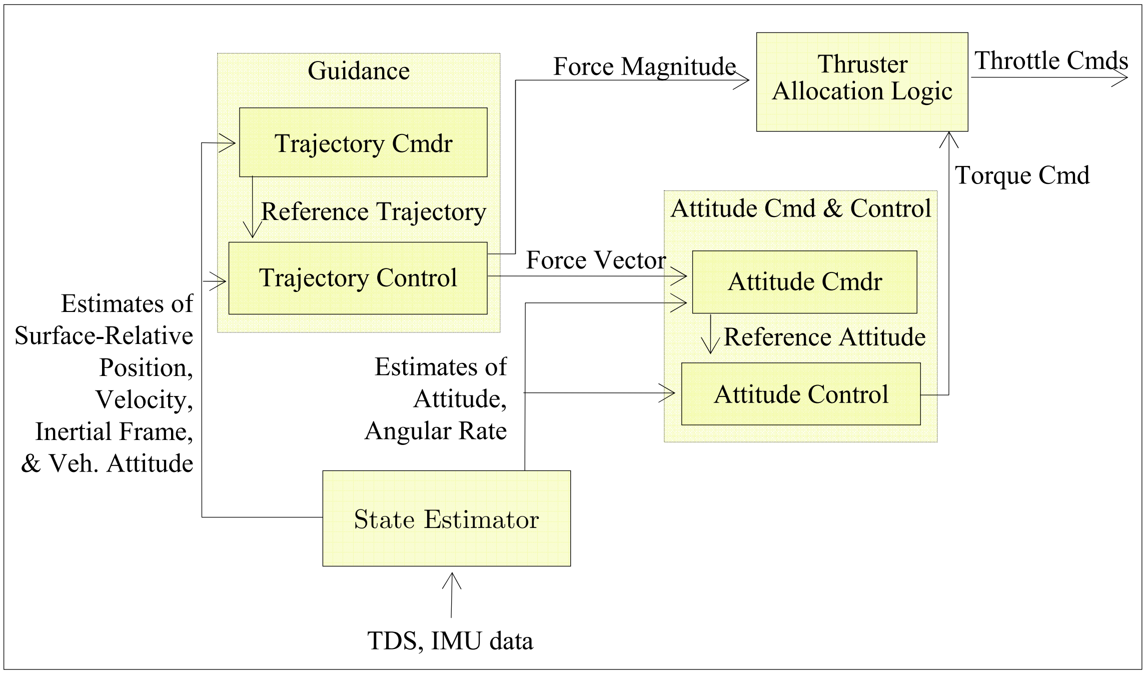

Figure 2. Guidance, Navigation, & Control Architecture

GN&C machinery is engaged when the backshell is separated from the PDV.

Guidance (Trajectory Command and Control) commands the back-shell separation so that terminal descent may begin at the appropriate altitude. Naturally, the TDS data (surface-relative altitude and velocity) must be available prior to this event. The principal function of Guidance is to establish a surface-relative reference trajectory (how this is done will be described in Section 4), and follow it by commanding an appropriate force vector. Estimates of surface relative position and velocity are needed to close the loop around the reference trajectory. Since the actuators (MLEs) are body fixed entities, the reference trajectory is followed by commanding an appropriate force vector. Internally, Guidance computes the desired force vector in Surface-Fixed Frame, which is then resolved in the Inertial Frame using the State Estimator provided Surface-Frame-to-Inertial-Frame transformation. Since this force is to be applied along the vehicle –Z axis, the vehicle attitude must be such that –Z axis is aligned with the Guidance-commanded force vector. The realization of the appropriate attitude is the responsibility of the Attitude Command and Control functions. Attitude Commander computes a reference attitude which allows the vehicle thrust axis (-Z Body axis) to be pointed in the direction of the inertial force vector commanded by Guidance. It profiles a turn to this attitude in the event that there is a large offset between the Guidance-desired attitude and the current attitude estimate. The reference attitude is passed on the Attitude Control function, which computes an appropriate torque value such that the errors between the reference attitude and rate, and the State Estimator provided estimates are minimized. The torque desired by the Attitude Controller and the magnitude of the force desired by Guidance are provided to the Thrust Allocation Logic function. The logic computes realizable throttle settings such that the resulting force and torque agrees as best as possible with the commanded values.

