

SYSTEMS FOR PINPOINT LANDING AT MARS

Aron A. Wolf, Claude Graves, Richard Powell, Wyatt Johnson

Mars landers have been able to land only within tens to hundreds of km of a target site due to uncertainties in approach navigation, atmospheric modeling, and vehicle aerodynamics; as well as due to map-tie error and wind drift. The Mars Science Laboratory mission will improve this to 5 – 10 km using optical navigation and entry guidance. To achieve “pinpoint landing” (within 100m) for future missions, ways of addressing the remaining error sources (approach navigation, wind drift and map-tie error) must be found. This work analyzes performance and cost trades for pinpoint landing systems, identifying the most promising areas for technology investment.

INTRODUCTION

Estimated landing accuracy of Mars landers to date has steadily improved, from ~150 km of the target for Mars Pathfinder to ~90 km for Mars Polar Lander to ~35 km for the Mars Exploration Rovers. Factors which contribute to these position uncertainties at landing include uncertainties in our ability to navigate the spacecraft to the desired entry point in the atmosphere, atmospheric modeling uncertainties, uncertainties in our ability to measure vehicle aerodynamic coefficients, map-tie error, and wind drift. The improvement is largely due to improved approach navigation. These landers all flew unguided ballistic trajectories during Entry, Descent, and Landing (EDL). This level of performance has been characterized as “Generation 1”.

The 2007 Phoenix and 2009 Mars Science Laboratory (MSL) mission will advance to “Generation 2”, delivering the lander to within ~5 – 10km of the target, by further improving navigational accuracy at entry with optical navigation and using hypersonic entry guidance to “fly out” the known error at entry as well as the uncertainties in atmospheric modeling and vehicle aerodynamics.

Ways to address the remaining error sources (map-tie error and wind drift) must be found to make the leap to “Generation 3”, or “pinpoint landing” capability, landing within 100m of the target. Pinpoint landing is necessary to reach targets of high science value, especially those within hundreds of meters of landing hazards.

This paper summarizes the results of a study done for NASA’s Mars program analyzing performance and cost trades for pinpoint landing systems in order to arrive at a

preliminary assessment of whether a pinpoint landing system can be developed for a reasonable cost, and to identify the most promising areas for technology investment.

A pinpoint landing system is a collection of elements that together are capable of landing a spacecraft within 100m of the target. The portion of the mission relevant to pinpoint landing can be divided into four phases: approach, hypersonic entry, the parachute phase, and powered descent.

The approach phase starts late in the interplanetary trajectory at the last propulsive maneuver before atmospheric entry and ends at entry interface (typically 125 km, for Mars).

The hypersonic entry phase starts at entry interface and ends at deployment of the parachute (or, if there is more than one parachute, the first parachute).

The parachute phase starts at deployment of the first parachute and ends at ignition of the descent engines. The conditions required for successful parachute deployment in the thin Martian atmosphere (high enough dynamic pressure and low enough mach number) are achieved at altitudes of typically 10 km or less above Mars' surface. Chute deployment must occur high enough to assure enough time for deceleration to a soft landing and completion of all required events before landing. This is especially difficult at higher landing site elevations and in seasons of low atmospheric density, where the atmosphere is thinnest and the demands on the system are greatest.

The powered descent phase starts at ignition of the descent engines and ends at touchdown. Terminal velocity on a parachute in the thin Martian atmosphere is significantly higher than in Earth's atmosphere. Descent rates are typically in the 50 m/s range at Mars with parachutes of reasonable size, requiring the use of thrust to slow the lander's descent enough to permit a successful soft landing.

Study assumptions

For the purposes of this study, it was assumed that the first mission using a Pinpoint Landing system launches in 2018, with a technology development cutoff in 2015. Also, landed mass is assumed to be 1000 kg (MSL - class).

SOURCES OF UNCERTAINTY

To design a system which can successfully accomplish pinpoint landing, it is necessary to identify all the uncertainties which can affect landing precision and to incorporate sufficient control authority to overcome system uncertainty in excess of the 100m requirement. Uncertainties can be classified into three categories:

Control uncertainty is miss distance from the target given perfect knowledge of position and velocity (i.e. the accuracy with which the vehicle can be targeted).

Knowledge uncertainty is the uncertainty in the ability to measure position and velocity. This depends on the method of measurement being used (Earth-based radio navigation, onboard sensors, etc.).

Delivery uncertainty is the actual miss distance from the target, which is a function of both knowledge and control uncertainty.

Estimates of the three types of uncertainties in each of the four mission phases (approach, hypersonic, parachute, and powered descent), and the control authority available to overcome them in each phase, are shown in Table 1. Factors contributing to these uncertainties are discussed below.

Approach phase

Approaching the planet, the trajectory is controlled with propulsive maneuvers. These are subject to maneuver execution errors (uncertainties in thruster force modeling, pointing, etc.) that are generally classified in terms of thrust magnitude and pointing, each having a “fixed” (constant) component and a component proportional to thrust magnitude. Control uncertainty measured at atmospheric entry is equal to execution error at the last maneuver before entry, mapped from the maneuver time to entry.

The spacecraft’s orbit is determined by taking radiometric and / or optical observations and processing them either separately or in combination using a navigation filter. Earth-based radio tracking produces Doppler, range, and Δ DOR (delta-differential one-way range) observations¹, all of which have been used previously during missions to Mars and other bodies. The presence of other spacecraft in orbit at Mars makes it possible to consider tracking between two or more spacecraft (i.e. observations of the lander from an orbiter or vice versa, provided each is equipped with suitable telecommunications equipment). These observations can significantly reduce knowledge uncertainties because they provide direct measurements of the approaching spacecraft relative to the target body and eliminate the need to upload state updates from Earth, allowing for data collection much closer to encounter.

Optical navigation also produces target-relative observations. An optical navigation camera has been developed for use on approach to Mars which utilizes imaging of Phobos and Deimos against a star background². High-resolution optical navigation will be demonstrated for the first time on the Mars Reconnaissance Orbiter mission to be launched in 2005. Another method of optical navigation involves tracking features on Mars in images taken as the spacecraft approaches the planet.

