

Copy 1 of 4

NASA CR-188331

CSDL-C-6303
**AUTONOMOUS HAZARD DETECTION
AND AVOIDANCE**

by
Homer Pien

June 1992

(NASA-CR-188331) AUTONOMOUS HAZARD
DETECTION AND AVOIDANCE Final
Report (Draper (Charles Stark)
Lab.) 53 p

N94-36439

Unclas

G3/61 0018530



The Charles Stark Draper Laboratory, Inc.
555 Technology Square, Cambridge, Massachusetts 02139-3563

CSDL-C-6303

AUTONOMOUS HAZARD DETECTION AND AVOIDANCE
- Final Report -

Homer Pien

June 1992



The Charles Stark Draper Laboratory, Inc.
555 Technology Square, Cambridge, Massachusetts 02139-3563

TABLE OF CONTENTS

1. INTRODUCTION	3
2. BACKGROUND	6
3. APPROACH	10
4. PROCEDURE	13
5. RESULTS	20
6. CONCLUSIONS	28
7. SUMMARY AND FUTURE WORK	30
8. ACKNOWLEDGMENT	32
9. REFERENCES	32
APPENDIX	
A. The Simulation Testbed	33
B. Processing Requirements	38
C. Issues and Lessons	41
D. Sample Images	44
DISTRIBUTION LIST	51

1. INTRODUCTION

During GFY'91, Draper Laboratory was awarded a task by NASA-JSC under contract number NAS9-¹⁸⁴²⁶~~18248~~ to study and evaluate the potential for achieving safe autonomous landings on Mars using an on-board Autonomous Hazard Detection and Avoidance (AHDA) system. This report describes the results of that study¹.

The AHDA task had four objectives, shown in Figure 1.1. The primary goal of the first year's effort is to develop the necessary tools for simulating a closed-loop AHDA landing, while an expected follow-on task for the second year would evaluate the performance of such a system. Due to unforeseen budget cutbacks, however, the follow-on task was canceled. This report summarizes the progress made during the first year, with primary emphasis on describing the tools developed during the year. Some cursory performance evaluation results, originally slated for the second year, are also presented.

- To demonstrate, via a closed-loop simulation, the ability to autonomously select safe landing sites, and the ability to maneuver to the selected site.
- To identify key issues in the development of AHDA systems.
- To produce strawman designs for AHDA sensors and algorithms.
- To perform initial trade studies leading to better understanding of the effect of sensor/terrain/viewing parameters on AHDA algorithm performance.

Figure 1.1. The AHDA Task Objectives

To achieve the objectives in Figure 1.1, a simulation testbed integrating terrain models, sensor models, hazard detection algorithms, site selection algorithms, and terminal descent guidance algorithm was produced. Since past studies have indicated the potential usefulness of optical sensors in hazard detection [Carmer, et. al., 1990], this study placed particular emphasis on modeling passive intensity and laser ranging sensors, as well as assessing their usefulness within a hazard detection system. Lastly, although many factors can influence the performance of an AHDA system, we have elected to assess the impact of several (presumably) dominant parameters on hazard detection. These parameters include terrain type, signal-to-noise ratio, resolution, and

¹ Some of the simulation tools used in this effort were developed under Draper Laboratory's Corporate Sponsored Research Program.

Intensity Image Sensor Parameters

- Terrain Type: Nominal, Rocky
- Spatial Resolution: 0.75 m/pix - 3.0 m/pix
- Signal-to-Noise Ratio: 20 dB - 7.5 dB
- Sun Azimuth Angle: 0 deg - 180 deg
- Sun Elevation Angle: 30 deg - 60 deg

Range Image Sensor Parameters

- Terrain Type: Nominal, Rocky
- Spatial Resolution: 0.25 m/pix - 0.75 m/pix
- Range Noise (1- σ): 0.0 m - 0.2 m

Figure 1.2. Parametric Analysis Variables and Values

sun angle. The range of values considered in this study for each parameter is shown in Figure 1.2.

Results of this study indicate that the use of an on-board hazard detection system can increase the probability of finding a safe landing site. Specifically, we analyzed a scenario in which the intensity sensor is first used to select initial candidate landing regions, and the laser ranger subsequently interrogates the candidate regions in order to select a final landing site. The primary conclusions of this sequential process, *based on a small sample size*, are shown in Figure 1.3.

The remainder of this report will expand upon the procedures used in the analyses, as well as the results obtained. Specifically, Section 2 reviews some background material, and Section 3 discusses the approach that was adopted in order to assess AHDA capabilities. Section 4 describes in some detail the procedures used in this study;

- The availability of safe landing sites depends greatly on number of hazards in the terrain.
- The intensity hazard detector can reduce the ranging laser search space by eliminating areas which appear too hazardous.
- An intensity sensor, operating (without the benefit of the ranging sensor) at 20 dB S/N and 0.75 m/pix resolution, can increase the probability of selecting a safe landing site by a factor of three on rocky terrain.
- With nominal terrain, the landing site selected from intensity images was indicated to be hazard-free by the simulation.
- A laser ranger (when used in conjunction with an intensity sensor) operating at 0.25 m/pix resolution and 0.1 m (1- σ) range noise selected sites that contained no hazards in the case examined in this study.

Figure 1.3. A Summary of the Results

Section 5 shows the results of this study, and Section 6 contains the conclusions. Section 7 contains the summary and a discussion of future work. Acknowledgments and references are given in Sections 8 and 9, respectively. Four appendices provide additional details regarding specific aspects of this study. Appendix A provides a detailed discussion (including derivations) of the various modules in the simulation test-bed. Appendix B shows some back-of-the-envelope calculations of the processing power necessary for real-time support of an AHDA system. Appendix C is an informal discussion of the issues and lessons learned during the development of the simulation test-bed. Lastly, some sample images used in the analysis are shown in Appendix D.

2. BACKGROUND

Future space exploration missions will require the ability to safely land vehicles -- without the aid of humans -- on planetary and lunar surfaces. There are two aspects to this requirement: the ability to perform safe landings on planetary surfaces, and the ability to achieve such landings autonomously. We address each of these problems in the following paragraphs.

To facilitate the analysis of autonomous landings on Mars, a surface model of Mars was created [Moore, 1987]. Three types of models are given: slope frequency distribution, rock size distribution, and mechanical properties for various materials. These models are based on Viking data, previous work, and experience. In particular, surface data from the Viking imaging camera, and data from the Viking infrared thermal mapper, are used to obtain rock size distributions. Similarly, radio echoes from the Martian surface provided slope distribution information. The rock size distributions are given for three cases: smooth, nominal, and rocky.

Using Moore's rock distribution of the Viking II landing site, it has been shown that blind (i.e., unguided) landings on Mars have an unacceptably high probability of failure [Engel, et. al., 1989]. In fact, it was shown that the Viking II lander had a 30% to 40% probability of failure, and a Mars Rover Sample Return (MRSR) vehicle (see below) can have a probability of failure as high as 9%, depending on terrain roughness [Engel, et. al., 1989]. These results have motivated the search for safe landing techniques.

The ability to achieve safe landings *autonomously* is motivated by the exploration missions currently envisioned by NASA. In particular, the Space Exploration Initiative (SEI) will require substantial support from robotic missions [Stafford, 1990]. Although some robotic missions may be accomplished telerobotically (i.e., via earth-based human control), the ~ 40 min round trip transmission delay between Mars and earth makes this option infeasible. Without a man in the loop, these robotic missions must land autonomously.

Two candidate approaches to achieving safe autonomous landings have been proposed: *precision landing* and *hazard detection and avoidance*. In the precision landing case, a set of navigation aids are used to guide the vehicle to the pre-selected safe landing site. In the hazard detection and avoidance case, the vehicle uses its on-board sensors to find a safe site in real-time, and the vehicle maneuvers to the selected site. The primary differences between these two approaches are the *a priori* knowledge requirements, the necessary navigation infrastructure, and the resulting landing accuracies. A vehicle employing AHDA requires very little *a priori* knowledge of the

terrain, and is expected to achieve a landing accuracy of 1 - 2 km with respect to the "nominal" landing spot. Because it needs very little in the way of infrastructure, a vehicle utilizing AHDA is an attractive approach for precursor robotic missions to achieve safe autonomous landings on Mars. The remainder of this report will focus exclusively on the AHDA approach.

A survey of sensors potentially capable of supporting hazard detection was performed by ERIM [Carmer, et. al., 1989]. In this study, sensor options including lasers, synthetic aperture radars, interferometric imaging techniques, and passive intensity sensors were considered. In each case, the ability to detect surface slopes as well as the presence of boulders was analyzed. One recommended sensor option for supporting hazard detection is a sensor suite comprised of laser ranging and passive intensity sensors. This combined system is necessary because, although laser rangefinders are capable of resolving hazards, the slow imaging rate associated with lasers precluded the examination of a sufficiently large landing area for the presence of hazards. This combined laser ranging / passive intensity sensor option is the one we believe to be the most reasonable, and is the option analyzed in this study.

When considering AHDA systems, one issue of significant concern is the availability of safe landing sites within the vehicle's divert footprint. Two separate studies, using different statistical assumptions, have addressed this concern. In one study, where rocks are assumed to cluster (modeled as a normal distribution) around uniformly-distributed cluster centers, a MRSR-style vehicle had between 10% - 50% probability of finding a safe landing site, depending on terrain roughness, pin-point landing accuracy, and degree of clustering [Pien, 1990]. In a separate study, the spatial distribution of rocks was modeled as a random-walk with various correlation distances. This study concluded that there are anywhere from a few to several tens of safe landing sites available within a 250 m x 250 m area [Hain, 1990]. Although both these studies analyzed the existence of landing sites with respect to boulders only, these studies also point to the availability of safe landing sites despite different statistical and modeling assumptions, provided the hazard detection system can find them.

In order to assess the ability of a laser ranger to support a hazard detection system, ERIM performed a 3-D laser simulation incorporating laser sensor models and image processing algorithms [Reiley, 1992]. The study concluded that, in order to have a high probability of correct hazard detection while retaining low false alarm rates, near ideal imaging conditions are necessary. These conditions include near-nadir viewing geometry, high signal-to-noise ratios, and high spatial resolution. Furthermore, the

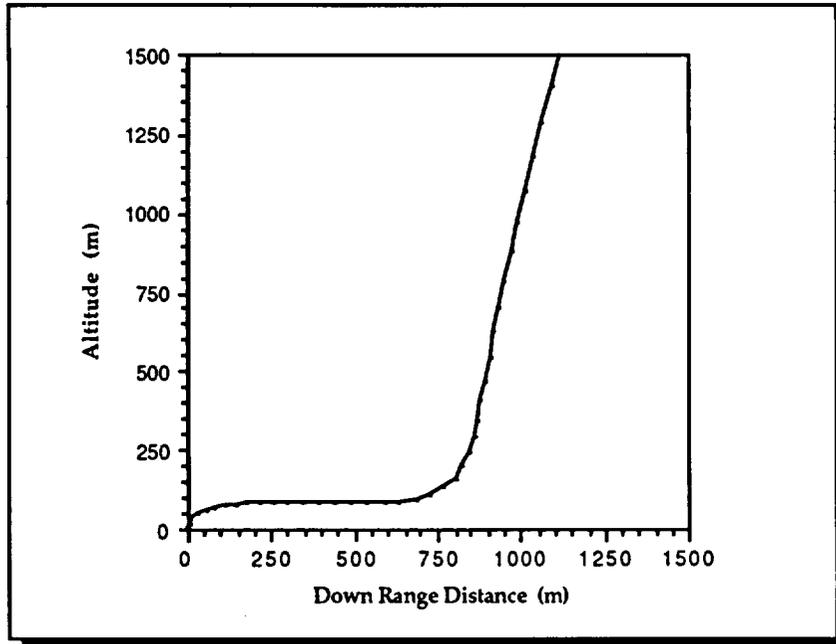


Figure 2.1. A Divert Trajectory Profile

study showed that the probability of false alarm increases dramatically when the imaging conditions deviated from near ideal.

The vehicle trajectory imposes several constraints on the AHDA system as well. Figure 2.1 shows the profile of a near fuel-optimal trajectory that diverted to a landing site approximately 800 m away from the nominal landing site (NLS). This trajectory illustrates the range of aspect angles (relative to the divert site) that may occur during a divert trajectory, reaching as low as 10 deg. Hence if near-nadir viewing geometries are required (as indicated by the ERIM study described above) then site selection must occur higher in the trajectory, the trajectory must be modified at the cost of additional fuel in order to enhanced the viewing geometry, and perhaps divert over a smaller footprint. Another way in which vehicle trajectory and dynamics affect the AHDA system is that the AHDA system must select a new landing site as early in the propulsive descent phase as possible (in fact, it would be desirable for site selection to be performed during the parachute descent phase). This is because the reachable footprint shrinks rapidly as a function of divert commit time, as shown in Figure 2.2 [Brown, 1989]. In particular, the origin of the axes in Figure 2.2 denotes the nominal landing site (NLS). The horizontal axis denotes along-range distance in meters, and the vertical denotes cross-range distances. Three contours indicating the divert footprint, or the "reachable area," are shown. The outermost contour is a full throttle divert from an altitude of 1.4 km at a velocity of 119 m/s; this divert was initiated at time $T = 0$. The middle contour shows

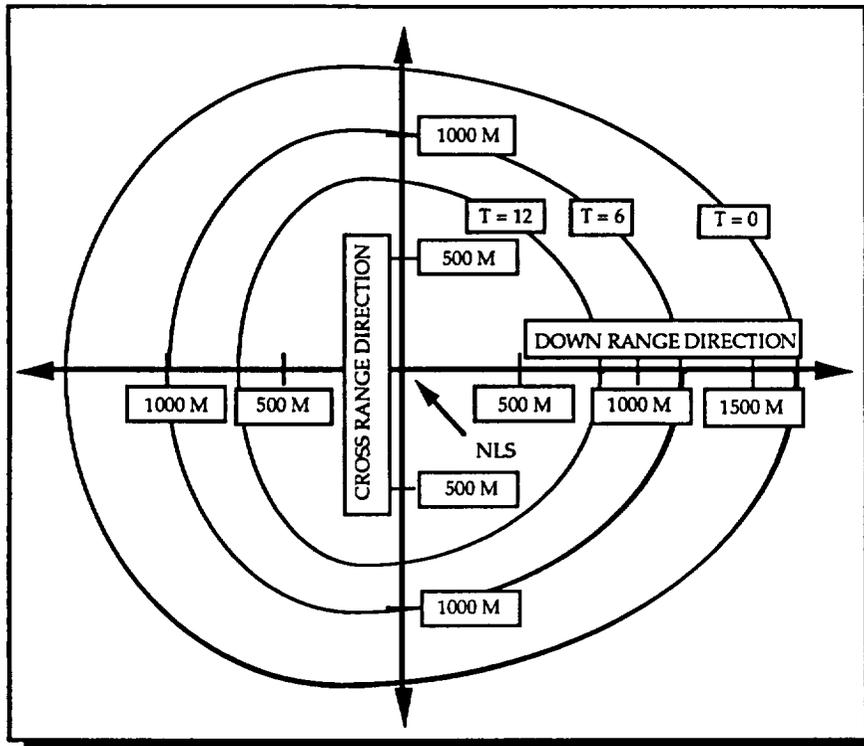


Figure 2.2. Divert Footprint as a Function of Divert Commit Time

the remaining footprint for a divert occurring 6 s ($T = 6$) later, and the innermost divert occurred at $T = 12$ s. As this figure indicates, in order to maximize the available footprint, divert site selection and divert commit must occur as early in the propulsive phase as possible. (Additional details can be found in [Pien, 1991].)

