JOURNAL OF SPACECRAFT AND ROCKETS
Vol. 40, No. 3, May–June 2003

Comparison of Obstacle Avoidance Strategies for Mars Landers

Andrew J. Sinclair*
Texas A&M University, College Station, Texas 77843
and
Norman G. Fitz-Coy†
University of Florida, Gainesville, Florida 32611

A novel autonomous obstacle avoidance controller for the safe landing of a Mars mission in hazardous terrain is developed. The controller uses iterative, closed-loop reactions based on feedback of the landing site terrain. The method was evaluated using a Monte Carlo simulation, including models of reference vehicle dynamics, landing site terrain, and a light detection and ranging sensor. An expected safe landing rate of approximately 60% was shown with no obstacle avoidance control. With the controller, success rates improved to 72.2% for a low control authority thruster design and 97.4% for a high control authority thruster design. The effectiveness of the novel strategy is compared to previous methods. Implications regarding the importance of control logic design and thruster design are inferred.

Nomenclature

C_D	=	drag coefficient
$\boldsymbol{F_D}$	=	aerodynamic drag, N
$[\boldsymbol{I}]$	=	body component inertia matrix, $\mathrm{kg \cdot m^2}$
j	=	map ring number
K_P	=	controller gain constant
m	=	vehicle mass, kg
n	=	number of map points designated as hazards
S	=	representative area, $\mathrm{m^2}$
v_{rel}	=	relative velocity magnitude, m/s
$\hat{\boldsymbol{v}}_{\mathrm{rel}}$	=	normalized relative velocity, m/s
z	=	map height, m
α	=	transverse map coordinate, rad
ρ	=	atmospheric density, $\mathrm{kg/m^3}$

Subscripts

i	=	map points
X, Y	=	high-authority engines

Introduction

ALTHOUGH methods for autonomously landing probes on the moons and planets of the solar system have been developed, autonomous obstacle avoidance is a current area of research and development. Autonomous obstacle avoidance for Mars landers refers to the capability of a lander to guide itself to a safe landing site while avoiding landing hazards such as boulders, craters, hills, and valleys. To date, the only lander to use obstacle avoidance has been the Apollo Lunar Excursion Module. This obstacle avoidance was accomplished using pilot control and provided the Apollo program with greater capability to explore interesting landing sites with improved mission safety. Current concepts for the human exploration of Mars call for split mission architectures that depend on the use of autonomous landers.[1–3] To provide such a mission with the same advantages enjoyed by the Apollo program, autonomous obstacle avoidance capability needs to be developed.

As with any new technology, there are a wide variety of design options that can be considered in the development of autonomous obstacle avoidance. Preliminary studies can serve to improve understanding of the overall design space before researchers commit to any one design option. The purpose of this paper is to analyze and compare several possible strategies for autonomous obstacle avoidance. A two-dimensional design space is considered with two different control logic designs and two different actuation thruster designs.

The first obstacle avoidance strategy to be considered uses control logic that steers the lander to a desired landing site and uses a thruster design that imparts high control authority. This method of autonomous obstacle avoidance is being developed by Johnson et al. at the Jet Propulsion Laboratory, California Institute of Technology, for the Mars Smart Lander mission.[4] This strategy mechanizes the approach of the Apollo landers. A light detection and ranging (LIDAR) sensor is used to scan the surface underneath the lander as it descends. This surface map is then analyzed for roughness and slope, and a desired landing site is selected. A guidance law is then used to steer the vehicle to this site. In their work, the guidance law calculated the current acceleration magnitude and direction based on the desired acceleration, jerk (the first derivative of acceleration), and snap (the second derivative of acceleration). When the acceleration magnitude was solved for, thruster saturation was not considered. Additionally, to solve for the acceleration direction, a vehicle with gimbaled or triaxial thruster sets must be used. Therefore, in this study, the vehicle was assumed to have been specifically designed for obstacle avoidance with high control authority and high maneuverability.

To evaluate their strategy, Johnson et al.[4] performed Monte Carlo simulations of the landing maneuver. The results of these simulations showed safe landing probability for the Mars Smart Lander related to the density of rocks included in their model of the Martian surface. For a 10% rock density, a 100% safe landing rate was achieved, whereas rock densities of 15 and 20% resulted in safe landing rates of 98 and 93%, respectively.

An obstacle avoidance strategy can be categorized by three main subsystems: sensors, controller, and actuators. The strategy, developed by Johnson et al.,[4] used a LIDAR sensor, a control approach to select and fly to a desired landing site (fly-to control), and high control authority actuators. LIDAR is a key technology for autonomous obstacle avoidance because it can be directly used to provide data regarding the terrain beneath a lander. It allows much simpler data analysis than photographic methods. However, other options exist in the selection of control logic and actuation thrusters. The current

Received 13 June 2002; revision received 6 January 2003; accepted for publication 7 January 2003. Copyright © 2003 by Andrew J. Sinclair and Norman G. Fitz-Coy. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the $10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0022-4650/03 $10.00 in correspondence with the CCC.

*Graduate Student, Department of Aerospace Engineering, 701 H. R. Bright Building TAMU-3141; sinclair@tamu.edu. Student Member AIAA.

†Associate Professor, Department of Mechanical and Aerospace Engineering, 231 Aerospace Building, P.O. Box 116250; nfc@ufl.edu. Member AIAA.

authors developed a novel controller that, instead of selecting and flying to a desired landing site, steers the vehicle away from hazardous obstacles in an iterative fashion (fly-away control).

The controller was implemented with both high and low control authority thruster designs. The two resulting obstacle avoidance strategies can be compared to the work of Johnson et al.[4] to gain a better understanding of the impact of control logic and thruster design on obstacle avoidance performance. The fourth option in the design space, a fly-to controller with low control authority, is not physically realizable because a controller can not be designed to guide the vehicle to an arbitrary landing site without sufficient authority to reach the target.