When a maneuver is required to correct the spacecraft’s targeting, navigation data are accumulated until a cutoff time TCO1 before the maneuver. The cutoff time is defined based on the time required for the ground-based operations teams to determine the orbit, design the maneuver to correct the flight path and transmit the required commands to the

Table 1: Sample System Uncertainties and Control Authority

Delivery errors in chute & powered descent phases are RSS of knowledge error and wind drift at 50 m/s wind speed (3.5km, for 70 sec. on chute; 1.0 km for 20 sec. on engines). Assumed that targeting can be biased to account for winds in excess of 50 m/s. One-time pitchdown before deploying chute allows increase in L/D from 0.18 to 0.25 without requalifying chute for deployment at increased angle of attack

spacecraft to execute the maneuver. Accumulation of tracking data typically continues after the cutoff time; however, only data accumulated up to the cutoff time are available for the design of the maneuver. Consequently, delivery uncertainty measured at atmospheric entry includes control uncertainty as well as orbit determination errors mapped from the tracking data cutoff time for design of the last maneuver to entry.

Although tracking data accumulated after the cutoff time for the design of the last maneuver cannot improve delivery uncertainty, it can be used to improve knowledge uncertainty at entry. A final knowledge update is generally planned as late as possible before entry to minimize knowledge uncertainty at entry. A second tracking data cutoff time TCO2 is set after the last maneuver and prior to the late update (several hours before entry) to allow time for computation of this late estimate of knowledge uncertainty on the ground. Knowledge uncertainty at entry then includes control uncertainty from the last maneuver, as well as orbit determination uncertainties at TCO2 mapped to entry. With this late update, knowledge uncertainty at entry is less than delivery uncertainty. Hypersonic entry guidance then “flies out” the difference between knowledge and delivery uncertainties, as discussed in the next section. Figure 1 conceptually illustrates delivery and knowledge errors vs. time, in the approach and hypersonic phases.

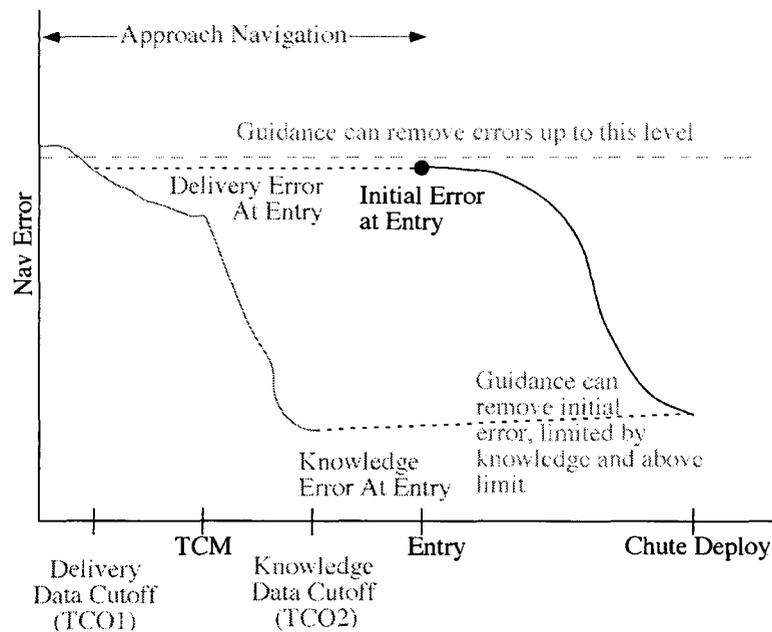


Figure 1 Uncertainties in Approach and Hypersonic Phases

In Mars missions to date, all processing of tracking observations has taken place on the ground. Developing the capability to process tracking observations autonomously onboard the lander has been considered. This could substantially reduce the time required for processing during the last critical hours on approach. The late update could then occur minutes (rather than hours) before entry, significantly improving knowledge at entry.

This capability would have to be validated appropriately to resolve all residual concerns about the operational risk involved.

Hypersonic entry phase

Uncertainties in the ability to navigate the spacecraft to the desired entry point, to model the Martian atmosphere, and to measure the aerodynamic characteristics of the entry vehicle dominate the hypersonic entry phase. All Mars landers to date, including the recently arrived MER mission³, have flown unguided entries in which no control of the trajectory to the desired target was exercised during atmospheric flight. Substantially improving the landing accuracy from that of previous missions, which is necessary to meet the requirements of Generation 2 or Generation 3, necessitates the use of guidance and control during the hypersonic entry phase.

The hypersonic entry guidance scheme developed for use on Apollo during Earth atmospheric entry has been modified for use at Mars, and is included in the baseline for the Generation 2 Mars Science Laboratory mission. The Apollo entry guidance also had the capability to fly a controlled skip-out of the earth's atmosphere with a subsequent reentry of the atmosphere and control to a targeted landing point. This part of the Apollo guidance was never used during the Apollo Program and is not needed for MSL. The MSL entry guidance scheme is a terminal point controller that uses a stored reference trajectory and sensitivity parameters to compute bank angle commands that modulate the aerodynamic lift vector to control the terminal state of the entry trajectory. The magnitude of the bank angle controls the in-plane trajectory and bank angle reversals control the cross-range trajectory until deceleration to about 900 m/s. The MSL entry guidance incorporates some changes for low speed flight that improve the guidance performance compared to the Apollo entry guidance. At speeds below about 900 m/s in the Martian atmosphere, there is insufficient aerodynamic lift to control down-range flight, but the effectiveness of the cross range trajectory control improves significantly. MSL guidance takes advantage of this, switching at 900 m/s to use the bank angle to control the crossrange trajectory so that the vehicle flies over the target position at parachute deployment. At the appropriate point, the parachute is deployed as a final energy management device for downrange trajectory control.