For the purposes of this study a MRSR-style trajectory and vehicle have been assumed [Gamber, 1990]. This MRSR trajectory necessitates the use of an aeroshell during the entry phase for aerobraking. At approximately 6 km altitude, a parachute is deployed. The parachute "yanks" the vehicle away from the aeroshell to begin the parachute descent phase. The parachute phase lasts for approximately 45 sec, at the end of which engines are ignited and the parachute is jettisoned to begin the propulsive (a.k.a. powered) descent phase. The vehicle stays on this propulsive phase for approximately 45 sec, and terminates by performing a soft landing on the surface. In the context of a hazard detection scenario, it is during the propulsive phase that the vehicle has the ability to maneuver toward the site selected by the hazard detection system. Hazards, in this study, refer to boulders or holes greater than 1 m in diameter, as well as slopes steeper than 15 deg over the baseline of the vehicle. The vehicle footprint, i.e., the diametric distance between footpads, is assumed to be 6 m in diameter, and the divert range is assumed to be approximately 1 km radius.

3. APPROACH

In order to address the complex integration issues that arise in an AHDA system, a closed-loop simulation test-bed was built. A high level view of the AHDA closed-loop simulation is shown in Figure 3.1. The simulation takes as inputs a simulated Martian terrain height map, a nominal trajectory, and a set of sensor parameters. The terrain simulation code, produced by JPL, fractally simulates the surface of the Martian terrain, and subsequently emulates various geological effects [Gaskell, 1989]. As output, the AHDA simulation produces the location of the final landing site, as well as the divert trajectory and the sequence of range and intensity images acquired by the vehicle during its descent. To produce the desired output, modules for simulating passive intensity images, laser range images, and relevant sensor effects were created. The simulated sensor images are then given to a hazard detector, which segments the image into safe and hazardous regions. A landing site selector subsequently selects an appropriate landing site on the basis of past information and the current sensor image. The position of the desired landing site is passed to the guidance module, which performs the necessary computations to maneuver the vehicle toward the selected site. The simulation also includes a Mars environment model and a vehicle model. Additional details concerning the simulation are given in Appendix A.

Two types of optical sensors have been selected for consideration in the AHDA system: *laser range* and *passive intensity* sensors. In a passive intensity sensor, a detector array receives light reflected off of the terrain, using the sun as the primary source of illumination. A laser ranger, on the other hand, raster scans the field-of-view and returns, for every detector element on the detector array, the line-of-sight (LOS) range between the detector plane and the corresponding point on the ground. The

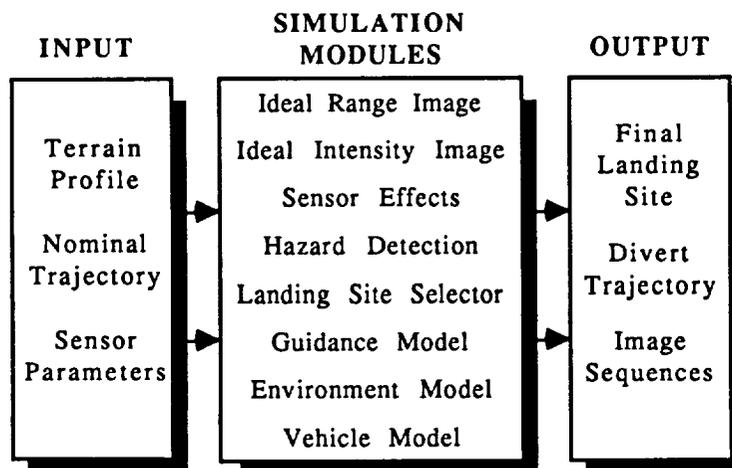


Figure 3.1. Simulation Test-Bed Overview

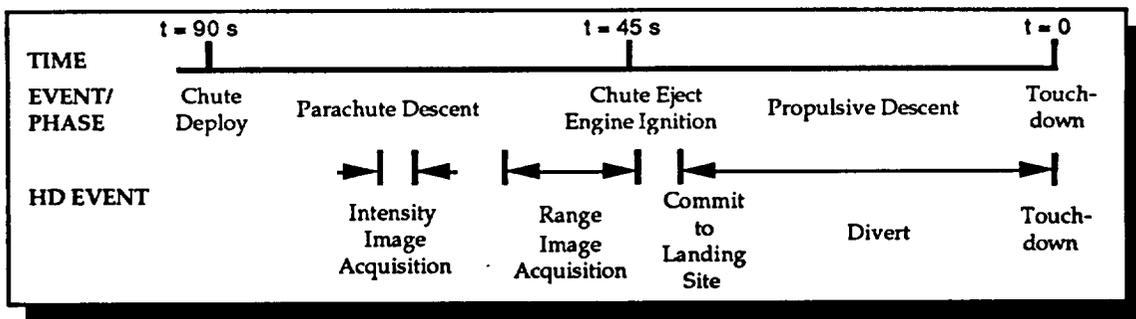


Figure 3.2. Hazard Detection Imaging Events

primary advantage of using a laser ranger is that it provides explicit three-dimensional shape information, whereas one must *infer* shapes when given only an intensity image. The disadvantages of using a laser, however, are complexity and power requirements.

Because of the explicit shape information contained in range images, it is highly desirable to use range information to facilitate hazard detection. Due to the slow imaging rates (as constrained by power consumption and signal-to-noise requirements), however, the utilization of a laser ranger makes sense only over small areas of the terrain, and from relatively close ranges. One alternative is to supplement the laser ranger with a high image-rate passive intensity sensor capable of "screening" potentially safe areas and subsequently reducing the laser search area. This is the approach we have adopted.

Note that the particular approach we have adopted -- that of utilizing intensity images for coarse-scale screening and laser range images for high resolution hazard detection -- reflects our bias towards considering boulders as the primary hazard. If slopes were the primary hazard, then we may adopt an approach in which lasers are first used for coarse-scale slope detection, and the remaining regions are subsequently examined by a high resolution intensity sensor.

The sequence of imaging events shown in Figure 3.2 was selected as the initial approach to analyzing a multi-sensor hazard detection system. *We emphasize that this is only one approach to imaging during vehicle descent; it should not be interpreted as the approach.* We assume that there is some way of reducing the oscillatory motion that occurs on the parachute phase, either via firing the reaction control jets or via the use of flap-type control surfaces. A number of intensity images can be obtained in rapid succession, say with Δt of 0.2 s between images. Assuming each intensity image is acquired at medium resolution (~ 2 m/pix), then at 512 x 512 pixels per image, each image will cover a field of view of approximately 1 km x 1 km. Thus within a small

number of images (on the scale of four), the entire footprint can be covered, within a very short time span (~ 1 s).

As each image is acquired, a contrast-based hazard detection algorithm is used to produce the n most promising landing areas. At the end of the intensity image acquisition process, a queue of n landing areas, in increasing order of contrast, is passed to the vehicle's mission planner. Assuming each 128 x 128 range image requires an image acquisition time of 1 s, then the mission planner will point the laser ranger at each of these n regions during the last n seconds of the parachute phase. (The range image acquisition takes place during the parachute phase so as to maximize the divert range during the propulsive descent phase, as well as to avoid the potentially adverse effects of plume and dust during the propulsive phase. Range image acquisition takes place at the *end* of the chute phase because this maximizes the oscillation damping period and minimizes the line-of-sight range, thus increasing the signal-to-noise ratio and improving the image quality.) The laser ranger will interrogate each region in priority order with 0.25 - 0.75 m/pix resolution. As soon as a sufficiently large region is found to be free of hazards, the process stops. If a large enough site cannot be found, then the vehicle will commit to the largest hazard-free site.

To perform hazard detection, two independent algorithms are used -- one operating on intensity images, the other operating on range images. In the intensity hazard detector case, a simple contrast detector is implemented. More specifically, a moving window is applied across the entire image. At every window location, the intensity standard deviation is calculated. The window location (or equivalently, the corresponding patch of terrain) with the smallest intensity standard deviation is assumed to be the location of the safest landing site. *We emphasize that this is merely an heuristic designed to eliminate regions with high concentrations of large boulders (hence producing greater amounts of bright and shadowed regions, and subsequently increasing the observed intensity variations); it is not meant to be a fool-proof hazard detection algorithm.* For the range hazard detector case, a topographic map is first constructed, from which point-wise slopes are computed and subsequently thresholded to produce safe/hazardous labels. The hazard detection algorithms are described in greater detail in Appendix A.

4. PROCEDURE

Although a conceptual scheme for implementing combined range/intensity hazard detection was described in the previous section, due to simulation limitations (memory requirements and processing time), some modifications were made in practice. We describe in this section the actual simulation parameters used in producing the results.

The coordinate system we chose is a planet-fixed system such that the X-axis points east, Y-axis points north, and the Z-axis denotes altitude. Additionally, the origin of the coordinate system is also used to define the lower left corner of the low-resolution terrain file (the JPL terrain generation program was designed such that all calls to the program must be referenced relative to the lower left corner of the low-resolution terrain file). Vehicle positions cited in this report will be in units of meters relative to this coordinate system

A simple parabolic trajectory was used to simulate the vehicle trajectory; Figure 4.1 shows a plot of the trajectory profile. In particular, a planar trajectory was used, and Figure 4.1 shows the XZ-plane of the trajectory, where the Z-component denotes the

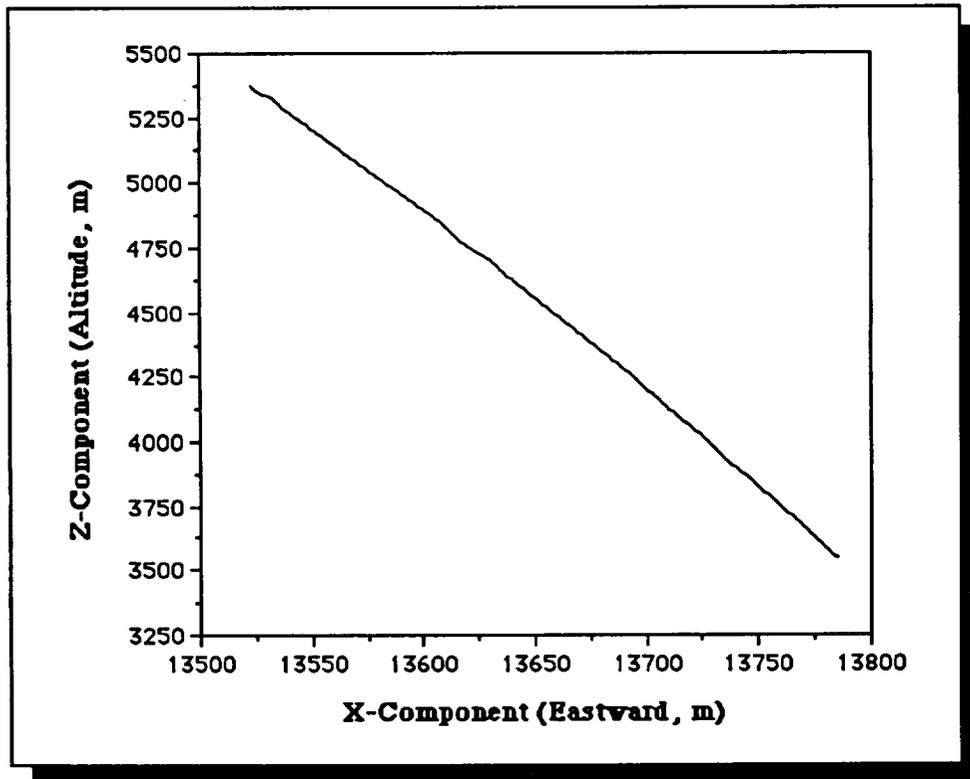


Figure 4.1. Image Sequence Trajectory Profile -- X- and Z-Components of the Vehicle Position Vector

altitude, and the X-component denotes the eastward component of the vehicle position vector; the Y-component was held near constant (see below). No parachute dynamics were simulated for this trajectory. The baseline set of imaging events, corresponding to the conceptual design described in the previous section, is shown in Table 4.1. The first five intensity images bracket the nominal landing site in a cross configuration (see Figure 4.2), and are obtained with a Δt of 0.2 s between images, at 2.0 m/pix resolution. Due to memory limitations, we were unable to make 512 x 512 intensity images, hence

Table 4.1 Image Sequence Trajectory and Image Acquisition Parameters.
Position vector components are in units of meters.