The next section of the paper describes the development of a simulation environment for the terminal descent phase of a Mars landing. The controller design and its implementation with high and low control authority thrusters are then discussed. The results of Monte Carlo simulation are covered. Finally, conclusions regarding the effectiveness of the strategies are drawn.

Simulation Development

Dynamics Model

For the current work, the conceptual Mars Ascent Vehicle (MAV), an autonomous lander described in the Mars Reference Mission (MRM), was chosen as the reference vehicle.[1,2] The MAV would land autonomously on Mars to be used by a human crew in leaving the surface of the planet. Although no humans would be onboard during landing, a success rate for this phase approaching 100% would still be necessary for overall mission success. Elements of the detailed design of the sensors, controller, and thruster architecture were based on the Viking landers. A simulation environment was developed for the terminal descent stage of a Mars landing. The terminal descent stage begins with parachute release and ends with landing on the surface. This phase of flight includes the range from 1 km to zero altitude and less than 100 m/s speed, with typical landing speeds of 1–3 m/s. Therefore, the local-vertical Mars-fixed frame was assumed to be an inertial reference frame.

A six-degree-of-freedom model of the lander dynamics was obtained using the Newton/Euler method. The forces included were control thrust from the engines, gravity, and aerodynamic drag. The only torque considered was the applied moment from the engines' thrust. The vehicle attitude was represented using Euler parameters. During simulation, these equations were integrated using a fourth-order Runge–Kutta integration scheme with a fixed time step of 0.01 s. This time step was chosen for simulation stability and reasonable numeric efficiency.

The engine model based on the MRM description used four RL10-class engines, with a maximum thrust of 93,400 N for each engine and specific impulses of 450 s. These four engines were the only thrusters included in the vehicle model and were used to produce all control force and control torque necessary to accomplish maneuvers.

This thruster design provided low control authority because performing lateral maneuvers required attitude maneuvers to orient the thrust vector. Each engine's force vector was oriented parallel to the positive vertical body axis. The engines were all modeled as being located 5 m out from the geometric/mass center of the lander in the positive and negative x and y body directions.

For the high control authority thruster design, gimbaled thrusters or triaxial thruster sets were used. In addition to the 373,600-N total maximum thrust in the vertical direction, the same thrust was allowed in both of the horizontal directions in both the positive and negative sense. Although such a design would have large propulsion system requirements, it is useful to look at such a scheme for the purpose of comparison to the previous work of Johnson et al.[4] The two-thruster designs are shown in Fig. 1.

For both high and low control authority designs, the initial mass of the lander was 48,100 kg (the 44,440-kg landed mass plus a 3660-kg propellant load). The mass was reduced according to fuel consumption by the engines. The MRM gives rough geometry specifications

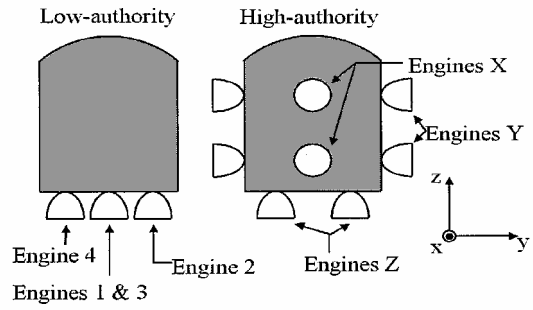

Fig. 1 High- and low-authority thruster designs.

for the MAV of $10 \times 10 \times 10$ m. Based on these dimensions, the following inertia tensor was used:

$$[\boldsymbol{I}] = m \begin{bmatrix} 10 & 0 & 0 \\ 0 & 10 & 0 \\ 0 & 0 & 10 \end{bmatrix} \tag{1}$$

The use of the mass as a scaling factor for the inertia assumes that as the mass changes due to fuel consumption, the mass distribution is unaffected. This assumption is equivalent to assuming that the fuel is always spread evenly throughout the vehicle and is always centered at the mass/geometric center of the lander. Also, no fuel sloshing has been modeled. Because the scope of the desired model was limited to altitudes of less than 1 km, gravitational acceleration was treated as a constant. This acceleration was -3.724 m/s^2 in the Z inertial direction. Aerodynamic drag was also included in the model. The following formula was used to calculate drag:

$$\boldsymbol{F_D} = -\tfrac{1}{2}\rho v_{\text{rel}}^2 C_D S \hat{\boldsymbol{v}}_{\text{rel}} \tag{2}$$

The following values were used for the model: $\rho = 0.017$ kg/m^3, $C_D = 2$, and $S = 78.5$ m^2. The relative velocity was computed from the vehicle velocity and wind speed.

Determination of Obstacle Size

Another model was developed to determine the size of terrain features, such as boulders, that could represent a landing hazard for the reference vehicle. The following scenario was considered. A lander is initially in a purely vertical descent. One landing leg then impacts a point on a raised terrain feature such as a boulder. The vehicle then rotates about that point until the opposite landing leg strikes the ground. The lander then rotates about the point on the ground.

The initial descent speed required for the lander to just reach the point where the center of mass was balanced over one landing leg was calculated for various heights of the terrain feature. Velocities greater than this descent speed would cause the lander to move through the equilibrium point and tip over. This calculation was performed by assuming angular momentum conservation (about the point of impact) through the impacts and conservation of energy in between them. The impacts were assumed to be perfectly plastic.

