

Mars Descent Imager (MARDI) on the Mars Polar Lander

M. C. Malin and M. A. Caplinger

Malin Space Science Systems, San Diego, California

M. H. Carr

U.S. Geological Survey, Menlo Park, California

S. Squyres, P. Thomas, and J. Veverka

Center for Radiophysics and Space Research, Cornell University, Ithaca, New York

Abstract. The Mars Descent Imager, or MARDI, experiment on the Mars Polar Lander (MPL) consists of a camera characterized by small physical size and mass ($\sim 6 \times 6 \times 12$ cm, including baffle; < 500 gm), low power requirements (< 2.5 W, including power supply losses), and high science performance (1000×1000 pixel, low noise). The intent of the investigation is to acquire nested images over a range of resolutions, from 8 m/pixel to better than 1 cm/pixel, during the roughly 2 min it takes the MPL to descend from 8 km to the surface under parachute and rocket-powered deceleration. Observational goals will include studies of (1) surface morphology (e.g., nature and distribution of landforms indicating past and present environmental processes); (2) local and regional geography (e.g., context for other lander instruments: precise location, detailed local relief); and (3) relationships to features seen in orbiter data. To accomplish these goals, MARDI will collect three types of images. Four small images (256×256 pixels) will be acquired on 0.5 s centers beginning 0.3 s before MPL's heatshield is jettisoned. Sixteen full-frame images (1024×1024 , circularly edited) will be acquired on 5.3 s centers thereafter. Just after backshell jettison but prior to the start of powered descent, a "best final nonpowered descent image" will be acquired. Five seconds after the start of powered descent, the camera will begin acquiring images on 4 s centers. Storage for as many as ten 800×800 pixel images is available during terminal descent. A number of spacecraft factors are likely to impact the quality of MARDI images, including substantial motion blur resulting from large rates of attitude variation during parachute descent and substantial rocket-engine-induced vibration during powered descent. In addition, the mounting location of the camera places the exhaust plume of the hydrazine engines prominently in the field of view.

1. Introduction

Among the most dramatic images returned from space over the past 40 years were those transmitted by the Ranger spacecraft and those filmed by the Apollo astronauts during their descents to the lunar surface. These images not only provided impressive views of the Moon but did so in a particularly memorable way. In addition to unambiguously telling where a spacecraft has landed, images acquired during descent to a planet's surface provide the public with a visual perspective of spectacular, often breathtaking beauty and excitement. The process of acquiring such images is simple, and the results are easily understood by all who see them. Descent imaging provides tangible results for early release to the public and engenders a sense of "being there" not usually available with planetary missions.

Descent imaging systems also provide a crucial link between orbiter and lander observations. They provide context for the lander data as a function of scale (resolution) and area. No other form of observation provides such context. Descent im-

aging addresses the common themes of the Surveyor Program: volatiles, life, and resources. It does this not by measurement of a particular elemental isotope or molecular species but by searching for landforms representative of surface processes that reflect the Martian environment, by providing other instruments the ability to reject or accept their own measurements on the basis of the general and specific geologic context of those observations, and by linking orbiter observations to those made by the lander. The primary scientific contributions of descent imaging are likely to be through serendipitous observation of primary and/or secondary landforms that show the action of specific processes (e.g., channeling, patterned ground, etc.) and through the constraints placed on relative age relationships of features seen at the landing site (e.g., features imaged at resolutions between millimeters and meters are likely to reflect surface processes that operate on time-scales of a few years to a few hundred million years).

The relative importance of descent imaging to science and engineering has evolved with time: the return on investment has, on occasion, been deemed less than for more highly valued measurements (e.g., the aforementioned chemical analyses). However, the Mars Surveyor program has acknowledged its importance, selecting descent imaging first as a Principal In-

Copyright 2001 by the American Geophysical Union.

Paper number 1999JE001144.
0148-0227/01/1999JE001144\$09.00

vestigator experiment and then as government-furnished equipment to be used jointly by a competitively selected facility science team (for science studies) and mission engineers (for mission planning and, potentially, landing hazard avoidance).