Before leaving the subject of the GN&C architecture, it is important to address the subject of altitude uncertainty and how MSL GN&C architecture intends to address it. Note that for MSL the intent is to deliver the Rover payload with a small nominal vertical and zero horizontal velocity. The precise touchdown location inside the landing ellipse is not pre-determined, rather it is established by the on-board algorithms. Terrain variation induced altimetry error is a problem that must be dealt with by every controlled landing application. These errors arise from two sources. The first is simply the range-dependant measurement error, which improves as altitude decreases. The second more troublesome element is a function of terrain variations, exacerbated by the fact that, in general, altimetry measurements made to initiate terminal descent may illuminate surface locations far removed from the eventual landing site. A large number of simulation were run to get a sense of the variation in this distance, i.e. the distance between the surface location illuminated by the TDS at back-shell separation altitudes and the Rover touchdown location. Simulations suggest that this distance is less than 500 m (3σ). MSL is therefore expected to be tolerant to terrain and slope variations on a scale of 500 m.

Vehicle horizontal velocity is removed first during the Powered Approach (PA) segment. Vehicle attitude at the end of the PA phase is constrained such that the TDS is looking straight down at some surface location. Following the PA segment the vehicle is forced to follow a controlled vertical descent trajectory to this location. A near-steady observation of this location by the TDS yields a relatively noise-free altimetry data. During the PA segment however, since the body-fixed TDS may be sweeping over a terrain on a 500 m scale the altimetry data can be noisy, changing as the illuminated terrain dictates. In order to be robust to altimetry noise, the PA segment reference trajectory is followed using a state estimate which ignores altimetry data. This surface relative state, initialized with the altimeter measurement made at the start of terminal descent, is propagated using IMU measurement and a gravity model. Altimetry-inclusive state estimate is used to follow the vertical reference trajectory after the conclusion of the PA phase. Since the PA altitude estimate may be in error by the terrain variations on a 500 m scale, the PA phase targets a terminal altitude which is biased (high) by some nominal design value.

4. GUIDANCE AND CONTROL DESIGN

Guidance has two distinct sub-functions: Reference Path generation (Trajectory Commander) and Reference Path Following (Trajectory Control). We will address the Trajectory Commander function first. The terminal descent trajectory is fixed. It is not recursively computed and is required to pass through three way points in the position, velocity space. The horizontal velocity is zero for all three. The three way points divide the terminal descent trajectory into four segments: the Powered Approach (PA), the Constant Velocity (CV), the Constant Deceleration (CD), and the Sky-Crane (SC) segments. Two purposes are served by the PA segment: the removal of the horizontal velocity, and realization of a prescribed out-of-plane distance. The latter is required in order to avoid landing in the proximity of the back-shell. Targeting a 300 m out-of-plane distance is sufficient to minimize this possibility when winds are taken into account. The reference trajectory is trivial (vertical path) for all but the PA segment, for which it satisfies a two-point boundary value problem in three dimensions. The vertical component, a constant deceleration path, satisfies the following boundary conditions:

$$z(0) = h_0,\ \dot{z} = w_0,\ z(T) = h_0 - \Delta h,\ \dot{z}(T) = w_T, \tag{1}$$