The results of this analysis are shown in Fig. 2. Figure 2 shows the maximum safe landing speed as a function of boulder size. The landing speeds of the reference vehicle and most landers in general is around 1–3 m/s. At these speeds, the MAV could survive landing on a boulder of 3–4 m in height. Given the approximations made in the preceding analysis, as well as the desire for a safety factor, boulders with a height greater than 2 m were designated as landing hazards. This hazard size is much larger than the rock size traditionally considered hazardous to landers (approximately 0.5 m); however, this large size is due in part to the large size of the MAV.

Terrain Model

A landing site terrain model was developed for a circular area with a 1-km diameter. Several terrain features can pose landing hazards to a Mars lander. These features include relatively small-scale features, such as boulders and craters, as well as large-scale features, such as the surface slope of hills and valleys. The terrain model used

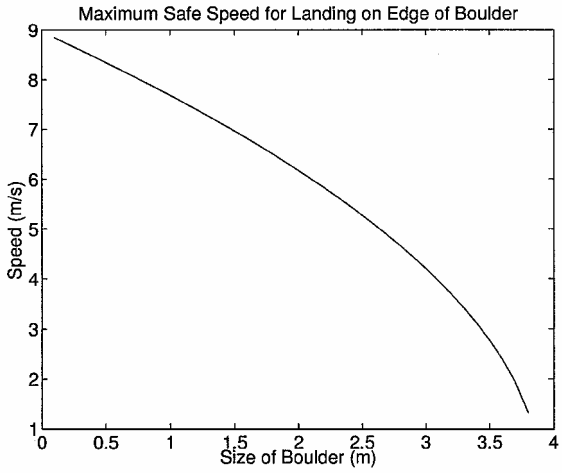

Fig. 2 Determination of obstacle size.

in the current work treated the surface as a flat plane populated with boulders.

The size and density of the boulders in the terrain model were based on data published in the collection "Scientific Results of the Viking Project" for the rock density vs rock size at the Viking landing sites on Mars and the Surveyor 7 and Apollo 12 landing sites on the moon.[5] These data show the similarity at all four sites for the number of rocks found within a given size range in a given area on the surface as a function of the rock size. The published data cover rocks of 0.1–1-m sizes. The current authors applied these data in forming a terrain model by applying a curve fit and then extrapolating to find the number of rocks expected in a 1-km-diam circle in 1-m rock-size diameter steps. A rock of diameter 30 m was the largest rock expected in the landing site. Although encountering a rock of this size on the surface of Mars is unlikely, it was included in the model to represent other large terrain features such as outcrops. Table 1 shows the results that were obtained.

To assemble the terrain model, rocks were randomly placed in the landing site according to the sizes and numbers shown in Table 1. The current authors made the assumption that the rocks were spherical in shape with 60% of their diameter above the surface. This assumption gave a height-to-diameter relationship for each rock. They were placed in the landing site so that no rocks overlapped the same area. The total group of rocks covered 9.89% of the landing site area. Of this, 2.48% was covered by the 1-m-diam rocks, 1.24% was covered by the 2-m-diam rocks, 0.83% was covered by the 3-m-diam rocks, and the remaining 5.34% was covered by the larger rocks.

Based on the assumed boulder height-to-diameter relationship, rocks with diameters of 4 m or greater were considered hazards. With approximately 5% hazard density, this landing zone is relatively clear of obstacles because it was based on the Viking Lander 1 and 2 and Surveyor 7 landing sites (all vehicles that landed successfully without obstacle avoidance capability). However, the large size of the reference vehicle makes landing safely more difficult. The summed area of the hazardous rocks and a 6-m apron (the radius of the lander's footprint) around each equals 34.4% of the landing zone. This calculation indicates an expected safe landing rate of 65.6% without obstacle avoidance control.

Sensor Model

The sensor suite modeled in the simulation was based on the Viking sensors. The Viking landers carried an inertial reference unit aided by a radar altimeter and the terminal descent and landing radar (a four-beam Doppler radar). In the current work, the controller was assumed to have perfect knowledge of the true altitude, velocity, attitude, and angular velocity. No sensor errors were modeled in these values.

To provide the controller input regarding the presence of landing hazards a LIDAR system was modeled. This model uses a system similar to the LIDAR discussed by Johnson et al.[4] The LIDAR

Table 1 Terrain model rock sizes

Diameter, m	Number of rocks
1	24,836
2	3,105
3	920
4	388
5	199
6	115
7	72
8	49
9	34
10	25
11	19
12	14
13	11
14	9
15	7
16	6
17	5
18	4
19	4
20	3
21	3
22	2
23	2
24	2
25	2
26	1
27	1
28	1
29	1
30	1

modeled in the current work performs range measurements using the time of flight of laser pulses. It uses a two-axis gimbaled optical system to scan over a region taking discrete measurements. The scanning pattern swept out a series of concentric rings. Eight rings are collected by skewing the sensor axis in steps through a 20-deg half-cone angle. Along each ring, 46 measurements are taken in the transverse direction. The number of rings and the number of measurements per ring were chosen based on the desired spatial resolution on the surface at particular heights.

Sweeping out each ring was modeled to take 0.01 s. Stepping from one ring to the next was also modeled to take 0.01 s. Because the integration time step was also equal to 0.01 s, one of these tasks occurred each time step. This model translates into a sensor capable of collecting measurements at 4600 Hz. This also translates into specifications for the slewing and stepping speeds of 100 rpm and 288 deg/s, respectively, and a complete scan time of 0.16 s. Additionally, a range of approximately 250 m was assumed. In the work of Stentz et al.,[6] a LIDAR system was built and tested that had a sample rate of 12 kHz, a spinning rate of 1800 rpm, a step rate of 120 deg/s, and a range of 200 m. The LIDAR proposed by Johnson et al.[4] used a sample rate of 10 kHz and a range of 500 m. Whereas that work did not specify scanning rates, it did call for scanning a similar angular region of 10×10 deg.