The Mars Surveyor 1998 (MS'98) Lander Descent Imager (Mars Descent Imager, or MARDI) consists of optics, a focal plane assembly (FPA), Data Acquisition System (DAS) electronics, and a power supply. It is characterized by small physical size ($\sim 5 \times 5 \times 9$ cm, 480 gm), low power requirements (< 3 W, including power supply losses), and high science performance (1000×1000 pixel, low noise images acquired every few seconds, and ultimate resolution better than 1 cm/pixel). MARDI will acquire ~ 30 images of mixed sizes and time sequencing, spanning 3 orders of magnitude in scale, during the time between heatshield jettison and spacecraft touchdown.

2. Scientific Objectives

2.1. Landing Site Context

Among the most important questions to be asked about a spacecraft sitting on a planetary surface is "Where is it?" Radiometric tracking [e.g., *Folkner et al.*, 1997] and orbit determination (both spacecraft-to-Earth and spacecraft-to-spacecraft) and integration of inertial reference system variations (accelerometers tied to inertial measurement units) provide answers to this question to varying degrees of accuracy: at best, these techniques can tell the position of the spacecraft (in inertial space) to perhaps a hundred meters. Tying these positions to the surface requires correlation with planetary geodetic networks, which are often inaccurate to hundreds if not thousands of meters [e.g., *Zeiler and Oberst*, 1999]. Surface imaging of features also visible from orbit can be used to pinpoint lander positions to a few tens of meters or better [e.g., *Golombek et al.*, 1997; *Oberst et al.*, 1999], provided that such features are found. However, if the orbiter image resolution is insufficient to see features visible to the lander, or the local, meter-scale relief is too great (so the lander cannot see very far), or the surface is relatively featureless, or the surface has many features but they all look the same, then the lander cannot be located, or its location may be misidentified.

The Viking Landers and Mars Pathfinder provide good examples of such circumstances. Through a combination of 20–40 m/pixel, relatively low-Sun orbiter photography, excellent radiometric tracking from Earth over a long period of time combined with good inertial position measurements during landing, and fortuitously landing near craters and hills large enough to be seen on the horizon in lander images, Pathfinder was located to within about 40–50 m [*Golombek et al.*, 1997; *Oberst et al.*, 1999]. However, despite good inertial position measurements during landing and good radiometric tracking data both during the descent and for a number of weeks thereafter, the homogeneously rugged local relief, nearly featureless horizon, and the lack of spatially variable landforms in the 40 m/pixel orbiter images defeated all attempts to determine the location of the VL-2 to better than 10 km [e.g., *Stooke*, 1998]. Recently, the location of VL-1, previously considered well established [*Morris et al.*, 1978; *Morris*, 1980], has been called into question owing to inconsistencies with the indisputable location of Pathfinder [*Stooke*, 1999; *Parker*, 1999; *Oberst et al.*, 1999]. Such uncertainties, raised decades after Viking landed, simply illustrate the degree of difficulty in identifying landing locations through indirect means. Direct observation of the landing site during the descent to the surface, starting at altitudes

that permit the images to be placed without question into the context of orbiter observations, eliminates all such uncertainty.

Why is it important to know "exactly" where a lander is located? The principal reason is context. It is necessary to determine if the locale is representative of the region and, indeed, of the planet. It is usually not possible, just from a lander's perspective, to tell the difference between what is visible and what is just over the horizon. The locale may be anomalous; this must be determined before general interpretations can be made. Knowing that local meteorology is affected by a nearby escarpment, or that the lander sits on ejecta from a nearby crater, is important both for local interpretation and for extending it farther afield. The context of relating lander observations to those seen from the orbiter is also important. The simplest, and most obvious, example is to place surface imaging into the context of orbiter images (e.g., extending and linking crater and boulder size/frequency relationships, extending surface observations of eolian bedform wavelength/amplitude/particle size attributes to larger scale, etc.). Other examples include relating color and/or albedo boundaries seen in orbiter data down to lander scales (which is particularly difficult to do from the surface owing to the extremely oblique viewing geometry of lander instruments [e.g., *Maki et al.*, 1999; *Johnson et al.*, 1999] and providing data to test models used to calculate rock abundance and other granulometric properties of the surface from thermal emission measurements.