Im #	t	Im Type	Res	Pos-x	Pos-y	Pos-z	Comment
1	1.05	Intensity	2.0	13523	11132	5373	Intensity
2	1.25	Intensity	2.0	13527	11132	5349	Image
3	1.45	Intensity	2.0	13532	11132	5324	Acquisition
4	1.65	Intensity	2.0	13536	11132	5300	Phase
5	1.85	Intensity	2.0	13540	11132	5275	-----
6	5.10	Range	0.7	13607	11133	4854	Range Im Acq Phase
7	5.70	Intensity	2.0	13619	11133	4772	-----
8	6.20	Intensity	2.0	13630	11133	4702	
9	6.70	Intensity	2.0	13639	11133	4632	Descent
10	7.20	Intensity	2.0	13650	11133	4561	Images
11	7.70	Intensity	2.0	13661	11133	4489	for
12	8.20	Intensity	2.0	13671	11133	4416	Error
13	8.70	Intensity	2.0	13681	11133	4342	Checking
14	9.20	Intensity	2.0	13692	11133	4267	and
15	9.70	Intensity	2.0	13702	11133	4191	Image
16	10.2	Intensity	2.0	13712	11133	4114	Sequence
17	10.7	Intensity	2.0	13723	11133	4037	Display
18	11.2	Intensity	2.0	13733	11133	3958	(No Feedback
19	11.7	Intensity	2.0	13744	11133	3878	Given to the
20	12.2	Intensity	2.0	13754	11133	3798	On-board
21	12.7	Intensity	2.0	13764	11133	3717	AHDA
22	13.2	Intensity	2.0	13775	11133	3634	System)
23	13.7	Intensity	2.0	13785	11133	3551	

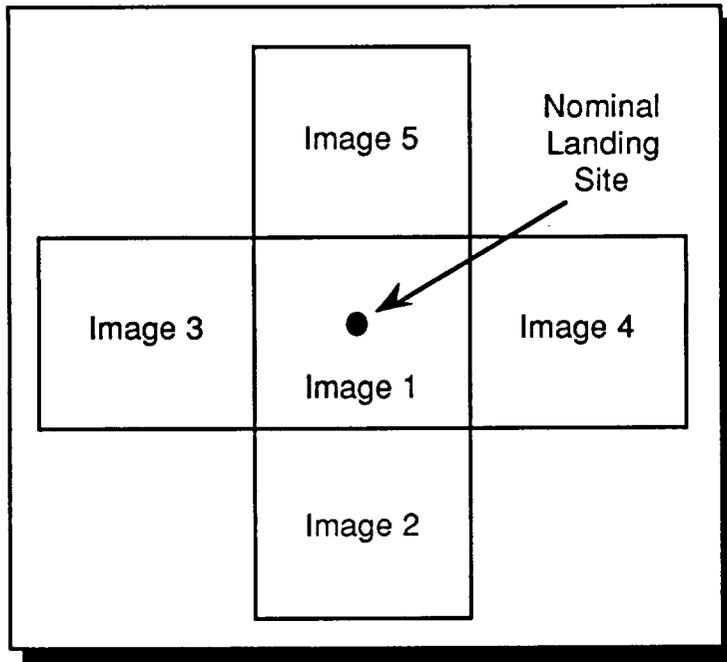


Figure 4.2. Scan Configuration of the First Five Intensity Images

each image in this sequence, be it range or intensity, is 128 x 128 pixels. At the end of the intensity image acquisition phase, a queue of minimum-contrast regions is created. We assume this process takes ≈ 2 s to complete. At the end of the queue-creation phase, the laser ranger is aimed at the best region on the queue, and proceeds to acquire a laser range image at high resolution -- 0.75 m/pix (laser pointing and laser image acquisition is assumed to take ≈ 1 s). In this particular sequence, the first region scanned by the laser was found to be safe, and thus the laser imaging phase is terminated. For the remainder of the sequence, intensity images of the new landing site are taken once every 0.5 s (so that a complete image sequence from the beginning of the hazard detection phase to near-touchdown can be obtained for error checking and display purposes), but no feedback is given to the AHDA system past the laser image acquisition phase.

Note that the sequence shown in Table 4.1 is an east-bound trajectory, since the X-component is increasing while the Y-component is (roughly) constant. This sequence started with a nominal landing site of (14,247 m, 11,134 m, 2,204 m) relative to the coordinate system origin; the selected safe landing site was offset from the nominal by (-24 m, -98 m, +32 m). That is, the selected safe site is approximately 100 m south of the nominal site. To make the simulation generic, the origin of the imaging coordinate

Table 4.2 Intensity and Range Image Viewing Geometries

Intensity Image Viewing Geometry	
• Line-of-sight Range:	4350 m
• Vehicle Elevation Angle:	77 deg
Range Image Viewing Geometry	
• Line-of-sight Range:	3850 m
• Vehicle Elevation Angle:	75 deg

system, and hence the low resolution terrain, can be placed at any location on the surface of Mars.

Because of the execution time associated with a single run of the closed-loop simulation, the effect of sensor degradations on hazard detection performance was considered from only two points along the trajectory -- one for range and one for intensity analysis. These two points are selected as follows. Although Table 4.1 shows 23 images, the selection of the (new) landing site is performed on the basis of the first 6 images -- the five intensity images and the one range image. In fact, since the range image is taken on the basis of the least contrast region selected from image 4, the entire site selection process is really performed on the basis of two images -- images 4 and 6². Thus, parametric analyses (i.e., the analysis of the effects of sensor degradations and illumination conditions on hazard detection algorithm performance) are performed over the same terrain patches as those in images 4 and 6. The imaging geometries of these two images are shown in Table 4.2; the baseline sensor parameters used in the parametric analyses are shown in Table 4.3.

An overview of the procedure used in the performance evaluation process is shown in Figure 4.2. Figure 4.2 is divided into two parts: the closed loop simulation steps are shown on the left, and the parametric analysis steps are shown on the right. The simulation starts with an initialization process, which includes inputting a set of parameters specifying viewing conditions and sensor parameters. After initialization, the simulation clock T is set to some t_0 , and the simulation begins. First, the guidance

²Note that, the fact that all decision making occurred on the basis of images 4 and 6 is an artifact of the particular trajectory and the nominal landing site we selected. Decisions based on different images will likely be the result if different trajectories or nominal landing sites were used.

Table 4.3 Baseline Intensity and Range Image Parameters

Baseline Intensity Image Parameters	
• Line-of-sight Range	4350 m
• Spatial Resolution	0.75 m/pix
• Signal-to-Noise Ratio	20 dB
• Vehicle Elevation Angle	77 deg
• Sun Azimuth Angle (wrt vehicle)	13 deg
• Sun Elevation Angle	45 deg
• Image Plane Size	128 x 128 pix
Baseline Range Image Parameters	
• Line-of-sight Range	3850 m
• Spatial Resolution	0.25 m/pix
• Range Noise (1- σ)	0.1 m
• Vehicle Elevation Angle	75 deg
• Image Plane Size	128 x 128 pixels

and navigation modules produce the vehicle position and line-of-sight (LOS) vectors for time $T = t_0$. From these vectors, the terrain point (the point on the terrain that will become the center of the sensor image) is computed. A high-resolution patch of terrain centered about the terrain point is generated using the sensor parameters input during the initialization step, and a sensor image is created. The hazard detection algorithms are executed next, and if the acquired image is a range image, then site selection is attempted. If a safe site is found, then the vehicle targeting point is updated to reflect the new landing site. If the acquired image is an intensity image, then a priority queue of the least contrast sites is maintained. In either case, the simulation clock is updated and the cycle repeats.

Since sensor degradation analyses are to be performed using the same viewing geometry and field-of-views as images 4 and 6 of the closed-loop simulation, the terrain corresponding to these field-of-views can be generated off-line. However, because the terrain simulation does not incorporate the appropriate Martian boulder distribution as described in [Moore, 1987], a module for imposing the appropriate rock model is necessary. (Note that this rock model is not added to the closed-loop simulation. This is because the rock model, as a module completely separated from the JPL terrain

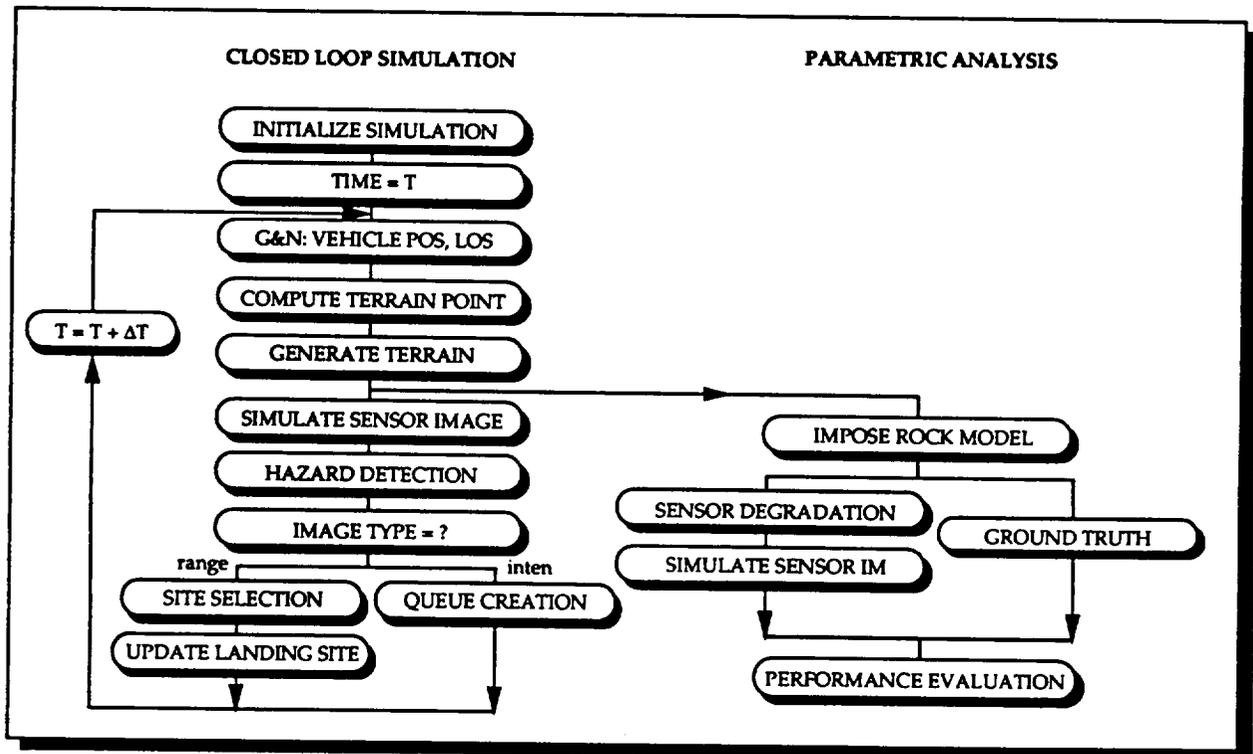


Figure 4.2. Overview of the Parametric Analysis Process

model, is not able to keep boulder locations consistent with resolution changes. Hence this rock distribution model can be used on static images only, and not as part of an image sequence.) Thus, after the creation of the terrain file, a rock distribution model is invoked, and subsequently sensor degradations used in the parametric analysis can be applied, and sensor images incorporating these degradations are created. Similarly, after creating the terrain, a ground truth image is created. The ground truth of an image is a pixel-by-pixel labeling of the image as to whether each pixel is hazardous or not. The ground truth image is obtained by taking a noise-less, nadir-viewed, high resolution range image of the terrain, and subsequently performing a range hazard detection on it. Having created the ground truth, performance evaluations on various degraded images can then be performed.

5. RESULTS

Using the sequence, trajectory, and procedures described in the previous section, analyses of the hazard detection system were performed. Two separate terrain cases were examined: nominal boulder size distribution and rocky boulder size distribution [Moore, 1987]. Furthermore, the performance of both the intensity and the range sensors (as well as the corresponding hazard detection algorithms) are independently analyzed. In particular, intensity hazard detection is used to extract regions that have minimal contrast³, and range hazard detection is used to extract candidate safe landing sites that are at least 5 m in radius⁴ (from the minimum contrast regions).