To model the measurement of each discrete scanning point, the position of the lander, the attitude of the lander, and the orientation of the sensor were compared to the terrain model. A ray–surface intersection was used to estimate the range from the sensor to the target point on the surface model.

Figure 3 shows a 20×20 m section of the surface array. Figure 4 shows an example of the LIDAR model output. This scan was taken with the lander at an altitude of approximately 50 m over the point (10 m, 10 m) in Fig. 3. The scan was taken while the lander was descending at approximately 10 m/s. The larger boulders in Fig. 3 can be noted near the center of Fig. 4

This method introduced two sources of error into the model. First, the sensor is assumed to always intersect with one of the grid points. The size of this error is related to the grid resolution. It is possible to use a linear interpolation scheme to solve for a point of intersection between the grid points, but this was not done for the sake of

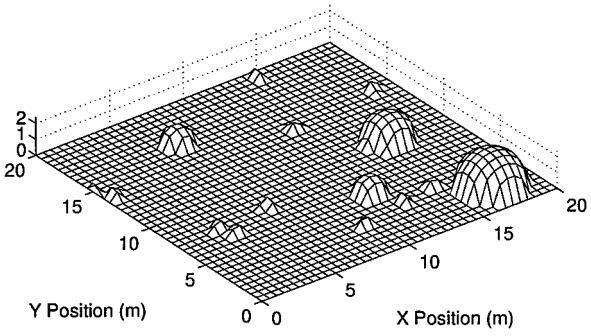

Fig. 3 Surface array for 20 × 20 m section of landing site.

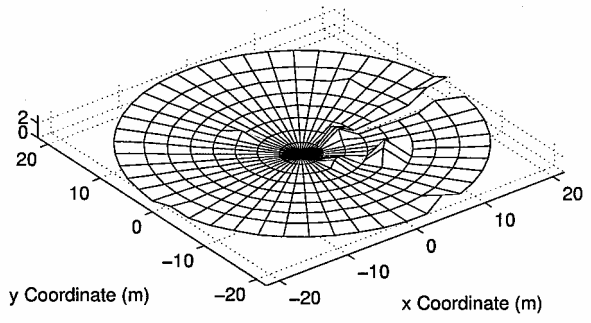

Fig. 4 Sensor model example output.

simplicity. The second error is the possibility for the search scheme to converge to the wrong point. This is possible if the angle of incidence of the sensor axis with the surface is low such that it intersects the surface in multiple locations or if the sensor is below the local horizon at some point on the array. Because of the sensor angles and heights that the model was used for, it was determined that this problem would not arise frequently enough to misshape the LIDAR terrain map significantly.

Two physical sources of error in LIDAR systems are beam divergence and scattering and pointing errors. Beam divergence and scattering result in noisy output, giving a spectrum of ranges for each pulse. Pointing errors are due to imperfections in the encoders used to aim the sensor beam. Neither of these errors were directly modeled in the model. However, the errors just described can be seen to be somewhat similar to the errors in the true system.

Fly-Away/Low-Authority Strategy

The first strategy to be implemented used a fly-away controller with the low control authority thruster design called for in the MRM.[1,2] The fly-away controller is based on the Viking Guidance System (VGS). The VGS contained laws designed to allow the landers to meet safe landing criteria on the landing vertical velocity, lateral velocity, attitude, and angular velocity. Their design is fully described by Holmberg et al.[7] These control laws were used, and the fly-away controller added another control law to meet criteria on the position of the landing site relative to any landing hazards.

The fly-away obstacle avoidance controller (OAC) used a feedback loop to drive the lander to a safe landing site. This guidance method is similar to the gradient field maps developed by Payton for autonomous robots.[8] Instead of selecting a safe landing site and flying toward it, the controller detects hazards and flys away from them.

A method was developed to identify hazards. The map of LIDAR range readings was converted from spherical coordinates (sensor gimbal angles and ranges) to Cartesian coordinates. Each point in the map was then compared to its neighboring points in the radial and transverse directions. If the value of that point in the map was different from the value of one of its neighbors by greater than 1 m, then that point was designated as a landing hazard. This method is demonstrated in Fig. 5. Figure 5 shows the analysis of the example output from Fig. 4. In this

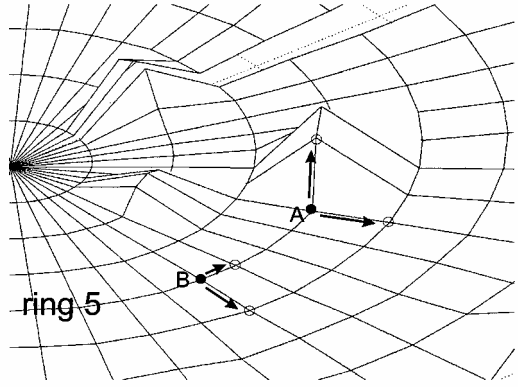

Fig. 5 Example: analysis of sensor model output.

analysis, point A was designated as an obstacle, whereas point B was not.

The determination of control effort was based on the identification of landing hazards as pitch,

$$K_P \sum_{i=1}^{n} \frac{z_i}{j_i} \cos \alpha_i \tag{3}$$

and yaw,

$$K_P \sum_{i=1}^{n} \frac{z_i}{j_i} \sin \alpha_i \tag{4}$$

A value of 10,000 was used for the gain constant K_P. The transverse coordinate was used to orient each term of the control effort radially opposite to the obstacle that it was associated with. The pitch and yaw terms are translated into throttle commands for each engine, as follows:

$$\text{Throttle}_1 = \text{Pitch} \tag{5}$$

$$\text{Throttle}_2 = \text{Yaw} \tag{6}$$

$$\text{Throttle}_3 = -\text{Pitch} \tag{7}$$

$$\text{Throttle}_4 = -\text{Yaw} \tag{8}$$

These attitude control commands were limited to ±15% because the remaining 85% of the throttle setting was reserved for the other controller portions. The analysis of the output shown in Fig. 4 generated pitch = 20.8 and yaw = −4.3. These commands would rotate the lander to reorient the engine thrust. This would drive the vehicle in the negative X and positive Y directions, away from the three large boulders.