Here T is the duration of the PA segment, h_0 is initial altitude, Δh is the altitude loss during the segment, w_0 is the initial descent rate, and w_T is the desired final vertical velocity. T, the duration of this segment is

$$T = -2\,\Delta h / (w_0 + w_T), \tag{2}$$

and it requires a vertical acceleration:

$$A_z = 2\,(w_0^2 - w_T^2) / \Delta h + g. \tag{3}$$

Note that $w_T < w_0 < 0$ is implied, Δh (> 0) is a function of initial velocities, w_T the targeted vertical velocity is a guidance parameter (-20 m/s currently), g is the acceleration due to gravity. Terminal descent start altitude h_0 is a function of initial velocities, and its determination will be described later on. Let coordinates x and y denote, respectively, the in-plane (or along-track) and out-of-plane (cross-track) positions of the vehicle. The boundary conditions to be satisfied by the two horizontal motion components are the following:

$$\begin{gathered} x(0) = x_0,\ \dot{x}(0) = u_0,\ \ddot{x}(0) = a_{x0}, \\ x(T) = x_0 + u_0 T/2,\ \dot{x}(T) = 0,\ \ddot{x}(T) = 0. \end{gathered} \tag{4}$$

$$\begin{gathered} y(0) = y_0,\ \dot{x}(0) = v_0,\ \ddot{x}(0) = a_{y0}, \\ y(T) = y_0 + d,\ \dot{y}(T) = 0,\ \ddot{y}(T) = 0. \end{gathered} \tag{5}$$

Note the enforcement of zero horizontal velocity and accelerations at the end point. A zero terminal acceleration forces the vehicle to reach the end-point in an upright attitude. The initial acceleration components a_{x0}, a_{yo} are initial attitude dependant. The desired along-track change in position is a function of initial horizontal velocity, and the choice $u_0T/2$ can be shown to be near-minimum fuel. The total out-of-plane displacement is d, the desired divert distance. The segment time T is fixed by the vertical channel considerations (eqn.(2)). A fifth order polynomial (the minimal polynomial) in time satisfies the boundary conditions in each of the two directions, i.e.

$$x = x_0 + u_0 t + a_{x0} t^2 + c_3 (t/T)^3 + c_4 (t/T)^4 + c_5 (t/T)^5. \tag{6}$$

Same form applies in the two directions (x, y) with different coefficients. In the along-track direction we have:

$$c_3 = -1.5 a_{x0} T^2 - u_0 T,\ c4 = 1.5 a_{x0} T^2 + 0.5 u_0 T,\ c_5 = -0.5 a_{x0} T^2. \tag{7}$$

The cross-track coefficients are as follows:

$$c_3 = -1.5 a_{y0} T^2 + 10d,\ c4 = 1.5 a_{y0} T^2 - 15d,\ c_5 = -0.5 a_{y0} T^2 + 6d. \tag{8}$$

Yet another constraint must be satisfied by the Powered Approach trajectory. It is the requirement that the acceleration required to track the path not exceed some prescribed value, e.g. 90% of the total available linear acceleration. The vertical acceleration component is constant, but the along- and cross-track acceleration components are cubic functions of time. Enforcement of this constraint therefore requires off-line computations for the range of expected initial velocities. The PA segment duration T is a linear function of Δh (eqn.(2)). Further, the peak horizontal acceleration components decrease as T is increased. Clearly there is an optimum Δh for which the peak required acceleration is exactly the desired value. Figure 3 makes this dependence clear. The vertical and horizontal velocities at the start of terminal descent are expected to be less than 110 m/s and 50 m/s (3σ), respectively. A 300 m out-of-plane divert, and an acceleration allocation of 90% of the maximum available is assumed in Figure 3.

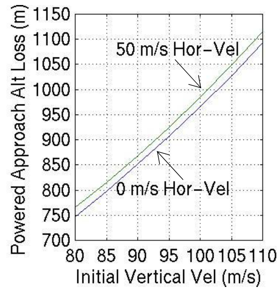

Figure 3. Δh as a Function of Initial Velocities

An example PA reference trajectory time history is shown in Figure 4 (normalized time). Note the constant deceleration z component, zero horizontal (x,y) velocities and acceleration components at end time. Note also the satisfaction of the acceleration constraint that the commanded acceleration magnitude does not exceed 90% of the maximum available acceleration (bottom right plot).