Figure 5.1 illustrates the relationship among the various regions selected by the two sensors during different parts of the decision making process. In particular, Figure 5.1a

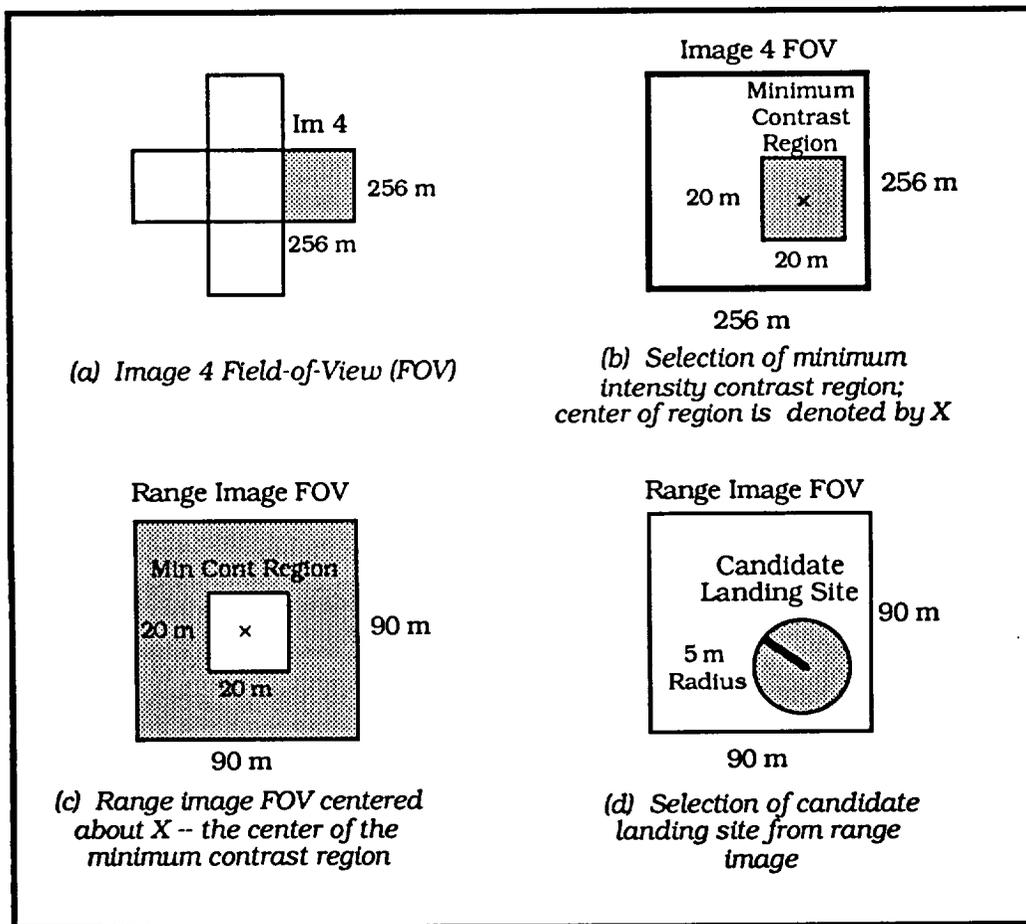


Figure 5.1. Relative Sizes of Terrain Regions Scanned

³ Each intensity region is 20 m x 20 m. The size of this region represents a heuristic upper bound of the error in the vehicle's ability to touchdown at a site selected during the parachute phase when no LOS range information is available.

⁴ A 5 m radius circle represents a heuristic upper bound of the landing error when range information is included as input to the landing site redesignation module.

shows the patch of terrain corresponding to image 4 being scanned as part of the intensity image acquisition process. In Figure 5.1b, a minimum contrast region is selected from the intensity image. Figure 5.1c shows the acquisition of a range image centered about the center of the minimum contrast region. After acquiring the range image, the hazard detection system attempts to select a 5 m radius candidate safe landing site; this process is shown in Figure 5.1d.

To assess the degree of hazard associated with the terrain, a large number of random landing sites are examined, where each landing site is defined to be a 3 m radius circle (i.e., vehicle footprint radius). More specifically, given any region of interest (which can be the 20 m x 20 m area selected by the intensity hazard detector, the 5 m radius circle selected by the range hazard detector, or the entire field-of-view), a suitable number of 3 m radius circles are (uniformly) randomly sampled within the region. If a random sample contains any hazardous pixels, then the corresponding sample is said to have "failed," otherwise its said to be "safe." The ratio of the number of safe samples to the total number of samples is an indication of the degree to which the region is considered safe (or hazardous), and is termed the "safeness" of the region.

Because the range sensor scans the patch of terrain centered about the least contrast intensity region, one would expect the range image field-of-view to be safer than the intensity field-of-view. Table 5.1 compares the safeness values for the intensity field-of-view versus the range field-of-view, for both the nominal and the rocky cases. In particular, note that the intensity contrast detector selected a region that is 100% safe in the nominal case.

Table 5.1. FOV Safeness as a Function of Terrain Model

	<u>Intensity FOV</u>	<u>Range FOV</u>
Nominal	94.6%	100.0%
Rocky	72.2%	82.6%

Results are shown in groups. Section 5.1 contains results of the intensity image analysis; Section 5.2 contains range image analysis results. In each case, results are shown separately for nominal and rocky terrains. In the intensity case, two performance numbers are given: because contrast is used as an heuristic indicator of terrain safeness, the first number (labeled Best Site) shows the safeness associated with the 20 m x 20 m region selected by the intensity contrast detector as having the *least* contrast. The second number, labeled Average, is the *average* safeness of the *five* regions selected as having the least intensity contrast. Note that because *safeness does not necessarily vary monotonically with intensity contrast*, since contrast is merely

an heuristic indicator of safeness, it is possible for the 5-region average safeness to be better than the safeness of the best site.

In the range hazard detection case, two numbers are also used to provide slightly different insights into algorithmic performance: the safeness and size of the selected site. Clearly, *the safeness and size of the selected site are not independent*, since the probability of finding a safe landing site depends on the size of the landing site. Therefore, both numbers are given as indicators of the quality of the selected site. If the site radius exceeds 5 m, then a safeness assessment is made. A selected site is considered too small if it is less than 5 m in radius, hence safeness is not computed for such sites.

5.1 INTENSITY IMAGE ANALYSIS RESULTS

Results of the intensity image analysis based on the *nominal* boulder size distribution are shown in Tables 5.2 - 5.5, where performance as a function of resolution degradation is shown in Table 5.2; performance as a function of S/N is shown in Table 5.3⁵; performance as a function of sun azimuth angles is shown in Table 5.4, and performance as a function of sun elevation angles is shown in Table 5.5⁶. For each degraded image, the image resolution, signal-to-noise ratio (S/N), sun azimuth angle, and sun elevation angle are shown. Additionally, the safeness of the minimum contrast region, as well as the averaged safeness of the 5 minimum contrast regions, are shown. The safeness corresponding to these regions should be compared to the safeness of the terrain corresponding to the entire field-of-view, which is 94.6% (see Table 5.1).

Tables 5.6 - 5.9 show the performance of the intensity hazard detector on *rocky* terrain. Table 5.6 shows performance as a function of resolution degradation; Table 5.7 shows performance as a function of S/N degradation; Table 5.8 shows performance as a function of sun azimuth angles, and Table 5.9 shows performance as a function of sun elevation angles. Again, performance numbers should be compared to 72.2%, which is the safeness of the terrain corresponding to the entire field-of-view (from Table 5.1).

⁵ Signal-to-noise, for this study, is defined as the ratio of the peak signal to the average noise.

⁶ The sun azimuth angle is defined in world coordinates. Thus, an azimuth of 0 deg has the sun aligned with the X-axis, which points eastward. The sun elevation angle is defined with respect to the azimuth. Thus an elevation angle of 135 deg when the azimuth is 0 deg is equivalent to an elevation angle of 45 deg when the azimuth is 180 deg.

Table 5.2. Intensity HD Performance on Nominal Terrain -- Resolution Degradations

Resol	S/N	Sun Az	Sun El	Best Site	Average
0.75 m/p	20 dB	269 deg	45 deg	100%	100.0%
1.5 m/p	20 dB	269 deg	45 deg	100%	97.8%
2.25 m/p	20 dB	269 deg	45 deg	100%	95.4%
3.0 m/p	20 dB	269 deg	45 deg	100%	96.8%

Table 5.3. Intensity HD Performance on Nominal Terrain -- S/N Degradations

Resol	S/N	Sun Az	Sun El	Best Site	Average
0.75 m/p	20 dB	269 deg	45 deg	100%	100.0%
0.75 m/p	15 dB	269 deg	45 deg	100%	100.0%
0.75 m/p	10 dB	269 deg	45 deg	100%	100.0%
0.75 m/p	7.5 dB	269 deg	45 deg	100%	100.0%

Table 5.4. Intensity HD Performance on Nominal Terrain -- Sun Azimuth Angle Variations

Resol	S/N	Sun Az	Sun El	Best Site	Average
0.75 m/p	15 dB	0 deg	45 deg	100%	100.0%
0.75 m/p	15 dB	45 deg	45 deg	100%	100.0%
0.75 m/p	15 dB	90 deg	45 deg	100%	100.0%
0.75 m/p	15 dB	135 deg	45 deg	100%	100.0%
0.75 m/p	15 dB	180 deg	45 deg	100%	100.0%

Table 5.5. Intensity HD Performance on Nominal Terrain -- Sun Elevation Angle Variations

Resol	S/N	Sun Az	Sun El	Best Site	Average
0.75 m/p	15 dB	45 deg	30 deg	100%	100.0%
0.75 m/p	15 dB	45 deg	45 deg	100%	100.0%
0.75 m/p	15 dB	45 deg	60 deg	100%	100.0%
0.75 m/p	15 dB	45 deg	120 deg	100%	97.4%
0.75 m/p	15 dB	45 deg	135 deg	100%	100.0%
0.75 m/p	15 dB	45 deg	150 deg	100%	100.0%

Table 5.6. Intensity HD Performance on Rocky Terrain -- Resolution Degradations

Resol	S/N	Sun Az	Sun El	Best Site	Average
0.75 m/p	20 dB	269 deg	45 deg	95%	91.0%
1.5 m/p	20 dB	269 deg	45 deg	94%	84.6%
2.25 m/p	20 dB	269 deg	45 deg	85%	68.0%
3.0 m/p	20 dB	269 deg	45 deg	100%	78.0%

Table 5.7. Intensity HD Performance on Rocky Terrain -- S/N Degradations

Resol	S/N	Sun Az	Sun El	Best Site	Average
0.75 m/p	20 dB	269 deg	45 deg	95%	91.0%
0.75 m/p	15 dB	269 deg	45 deg	95%	84.6%
0.75 m/p	10 dB	269 deg	45 deg	95%	82.4%
0.75 m/p	7.5 dB	269 deg	45 deg	95%	89.4%

Table 5.8. Intensity HD Performance on Rocky Terrain -- Sun Azimuth Angle Variations

Resol	S/N	Sun Az	Sun El	Best Site	Average
0.75 m/p	15 dB	0 deg	45 deg	95%	88.6%
0.75 m/p	15 dB	45 deg	45 deg	100%	88.6%
0.75 m/p	15 dB	90 deg	45 deg	100%	87.2%
0.75 m/p	15 dB	135 deg	45 deg	95%	87.8%
0.75 m/p	15 dB	180 deg	45 deg	95%	96.8%

Table 5.9. Intensity HD Performance on Rocky Terrain -- Sun Elevation Angle Variations

Resol	S/N	Sun Az	Sun El	Best Site	Average
0.75 m/p	15 dB	45 deg	30 deg	95%	84.8%
0.75 m/p	15 dB	45 deg	45 deg	100%	88.6%
0.75 m/p	15 dB	45 deg	60 deg	99%	88.6%
0.75 m/p	15 dB	45 deg	120 deg	95%	88.8%
0.75 m/p	15 dB	45 deg	135 deg	95%	85.6%
0.75 m/p	15 dB	45 deg	150 deg	95%	83.6%

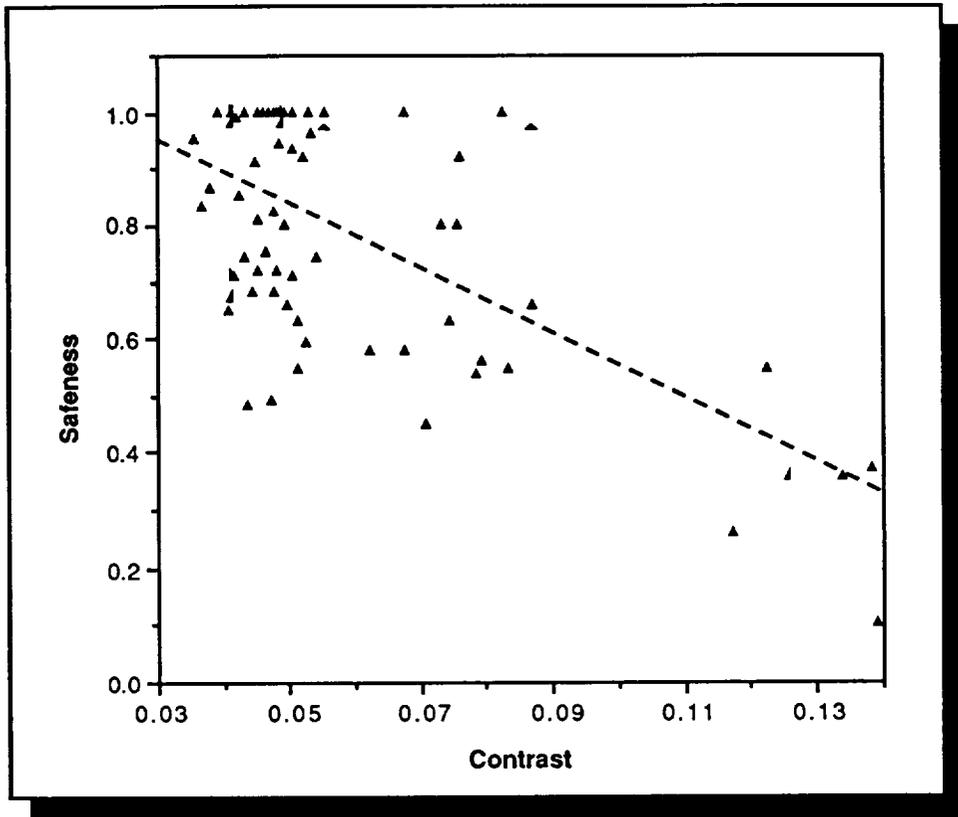


Figure 5.2. The Correlation Between Contrast and Safeness

Recall in general, contrast will not vary monotonically with safeness. Figure 5.2 shows the relationship between contrast and the safeness of each region (i.e., for each 3 m radius random landing attempted), for the case of rocky terrain, 0.75 m/pix resolution, and 20 dB S/N. That is, Figure 5.2 corresponds to the case shown in the first row of Table 5.6, in the sense that the Best Site value of 95% corresponds to the leftmost point (i.e., the least contrast point) in Figure 5.2, and the Average value of 91% is the average of the 5 leftmost points in Figure 5.2. Lastly, a dashed line indicating the least-squares fit is shown. The correlation coefficient of the data is 0.417, which, with degree of freedom > 70, is 99.9% significant. This implies the null hypothesis (i.e., the hypothesis that the observed correlation is due purely to chance) can be rejected, and the observed correlation between safeness and contrast is statistically significant.