Because of the low control authority of the thrusters chosen, producing lateral movement required rotating the spacecraft to give the thrust vector lateral components, waiting for lateral velocity to build up, rotating in the opposite direction to cancel the forward thrust, waiting for lateral velocity to die out, and, finally, rotating back to the initial attitude. It was decided to use a controller that directly used this sequence of maneuvers.

The OAC was activated at an altitude of 200 m and the VGS attitude controller was turned off. If the sensor map indicated that there was a hazardous terrain feature directly below the lander by analyzing the first two rings of the map, then the OAC was initialized as described by analyzing the entire sensor map. The OAC was then fixed to this initial value for 0.5 s. The controller then entered a coast phase with the OAC turned off for another 0.5 s. Finally, the VGS attitude controller was turned back on to reestablish the lateral velocity, attitude, and angular velocity states desired for landing.

Fly-Away/High-Authority Strategy

A second strategy was implemented using a fly-away controller with the hypothetical high control authority thruster design

described in an earlier section. The controller used in this strategy is similar to the one used in the fly-away/low-authority strategy; however, the different actuator dynamics called for a different implementation. Because the high control authority thrusters can produce thrust in any arbitrary direction, lateral maneuvers do not first require an attitude maneuver to orient the thrust vector. This high authority allowed the fly-away controller to be implemented in a continuous fashion running at a high control rate (as opposed to the single maneuver used in the preceding implementation). Changes in the baseline Viking-derived control laws were also implemented to account for the decoupling in the rotational and translational dynamics.

The OAC used a concept similar to the preceding strategy, but a different implementation. The OAC operated between altitudes of 150 and 20 m. As the LIDAR sensor provided new terrain maps, these were analyzed in the same manner as before. The control law was similar to that used in the preceding strategy, but some modifications were made to adapt it to the new implementation:

$$\text{Throttle}_X = K_P \sum_{i=1}^{n} \frac{z_i}{j_i^2} \cos\alpha_i \tag{9}$$

$$\text{Throttle}_Y = K_P \sum_{i=1}^{n} \frac{z_i}{j_i^2} \sin\alpha_i \tag{10}$$

The gain constant K_P was set to a value of 50. Additionally, if the lander was directly over a rock, then the control effort was multiplied by a gain of 100.

Monte Carlo Simulation

To evaluate the performance of these obstacle avoidance strategies, Monte Carlo simulations were carried out. Four simulations were performed. Each strategy was executed with the obstacle avoidance turned off as a baseline and then with the obstacle avoidance enabled: fly-away/low-authority strategy with the obstacle avoidance turned off, fly-away/low-authority with the obstacle avoidance enabled, fly-away/high-authority strategy with the obstacle avoidance turned off, and fly-away/high-authority with the obstacle avoidance enabled.

Each simulation consisted of 500 cases. All Monte Carlo cases were initialized with an altitude of 1000 m and zero angular velocity. A set of 500 Monte Carlo parameters was created for the initial lander velocity and attitude and the wind velocity. However, for the fly-away/high-authority cases, the initial attitude Euler parameters were set to $[0\,0\,0\,1]^T$ because this strategy did not involve attitude maneuvers. The selection parameters used for choosing these values are shown in Table 2.

Table 2 Monte Carlo parameter selection

Parameter	Distribution	Values
Velocity magnitude	Gaussian	Mean 60 m/s, standard deviation 3.50 m/s
Velocity elevation	Gaussian	Mean 10.5 deg, standard deviation 4.49 deg
Velocity azimuth	Uniform	0–360 deg
Attitude eigenangle	Gaussian	Mean 0 deg, standard deviation 11.63 deg
Attitude eigenaxis	Uniform	0–360 deg
Wind speed	Gaussian	Mean 25.5 m/s, standard deviation 10.90 m/s
Wind direction	Uniform	0–360 deg

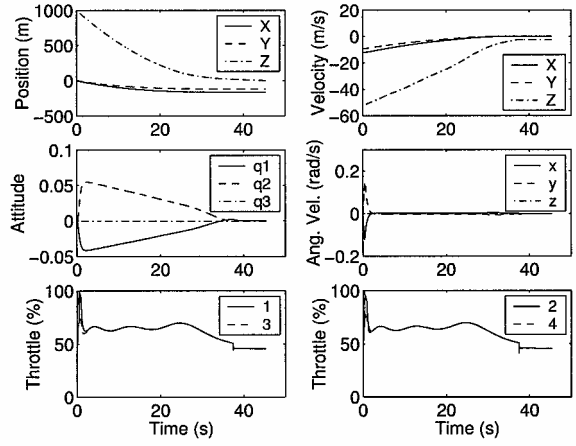

Fig. 6 Example: low control authority baseline simulation.

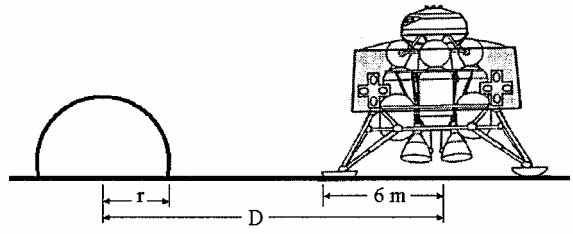

Fig. 7 Landing position hazard analysis.