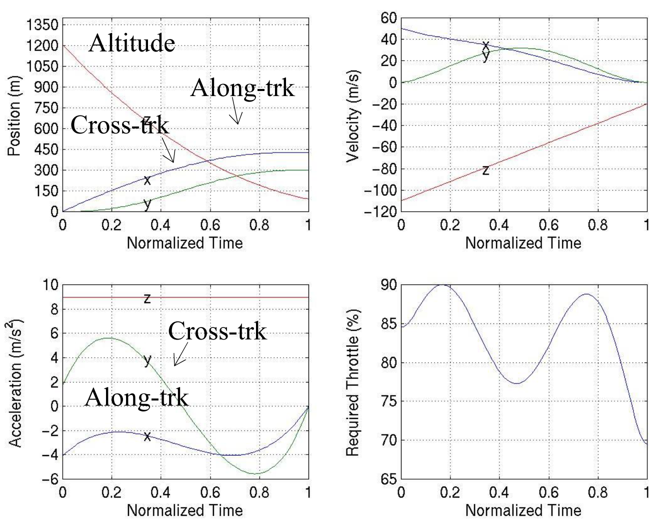

Figure 4. Example Powered Approach Ref. Trajectory

Before coming back to the question of terminal descent start altitude, a brief discussion of the remaining descent segments is in order. These are, in chronological order, the Constant Velocity (CV), the Constant Deceleration (CD), and the Sky-Crane (SC) segments. All three segments follow a vertical trajectory. The SC trajectory has a constant vertical velocity v_{SC} and is entered at a designated altitude of h_{SC}. The Rover separation from the Descent Stage and Rover touchdown are realized during this phase. Immediately before the SC segment lies the Constant Deceleration (CD) segment, which serves the purpose of slowing the vehicle down from an initial vertical velocity of w_T to v_{SC}, the velocity during the SC phase. A constant vertical deceleration is exercised to realize this change in velocity. The Constant Velocity (CV) segment bridges the gap between the CD and PA segments. The CV segment trajectory is also a vertical path traversed at a constant velocity of w_T. The CV segment allows the GN&C system to autonomously correct the altitude errors introduced at the start of the PA segment. It is for this purpose that the CV start altitude is biased high by a nominal design value. Some fuel inefficiency is inherent with this approach, but it is deemed acceptable for it provides significant robustness to altimetry errors and terrain-induced altimetry noise. The required additional fuel is budgeted in the total propellant allocation,.

Returning now to the subject of ignition altitude (or the Powered Approach start altitude), it can be expressed as the following sum:

$$h_{IGN} = h_{SC} + (w_T^2 - v_{SC}^2) / \{2(A - g)\} + \delta h + \Delta h, \qquad (9)$$

where h_{SC} is the SC start altitude, the second term on the right is the altitude loss during the CD phase, δh is the altimetry error allocation, and Δh, the altitude loss during the PA phase, is a function if initial velocities (Figure 3); A is the peak acceleration allocation (90% of the maximum available acceleration) and g is the acceleration due to gravity. A suitable parameterization of the curves shown in Figure 3 is all that is needed to establish the ignition altitude. Note that the variations shown are not particularly sensitive to horizontal velocities. Several possibilities exist here. A particularly simple parameterization is one which is independent of velocities, e.g. Δh = constant ≈ 1100 m. A more fuel efficient parameterization is a straight line which bounds, from above, the two curves shown in Figure 3. The vehicle has sufficient propellant to fly either of the two extremes. Exactly which parameterization will be flown on MSL is not determined at this time, but it is expected that it will lie somewhere in between the two extremes mentioned here. PA phase is initiated when the following condition is satisfied by the estimated altitude;