Performance of MSL entry guidance is shown in Figure 2. This figure shows the results of a Monte Carlo simulation of the MSL flight from the entry interface until parachute deployment for 2000 simulated entry profiles. These simulations include the effects of delivery and knowledge uncertainty at the entry interface using radio and optical tracking, variations in the Mars atmosphere, uncertainty in the aerodynamic characteristics of the entry vehicle, and navigation uncertainties during atmospheric flight for an entry vehicle that has sufficient aerodynamic maneuverability to compensate for the these trajectory dispersions. The scatter plot in the upper left of this figure shows "navigated" down-range and cross-range position (i.e. position estimated onboard) relative to the target position at parachute deployment. Control uncertainty in downrange and crossrange position at parachute deployment is only 26 m (3σ); however, as illustrated in the lower right plot, a delivery uncertainty of about +/- ~3 km remains in

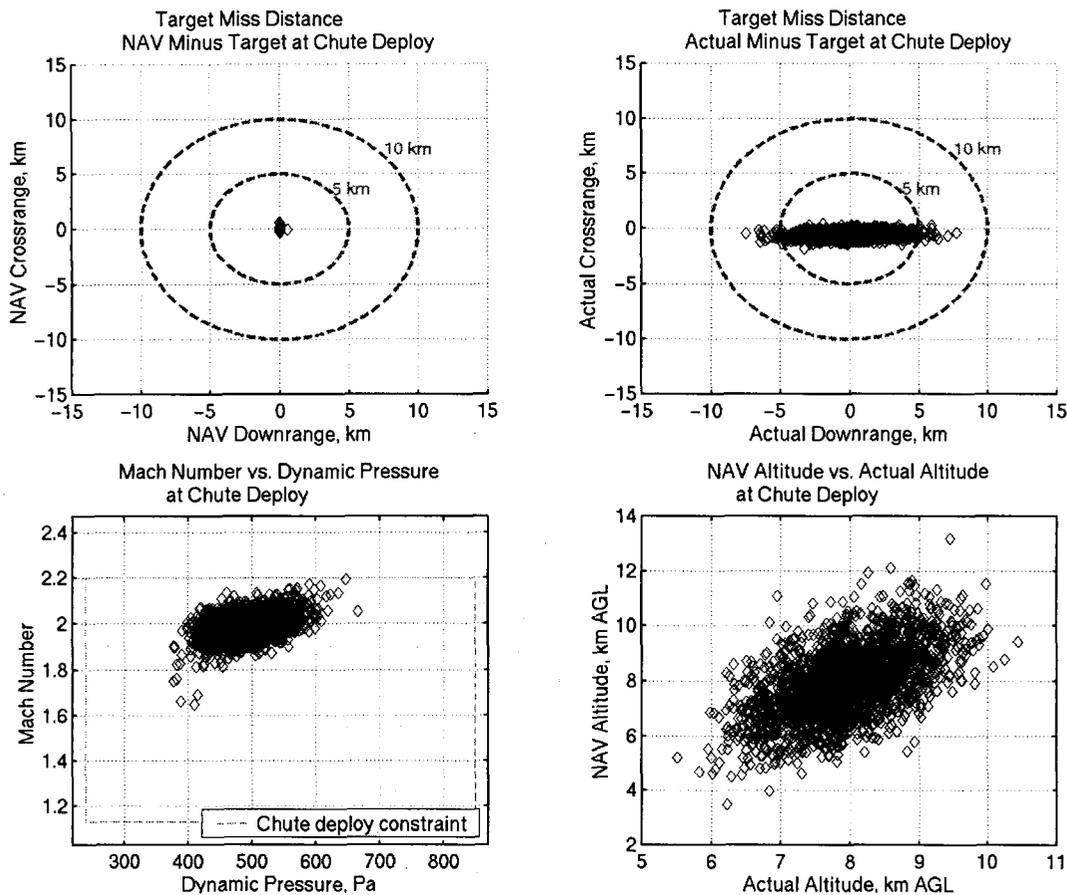


Figure 2 Monte Carlo results for a typical MSL EDL simulation

geometric altitude. The entry guidance controls to density altitude at parachute deployment rather than to a geometric altitude, just as an airplane controls to pressure altitude, and this large geometric altitude delivery uncertainty is caused primarily by variations in the Mars atmosphere and in the aerodynamic drag coefficient.

Onboard guidance can only correct errors it knows about, and as a result, real-world delivery errors perpendicular to the radius vector at parachute deployment are roughly equal to the knowledge error at this point. This is illustrated in the scatter plot in the upper right of Figure 2 that shows a 3σ delivery error of 6.50 km. The corresponding 3σ knowledge uncertainty is 6.48 km, illustrating that knowledge error dominates delivery error at parachute deployment (provided that the entry vehicle has sufficient aerodynamic control authority to compensate for the delivery error at entry, as is the case here). The plot of Mach number vs. dynamic pressure in the lower left corner illustrates the ability to control the MSL entry trajectory to the parachute deployment envelope illustrated by the dashed line. Potential enhancements to MSL entry guidance are discussed in a later section of this paper.

Parachute phase

Wind drift is the dominant error source in the chute phase. As the lander slows to terminal velocity on the parachute, winds begin to significantly influence its flight path. Winds from the surface to ~10km are not well characterized, and variations in both direction and magnitude can be large. Kass⁴ estimates the 1σ wind velocity over the entire planet to be in the range of 20 - 30 m/s; 3σ velocity over the whole planet is in the 70 m/s range. At these speeds, winds during descent on the chute (typical duration ~100 sec) can carry the lander several km from the intended target.

Sufficiently predictable winds would allow targeting of chute deployment upwind of the desired landing site, letting wind drift the lander to the target. However, low-level Martian winds are not well characterized. Predictability of both wind velocity and direction varies with topography. Both direction and velocity are more predictable, and wind speed is greater more of the time, in more mountainous regions. However, work beyond the scope of this study is required to quantify the predictability of low-level winds under various conditions. For this study, wind velocity was taken to be constant at 50 m/s, with no directional preference. Winds in excess of 50 m/s were assumed to be predictable, and could be compensated for by targeting chute deploy to occur upwind of the intended landing site.

Two design approaches may be taken in the chute phase. One is to use a steerable parachute or other aerodynamic decelerator to counter wind drift, with a sensor capable of quickly detecting drift. The other is to use a traditional non-steerable chute (which entails accepting the wind drift on the chute) and to make the necessary correction in the powered descent phase using added propulsion system capability (fuel, engines) to fly back to the target. It is also possible to mix the two approaches, providing a steerable chute with enough control authority to counteract some, but not all, of the wind drift with some augmentation of the propulsion system to counteract the rest.