Fly-Away/Low-Authority Results

The baseline simulations were performed to verify the operation of the Viking-style controller and to establish the failure rate that could be expected if no obstacle avoidance control was used. Figure 6 shows the time history of the vehicle states and the engine throttle settings from one of the baseline simulation cases. In this example, the vehicle was initialized with negative velocities in both the X and Y directions. In the first few seconds of the flight, the engines were differentially throttled to cancel out this lateral velocity. Relatively large angular velocities were produced, and the attitude rotation began to build up. Then, through the first 30 s, the lateral velocity decayed due to the engine thrust. As this happened, the lander slowly rotated back to the vertical. During this time, the engines were also throttled to track the desired altitude/velocity contour. Between 35 and 40 s into the flight, the final approach controller was activated. The lander descended at constant velocity, and the attitude converged to an upright orientation. The flight ended when the vertical position reached zero.

The final vehicle state and mass were stored for each case. These were then analyzed to determine if the vehicle landed successfully. The criteria for a safe landing were divided into two categories: vehicle state criteria and landing hazard criteria. The following vehicle state criteria were used: vertical velocity 2.5 ± 0.9 m/s, lateral velocity ± 1.25 m/s, eigenangle ± 5 deg, angular velocity ± 7 deg/s, and mass $> 44,440$ kg. To determine whether the landing hazard criteria were met, the vehicle's landing position was compared to the rocks in the terrain model. The landing position was compared to the location of the each of the landing hazards as shown in Fig. 7. If the distance from the center of the vehicle to the center of any rock of height greater than 2 m was less than sum of the radius of that rock and the radius of the vehicle (6 m), then that case was designated as landing in a hazardous site.

Figure 8a shows the analysis of the baseline simulation results. The position of each landing site is shown along with its designation for meeting the landing hazard criterion. All 500 of the simulation cases met the vehicle state criterion. There were 314 cases that met the landing hazard criteria, whereas 186 landed too close to one or more of the hazardous rocks. Therefore, 62.8% landed in safe sites. This roughly corresponds with the percentage of the safe area in the landing zone.

The second group of Monte Carlo simulations was performed with the OAC enabled to judge its effectiveness. Figure 9 shows the results from one of these simulations. This case used the same Monte Carlo parameters as the simulation case shown in Fig. 6. The vehicle states and throttle histories are similar to the baseline case except that, approximately 20–25 s into the flight, the obstacle avoidance

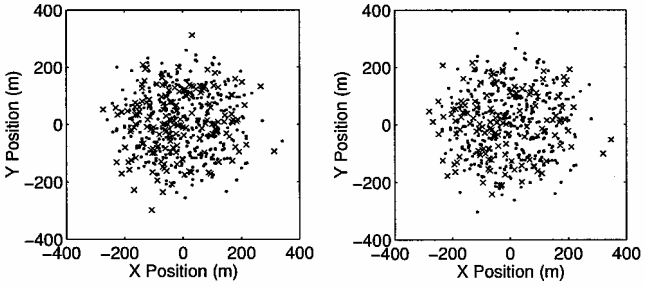

a) Without obstacle avoidance b) With obstacle avoidance

Fig. 8 Landing positions with low control authority: •, safe and ×, hazard.

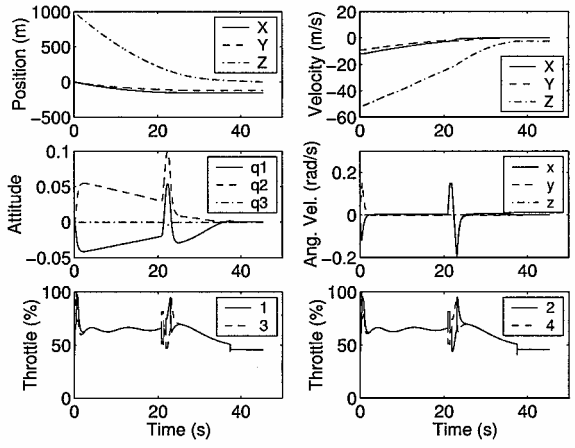

Fig. 9 Example: low control authority obstacle avoidance simulation.

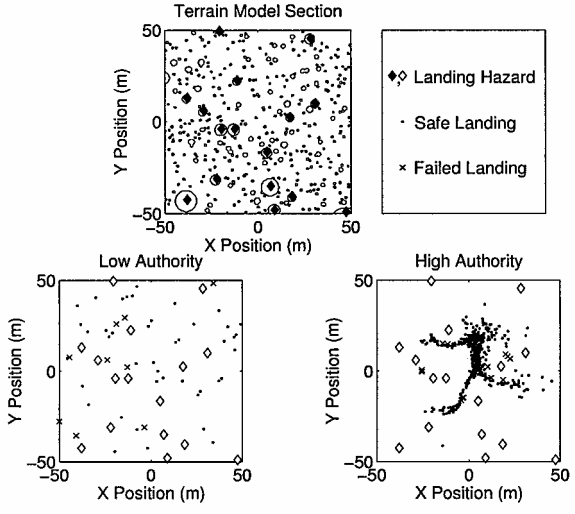

Fig. 10 Landing positions relative to landing hazards.

maneuver was performed. At this point, the lander performed an attitude maneuver. This skewing of the thrust vector caused blips in the lateral velocities. After a period of 1 s, the lander began to rotate back and the lateral velocities began to decay. The result of this was a change in the final position from the baseline case to the obstacle avoidance case.