$$h_0 \le h_{IGN}. \qquad (10)$$

Besides generating the reference trajectory, Guidance also commands a force vector to compel the vehicle to track the reference path with small errors. This is the Trajectory Control function depicted in Figure 2. Let the triple $\{p_{ref}, v_{ref}, a_{ref}\}$ denote the position, velocity and acceleration associated with the reference path which is a function only of time. The reference path is prescribed in the surface relative frame. In order for the vehicle to provide the reference acceleration value, it must align vehicle –Z axis with the reference path acceleration vector and provide a thrust consistent with the reference acceleration magnitude. In other words there exists a reference direction (direction of a_{ref}) and a reference magnitude ($|a_{ref}|$). Removal of path following errors required small corrections to this direction and magnitude. Attitude Commander computes a reference attitude. This is the attitude which aligns the DS –Z axis with the Guidance-commanded force vector. When a large offset exists between this attitude and the estimated attitude, a profiled turn is commanded. A preferred roll orientation does not exist, therefore the desired roll angle is one which minimizes turning about the vehicle Z axis.

The last function to be executed before completing one pass through the terminal descent GN&C logic is the Thrust Allocation Logic. Inputs to this function are the 3-axis torque commanded by Attitude Controller, and the force magnitude commanded by Guidance. The force is assumed applied along the DS –Z axis. The logic computes an appropriate set of 8 MLE throttle values such that the force and torque commands are realized as closely as possible while satisfying the realizability (minimum thrust ≤ thrust command ≤ maximum thrust) constraint. Thrust Allocation Logic solves a Quadratic Programming Problem [4]. To assure a definite computational time the function is exited after a fixed number of iterations. The algorithm is an efficient steepest-descent algorithm, suitable for a real-time implementation.

As planned, MSL does not have a direct means of sensing Rover touchdown; a software logic is employed instead. Consider what happens as the DS and Rover, separated by the triple-bridle continue to descend at a constant rate during the SC segment of the mission. Actually it is the DS which is forced to follow this constant velocity vertical trajectory. It is forced to remain on this path until the touchdown has been confirmed. As the system continues its descent, at some point in time the Rover will make contact with the surface, eventually causing the bridle to go slack as long as the DS can be made to continue on its downwards constant rate motion. A continuation of the DS motion requires less force. It is this change in the force command which is monitored (with a persistence check) to declare touchdown. Upon receiving this indication, a bridle cut is performed to sever the DS – Rover connection. The DS, in order to land at a far enough distance from the Rover, throttles up while executing a turn to place the velocity vector at an optimum angle with respect to the local horizontal. All engines are shut down after a fixed time interval, and the DS coasts eventually crashing on the surface. The desire is to not let this happen within 200 m of the Rover. This last phase of the mission which disposes of the DS is referred to as the Fly-Away phase. In order to realize an acceptable minimum thrust/weight ratio following touchdown, four of eight MLEs are permanently shutdown at the start of the SC phase. The 3-axis attitude control capability is preserved, but the linear acceleration capability is cut in half. The diminished capability however is more than adequate to execute the Fly-Away maneuver.

5. SIMULATIONS

Several simulations are discussed next. The PDV is initially at about 2 km altitude, descending at 110 m/s with a 50 m/s horizontal velocity. Surface-relative motions are depicted in Figure 5. Large angular rates at the beginning represent attitude motions while the parachute is still attached. Powered Approach starts at about the 7 second mark and concludes at 25 sec. Note that the descent rate goes to 20 m/s at the 25 sec mark and off-nadir motions remain small thereafter. SC starts when the vertical velocity reaches the nominal SC velocity slightly before the 30 sec mark. Fly-away maneuver starts at t = 45 sec. Figure 6 depicts control errors in this case. Attitude errors remain well below 5° most of the time. Position and velocity control errors are well below 20 cm, 50 cm/s in this case.