Map-tie error is the uncertainty in our ability to predict the position of a feature on the surface of Mars in inertial space. The accuracy with which we can locate a surface feature in inertial space is worse than the accuracy to which we can locate features with respect to each other on the Martian surface. Since images taken from orbit are used to locate features in inertial space, our ability to locate features in inertial space depends on instrument pointing, orbit determination, and other errors associated with the images taken from orbit. Gaskell⁵ has reduced map-tie error to ~500m – 1km for much of the Martian surface. Map-tie error is projected to be in the ~100m range by 2018. Although this is a notable improvement, map-tie error must be reliably reduced to less than 100m in order to meet the requirement to land within 100m of the target. This makes target-relative navigation necessary during the final descent, starting in the chute phase in order to determine the desired direction of flight in the powered descent phase before the start of powered descent.

Powered descent phase

If wind drift has not been compensated for in the parachute phase, the lander must fly back to the target under power. This requires sufficient control authority with low enough control uncertainty to overcome the accumulated wind drift of up to several km. A large increase in control authority (provided by fuel and possibly other propulsion system components) is required over MSL, which carries only enough fuel to dissipate vertical speed and to maneuver ~100m to avoid surface hazards.

PERFORMANCE AND COST OF CANDIDATE TECHNOLOGIES:

Preliminary estimates of cost to develop several candidate technologies for Pinpoint Landing were made as part of this work. Table 2 summarizes these cost estimates. Factors contributing to cost are discussed in the subsections below.

Table 2
Cost estimates for candidate Pinpoint Landing technologies

System Element	Assumed cost (\$M)	Comments
S/C-S/C UHF doppler, onboard processing from E-2hrs to touchdown	\$5.0	UHF LGA Upgrade + Onboard S/W + OPS + Entry Guidance Interface + OPS + Test
Terminal descent imager w/ feature tracking	\$7.0	Descent camera + processor
Requalify chute to deploy at higher angle of attack (needed for bank modulation if $L/D > 0.18$), or behind Apollo aeroshell (needed for $0.25 < L/D < 0.3$)	\$75.0	Chute qual costs assumed essentially same as qual for higher mach
Bank + AOA modulation using an aero control surface (e.g. flap or tab), $L/D 0.18 - 0.25$	\$70.0	No need to requalify chute. More accurate control of chute deploy alt.
Bank + AOA modulation using movable mass inside aeroshell, $L/D 0.18 - 0.25$	\$10.0	No need to requalify chute. More accurate control of chute deploy alt.
Bank modulation + one-time AOA change using jettisonable mass, $L/D 0.18 - 0.25$	\$1.0	No need to requalify chute. Not expected to provide control of chute deploy alt.
New aeroshell for $L/D > 0.3$, with increased volume	\$320.0	
Guided subsonic chute	\$10.0	Analogy to MER chute development
MSL pwred desc guidance + added prop.	\$3.0	
New pwred desc guidance + added prop.	\$5.5	

Sensors

The knowledge uncertainties shown in Table 1 are achieved with either onboard sensors or radiometric tracking. The navigation strategy that minimizes the number of sensors uses radiometric tracking and optical navigation on approach, with an IMU driving hypersonic entry guidance after atmospheric entry and a descent sensor capable of terrain-relative navigation after the heatshield is jettisoned. Terrain-relative navigation is needed in the chute and powered descent phases (starting at ~10 km) in order to eliminate map-tie error. Radiometric observations (even those from spacecraft to spacecraft) are subject to map-tie error because they can be interpreted only in inertial coordinate systems and not in target-relative coordinate systems.

Cameras which image in visible wavelengths can be used to sense terrain at low altitudes, with target-relative position computed by comparing the images with predicted images produced on the ground and uplinked to the lander prior to EDL. As part of preparation for EDL, images can be made on the ground which closely resemble the appearance of terrain by constructing a digital terrain map from images previously taken from orbit, and simulating the lighting which will be encountered at landing⁵. Use of a visible camera constrains the mission to daylight landing when the sun is between approximately 15 – 70 deg. above the horizon⁶.

Phased array terrain radar and lidar are capable of accumulating 3-dimensional images of the terrain and comparing these to stored onboard digital terrain maps of the target area prepared on the ground. A phased array terrain radar is currently planned for use on MSL for hazard avoidance. Development of this instrument for MSL considerably reduces the cost of development for pinpoint landing, although some algorithm development is necessary to add pinpoint landing capability to algorithms originally developed for hazard avoidance only. A preliminary cost estimate of ~\$6M was made for building a copy of the MSL radar, assuming test equipment is kept together from the MSL build. A preliminary estimate to rebuild the MSL radar, reassemble the test equipment, and test the radar is on the order of ~\$10M.

Additional sensors / sensing methods add system robustness and can improve knowledge uncertainty as shown in Table 1. Spacecraft-to-spacecraft radiometric observations can be made, and / or images can be collected, from approach to landing (although imaging through the heatshield requires special accommodations⁷).

Enhancements to entry guidance and control

The MSL mission baseline includes entry guidance with bank-only modulation (no direct drag modulation or modulation of angle-of- attack). As previously discussed, EDL simulations show that when bank-only modulation is used in the hypersonic entry phase, the 3σ uncertainty in chute deployment altitude is ~ +/- 3 km, primarily due to uncertainties in atmospheric density and vehicle drag coefficient. This means that the nominal chute deployment altitude must be chosen to provide 3 km of margin above and below the nominal.

Using limited angle-of-attack modulation in addition to bank angle modulation offers two advantages of particular interest to this study, compared to using only bank modulation:

- It allows higher angle-of-attack (and therefore higher L/D) with the MSL aeroshell configuration while maintaining the angle-of-attack at less than the current 15 deg. limit at parachute deployment.
- It has the potential to improve control of geometric altitude at parachute deployment when combined with good knowledge of geometric altitude early during the entry flight.

It also increases responsiveness in controlling aerodynamic drag and provides additional modulation of the in-plane L/D. This is especially important for landing at high altitude and during seasons that have low atmospheric density. It also allows an increase in landed mass without requiring an increase in the aeroshell size. The Space Shuttle Orbiter successfully uses both bank angle and angle-of-attack modulation for entry trajectory control, but the implementation logic would be different for the Mars EDL to enable better control of geometric altitude at parachute deployment. The concept for controlling geometric altitude with angle-of-attack modulation has been defined analytically but has not been verified or quantified in a closed-loop simulation.