5.2 RANGE IMAGE ANALYSIS RESULTS

Although intensity image analyses were performed for nominal and rocky terrain cases, range hazard detection performance is shown for the rocky terrain case only. This is because the range field of view of the nominal case *for the particular trajectory and nominal landing site we selected*, having already passed the intensity hazard

detector, contained no hazards, and thus *all* regions examined were safe. This was shown in Table 5.1.

Table 5.10 shows the effect of range noise and resolution on the range hazard detector *for the rocky terrain case*. Two numbers are shown in each cell -- the size of the selected site is shown as the top number, and the safeness of the 5 m region about the center of the selected site is shown as a percentage in parenthesis. A site is selected when it is the largest circle within the field-of-view that is free of any (detected) hazards. When the selected site is less than 5 m in radius, region safeness is not calculated.

It should be noted that the values shown in Table 5.10 are obtained from the range hazard detection algorithm threshold set at 0.9 m hazards, instead of 1 m hazard (see Appendix A for a detailed description of the hazard detection algorithm and threshold selection). This is because we felt that, in an operational system, tolerance would have to be built into the hazard detection system. Figure 5.3 shows, for the 0.25 m/pix case, the differences in the radii of the selected landing sites that occur when the hazard detection system is set for 0.6 m hazards instead of 0.9 m.

It is worthwhile describing the errors that occurred on the 0.5 m/pix case of Table 5.10. In this case, there was a small hazard (with only 1 hazardous pixel in the ground truth) that was missed. Subsequently, the selected site contained this pixel. The effects of this missed detection are 1) unsafe landing areas exist within the selected region, and 2) the radius of the selected region became larger (14 m from 10.5 m). It turns out that although the 0.75 m/pix case also failed to locate the same hazard, because the site selector chose a different landing region, no unsafe landing areas existed within this region. When the hazard detection threshold is lowered to 0.6 m, the 0.5 m/pix case was able to designate the hazardous area correctly, and subsequently selected a landing site that was safe (see Table 5.11).

Table 5.10. Range Hazard Detection Performance on Rocky Terrain

Resolution	Range Noise (m)				
	0.00	0.05	0.10	0.15	0.20
0.25 m/pix	10.5 m (100%)	10.5 m (100%)	10.5 m (100%)	2.3 m (-)	1.5 m (-)
0.50 m/pix	14.0 m (46.7%)	14.0 m (40.0%)	10.5 m (100%)	3.5 m (-)	2.9 m (-)
0.75 m/pix	10.5 m (100%)	10.5 m (100%)	8.9 m (100%)	4.5 m (-)	2.3 m (-)

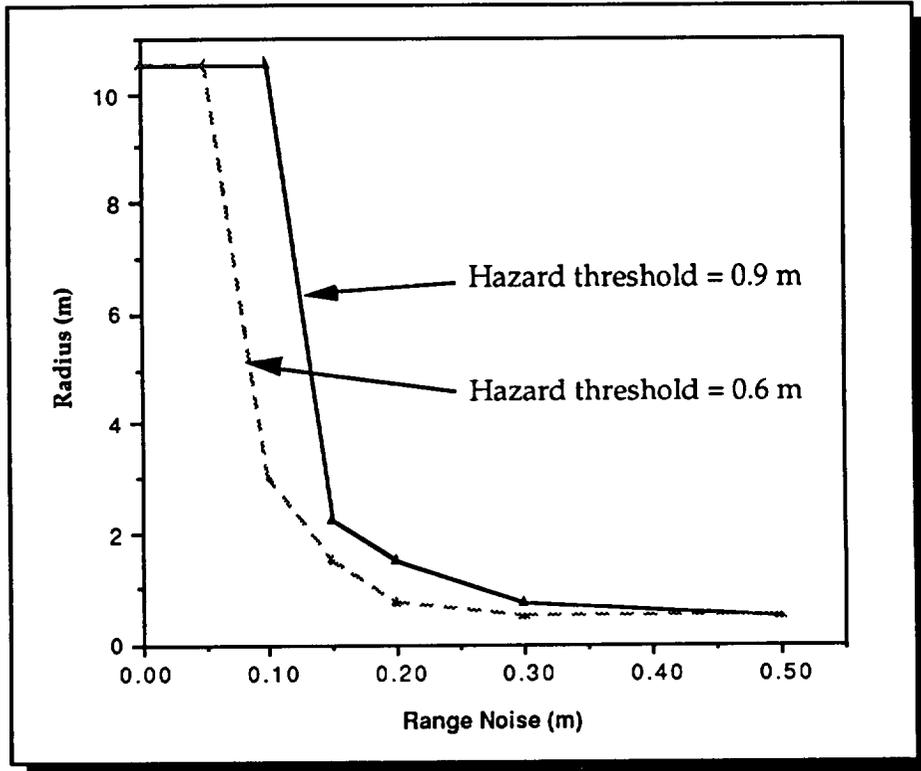


Figure 5.3. The Effect of Varying Hazard Detection Threshold on the Size of the Selected Landing Site Radius.

Table 5.11. The Effect of Hazard Threshold on Site Selection for the Noiseless Case

Resolution	Threshold = 0.6 m	Threshold = 0.9 m
0.25 m/pix	10.5 m (100%)	10.5 m (100%)
0.5 m/pix	10.0 m (100%)	14.0 m (46.7%)
0.75 m/pix	10.5 m (100%)	10.5 m (100%)

6. CONCLUSIONS

Several conclusions can be made from the results shown in the previous section. We start with the intensity hazard detector.

From Table 5.1, it is apparent that using the intensity hazard detector as a screener can provide a useful reduction of search space for the laser ranger. In fact, on nominal terrain, the intensity hazard detector was able to select regions that contained no hazardous pixels. Even on rocky terrain, the utilization of the intensity sensor in isolation provided good performance. For example, operating the intensity sensor at 0.75 m/pix resolution at 20 dB S/N, the intensity hazard detector selected regions that were consistently above 90% safe, from terrain that was only 72.2% safe. More specifically, the first row of Table 5.6 shows a case in which the average safeness of the five sites with the least contrast is 91.0%, which translates to a *three-fold reduction in the probability of selecting hazardous landing sites* (from 27.8% to 9%). When considering only the "best site" probabilities, the reductions are even more dramatic. It has also been shown that the intensity contrast is positively correlated to the safeness of each contrast region (Figure 5.3). Thus, even without the utilization of a laser ranger as the second stage of hazard detection, using contrast as an indicator of terrain safeness is superior to random (blind) landings.

As can be seen from Table 5.6, resolution has a noticeable effect on the performance of the intensity hazard detector. The signal-to-noise ratio (Table 5.7), however, seems to have a considerably less significant impact, presumably because the noise is "averaged over" when computing the contrast of a region. The effect of sunlight angles (Tables 5.8, 5.9) on performance is less clear. Previous studies have indicated that a sun elevation angle of near 90° creates problems for the hazard detection system [Pien, 1990]. When the range of elevation angles is restricted to between 30° to 60°, however, the impact is not noticeable. Similarly, no clear relationship between the sun azimuth angle and performance was observed.

In the range hazard detection case, results are easier to interpret. As shown in Table 5.10, when resolution is coarser than 0.25 m/pix, missed detections occur. When the range noise exceeds 0.1 m, the algorithm has noticeable difficulties extracting safe sites of appropriate sizes. Alternatively, one can decrease the hazard detection threshold so as that the algorithm labels smaller objects as hazards. This approach can be used to increase the performance of the system, but at the expense of decreasing the size of the selected sites.

Table 5.12. Strawman Sensor Design Parameters

Intensity Image Sensor Parameters	
• Spatial Resolution:	0.75 m/pix
• Signal-to-Noise:	15 - 20 dB
Range Image Sensor Parameters	
• Spatial Resolution:	0.25 m/pix
• Range Noise (1-s):	0.1 m

Based the above results and conclusions, a strawman set of sensor parameters can be formulated, as shown in Table 5.12.

Finally, a detailed calculation showing the processing requirements of this AHDA system is shown in Appendix B. The calculations show that, compared to an off-the-shelf Sun workstation, the necessary *intensity* image processing was easily achievable. For *range* image processing the worst case analysis showed that slightly less than 1 s was needed for the processing, whereas 1 s was allocated to the range image processing in our image sequence. We therefore believe the AHDA system, as described, does not pose any severe requirements on the processing power of the on-board computer system.

7. SUMMARY AND FUTURE WORK

In summary, the ability to achieve safe autonomous landings on Mars by utilizing an Autonomous Hazard Detection and Avoidance system was demonstrated *for a very small sample size* via the Autonomous Landing Simulation Testbed. In particular, of the cases examined, the simulation showed that improvements in the ability to select safe landing sites can be achieved by using an intensity contrast detector when compared to a blind landing system. Furthermore, when the intensity sensor is operated in conjunction with a high resolution laser ranger, the simulation indicated that the AHDA system is capable of selecting final landing sites that contain no hazards. The effect of various imaging parameters on the performance of the hazard detection system was evaluated, and a strawman sensor design was provided. Lastly, it was also shown that the processing power necessary for supporting such an AHDA system is well within the capabilities of current off-the-shelf computing hardware.

Despite the encouraging results, there are several weaknesses within the simulation. These are topics we plan to address as part of our future work, and we summarize them below.

First, the results shown were derived from a very small sample size. Specifically, the results were obtained from a single baseline trajectory over a single patch of terrain (although the terrain was modified by two different boulder models). Similar results were obtained when the simulation processed trajectories over other landing sites, but these other landings were not analyzed to the same extent and their results were not incorporated into this report.

Second, although a slope hazard detector was used for producing the results in this report, it should be improved. This is because the slope hazard detector as described in Appendix A is simply an extension of the boulder hazard detector, and hence is not really designed for slope detection. A more robust slope detector should be developed.

Third, it should be noted that the simulation testbed was not intended to be a high fidelity simulation. Specifically, no parachute nor touchdown dynamics were modeled. Additionally, the sensor simulations produced "normalized" outputs in the sense that the received power and noise is measured relative to the peak signal, and thus the simulation did not incorporate "real" system parameters such as power, photon noise, etc. Furthermore, Poisson intensity shot noise, range ambiguity functions, and range quantization were turned off in order to reduce the number of parameters in the study (see Appendix A). Lastly, the terrain simulation produced artifacts which, at times,

were clearly visible. A more robust terrain generation program is undoubtedly needed in the future.

Fourth, the work described in this report focused exclusively on one trajectory type. Alternative trajectories, such as those in which the vehicle enters the Martian atmosphere at a considerably steeper angle, and ignites the engines at higher altitudes (see [TRW, 1990]), should also be considered.

Lastly, because of interest in autonomous landings on the lunar surface, the Autonomous Landing Simulation Testbed should be modified so that it can incorporate boulder and crater distribution models appropriate for the lunar surface.

8. ACKNOWLEDGMENT

We gratefully acknowledge the support of NASA-JSC and Draper Laboratory's Corporate Sponsored Research Program. Additionally, the support of Tim Brand and Ken Baker (JSC) are appreciated.

The Autonomous Landing Simulation Testbed contains the work of many people. The sensor simulation, hazard detection, and site selection modules were written by David Oh, Todd Perlmutter, and Eric Flatt. The guidance and navigation modules were written by Ken Spratlin, Tony Bogner, and Tom Fill. Finally, Ken Spratlin, Tony Bogner, and Eric Flatt were responsible for the software integration.

Lastly, we wish to thank Judy Miller, Tim Brand, and Ken Baker for reading earlier drafts of this report; Bob Gaskell (JPL) for providing the terrain simulation program, and Duane Carmer and Mike Reiley (both of ERIM) for providing information on sensor capabilities early in the Autonomous Landing program.

9. REFERENCES

- Carmer, D., M. Trichel, M. Eismann, K. Gleichman, Y. Morita, P. Tchoryk, *Final Report on the Sensor Trade Study: Autonomous Hazard Detection and Avoidance*, ERIM 213400-50-F, July 1990.
- Engel, A., T. Bogner, T. Dierlam, and K. Spratlin, *Autonomous Landing Study -- Probability of Safe Landing*, Draper Laboratory MRSR-89-018, July 1989.
- Gamber, R. T., *Mars Rover Sample Return: Delivery and Return Study Interim Final Report*, Martin Marietta Astronautics Group, May 1990.
- Gaskell, R., *Preliminary Simulated Martian Surfaces*, JPL Engineering Memorandum 314-467, 19 July 1989.
- Hain, R., *Safe Landing Area Availability on the Surface of Mars*, Draper Laboratory LMI-90-32, June 1990.
- Moore, H., *Preliminary Mars Surface Models*, 20 April 1987.
- Pien, H., *Monte Carlo Analysis of a Simple Hazard Detection System*, Draper Laboratory LMI-91-008, May 1990.
- Pien, H., "Autonomous Hazard Detection and Avoidance for Mars Exploration" in *AIAA Computing in Aerospace 8*, AIAA, Washington D. C., 1991.
- Reiley, M., D. C. Carmer, W. F. Pont, and C. D. Lake, *Autonomous Hazard Detection: 3-D Laser Radar Simulation and Parametric Study*, ERIM 216900-13-F, Feb 1992.
- Stafford, T., *America at the Threshold: America's Space Exploration Initiative*, Report of the Synthesis Group on America's Space Exploration Initiative, May 1991.
- TRW, *Mars Rover Sample Return: Delivery and Return Study -- Interim Report*, TRW Space and Technology Group, Redondo Beach CA, June 1990.