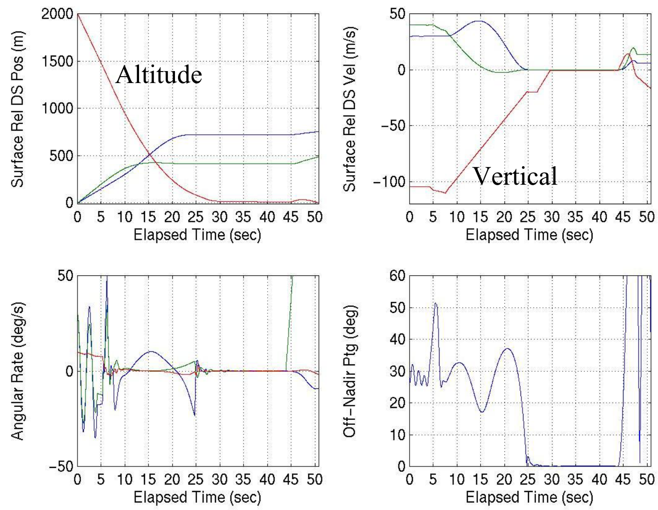

Figure 5. Surface Relative Motions of the DS

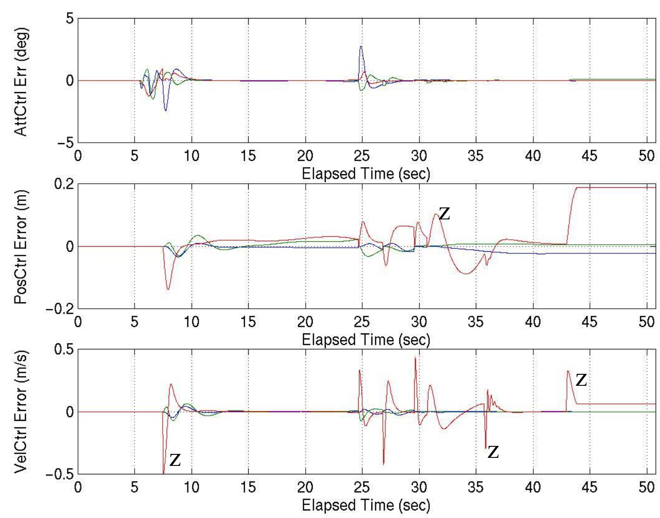

Figure 6. Control Errors Time History

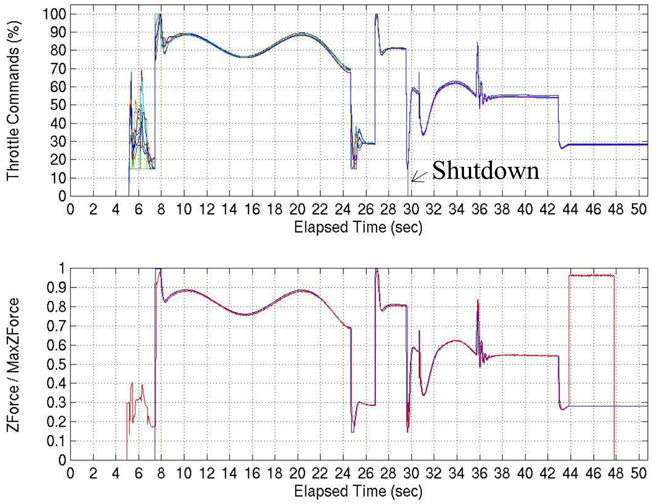

Figure 7. Throttle Commands

Lastly in Figure 7 we depict throttle commands (top) and the –Z force commanded by Guidance. Momentary saturation of the MLEs is perfectly acceptable. Note that control error behavior during saturation events remains benign. All eight MLE commands are shown. Four of the eight are shutdown (commanded thrust goes to zero) at the start of the SC phase at the 30 sec mark. Note also the throttling down of the 4 MLEs in response to Rover touchdown at t = 43 sec. Throttling up immediately thereafter denotes the start of Fly-Away with a final shutdown 4 seconds later at the 48 sec mark.

6. CONCLUSIONS

Design of the terminal descent Guidance, Navigation, and Control algorithms that would be used to land the MSL rover is presented here. A novel Sky-Crane approach would be used to gently land the MSL Rover. Yet another innovation is terrain-accommodating guidance. Several simulations have been carried out to gain confidence in the proposed approach. It has been demonstrated that the proposed approach is indeed robust to terrain variations and would perform as intended.