Descent imaging can also provide a context for operations after landing. For example, the final images can cover the area around the lander out to several tens of meters or more at spatial scales of a few centimeters. Such images can be used to plan sampling activities and/or rover traverses, both initially before surface imaging and complimentary to those data once they are received. The easily interpreted, overhead perspective provides such planning activities considerable speed and flexibility. Advanced techniques in computer graphics and data visualization have been used to merge lander images with distance measurements [*Stoker et al.*, 1999], derived from stereoscopic images or laser rangefinding, in efforts to mimic the overhead perspective. However, the inability to see surfaces hidden from direct view from the lander perspective is an essentially fatal flaw in such efforts. The simplest, most comprehensive way to achieve overhead viewing is from a descent camera.

2.2. Science Studies From Descent Imaging

The scale at which processes modify a planet's surface is dependent on the vigor of the processes and the timescales over which they act. On Mars the vigor of environmental processes appears to have varied substantially: contemporary phenomena (or at least those that have actually been observed) appear to be relatively weak (e.g., wind transport of dust and sand), while ancient processes appear to have been much more vigorous (e.g., channel formation by catastrophic flood). However, some areas, for example, the layered materials within Candor Chasma, are devoid of impact craters, implying an erosional process of substantial capability of which there is little other evidence. The true nature of Martian erosional processes is just now being deciphered with the help of high-resolution images.

For the purpose of this discussion, dependent on resolutions not yet achieved on Mars, relationships between scale and time as seen on Earth can be used to illustrate the study of temporal




Figure 1a. Portion of Landsat image showing Antarctic ice-free valleys at 80 m/pixel. The box marks the location of Figure 1b.

relationships with descent images. Figures 1, 2, and 3 present a suite of images representative of a descent-imaging sequence (a description of the planned MS'98 operational scenario is presented later).

Figure 1 shows the relationship between a typical orbiter image (in this case, Figure 1a shows an 80×80 km portion of a Landsat frame at 80 m/pixel) and the first image of a descent sequence (Figure 1b, a portion of an aerial photograph). At a resolution of 7.5 m/pixel and covering an area ~ 8 km on a side, the location of the descent image is reasonably visible in the orbiter data. Images taken early in a descent provide both this crucial link to orbiter observations and the context for all subsequent frames.

In the specific case of the images in these illustrations, the increase in resolution from 80 to 8 m/pixel spans several im-

portant transitions in geomorphic interpretation. That the area is mountainous, with glaciers moving down various topographic gradients, can be seen in the orbiter data. Details of, for example, the glacial flow cannot be seen at this resolution but become obvious in the 8 m/pixel data. Note the snout of the Taylor Glacier (bottom center, image on right), which shows ablation pitting along medial streamlines of shear and morainal debris. Note, too, details of the mountain walls, including mass movements and evidence of liquid flow (e.g., small stream channels).

Figure 2 shows a sequence of six images at successively higher resolutions of 5, 2.5, 1.2, 0.6, 0.3, and 0.16 m/pixel, equivalent to images taken roughly 65, 40, 26, 18, 13, and 9 s prior to landing. As resolution improves, the landscape changes dramatically. At low resolution the scene shows large



Figure 1b. Aerial photograph of area outlined in Landsat image, at 7.5 m/pixel, also indicating location of nested descent images. Note the comparison of landforms between the Landsat and aerial photograph at a scale ratio of roughly 10:1.

diversity of landforms, while at high resolution the scene is much simpler. However, this trend should not be interpreted to mean that high-resolution images are not of interest. Rather, the type of information derived from images of such resolution is different from that derived from lower-resolution data. Surficial sedimentary deposits can just be discriminated at ~ 1.2 m/pixel and become more prominent at resolutions of 0.3 m/pixel and better. The size/frequency and spatial distribution of these materials reflect the processes that formed the surface (and the materials of which the surface is composed), as well as the processes that have modified the surface and transported and deposited materials. The size/frequency of the boulders indicates, independently of any other association, that these materials are glacial debris, and the areal distribution indicates

both the direction of flow of the glacier and the nature of the interaction of the bedrock with the glacier. In this case, the motion was from left to right (seen by the location of the boulder field), and the glacier was wet-based (shown by the nature and caliber of debris in the ground moraine on which the boulders sit). Given that these glaciers in the present environment are dry-based, the debris is thus remnant from an earlier period.