Angle-of-attack modulation can be implemented in several ways. Use of propulsive thrusters to provide a torque to overcome the aerodynamic characteristics is expected to cause significant propellant consumption and may result in the need for larger thrusters. A one-time pitch down maneuver just before parachute deployment using an inexpensive jettisonable mass could be used but may not provide the modulation needed to control geometric altitude. Vehicle center-of-mass control or movable aerodynamic control surfaces are also possibilities for providing the angle-of-attack modulation and should be investigated. Angle-of-attack modulation complicates the vehicle design and is an immature concept for Mars EDL, but it should be assessed as a way for improving the Mars EDL capability for high altitude landing with low atmospheric density and for improving the landed mass capability. All of the above concepts need to be evaluated in more detail if angle-of-attack modulation is used.

The vehicle L/D is determined by the angle-of-attack that is controlled by offsetting the center-of-mass. Practical limitations on moving the center-of-mass limit the angle-of-attack to the value corresponding to an L/D of ~ 0.24 for the Viking aeroshell. However, this value of angle of attack is greater than the current limit of 15 deg. for which parachute deployment has been qualified. The Apollo aeroshell (not yet used at Mars) can achieve a maximum L/D of 0.30, for which the angle of attack is also greater than the 15 deg. parachute deployment limit. For L/D values which require angle of attack in excess of 15 deg., the parachute must be qualified for deployment at the higher angle of attack, or angle of attack must be reduced before chute deployment. This can be accomplished with any of the angle-of-attack modulation methods discussed above, of which the least expensive and complex is a mass jettison leading to a one-time pitch-down immediately before chute deployment.

Aerodynamic decelerators

Wind drift, the dominant error source remaining after chute deployment, can be countered using either steerable aerodynamic decelerators in the chute phase (parafoils, steerable parachutes), or extra propellant for extended powered flight to the target in the powered descent phase, or a combination of the two.

Preliminary analysis raised serious doubts about the feasibility of parafoils, which have never been flown at Mars entry-like flight conditions. The large (~ 700 m²) parafoil developed as a prototype for the X-38 Crew Rescue Vehicle can be taken as a point of

comparison^{8,9}. The X-38 parafoil deploys at a dynamic pressure of 480 Pa on Earth, which equates to 283 m/s, or Mach 1.2, in Mars' thin atmosphere. This study found no evidence of previous experience in deployment or flight of parafoils at supersonic speeds. Mach 0.5 at Mars corresponds to a dynamic pressure of ~81.3 Pa, much lower than required to deploy the X-38.

Even if the deployment problem can be solved, cruise flight at Mars requires either a large parafoil or high cruise speeds. The X-38 was flown at a minimum dynamic pressure of 278 Pa (216 m/s or Mach 0.91 at Mars' surface). Cruise flight at Mach 0.5 (118 m/s, still much greater than the typical 50m/s terminal velocity on a parachute) requires a parafoil of 155 m². An additional complication that applies to parafoils, and any other system requiring turning flight for maneuvering, is low turn rate (~2 deg / sec)¹⁰.

Steerable parachutes appear to be a technically feasible alternative. Dellicker et al¹¹ have successfully demonstrated an autonomous steering mechanism with a flat, circular parachute, achieving an L/D of 0.8. In this system, actuators connected to each of four parachute risers steer the chute by pulling on the risers. Accuracies of 70 meters Circular Error Probability (CEP) were demonstrated with fifteen successful fully autonomous airdrops. An effort estimated to be approximately equal in scope to that required for the development of the parachutes for the Mars Exploration Rover mission (total cost ~\$10M) would be required to develop a steerable chute system for use at Mars. No detailed cost estimate for this development was made as part of this study.

The lowest cost alternative is a non-steerable two-parachute system similar to that used on MSL. This system consists of supersonic chute of the disk-gap-band design used in all NASA Mars landers since Viking, and a subsonic chute of a design to be developed for MSL. This system does nothing to counteract wind drift. The lander drifts up to several km with the wind depending on wind velocity, and must fly back to the target under power.

An expensive series of Balloon-Launched Decelerator Tests (BLDT) were conducted for Viking in 1972 to flight-qualify a parachute of the disk-gap-band design. Because of the prohibitive expense of repeating these tests, parachutes for all of NASA's Mars landers since Viking have been designed to deploy within mach and dynamic pressure limits defined as a result of those tests and subsequent analysis. For MSL, the limits are Mach 1.13 to 2.2 and dynamic pressure between 240 to 850 Pa. These limits are sometimes referred to as the "parachute deployment box".

Enhancements to terminal guidance and control

With the expected approach navigation performance, MSL's hypersonic entry guidance is capable of delivery to ~6.5 x 2 km (99.7%) at chute deploy. Therefore, even when wind drift on the non-steerable MSL parachute is accounted for, no action is necessary in powered descent to meet the Generation 2 precision landing requirement to land within a ~10 x 5 km ellipse centered on a preselected site. The principal role of powered descent guidance on MSL is to decelerate the lander for a soft landing on the surface.

MSL does carry enough fuel for about 100m of maneuvering over the surface for avoidance of hazards in the landing zone.. The powered descent guidance scheme planned for use on MSL¹² is a variant of the Apollo scheme used for lunar landing. Like the Apollo lunar landing scheme, it is capable of accepting a redesignation of the target at any time in the descent. However, constraints are applied to prevent the lander from selecting targets which it cannot reach with the amount of propellant remaining. MSL powered descent guidance calculates the landing site reachable with minimum propellant expenditure on the trajectory after jettisoning of the chute, and uses that site as the target unless a new target is designated to meet hazard avoidance requirements.

A terminal guidance algorithm solves a two-point boundary value problem (TPBVP) to obtain a trajectory that starts at an initial (i.e., the current) state and finishes at some desirable final state. For MSL, the final state includes position, velocity, and acceleration components. The horizontal velocity and acceleration must be zero to land vertically, in an upright attitude.