The results of each case were stored and analyzed in the same manner as the baseline simulations. The position of each landing site is shown in Fig. 8b. A 100 × 100 m section of the landing zone is shown in Fig. 10 with the local landing hazards labeled with diamonds. As with the baseline cases, all 500 of the obstacle avoidance cases satisfied the vehicle state criterion. The OAC had a negligible effect on fuel consumption. On average, the obstacle avoidance cases

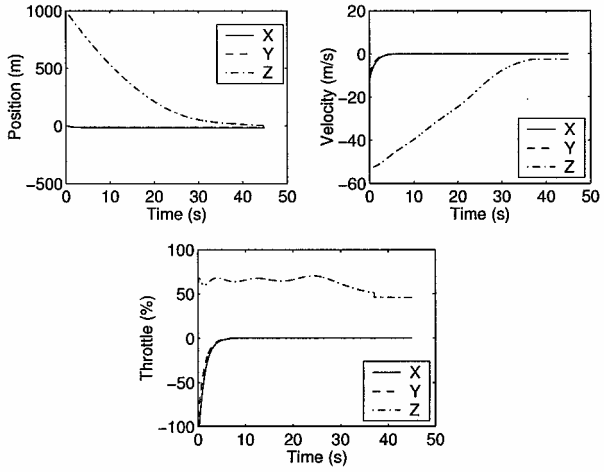

Fig. 11 Example: high control authority baseline simulation.

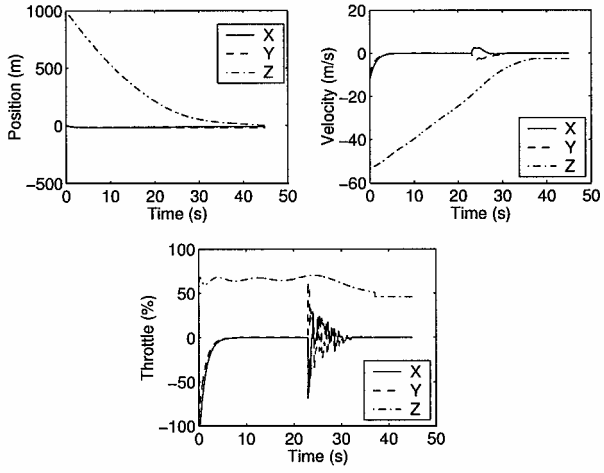

Fig. 12 Example: high control authority obstacle avoidance simulation.

consumed only one additional kilogram of fuel (out of the 3660-kg fuel load) as compared to the baseline cases. Of the obstacle avoidance cases, though, 361 met the landing hazard criteria, whereas 139 landed too close to one or more of the hazardous rocks. These results translate to a 72.2% success rate, or a 9.4% improvement in landing success.

Fly-Away/High-Authority Results

To evaluate the fly-away/high-authority strategy, a baseline simulation was again conducted with the obstacle avoidance disabled. Figure 11 shows an example simulation case. With the high-authority thruster design, attitude maneuvers are not required; therefore, the rotational states are not shown. These results are similar to the low-authority baseline, except that the higher control authority allows quicker convergence of the lateral velocities.

Figure 12 shows the results with the OAC enabled. The vehicle states and throttle histories are similar to the baseline case, except that approximately 20–30 s into the flight the lateral maneuvering is performed. The analysis of the LIDAR data leads to modulation of the X and Y components of thrust. This causes blips in the lateral velocities and a change in the final position from the baseline case to the obstacle avoidance case.

The final landing positions for the high control authority simulations are shown in Fig. 13. The landing positions with obstacle avoidance enabled are compared with the local landing hazards in Fig. 10. Because of the high-control authority, the landing positions are confined in a smaller area than the low-authority cases. This portion of the terrain model includes 16 landing hazards (labeled

Table 3 Safe landing rate summary

Thruster design	Fly-to control, %	Fly-away control, %
Low authority	Not tested	72.2
High authority	100	97.4

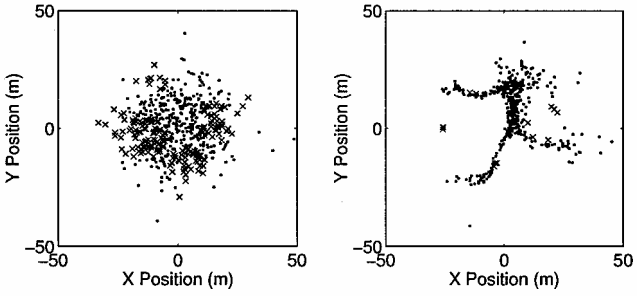

a) Without obstacle avoidance **b) With obstacle avoidance**

Fig. 13 Landing positions with high control authority: •, safe and ×, hazard.

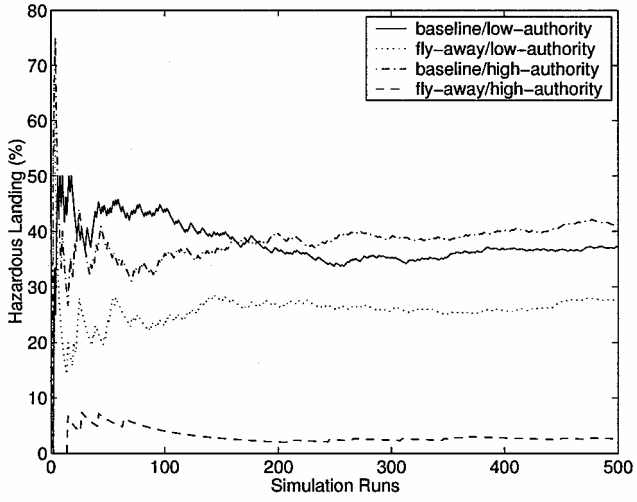

Fig. 14 Monte Carlo convergence of failure rate.

with diamonds) that lie in the plotted area, as well as many smaller rocks.

For both the baseline and obstacle avoidance simulations, all 500 of the simulation cases met the vehicle state criterion. For the baseline simulation cases, 205 landed too close to one or more of the hazardous rocks. Therefore, 59% landed in safe sites, similar to the low-authority baseline case. However, with the obstacle avoidance control turned on, 487 met the landing hazard criteria, whereas only 13 landed too close to one or more of the hazardous rocks. These results translate to a 97.4% success rate, or a 38.4% improvement in landing success. A summary of the success rates for the three strategies considered in this paper is shown in Table 3.