Two “serendipitous” observations can be seen in these images, relating to liquid water. First, the dark stains in some valley wall and floor channels indicate that liquid water was flowing on the surface in the very recent past. Indeed, using relatively simple calculations reflecting the seasonal heating high evaporation rates in the Antarctic ice-free valleys, it is




Figure 2. Sequence of images at a nesting scale ratio of 2:1 showing changes in landforms as a function of scale. Images shown are at scales of 5, 2.5, 1.2, 0.6, 0.3, and 0.16 m/pixel, equivalent to images acquired roughly 65, 40, 26, 18, 13, and 9 s prior to landing.



Figure 3. View of area shown in last two images of Figure 2, showing advantages of oblique viewing. Arrows point to a circle of small stones 7 m in diameter, used as a landing site for helicopters.

possible to show that the moisture is only a few weeks old. Second, the dark band around the ice-covered lake can be seen in several of the higher-resolution images to be a liquid water moat, as indicated by, for example, wind ripples and the reflecting properties of the surface. Again, relatively simple calculations suggest that such moats are ephemeral.

Figure 3 shows an image taken from ~ 80 m above the surface but looking obliquely (the near field is viewed at an emission angle of 30° , and the far field is viewed at $\sim 75^\circ$). The advantages of oblique viewing, in particular, the ability to gain from a single image some knowledge of subtle relief in the scene and to provide a more familiar view, are clearly evident in this image. Note in the bottom center foreground a helicopter landing circle ~ 7 m in diameter (white arrows).

One of the more interesting observations that can be made

from this sequence of images is that the scene content varies dramatically with scale. This is *prima facie* evidence against the concept of scale invariance in nature, that is, that nature is not fractal-like. Mathematicians [e.g., *Mandelbrot*, 1967, 1982], geophysicists [e.g., *Turcotte*, 1989, 1991], and hydrologists/geomorphologists [*Scheidegger*, 1970; *Tarboton et al.*, 1988; *Xu et al.*, 1993] have argued, through various examples, that self-similarity is a fundamental attribute of geology. Others contend that the types and style of geologic processes and materials clearly vary with scale (i.e., the mechanisms responsible for breaking individual grains of sand are very different from those responsible for the shape of river valleys). The sequence of images shown here attests to this view. There are clearly several points in the continuum of scales where the surface takes on distinctly different properties; the last two frames in

Table 1. Science Requirements for Descent Imaging

	Requirement
Resolution (highest)	~2 cm/pixel
Landing site visibility	must be seen in last frame; desire to see throughout descent
Field of view	landing site visible in last frame that covers ≥ 10 m at 2 cm/pixel
Nesting scale ratio	better than 5:1 ($\leq 2:1$ goal)
MTF at Nyquist	>0.10
SNR	$\leq 20:1$, for albedo = 0.10, at aphelion, with illumination angle (i) $\leq 75^\circ$ (Sun elevation $\geq 15^\circ$)
Photometry	5% relative (within an image), 10% absolute (between images)
Images returned	~8–16
Spectral response	500–800 nm

Figure 2 are very different from the first two frames. To the extent that these surfaces reflect the different processes and materials, an analogous sequence on Mars will provide considerable insight into similarities to and differences with terrestrial conditions.

3. Investigation Approach

3.1. Science Requirements

Table 1 outlines the basic capabilities of the MS'98 MARDI experiment, including science requirements based on either science or context arguments and mission constraints. Some of these capabilities are discussed in more detail below.

3.1.1. Resolution. A resolution goal of 2 cm/pixel was deemed necessary to adequately characterize the local environment of the lander. This resolution was also both scientifically and technically achievable. Resolution must be traded against field of view (FOV); the larger the requisite FOV, the poorer the spatial resolution. Resolution is also affected by the altitude at which imaging occurs (a lowest “good” image altitude of 10 m seems reasonable) and by mechanical attributes of the spacecraft, primarily vibration and the rate of attitude excursions.