MSL solves this TPBVP by computing a quadratic acceleration profile for each of the three axes in the following form:

$$a(t) = C_0 + C_1t + C_2t^2$$

Integration yields the velocity and position components, and the unknown coefficients C_0 , C_1 , and C_2 are solved for in terms of the known boundary conditions and the descent time. The descent time is chosen such that the vertical acceleration is linear with time ($C_2=0$) in order to reduce the computational complexity. Since MSL does not require much horizontal thrusting, a linear vertical acceleration profile is nearly mass-optimal for short descents.

However, for a target re-designation that is far away, a significant amount of horizontal thrust is required and optimization of the horizontal channel cannot be done independently of the vertical channel. A linear vertical acceleration profile (and the corresponding descent time) is no longer mass-optimal.

An optimal trajectory can be constructed using the same basic concept as the Apollo/MSL algorithm, but using a higher-degree polynomial. In our work, we used an 11th degree polynomial (a balance between compute time and optimality of solution) to represent the position profile for each axis.

$$r(t) = \sum_{i=0}^{11} C_i h_i \left(\frac{t}{T} \right)$$

where $r(t)$ denotes the position profile, C_i are the unknown coefficients, h_i are linearly independent polynomials of maximum degree 11 (and chosen to be orthogonal for numerical issues), and T is the descent time. This profile is differentiated 5 times to yield velocity, acceleration, jerk, snap, and crackle and is evaluated at both the initial and final boundary conditions. Of these 12 boundary conditions (per axis), the TPBVP only has

constraints on 5 of them (4 for the vertical axis). The remaining are optimization variables, as is the descent time, yielding a total of 23 optimization variables.

Using a parametric optimizer, the optimal trajectory is computed that also satisfies the following constraints:

- Maximum thrust < 80% capacity (robustness)
- Minimum thrust > 30% capacity (MSL limitation)
- Altitude > 0

Each of these trajectories takes about an hour of CPU time to optimize. To be suitable for use onboard the lander in realtime, a guidance algorithm which can closely approximate mass-optimal performance with much faster computation time must be developed.

The optimal trajectory derived by Lawden is inapplicable for this problem because it does not allow for our last 2 constraints, nor does it allow a specified final thrust level (which is necessary to land vertically and upright).

EXAMINATION OF SYSTEM-LEVEL TRADES

Using thrust to compensate for wind drift on the chute may necessitate carrying significant amounts of extra propellant, in addition to extra propulsion system dry mass (tanks, possibly more engines). Since the diameter of the aeroshell is fixed by the diameter of the shroud of the launch vehicle, any increase in entry mass increases the ballistic coefficient of the vehicle ($= \text{mass} / C_d * \text{area}$). Increasing ballistic coefficient decreases parachute deployment altitude. Increasing the L/D of the vehicle increases parachute deployment altitude. The following two rules of thumb, based on experience with EDL simulations¹³, capture these relationships:

- Rule of thumb #1: A 10% increase in entry mass results in a 1 km decrease in the altitude at which the supersonic chute is deployed
- Rule of thumb #2: increasing L/D by 35% raises chute deployment altitude 1km

In order to examine these trades in a quantitative (but preliminary) way, a low-fidelity software simulation of the parachute and powered descent phases was created. The following assumptions are used in the simulation:

- Two chutes are used in the parachute phase, one supersonic and the other subsonic (which is assumed either steerable or non-steerable). Descent rate at terminal velocity on the subsonic chute is 50 m/s.
- The steerable chute is assumed to have an L/D of 0.8
- The initial state vector at deployment of the supersonic chute is derived by applying the above two rules of thumb to a typical MSL trajectory used as a reference (for which supersonic chute deployment altitude is 10 km).
- Mass properties are those shown in Table 3 - approximately 2100 kg of dry mass at entry is needed to land 1000 kg on the surface of Mars.

- Time on the supersonic chute is fixed at 22.6 sec. (derived from the sample MSL case) to slow to near terminal velocity
- Time on subsonic chute must be ≥ 20 sec. to slow to near terminal velocity
- Wind velocity is assumed constant at 50 m/s, carrying the lander away from the target.
- Atmospheric drag is not modeled during the powered descent phase
- Landing site elevation is 2.5 km above the MOLA reference ellipsoid (a requirement which allows accessibility to 90% of Mars' surface)

Table 3
MASS BREAKDOWN FOR MSL REFERENCE CASE

Supersonic chute	21
Subsonic chute	45
Heatshield	347
Backshell	167
"Skycrane" descent stage (dry)	505
Landed	1000
Total dry mass (kg)	2085
RCS propellant	15
Terminal descent prop	300
Total wet mass (kg)	2400

For a given combination of entry mass, chute deployment altitude, and time on subsonic chute, the simulation integrates a trajectory from supersonic chute deploy to subsonic chute jettison. The final state in the chute phase includes position, velocity, and mass.

In a separate set of computations, the simulation builds a powered descent trajectory using an initial position and velocity from the end of the chute phase, for which final propellant mass is 0 kg. Wet mass at the start of powered descent is determined from this computation using the dry masses in Table 3. If the difference between the mass at the end of the chute phase (determined from the chute phase computation described above) and the mass at the start of the powered descent phase is positive, then the case is marked as feasible. Otherwise, it is marked as infeasible.

Table 4 shows some relevant parameters describing a single simulation case (2500 kg entry mass, optimal powered descent guidance, nonsteerable chute, 400m delivery error at supersonic chute deploy). Supersonic chute deployment altitude, calculated from the first rule of thumb, is 9.5 km above the MOLA reference ellipsoid (7.5 km above the surface). At this point, the flight path angle is still shallow (~20 deg); the trajectory has not "bent over" to the vertical yet. Deployment occurs 7.2 km before reaching the target so the lander will be over the target when terminal velocity is reached and the trajectory is nearly vertical. Subsonic chute deployment occurs nearly over the target. In 74.6 sec on the two chutes, the lander drifts 3.5 km downwind in the 50 m/s winds. The engine is ignited at 4.1 km above MOLA (1.6 km above the surface), and after a 66-sec. powered descent, the lander touches down at a site whose elevation is 2.5 km above MOLA.