ACKNOWLEDGMENT

The design and development described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

REFERENCES

[1] A. Klumpp, "Apollo Lunar Descent Guidance," Automatica, Vol. 10, Paragon Press, 1974, pp. 133–146

[2] R.N. Ingoldby, "Guidance and Control System Design of the Viking Planetary Lander," Journal of Guidance and Control, Vol. 1, No. 3, May-June 1978, pp. 189-196

[3] E.C. Wong, G. Singh, J.P. Masciarelli, "Guidance and Control Design for Hazard Avoidance and Safe Landing on Mars," Journal of Spacecrafts and Rockets, Vol. 43, No. 2, March-April 2006, pp.378-384.

[4] G. Singh, "Quadratic Programming for Allocating Control Effort." NPO-40592, NASA Tech Briefs, Vol. 29, No. 8, August 2005, *www.techbriefs.com*

BIOGRAPHY

Gurkirpal Singh is a Principal Engineer in the Guidance and Control Analysis Group. He received his M.S. and Ph.D. in Aerospace Engineering from The University of Michigan in 1984 and 1988, respectively. He has been with the Jet Propulsion Laboratory since 1989. At JPL he has been primarily involved in the design and development of spacecraft guidance and control algorithms for several projects including Galileo, Cassini, Mars Pathfinder, Deep-Space 1, and Mars Exploration Rovers. His areas of expertise include 6-dof guidance and control, multibody dynamics, and inertial vector propagation. He has published in the areas of flexible spacecraft slewing control, real-time simulation of multi-body spacecraft dynamics, pointing constraint avoidance, precision landing guidance, and formation-flying path-planning guidance and control. He is currently developing terminal descent guidance and control algorithms for the Mars Science Laboratory Project.

Miguel San Martin received his B.S. Degree in Electrical Engineering with honors from Syracuse University in 1982, and his M.S. Degree in Aeronautics and Astronautics Engineering from the Massachusetts Institute of Technology in 1985. He joined the Jet Propulsion Laboratory in 1985. His area of interest is the analysis, design, implementation, and testing of spacecraft articulation and attitude control systems, with an emphasis on applied estimation theory. He has participated in several flight projects and has been a member of numerous flight anomaly tiger teams. He was the designer of the Cassini spacecraft Attitude Estimator, the TOPEX/Poseidon Altimeter pointing calibration ground software, and was the technical lead for the Mars Pathfinder Attitude Control Subsystem flight software. More recently he was the Guidance and Control System Manager and Chief Engineer for the the Mars Exploration Rover Project, that successfully landed the Spirit and Opportunity rovers in January 2005. He is currently the Guidance Navigation and Control System Chief Engineer for the Mars Science Laboratory project.

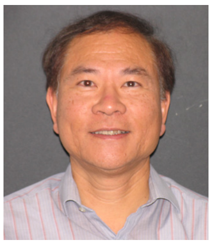

Edward C. Wong is a Principal Engineer and a Technical Group Leader for spacecraft guidance and control design and analysis for flight projects. He received his B.S. degree in Engineering Physics from Oregon State University, his M.S. degree in Electrical Engineering from Brown University, and his doctorate degree in Dynamics Systems Control from UCLA. He joined the Jet Propulsion Laboratory in 1977. He has worked on many JPL projects such as the Voyager, Galileo, Cassini, Space Station Payload Pointing Systems, Space Infra-Red Telescope Facility, Shuttle Radar Topography Mission, Mars Rendezvous and Sample Return, the New Millennium ST-6 Autonomous Rendezvous Experiment in Earth Orbit, and Mars Science Laboratory mission. He is primarily responsible for the design, development, and validation of control algorithms in both flight and ground software for these projects. His primary areas of interest are spacecraft attitude control and attitude determination, system identification, dynamics and modeling, formation fling, and autonomous rendezvous and spacecraft control.