3.1.2. Landing site visibility and field of view. The desire to view the landing site throughout the descent places severe requirements on the FOV, as does the areal coverage and resolution of the last frame (10 m at 1–2 cm/pixel). Despite a rapid, near-vertical descent (which would permit a high spatial resolution and small FOV), large attitude excursions of up to $\pm 30^\circ$ while descending beneath the parachute is the primary factor driving a large FOV. A large FOV also provides the opportunity for oblique viewing geometry over part of the field that can be used to acquire stereoscopic coverage as well as to provide a visual impression of landscape relief.

3.1.3. Nesting coverage. Nesting coverage is the way successive images relate to one another and to the entire suite of images. Full nesting means that all higher-resolution images are fully contained within the lowest resolution image (Figure 1b). This provides a direct view of a small, localized area (including the landing site) at all resolutions down to the surface and occurs when, for example, the descent is entirely vertical. The rapid Mars Surveyor 1998 Lander descent could permit full nesting, but the large attitude excursions during parachute descent make this very unlikely. Some nesting is expected.

3.1.4. Nesting scale ratio. The nesting scale, or resolution ratio, is determined at the beginning of descent imaging by the altitude at which imaging starts and by the descent rate and at the end of the descent by the descent rate, the minimum time

between frames, and the size of each image. An optimized solution would have a variable nesting ratio, triggered by altitude. However, constraints imposed by project management prevented an optimized descent scenario. Rather, two fixed intervals between images (one for parachute descent, one for powered descent) were implemented, which yield variable nesting scale ratios that range from a low (at high altitude) of about 1.1:1 to a high of 4:1 just prior to landing.

3.2. Known Issues Affecting Science Return From MARDI

3.2.1. Rocket plumes. MARDI is fastened to a cruise rocket engine motor (REM) mounting bracket. Its 66° FOV is canted 22° outboard from the centerline (thrust or descent velocity) axis (Figure 4). The primary reason for this cant is that much of the view in the descent direction will be through the overlapping plumes of the descent engines. The optical distortion of the plume due to temperature and density differences may be quite extreme, and there is a possibility that an ammonium bicarbonate contrail will be produced as cooling or incompletely catalyzed hydrazine products interact with atmospheric CO_2 . Canting the instrument is an attempt to acquire some undistorted imaging during powered descent.

3.2.2. Image smear. The camera will also experience rapid body angular rates during parachute descent and substantial vibration during powered descent. The maximum composite effect of these motions is about 3–5 pixels of motion blur (smear) for a 4 ms exposure. At this short exposure time, detector smear (created by light impinging on the inefficient interline lightshields of the detector during electronic transfer of an exposed image off the detector) contributes a moderately large fraction ($>20\%$ of total signal) of noncoherent, spatially varying brightness patterns to each image. The MS'98 Project chose to expend no effort or resources to mitigate these problems. Taken together, solid body and vibrational motion blur and detector smear may substantially reduce spatial resolution and image fidelity throughout the descent.

3.2.3. Deterministic acquisition sequence. A major limitation imposed on the MS'98 descent imaging system is the MS'98 Project requirement that all image acquisitions be “pre-determined” (or “deterministic”) with respect to the entry, descent, and landing (EDL) sequence. At the time this decision was made, this meant that altitude could not be used as a trigger for imaging sequences. To attempt to acquire adequate overlap in resolution and coverage, images are planned to be taken continuously at fixed intervals. Two different intervals (one during parachute descent and one during powered descent) are used to try to accommodate the dramatic variation in descent rate as the spacecraft approaches the surface.




Figure 4. View of Mars Polar Lander in its descent configuration, showing the location of MARDI and the orientation of the MARDI FOV.




Figure 5. Imaging scenario for nominal MS'98 MPL descent.

3.2.4. Other factors. There are two additional limitations imposed by the Lander on the number and spacing of descent images:

1. The spacecraft command and data handling (C&DH) system and implementation of the RS-422 protocol dictate the maximum transfer rate of data from the camera to the Lander. The 1 Mbps transfer rate is affected by both bus and central processing unit (CPU) utilization (and hence subject to Project-mandated margin requirements). Thus, full-frame, raw, uncompressed 8-bit images (1000^2 pixels) cannot be taken faster than one every 8 s.