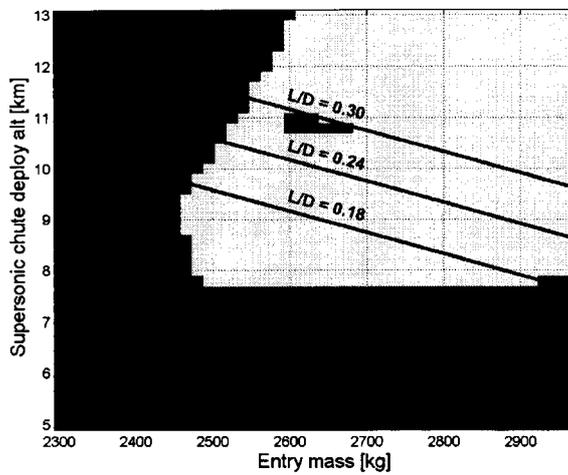
Table 4
RESULTS FROM A SAMPLE SIMULATION CASE

Event	Time since chute deploy (sec)	Altitude (km above MOLA)	Mass (kg)	Dist. from target (km)
Supersonic chute deploy	0.00	9.50	2500.00	-7.23
Subsonic chute deploy	22.60	7.35	2479.00	-0.79
Engine ignition	74.60	4.14	1920.00	3.53
Touchdown	140.59	2.50	1000.00	0.00

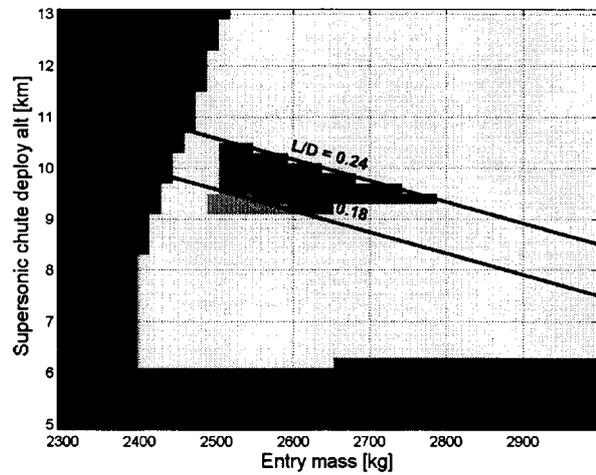
A variety of cases were run in this simulation to examine system trades. Four “case types” (MSL or optimal powered descent guidance, for either steerable or nonsteerable subsonic chutes) were run for delivery errors at supersonic chute deploy of 400, 1000, and 1500m. Results of the cases with 400m delivery error are shown in Figure 3 as plots of supersonic chute deployment altitude vs. entry mass. In these plots, blue lines show nominal supersonic chute deployment altitude, computed using the two rules of thumb above. Red regions show combinations of supersonic chute deployment altitude and entry mass for which the lander cannot reach the target, either due to insufficient propellant, insufficient thrust, or both.

Green and yellow regions show combinations of chute deployment altitude and entry mass for which the lander successfully reaches the target. The difference between the green and yellow regions traces back to the fact that if bank-only modulation is used in the hypersonic phase, variations in atmospheric density can cause the actual supersonic chute deployment altitude to vary up to ± 3 km from the nominal value. This means that to assure successful pinpoint landing for a given entry mass, it is necessary to be able to reach the target with the available propellant not only for chute deployment at the nominal altitude, but also for deployment 3km higher and 3km lower than the nominal. The green regions in the plots satisfy this requirement, but the yellow regions do not. Three shades of green are used which apply to three vehicle L/D values (lightest green: 0.18, darker: 0.24, and darkest: 0.30).

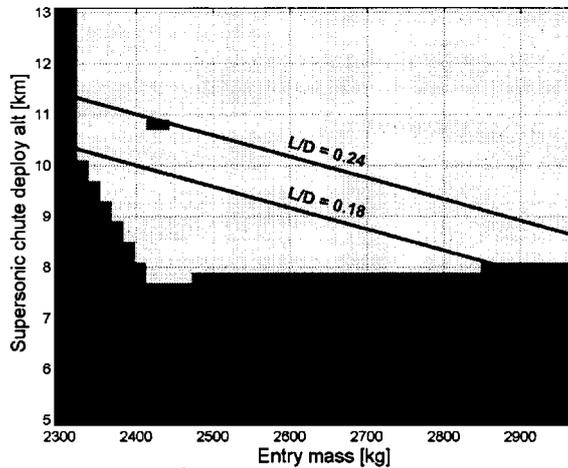
For example, Figure 3c (optimal guidance, nonsteerable chute) shows that for entry mass < 2400 kg, all points are in the red region, indicating that there is no chute deployment altitude for which the lander can reach the target. For a vehicle L/D of 0.18 and an entry mass of 2450 kg, the chute deployment altitude predicted by the first rule of thumb (shown on the lower blue line) is 9.8 km, in the yellow region. From this altitude, under the assumptions listed previously, the lander is capable of reaching the target. However, if bank-only modulation is used, the ± 3 km uncertainty in chute deployment altitude means that the chute may be deployed as high as 12.8 km (in the red region, meaning that the target cannot be reached) or as low as 6.8 km (in the yellow region, meaning that the target can be reached). So, mission success is not assured throughout the expected range



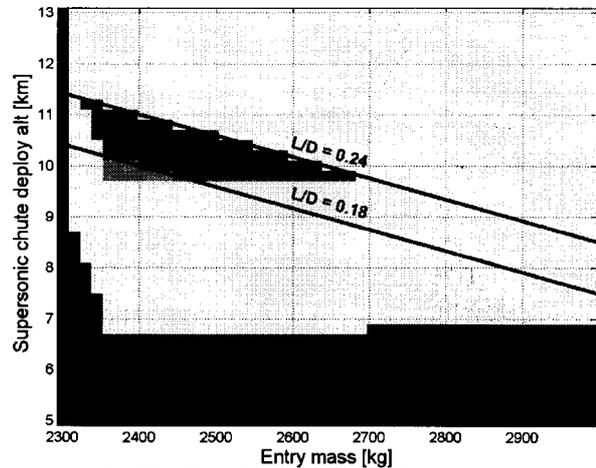
a) MSL powered descent guidance. nonsteerable chute.



c) Optimal powered descent guidance. nonsteerable chute.



b) MSL powered descent guidance. steerable chute



d) Optimal powered descent guidance. steerable chute.