APPENDIX A. The Simulation Testbed

To simulate images, a sensor LOS direction is first calculated as a function of vehicle orientation and sensor tilt angles. By combining the LOS vector with the vehicle's trajectory, instantaneous FOVs can be computed. The ideal image and sensor simulations are shown in Figure A1. To produce range images, z-buffers are used to facilitate shadowing (i.e., occlusion). To produce intensity images, z-buffers are used in conjunction with a Lambertian scattering model [Foley and Van Dam, 1982]. To simulate sensor effects, a convolution with the Airy function (if the sensor is diffraction limited) and a detector-size rectangle function is performed in order to simulate both the circular lens and rectangular detector response functions. Noise is subsequently added (Gaussian noise for range images and Poisson noise for intensity), followed by

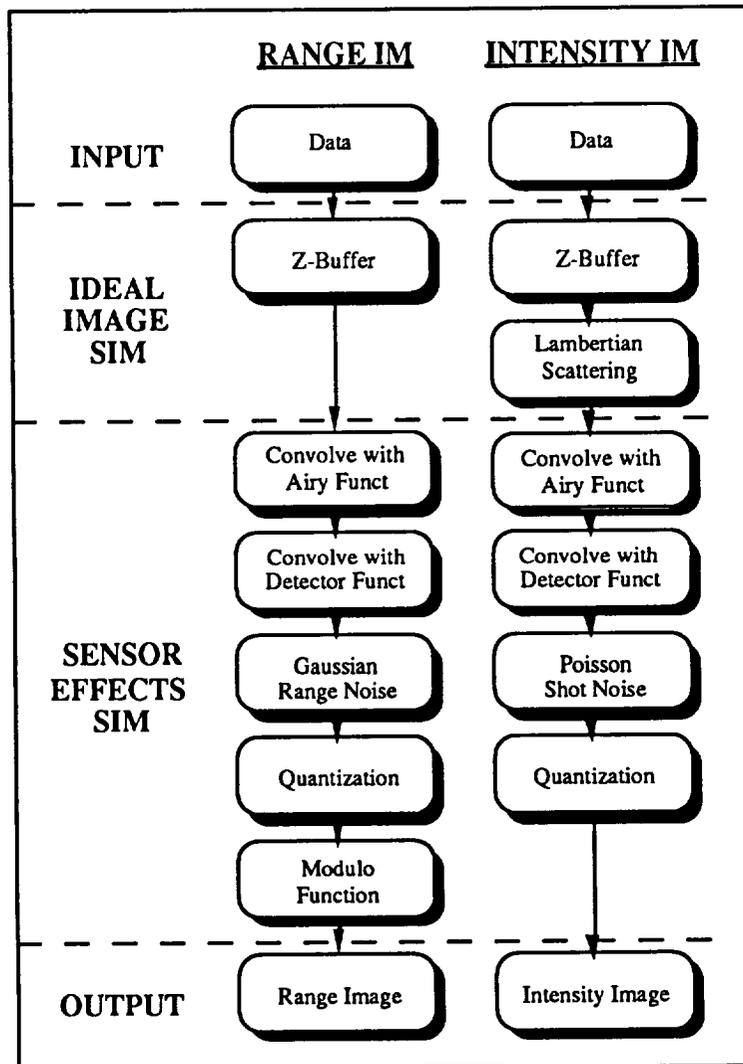


Figure A1. Image Simulation Modules

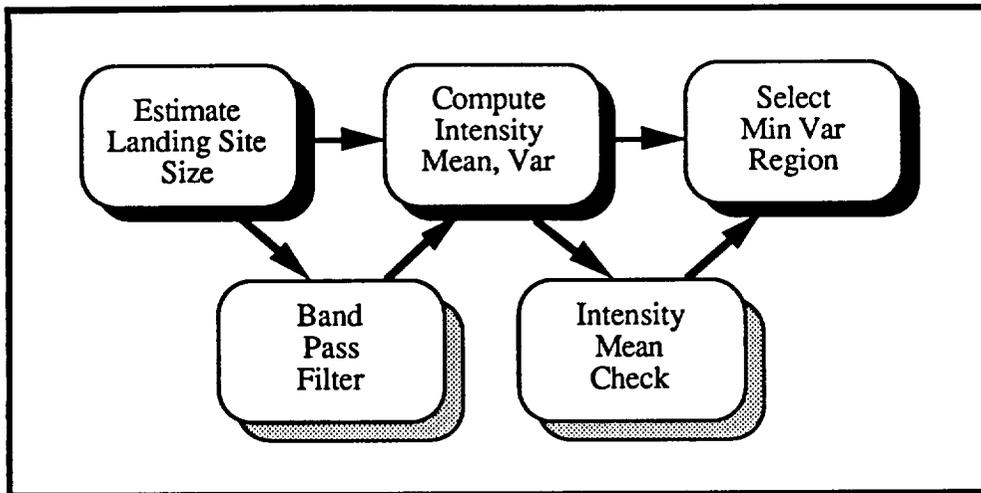


Figure A2. An Intensity-Based Hazard Detector

a quantization step. Lastly, a modulo function is performed on range images, if necessary, for simulating range ambiguities. For the intensity images used in the parametric studies, Gaussian noise was used for computational efficiency reasons. For the range images, the quantization and range ambiguity simulation steps were eliminated in order to reduce the number of study parameters.

Currently, hazard detection for range and intensity images is performed independently. Several algorithms have been implemented for hazard detection based on intensity images. The primary algorithm measures the intensity variation of each region within the image and subsequently selects the regions with small intensity variations as candidate landing sites (see Figure A2). More specifically, if, from some point along the trajectory, the vehicle selects a spot on the ground as the landing site, a heuristic is used to estimate the vehicle's ability to land at that selected spot (as a function of slant range distance to the selected spot). This heuristic bound on the landing accuracy defines the size of the hazard-free region that must be found. To find regions with small intensity variations, a sliding window of size determined by the above heuristic is moved across the intensity image (windows are 50% overlapped). For each window, the intensity mean and variance are calculated and stored. After this process has been performed over the entire image, the windows with the smallest intensity variance are chosen as candidate landing sites. In practice, since intensity images are acquired within a short interval, a constant window size of approximately 20 m x 20 m is used.

The incorporation of two additional steps (shown as the gray boxes in Figure A2) were also examined. First, the image undergoes a band-pass filtering step that removes intensity variations introduced by objects too large (i.e., low frequency components) or too small (i.e., high frequency components) to be of concern to the hazard detector. The band-pass filter used is the difference of two Butterworth low-pass filters [Jain, 1989]. One side-effect of using the band-pass filter is that it eliminates the need for a separate noise reduction low-pass filter. The second step utilized is a check on the mean intensity of every window in order to eliminate the possibility of selecting a totally shadowed region (i.e., a uniformly dark region) as a candidate landing site. That is, the mean intensity of a window must exceed some threshold in order to be selected as a candidate site. In practice, the improvement in performance achieved via the use of band-pass filters were negligible, and hence was not used as part of the parametric analyses. Additionally, a threshold of 0.1 was imposed on the intensity mean in order to eliminate the possibility of selecting completely shadowed regions as landing sites.

For range images, a preprocessing step is first performed prior to invoking the range hazard detector. This preprocessing consists of: 1) range ambiguity removal, 2) moving average noise filter, and 3) rotation to a top-down viewing perspective. The purpose of the rotation is to obtain an aspect-independent representation of the terrain, corresponding roughly to a topographic or height profile of the terrain. More specifically, assuming the image is acquired from a camera with a roll angle of θ (i.e., θ is the angle between the horizontal axis of the image plane and the horizon) and an elevation angle of α (measured from the horizontal), the rotations necessary to form the topographic perspective are: 1) a roll-removal rotation, and 2) an elevation-removal rotation. These rotations, for point (x, y, z) , are:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix}^T \begin{pmatrix} \cos \theta & \sin \theta & 0 \\ -\sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \phi & \sin \phi \\ 0 & -\sin \phi & \cos \phi \end{pmatrix}$$

where $\phi = \pi/2 - \alpha$. The first rotation removes the roll angle, and the second rotation removes the elevation angle.

The AHDA simulation currently utilizes a simple but brute force range image hazard detection algorithm: the algorithm measures the point-wise slope between pixels of the rotated range image to determine if each slope exceeds the hazard threshold. In other words, the hazard definitions for boulders/holes and slopes are used to impose upper bounds on the point-wise slopes (as well as the associated spatial extents) that can occur between two pixels in order for the region between these two pixels to be considered

hazard-free. The 1 m diameter boulder/hole hazard definition, for example, imposes the point-wise slope threshold of

$$\xi = \tan^{-1} (1 / 0.5),$$

i.e., if any two pixels (within 1 m of each other) of the rotated range image produce a point-wise slope exceeding ξ , then the region between these two pixels is considered hazardous by the boulder/hole hazard definition. By repeating this procedure for the slope hazard definition, every pixel in the rotated image can be labeled with one of four pixel labels: safe, slope hazard, boulder/hole hazard, or occluded. The last category is used to denote shadowed (occluded) pixels in the sensor range image, which do not have a corresponding height value after rotation.

To perform site selection, one of two algorithms is invoked, depending on the number of hazardous pixels. The site selection algorithm assumes a binary image as input, where the binary image indicates the presence or absence of hazards at every pixel of the image. When the number of hazardous pixels is large, the size of the safe landing site is small, in which case an exhaustive search for the largest safe landing site is possible. This is accomplished by examining every pixel in the image, and if the pixel is safe, then the algorithm recursively checks to see if the neighbors of the pixel are safe, and so on. Intuitively, this algorithm "grows" a circle of safe pixels around an initial pixel, until the circle encounters a hazardous pixel.

The circle growing algorithm in fact performs its search a bit more efficiently. Instead of growing a circle one pixel at a time, it first estimates an upper bound on the size (radius in pixels) of the largest circle (denote this by n), and subsequently performs a binary search for the largest safe circle. The binary search reduces the number of circles to be checked from n to $\log_2 n$.

When the number of hazardous pixels is not as large, an algorithm known as the Largest Empty Circle [Preparata and Shamos, 1985] is used. The algorithm makes use of the Voronoi diagram, and can be computed in $O(n \log_2 n)$ time⁷, where n denotes the number of hazardous pixels (actually, n denotes the number of hazardous pixels plus 4; the 4 additional pixels correspond to the four corners of the image, which form the convex hull within which the search for the largest empty circle is performed). Because of the complex nature of the algorithm, we refer the interested reader to the reference.

⁷ Intuitively, an algorithm is $O(n \log n)$ if it requires on the order of $n \log n$ processing steps. A formal definition of this notation can be found in [Knuth, 1973].

REFERENCES

- Foley, J. D. and A. Van Dam, *Fundamentals of Interactive Computer Graphics*, Addison-Wesley, Reading MA, 1982.
- Jain, A. K., *Fundamentals of Digital Image Processing*, Prentice-Hall, Englewood Cliffs NJ, 1989.
- Knuth, D. E., *The Art of Computer Programming*, vol 1, 2nd ed., Addison-Wesley, Reading MA, 1973.
- Preparata, F. and M. Shamos, *Computational Geometry -- An Introduction*, Springer-Verlag, NY, 1985.

APPENDIX B. PROCESSING REQUIREMENTS

Processing can be divided into five stages: contrast computation and queue merging for intensity images; image rotation, hazard detection and site selection for range images. For timing purposes, we assume an on-board serial processor operating at 5.6 million floating point operations per second⁸ (i.e., 5.6 MFLOPS). Furthermore, we will assume reasonable but worst case scenarios in order to bound the processing requirements. We will assume the imaging strategy described in Sections 3 and 4.

To compute the contrast of an image, we assume a 128 x 128 high resolution (0.5 m/pix) intensity image, which will maximize the size of the neighborhood over which the mean and standard deviation is to be obtained. At this resolution, a 20 m x 20 m region spans 40 x 40 pixels, and with 50% overlap, there are 7 x 7 such regions. To compute the mean of a region, it requires only 1 pass through the region. To compute the variance, it takes 3 floating point operation per pixel. Thus, the contrast of a region of size n can be computed in $4n$ floating point operations. To keep track of the 10 best regions within the image will require a number of operations on the order of the square of the number of regions (i.e., $\approx 50^2$). The total number of operations necessary to compute the mean and intensity of an image, and subsequently returning a queue of the 10 best sites, is therefore less than 320,000, which can be performed within the 0.2 s processing time allocated per image when assuming a 5.6 MFLOPS processor.

To combine the 5 priority queues (each queue comprised of the 10 least-contrast regions from each intensity image) into a single priority queue at the end of the intensity image acquisition phase requires a number of operations less than the square of the total number of items. Thus, the merging of the 5 queues into one requires a trivial amount of time.

To analyze the number of operations per *range* image, we assume the acquisition of a 128 x 128 range image at 0.25 m/pix resolution.

To rotate an image obtained from some aspect angle to a top down perspective, we multiply the position of a pixel by the corresponding 3 x 3 rotation matrix. This rotation requires 9 (multiplication) operations, at a total cost of 150,000 floating point operations (FPOs). To perform boulder hazard detection, the algorithm checks every pair of points within a 1 m x 1 m (4 x 4 pixels) region about every pixel. Thus, boulder detection examines 136 points per pixel, and with 1 floating point multiply to

⁸ Sun SparcStation 2 workstation operates at 5.691 MFLOPS single precision arithmetic [Wilson, 1991]. For comparison, a SKYBolt accelerator board utilizing an i860 arithmetic processor and 64 MB static RAM operates at 80 MFLOPS single precision [Sky Computers, 1991].

determine the slope between any pair of points, the entire boulder hazard detection process takes 2.2 million floating point operations (MFPOs). Slope detection, depending on the specific algorithm, can be substantially faster. In particular, a moving-ball surface smoothing operation, operating on a 4 x 4 neighborhood, will be an order of magnitude faster. Thus, we will neglect the boulder processing requirements in the analysis of the hazard detection algorithm.