2. Only 100 Mbits of storage memory have been allocated to descent imaging. This is equivalent to 100 s of camera readout. If the descent takes longer than 100 s, the camera overwrites the memory.

Several factors mitigate the first of these limitations. Vignetting of the FOV at 65.9° permits editing the roughly square image to an inscribed circle ($\sim 30\%$ savings). Real-time, lossless compression by the camera can reduce the effective transfer volume by as much as an additional 30%. Subsecting the image (for example, taking only an 800×800 portion of the detector) can also reduce the data volume. The memory limitation is addressed by a sieving algorithm that separately preserves the first and last powered descent and parachute descent images while sampling intervals between these bounding images.

3.3. Observational Approach/Operational Scenario

Figure 5 shows a representative MS'98 descent profile at two scales. The right hand graph shows the profile from parachute deployment to touchdown in 2 s increments (pluses), with image acquisitions indicated by solid dots. The left image is an enlargement of the terminal descent. The total descent can

take from as short as 80 to as long as 180 s, depending on earlier entry and descent factors. Powered descent, using the hydrazine monopropellant engines, can last up to 40 s, although virtually all simulation scenarios limit powered descent to a window between 32 and 38 s in duration.

Descent imaging activities begin ~ 2.5 hours before entry, when the EDL sequence powers on the MARDI survival heater to bring the camera to operating temperature. The heater is turned off 24 min before entry. At entry minus 15 min, the MARDI software task within the spacecraft CPU begins executing. Thirty seconds before parachute deployment, the instrument is powered and the spacecraft CPU boots the camera digital signal processor (DSP) software. Imaging begins 0.3 s prior to heatshield release (9.7 s after parachute deployment) and continues every 0.5 s through the jettison event. These subsected images (256 pixels square) show the heatshield interior surface for calibration purposes. Surface observations begin with a full-field image ~ 2 s after heatshield jettison (altitude ~ 7 km). Images at that altitude will cover just under 75 km^2 at a resolution of 8.5 m/pixel. Images are acquired every 5.3 s thereafter. Each full-frame (1000×1000 pixel) image is circularly edited (the optics vignette at the corners of the square detector, and circular editing represents a savings of $\sim 30\%$ in data storage) and losslessly compressed. Immediately after parachute/aftshell jettison and before the beginning of powered descent, an image is acquired; this image constitutes the last full-frame acquired during unpowered descent. If parachute descent takes longer than 85 s, a sieving algorithm will be employed to control overwriting image storage memory in such a way as to preserve images that effectively sample the entire range of scale intervals. A maximum of 16

Table 2. Mars Surveyor 1998 Descent Imaging Scenario

Time to TD	Altitude	Range	Resolution, m	Image Size
107.30	6436.24	6786.95	8.05	216
106.80	6400.84	6651.34	8.00	216
106.30	6366.23	6521.46	7.96	216
105.80	6332.37	6396.95	7.92	216
105.30	6299.21	6277.48	7.87	1000
100.00	5978.94	5250.72	7.47	1000
94.70	5680.40	4501.94	7.10	1000
89.40	5370.85	3899.99	6.71	1000
84.10	5039.91	3390.06	6.30	1000
78.80	4688.89	2950.91	5.86	1000
73.50	4321.22	2570.86	5.40	1000
68.20	3939.26	2237.86	4.92	1000
62.90	3545.83	1938.52	4.43	1000
57.60	3145.59	1661.01	3.93	1000
52.30	2742.90	1397.57	3.43	1000
47.00	2337.82	1144.34	2.92	1000
41.70	1936.17	898.56	2.42	1000
36.40	1528.20	656.39	1.91	1000
31.10	1115.06	418.89	1.39	800
27.10	821.98	259.75	1.03	800
23.10	570.41	165.00	0.71	800
19.10	366.95	113.86	0.46	800
15.10	209.11	57.20	0.26	800
11.10	96.77	23.33	0.12	800
7.10	32.76	7.28	0.04	800
3.10	9.17	2.97	0.01	800

parachute descent images can be stored. Five seconds after powered descent begins, imaging continues on a one per 4 s schedule that continues until touchdown. These images are 800×800 pixels, square, and, again, losslessly compressed. If powered descent takes longer than 40 s, the same sieving algorithm is employed to control overwriting powered descent images stored in memory. A maximum of 10 powered descent images can be stored.