Figure 3 Comparison of MSL and optimal powered descent guidance for steerable and nonsteerable subsonic chutes and 400m delivery error at chute deploy. Altitudes shown are above MOLA, for landing 2.5 km above MOLA.

of chute deployment altitudes at an entry mass of 2450 kg if bank-only modulation is used.

If bank and angle-of-attack modulation is used in the hypersonic phase (as previously discussed), it is reasonable to anticipate a significant reduction in the uncertainty in chute deployment altitude (although simulations were not done in this study to quantify the reduction). If the addition of angle-of-attack modulation reduces the uncertainty to ± 1 km, successful pinpoint landing can be achieved for an entry mass of 2450 kg. Addition

of angle-of-attack modulation can be costly, however, which does not necessarily make this a desirable trade.

If entry mass is increased to 2500 kg, the expected nominal chute deployment altitude decreases to ~9.5 km. At this entry mass, it is possible to reach the target for deployments at altitudes between 12.5 and 6.5 km (+/- 3km of the nominal). Therefore, pinpoint landing can be successfully accomplished with bank-only modulation, and the (2500 kg, 9.5 km) point lies in a green region.

Table 5 summarizes the results of all the simulations, sorted in order of increasing cost (using the preliminary cost assumptions of Table X). Under the assumptions used in this study, pinpoint landing appears to be feasible for all the candidate systems shown in the table; however there is a significant cost differential between the least and most expensive. Development of some of the more expensive technologies (aerodynamic surfaces for angle-of-attack modulation, qualification of chutes for greater angle of attack at deployment) may be desirable or necessary for heavier landers; however, at first blush, they do not appear to be needed for an MSL-class pinpoint landing mission. The steerable parachute saves ~100 kg entry mass; however, that savings does not appear to be critical to the success of the mission if the assumptions made here are correct.

Table 5.
Summary of Simulation Results

Steerable

CONCLUSIONS

The biggest obstacle to landing within 100m of a desired target is uncertainty in low-level winds that can drift the lander several km from the target in the parachute phase. The results of this preliminary study show that under the set of assumptions discussed here, development of a powered descent guidance algorithm which optimizes propellant expenditure (or nearly so) appears to be the most cost-effective technology development of several candidate technologies considered. Steerable subsonic parachutes appear to

improve performance, but probably not enough to offset their development cost, especially given the acceptable performance of systems with nonsteerable chutes when used with optimal powered descent guidance. In addition, a sensor for target-relative navigation is needed at low altitudes to eliminate map-tie error.

Refinement of this preliminary work is warranted. Effects of variations in several parameters important to the design of future pinpoint landing missions have not been investigated here (e.g. entry speed, time of arrival in the Martian atmosphere's pressure cycle, wind speeds other than the 50 m/s value used here, landing site elevation, landed mass, etc.). A design study for a steerable parachute could verify or alter some of the assumptions regarding its cost and performance over a wider range of these variables. Also, the validity of the two "rules of thumb" has not been tested over the whole range of these parameters. Higher fidelity cost estimates for system components are necessary to more accurately ascertain the cost of developing pinpoint landing capability. Mesoscale modeling of winds for selected types of sites could aid considerably in developing a more informed understanding of the effect of wind on the design of pinpoint landing systems.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the contributions and advice of Dan Burkhart, Brian Rush, Joe Guinn, and Betty Quintanilla in the area of approach navigation, Christian Liebe on sensors, Ron Sostaric on aerodynamic decelerators, David Kass on low-level Martian winds, and in particular Gil Carman in the area of entry guidance and control.

REFERENCES

1. Thornton, Catherine L. and Border, James S., "Radiometric Tracking Techniques for Deep-Space Navigation," Monograph 1, Deep-Space Communications and Navigation Series, Joseph H. Yuen, Editor-in-Chief, JPL Publication 00-11, October 2000
2. Rush, B., Bhaskaran, S., and Synnott, S., "Improving Mars Approach Navigation Using Optical Data," paper # AAS-01-412 presented at the AAS/AIAA Astrodynamics Specialist Conference, 30 July - 2 August, 2001, Quebec City, Canada
3. MER Project - Entry, Descent, and Landing Performance Assessment Report (Preliminary), JPL D-19688, April 5, 2002
4. David Kass, personal interview, July 2003
5. Gaskell, R.W., "Automated Landmark Identification for Spacecraft Navigation", presented at the AIAA/AAS Astrodynamics Conference, 8/2001, Quebec City, Canada
6. Cheng, Y., Johnson, A., Matthies, L.H., and Wolf, A., "Efficient Passive Image-Based Hazard Detection for Safe Landing on Mars," paper # AM028, The 6th International Symposium on Artificial Intelligence, Robotics and Automation in Space, Montreal, Canada, June 18 - 22, 2001
7. Au, R. H., "Optical Window Materials for Hypersonic Flow," *SPIE Window and Dome Technologies and Materials* vol 1112 (1989), pp. 330-339
8. Madsen, C.M. and Cerimele, C.J., "Flight Performance, Aerodynamics, and Simulation Development for the X-38 Parafoil Test Program," paper # AIAA-2003-2108 presented at the 17th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar, 19 - 22 May, 2003, Monterey, CA
9. Strahan, A., "Testing of Parafoil Autonomous Guidance, Navigation & Control for X-38," paper # AIAA-2003-2115, presented at the 17th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar, 19 - 22 May, 2003, Monterey, CA

10. Lingard, J.S., Viviani, G.L., "Ram-Air Parachute System Design," 13th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar, 15 May, 1995, Clearwater Beach, FL
11. Dellicker, S., Benney, Rl., LeMoine, D., Brown, G., and Gilles, B., "Steering a Flat Circular Parachute? They Said It Couldn't Be Done," paper # AIAA-2003-2101 presented at the 17th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar, 19 – 22 May, 2003, Monterey, CA
12. E.C. Wong, E. C., G. Singh, and J.P. Masciarelli, "Autonomous Guidance and Control Design for Hazard Avoidance and Safe Landing on Mars", paper # AIAA-02-46, AIAA Atmospheric Flight Mechanics Conference, Monterey, California, August 5-8, 2002
13. Gil Carman, personal interview, September, 2003