To perform site selection, one of two algorithms is used, depending on the number of hazardous pixels in the image. We assume for the sake of analysis that the number of pixels in a 5 m radius circle is the switchover point at which the decision of which of the two algorithms to invoke is made. At 0.25 m/pix, a 5 m radius circle contains 1257 pixels.

When the number of safe pixels is less than 1257, a binary search algorithm for the largest circle of safe pixels is used. Since the worst case scenario occurs when the safe pixels are contiguous (thus forming the largest possible circle), we will assume this to be the case. For each radius r , the number of pixels that need to be scanned is $(2r)^2 = 4r^2$. For any pixel (x,y) within the circle, the average distance between (x,y) to the rim of the circle is less than 4 pixels (more pixels are near the rim than near the center). Thus, each of the 1257 points must check, on average, 64 pixels (i.e., 4×4^2), for a total of 80,500 FPOs. There are also $128^2 - 1257$ hazardous pixels that must be scanned. Thus in total, this stage takes 100,000 FPOs.

When the number of safe pixels is greater than 1257, we use the Largest Empty Circle algorithm [Preparata and Shamos, 1985] to find a circle that doesn't contain any hazardous pixels. This algorithm makes three $n \log_2 n$ passes through the data, where n denotes the number of hazardous pixels. In our case, n is bounded by $128^2 - 1257 = 15,127$. Thus, the total number of operation is $3 \times 212,000$, or 636,000 FPOs. Note, however, that because of the complexity of the algorithm, precise number of operations is hard to come by; we have used instead an order-of-magnitude analysis to derive the 636,000 FPOs. Although order-of-magnitude analysis does not bound the coefficient in front of the $n \log_2 n$ term, empirically, this coefficient is somewhere between 2 and 3. Assuming the worst case of a coefficient of 3, the Largest Empty Circle algorithm requires 2 MFPOs. Combining the 2.2 MFPOs for hazard detection with the 2 MFPOs for site selection shows that the range image processing requires less than 4.5 MFPOs.

In summary, some worst case assumptions have been made in order to bound the number of floating point operations necessary to support on-board hazard detection and site selection. For the intensity image processing case, our imaging scenario

allotted ample time. For the range image processing case, our imaging scenario allotted 1 s for the processing of each range image. With the number of operations bounded by 4.5 MFPOs, it can be seen that the necessary processing can be performed within 1 s of a CPU rated at 5.6 MFPO per second. Thus, we believe the AHDA system, as described, does not pose any severe processing demands.

REFERENCES

- Wilson, D., "Tested Mettle" in *UNIX Review*, 5/1991.
SKY Computers, Inc., *SKYBolt Hardware Reference*, 1/1991.
Preparata, F. and M. Shamos, *Computational Geometry -- An Introduction*, Springer-Verlag, NY, 1985.

APPENDIX C. ISSUES AND LESSONS

In this appendix we enumerate, in no particular order, a set of issues or lessons learned as part of the simulation testbed development and analysis process.

INTEGRATION

The integration aspect of the simulation was probably the most time consuming. There are several reasons for this. First, the mixture of languages provided enormous headaches. The terrain simulation code was written in Fortran, the guidance and navigation code was written in Ada, and the sensor simulation and hazard detection modules were written in C. Interface problems also necessitated wrapping the entire simulation in a dummy Ada shell so as to appease the Ada compiler. Because these pieces of code were written independently, the effort involved in trying to translate them into a single language, along with continuous debugging and modifications, would have been enormous. The second reason is that the terrain simulation was written on a different platform -- the DEC VAX system instead of the Sun workstations, and differences in their architecture introduced very subtle bugs that took considerable effort to resolve. Third, because the simulation incorporated modules developed independently, and more importantly, developed by different individuals, it was difficult not to treat these modules as black boxes. Although conceptually, one would like to build large software systems that are highly independent, the tight coupling necessary for a closed-loop system necessitated software synchronization at many levels. This coupling introduced significant amount of complexity into the simulation.

SIMULATION HARDWARE REQUIREMENTS

The simulation testbed is large. Because of the terrain and sensor simulations, the closed-loop simulation requires an almost obscene amount of hardware to run. Specifically, running on a Sun SparcStation1, the system required 64MB of RAM, 150MB of swap, and 100MB of hard disk storage to hold an image sequence. Even then, an image sequence almost always run out of memory before the end of the sequence (within a single UNIX process, no garbage collection is performed on swap until the process terminates). Thus, a sequence of 40 images typically requires 2 to 3 separate runs, each one picking up where the previous left off, in order to complete.

SENSOR SIMULATIONS

The intent of the sensor simulation part of the closed-loop simulation was to provide some "reasonable" images with which we can begin to devise hazard detection strategies. As such, there are several ways in which the sensor simulation did not incorporate as much fidelity as it could have. First, as was mentioned in the body of the report, no "real" units were used in producing the output. The intensity image outputs the size of the projected unit area that is reflected back to the sensor plane, assuming a Lambertian scattering model. Thus the intensity sensor image produced outputs within the range of 0.0 to 1.0. Although the noise model was initially envisioned to be a Poisson model, because Poisson is a discrete process, the quantization level of the noise became another parameter in the simulation. For the results described in the body of the report, therefore, we used a Gaussian model instead. The laser ranger simulator emulated a leading edge range detector -- one which returns the line-of-sight range to the nearest point of the terrain within the instantaneous field-of-view. In practice, if one included power measurements, one would perform a leading edge range detection by making sure the returned signal exceeded some threshold, whereas in our case, we assumed a model in which the surface geometry did not affect the measured range. Lastly, the issue of whether Lambertian scattering is the correct model for terrain was ignored.

TERRAIN MODEL

Because our ability to select safe landing sites depends greatly on the terrain, the terrain model quickly became a vital part of the simulation. There are several ways in which the terrain generation code can be improved, and we suggest some of them here⁹. First, although the terrain is visually realistic at coarse resolutions, simulation artifacts are clearly visible at high resolutions. Second, processes that can produce channels and other dominant geological features are clearly absent. Third, the incorporation of realistic crater and boulder size and location distributions is highly desirable. Fourth, the computational aspect of the terrain simulation adds a significant amount to the processing time, hence a reduction in the processing time would be highly desirable.

⁹ We are currently using version 1 of the terrain simulation code. We believe some of the suggested modifications have already been implemented (and distributed) in version 2 of the code.

VEHICLE, TRAJECTORY, AND TERMINAL GUIDANCE

Without a vehicle design, it becomes almost pointless to perform a high fidelity vehicle simulation. The vehicle is thus implemented as a point mass, and the simulation can aim the sensor anywhere at anytime. The next level of fidelity would presumably incorporate a vehicle body and some constraints on landing site visibility. The incorporation of a vehicle model at this level of fidelity should probably be accompanied by a parachute dynamics model.

An interesting fact that arose as part of our testing of the closed loop simulation is the need to incorporate radar altimeter models into the simulation. In particular, there was one run in which the slope hazard detector was (intentionally) turned off, and the vehicle selected a landing site that was 1000 m higher and 400 m up-range (i.e., closer to the vehicle) than the nominal landing site. Without an altimeter model, the ignition of the engines were synchronized to a timeline measured with respect to the nominal landing site. Thus, when the clock showed that the vehicle was approximately 1400 m above the nominal landing site and that the engines should be ignited, the vehicle was in fact only 400 m above the (newly) selected landing site, and the vehicle subsequently crashed (i.e., ran into the ground with a significant velocity). The incorporation of an altimeter model, however, implies a significant increase in simulation complexity because the altimeter would have to be coupled to the terrain file.

APPENDIX D. SAMPLE IMAGES

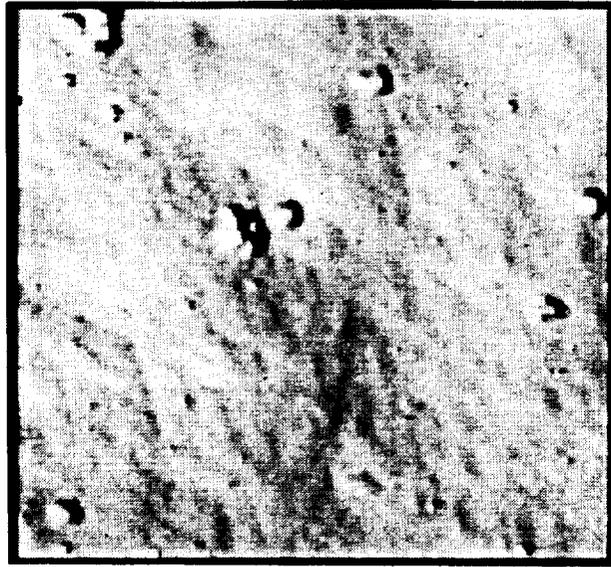
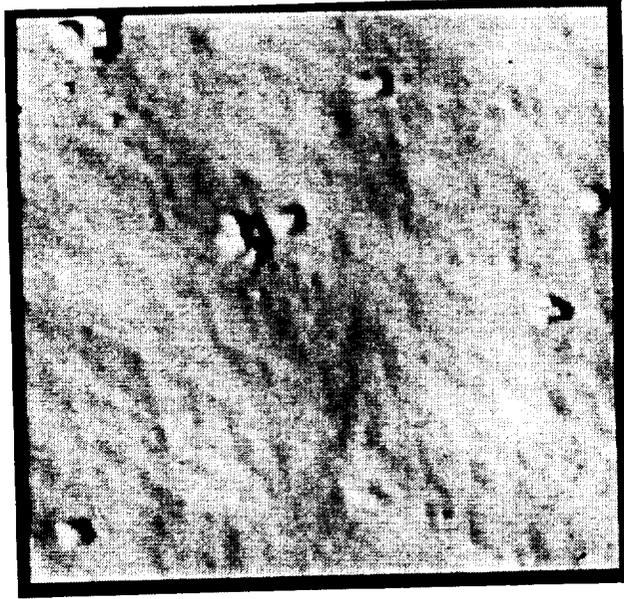
Representative images used in the parametric studies are shown in this appendix.

Figure D1 shows two intensity images with terrain generated using the JPL terrain simulation program version 1. Each image is 128 x 128 pixels, with a pixel resolution of 0.75 m/pix and an elevation angle of 77 deg. The top image shows the nominal boulder distribution case, whereas the bottom image shows the rocky case. Note that several "dominant" boulders have remained stationary within these two images. This is because version 1 of the terrain generation program can only output terrain with a fixed number of pixels, and when a field-of-view of some specific size is requested, the resulting terrain may not have sufficient pixel resolution, and hence may require interpolation. This is the case seen in Figure D1. The (large) boulders seen in the image were of the appropriate resolution prior to interpolation, and no smaller boulders were needed. After interpolation, however, there's a noticeable lack of small boulders. The process described in Section 4 was used to superimpose the appropriate boulder model.

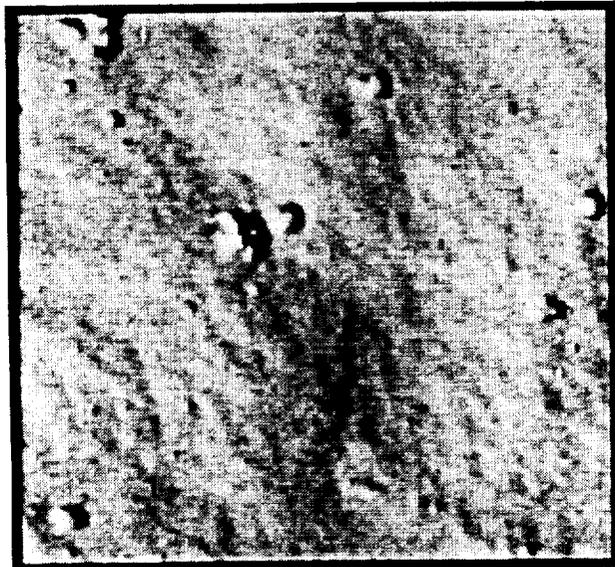
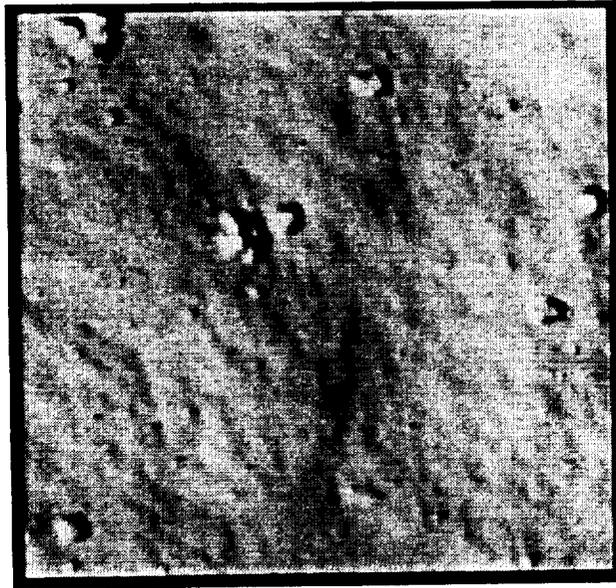
Figures D2 and D3 show the effects of signal-to-noise on an image. The top image in Figure D2 shows the intensity image at 20 dB, bottom image of Figure D2 shows the same image at 15 dB. The top image in Figure D3 shows the same image at 10 dB, and the bottom image in Figure D3 shows it at 5 dB.

Figures D4 and D5 show the effects of resolution on an image. Figure D4 shows the intensity image at 0.75 m/pix (top) and 1.5 m/pix (bottom); Figure D5 shows the same image at 2.25 m/pix (top) and 3.0 m/pix (bottom).