3.4. Investigation Methods and Procedures/Anticipated Results

Upon receipt on Earth, the images will be decompressed and preflight photometric and geometric corrections will be applied. Science analyses will include extraction of relief from stereo images (created as the lander moves laterally during the descent) and production of maps of the landing site in “near real time” in support of lander operations. A highlight of the data processing will be the recreation of the descent in animated form.

Under nominal circumstances, and limited by the available storage volume, a maximum of 30 images will be acquired from

altitudes below 8 km: four 256^2 images taken on 0.5 s centers as the heatshield is jettisoned, sixteen 1000^2 , circularly edited frames every 5.3 s during parachute descent, and ten 800^2 images taken on 4 s centers until touchdown. Table 2 provides representative information about these images from an example descent scenario. In general, the first image will cover about 8–9 km on a side at a resolution of about 8–9 m/pixel, and the typical last image (acquired 3–5 s before landing) will cover an area just under 9 m across at 0.9 cm/pixel. Note that all images acquired before powered descent are at the same or lower resolution than can be acquired by the Mars Orbiter Camera (1.4 m/pixel), while those taken during powered descent will progressively exceed this resolution by as much as 2 orders of magnitude.

In addition to the individual images, derived information will include (1) detailed geologic, geomorphic, and traverse planning maps of the landing site, ranging in scale from 1:24 to 1:30 K; (2) relief maps of the landing site, at scales from 1:14 to 1:80 K, with vertical resolution ranging between 4 m and 2 cm; and (3) time-sequential “real-time” descent animation, showing the 60–100 s descent to the surface.

4. Instrument Description

Science requirements and mission constraints were combined into a set of functional requirements to which the MARDI was designed (Table 3). Table 4, the Instrument Parameter Table, provides a summary of the results of this effort. More details of the design are given below. A photograph of the final instrument is shown in Figure 6.

4.1. Optics

The MARDI optics provide a nominal FOV of 65.9° across the 9 mm square photoactive area of the charge-coupled device (CCD) detector. The lens has an effective focal length of 7.01 mm, has a focal ratio of $f/5.6$, and weighs ~ 100 g. The lens has nine elements, including two cemented doublets. All optical surfaces are spheres except for a weak asphere immediately after the aperture stop, which deviates <3 waves at 633 nm from spherical.

The optical performance properties are weighted for Martian spectral radiance in the 500–800 nm band, although no band-pass filter is included in the optical path. The lens is designed to provide high modulation transfer function (MTF) performance over the entire field. Testing of the flight optics at 633 nm showed an MTF at the Nyquist frequency of the detector (56 line pairs per millimeter) of $\sim 80\%$, even at the edges of the field. The broadband performance is somewhat lower.

Table 3. MARDI Functional Requirements/Performance

Functional Requirements/Performance	
Field of view	65.9°
Exposure time (smear limit)	5 ms (1 pixel)
Focal ratio	$f/2$
Frame rate	0.5 frames/s (2 s/frame)
Signal-to-noise ratio	$\geq 20:1$, for Albedo = 0.10, at aphelion, with illumination angle—(i) $\leq 75^\circ$ (Sun elevation $\geq 15^\circ$)
Signal capacity	50 e- noise floor (5 ms, $A = 0.10$, $i \leq 75^\circ$, at aphelion)—30,000 e- full well ($A = 0.50$, $i = 0^\circ$, at perihelion)
Modulation transfer function	>0.10 @ Nyquist
Nesting ratio	2:1 or better
Image compression	2:1 realtime (<2 s) in DSP; 10:1 nonrealtime in S/C CPU