The appropriateness of using 500 Monte Carlo simulations can now be analyzed by looking at the convergence of the predicted failure rate. Figure 14 shows the failure rate that would have been predicted if the Monte Carlo simulation had used fewer than 500 simulations. Figure 14 shows that the predicted failure rates converged by when 500 simulations were used, and, therefore, some confidence can be given as to the precision of these rates.

Conclusions

The baseline simulations predicted safe landing rates of approximately 60% if no obstacle avoidance strategy were used. The three strategies that have been evaluated all showed improvements on the success rate. The fly-away/low-authority obstacle avoidance strategy had a success rate of 72.2%, whereas the fly-away/high-authority strategy had a 97.4% success rate. For a similar terrain model, the fly-to/high-authority strategy of Johnson et al.[4] achieved a 100% success rate for a smaller vehicle. These simulation results allow several conclusions to be drawn regarding the individual strategies and the relative importance of actuator and controller designs.

Although the fly-away/low-authority strategy had the lowest safe landing rate, it provided a 9.4% improvement in the expected success rate for the landing of the reference vehicle. Additionally, the impact of this obstacle avoidance system on the vehicle design would be small, only requiring the addition of the LIDAR sensor system and the obstacle avoidance software. Clearly, this improvement in mission safety would be well worth the required tradeoff in system cost and complexity.

Analysis of the results from the fly-away/high-authority strategy lead to conclusions about its effectiveness. First, the simulation cases that resulted in hazardous landings generally occurred with hazards in the positive X direction relative to the landing site. These hazardous landings were probably a result of the LIDAR sensor model that was used. Sensor scans began along the positive X axis and swept in a counterclockwise direction. The discontinuity along the initial and final direction left the controller partially blinded to obstacles in the positive X direction. In the design of an actual flight sensor, careful attention would be needed to provide the system with a full field of view.

Second, it was shown that the hazardous landings did not occur at the largest hazards in the local area, but near the smaller hazards. The trajectories of the failed cases often showed the OAC initialized with the lander above a location near one of the larger hazards. This prompted the controller to perform large lateral maneuvers. However, as the vehicle rapidly flew over the terrain, it would occasionally reach an area with a new hazard and not have time to react before landing. This tradeoff represents a drawback of the fly-away control strategy. Careful tuning of the gain must be performed to give the vehicle sufficient maneuverability to avoid large obstacles. At the same time, it must not maneuver into new areas without sufficient time to landing for the feedback system to converge on a safe landing site. Further improvements in the fly-away/high-authority strategy could probably have been shown by further tuning of the controller to the specific local terrain. However, this would not have represented the general capability of the strategy.

Clearly, the safe landing rates of 72.7 or 97.4% for the fly-away strategies would not be acceptable for even an autonomous lander phase of a human mission. These results indicate that the fly-to strategy would be a more fruitful area for future research, provided that the high success rate can be maintained for a vehicle such as the MAV. The current results also indicate a strong dependence of landing success on control authority. An important aspect of the future development of the fly-to strategy will be to maintain the safe landing rate near 100% as more realistic engine models are introduced. Additionally, providing high-control authority in the form of gimbaled thrusters or triaxial thruster sets will be of paramount importance in the design of a vehicle such as the MAV if autonomous obstacle avoidance is to be achieved.

References

[1]Hoffman, S. J., and Kaplan, D. L. (eds.), *Human Exploration of Mars: The Reference Mission of the NASA Mars Exploration Study Team*, NASA SP-6107, July 1997, URL: http://exploration.jsc.nasa.gov/marsref/contents.html [cited 19 Jan. 2002].

[2]Drake, B. G. (eds.), *Reference Mission Version 3.0 Addendum to Human Exploration of Mars: The Reference Mission of the NASA Mars Exploration Study Team*, Exploration Office Document EX13-98-036, SP-6107-ADD, June 1998, URL: http://exploration.jsc.nasa.gov/marsref/contents.html [cited 19 Jan. 2002].

[3]Donahue, B. B., and Cupples, M. L., "Comparative Analysis of Current NASA Human Mars Mission Architectures," *Journal of Spacecraft and Rockets*, Vol. 38, No. 5, 2001, pp. 745–751.

[4]Johnson, A., Klumpp, A., Collier, J., and Wolf, A., "LIDAR-based Hazard Avoidance for Safe Landing on Mars," *Spaceflight Mechanics 2001*, Vol. 108, Advances in the Astronautical Sciences, Univelt, San Diego, CA, 2001, pp. 323–338; also American Astronautical Society, Paper AAS 01-120, 2001.

[5]Binder, A. B., Arvidson, R. E., Guinness, E. A., Jones, K. L., Morris, E. C., Mutch, T. A., Peieri, D. C., and Sagan, C., "The Geology of the Viking Lander 1 Site," *Journal of Geophysical Research*, Vol. 82, No. 28, 1977, pp. 4439–4451.

[6]Stentz, A., Bares, J., Singh, S., and Rowe, P., "A Robotic Excavator for Autonomous Truck Loading," *Autonomous Robots*, Vol. 7, No. 2, 1999, pp. 175–186.

[7]Holmberg, N. A., Faust, R. P., and Holt, H. M., "Viking '75 Spacecraft Design and Test Summary," NASA Reference Publ. 1027, 1980.

[8]Payton, D. W., "Internalized Plans: A Representation for Action Resources," *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, edited by P. Maes, MIT Press, Cambridge, MA, 1990, pp. 89–103.

D. B. Spencer
Associate Editor