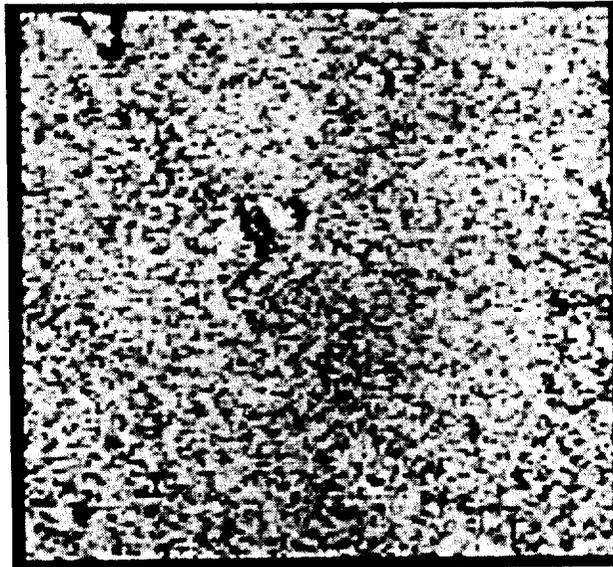
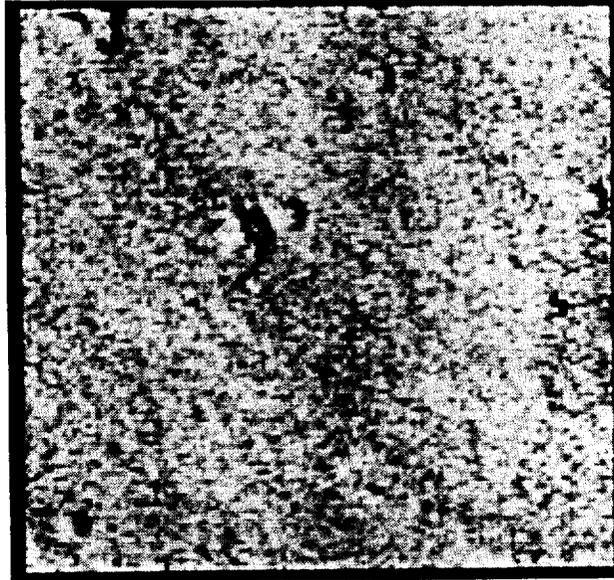
Figure D6 shows an intensity-range pair generated via version 2 of the JPL terrain generation program. The top image shows the intensity image with a nadir (top-down) viewing perspective, with solar illumination from the upper left at an elevation angle of 30 deg, and a pixel resolution of 0.4 m/pix. The bottom image shows the corresponding range image, with bright intensity indicating a shorter line-of-sight range. We are currently integrating this version of the terrain generation code into the AHDA simulation so that we can determine the performance of the hazard detection system on this (more realistic) terrain.



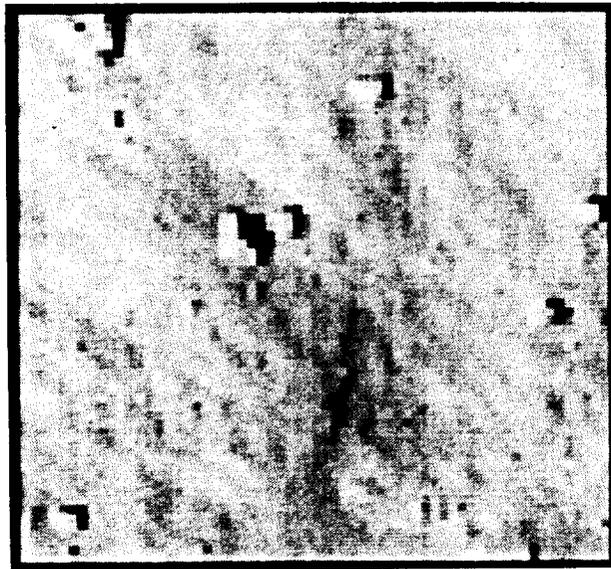
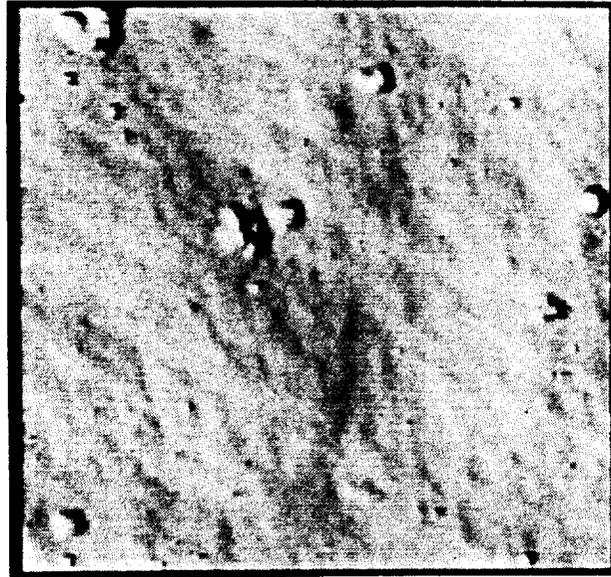
*Figure D1. Sample Intensity Images.
(Top) Nominal Terrain. (Bottom) Rocky Terrain.*



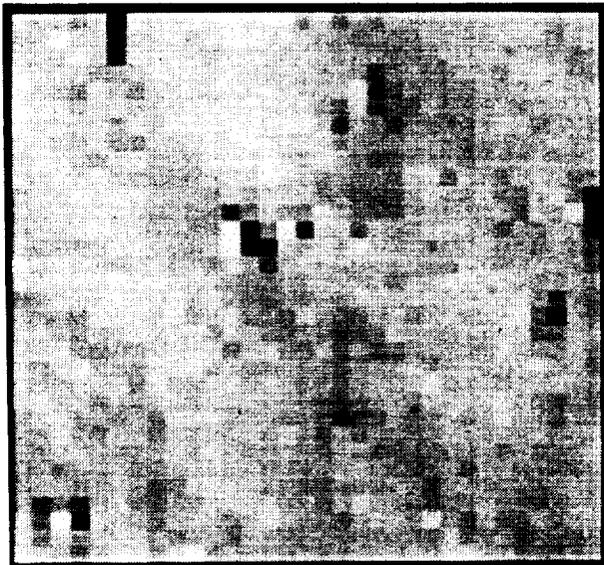
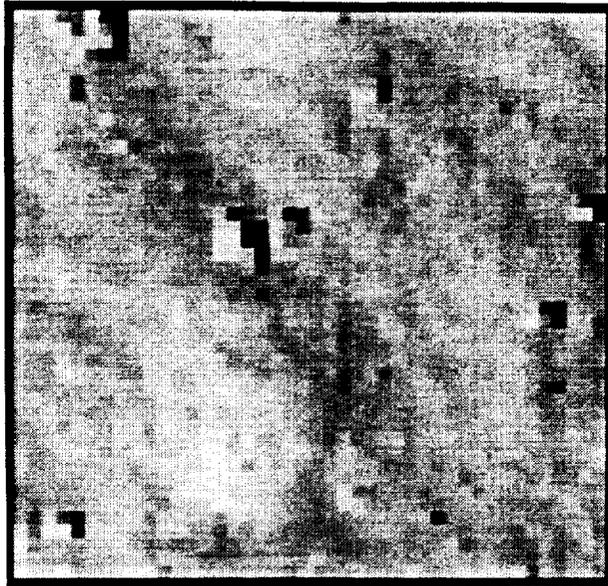
*Figure D2. The Effect of Signal-to-Noise on an Intensity Image.
(Top) 20 dB. (Bottom) 15 dB.*



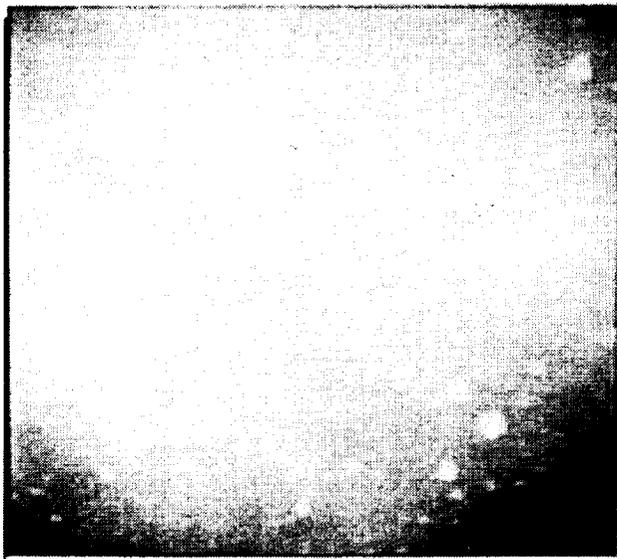
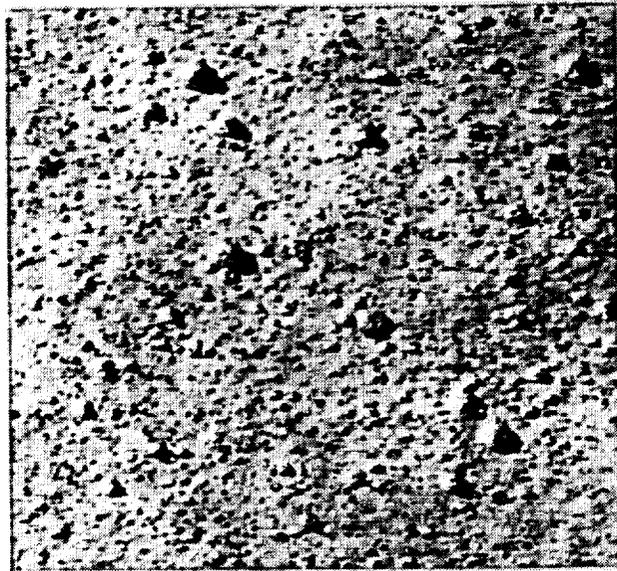
*Figure D3. The Effect of Signal-to-Noise on an Intensity Image.
(Top) 10 dB. (Bottom) 5 dB.*



*Figure D4. The Effect of Resolution on an Intensity Image.
(Top) 0.75 m/pix. (Bottom) 1.5 m/pix.*



*Figure D5. The Effect of Resolution on an Intensity Image.
(Top) 2.25m/pix. (Bottom) 3.0 m/pix.*



*Figure D6. Images Simulated Using Version 2 of JPL's Terrain Simulation Code.
(Top) Intensity Image. (Bottom) Range Image.*

DISTRIBUTION LIST

<u>RECIPIENT</u>	<u>COPIES</u>	<u>RECIPIENT</u>	<u>COPIES</u>
<i>NASA Johnson Space Center Houston, TX 77058</i>		<i>Lockheed Missiles & Space Co. 1111 Lockheed Way Sunnyvale, CA 94089</i>	
• EE/J. Lamoreux	1	• S. Smith	1
• EE/R. D. Juday	1		
• EE/W. X. Culpepper	1	<i>LTV Missiles & Electronics Group P. O. Box 65003 Dallas, TX 75265-0003</i>	
• EG111/A. Bordano	1	• MS WT-39/J. D. Billingsley	1
• EG2/G. G. McSwain	1		
• EG231/N. Smith	1	<i>LTV Aerospace & Defense Co. P. O. Box 65003 Dallas, TX 75265-0003</i>	
• EG42/T. J. Blucker	1	• M/S WT-07/Dr. L. G. Minor	1
• ER1/N. H. Chaffee	1		
• ER1/J. D. Erickson	1	<i>Computer & Vision Research Center The University of Texas at Austin Austin, TX 78712</i>	
• ER2/K. J. Healey	1	• Dr. J. K. Aggarwal	1
• ER2/T. W. Pendleton	1		
• ER2/K. Baker	4	<i>Dept. of Electrical Engineering University of Texas at Arlington Box 19016, URA Station Arlington, TX 76019</i>	
• ET2/J. Stecklein	1	• Dr. V. Devarajan	1
• ET4/D. Long	1		
• ET4/M. A. Tigges	1	<i>General Dynamics Space Systems P. O. Box 85990 San Diego, CA 92186-5990</i>	
• IE3/S. Bailey	1	• MS 24/E. Jones	1
• JM2/Technical Library	4		
• XE/G. Barton	1	<i>McDonnell Douglas Electronic Systems P. O. Box 516 St. Louis, MO 63166-0516</i>	
• XI/D. Sauchier	1	• MS 111-1161/E. M. Tribula	1
<i>NASA Ames Research Center Moffett Field, CA 94053</i>		<i>Advanced Systems Development Sandia National Laboratories Albuquerque, NM 87185</i>	
• MS 210-9/D. Denery	1	• M. A. Grohman	1
<i>Jet Propulsion Laboratory 4800 Oak Grove Drive Pasadena, CA 91109</i>		<i>ERIM/Electro-Optics Systems Lab. P. O. Box 8618 Ann Arbor, MI 48107-8618</i>	
• MS 264-276/R. Bourke	1	• D. C. Carmer	1
• MS 300-329/T. H. Chao	1	• M. Reiley	1
• MS301-125L/R. Gaskell	1		
		<i>ERIM 1101 Wilson Blvd., Suite 1100 Arlington, VA 22209-2403</i>	
<i>Martin Marietta Aerospace P. O. Box 179 Denver, CO 80201</i>		• M. C. Trichel	1
• MS H3482/J. Cuseo	1		
<i>LESC P. O. Box 58561 Houston, TX 77258-8561</i>			
• MS C19/E. Hack	1		
<i>TRW/Space Technology Group One Space Park Drive Redondo Beach, CA 90278</i>			
• MS 154-1878/D. Stager	1		

DRAPER LABORATORY INTERNAL DISTRIBUTION LIST

<u>Recipient</u>	<u>Copies</u>	<u>MS</u>
M. Adams	1	4E
B. Ahn	1	03
A. Bogner	1	2B
T. Brand	2	2B
J. Deyst	1	1C
D. Eyring	1	7C
T. Fill	1	2B
J. Gilmore	1	63
R. Hain	1	2B
M. Johnston	1	4A
J. Miller	1	2B
P. Motyka	1	4C
J. Negro	1	3A
H. Pien	5	4E
D. Shazeer	1	4E
K. Spratlin	1	2B
Technical Information Center	1	74