Table 4. Mars Surveyor 1998 MARDI Instrument Parameters

Parameter	Description
Mass and Volume	480 gm, $6 \times 8 \times 12$ cm
Power	3.5 W imaging
Optics	9-element refractive
FOV	65.9°
IFOV	1.25 mrad (7.5 m at 6 km, 12.5 cm at 100 m altitude)
<i>f</i> /ratio	<i>f</i> /5.6
Focal length	7.01 mm
Detector	Kodak KAI-1001: 1024×1024 (1018×1008 photoactive), 9 μ m pixels, 20% fill factor, interline transfer electronic shuttering
Noise	50 e-
Full-well	40,000 e-
Exposure time	4 msec
SNR	60:1 (albedo 0.1, $i = 70^\circ$)
Band pass	Panchromatic
Electronics	
Microprocessor	Motorola DSP 56166 at 60 MHz, 4096 bytes program RAM, 4096 bytes data RAM
Communications	Synchronous port at 1 Mbps/port Signals: RS-422 receive, transmit, clock, frame synch
Frame time	0.5 s for heatshield jettison, 5.3 s during parachute descent, 4.0 s during powered descent
Digitization/compression	12-bit ADC, square root encoded to 8 bits 1.8:1 lossless Huffman first-difference (real time in DSP)
Data return	≤ 30 images: four 256^2 of heatshield jettison, \leq sixteen 1024^2 circularly edited during parachute descent, \leq ten 800^2 during powered descent

The lens barrel and Sun shade are made of anodized aluminum. The design is athermalized to maintain focus over the expected temperature range. The optics are ~ 3.7 cm in diameter, including the Sun shade. When integrated with the MARDI electronics, the optics contribute 6 cm to the overall length.

4.2. Electronics

The low-mass camera system developed by Malin Space Science System for the Mars Surveyor 1998 missions is a logical development of previous designs for spaceborne CCD imaging systems. It uses the advances in sensor and computer technology realized in the 1990s to replace dedicated hardware with a flexible, software-driven design using a digital signal processor (DSP) and a large electronically shuttered CCD. A block diagram of the electronics is shown in Figure 7.

4.2.1. Detector. The heart of MARDI is a Kodak KAI-1001 CCD. This detector has 1024×1024 9-micron pixels (1018×1008 photoactive), and uses interline transfer to implement electronic shuttering. The KAI-1001's fill factor of 20% causes its quantum efficiency to be low, especially redward of 700 nm. These effects combine to make the detector comparatively poor in sensitivity, but the optics are sufficiently fast to compensate, allowing this compact CCD to be used.

Two capabilities of the detector can be used to reduce the raw data rate from the camera. First, on-chip summing can be used to add adjacent lines together in the charge domain under conditions of low signal. Second, the CCD's fast-dump feature can be used to read out only selected portions of the detector. Although the first capability is not used during descent imaging, the second is used to acquire both the heatshield deployment and powered descent images.

4.2.2. Clock generation. The CCD requires seven clock signals: a two-phase vertical clock (V1/V2), a two-phase horizontal clock (H1/H2), a substate clear clock (S), a reset clock (R), and a fast-dump clock (F). In addition, the analog-to-digital converter (ADC) requires a logic-level convert clock (C). The H1/H2 and R clocks are "fast" (maximum of ~ 3 MHz), while the V1/V2, S, and F clocks are slow (maximum of ~ 20 KHz).

The clocks are generated at logic levels by the DSP and are switched at the voltages required by the CCD using discrete P and N-MOSFET transistor pairs. There are two switch timing configurations: one for "fast" clocks, utilizing discrete HC-family logic to synthesize fast one-shots, and one for "slow" clocks, which use 74HC4538 integrated circuits (ICs).

4.2.3. Analog signal processing. The output signal from the CCD is alternating current (AC) coupled and then amplified by a single AD8011 opamp. To bring the amplified signal into the 0–5 V range accepted by the ADC, a fixed direct current (DC) offset is applied prior to amplification. The resulting AC-coupled signal encompasses half the ADC range: 1.25–3.75 V in the maximum signal condition of all pixels at full-well and 0–2.5 V in the case where a small number of pixels are at full-well and the rest are near reset.

4.2.4. Digitization. The amplified CCD signal is digitized by an Analog Devices AD1672 12-bit ADC running at its maximum rate of 3 million samples per second. For each pixel, both reset and video levels are digitized and then subtracted in the digital domain to perform correlated double sampling (CDS). Because of the dual sampling, the fastest pixel transfer time is 667 ns.

4.2.5. System speed. Image exposure time is 4 ms; transfer of charge from the photoactive sites to the vertical registers occurs in $< 8 \mu$ s. The full-frame maximum detector readout rate of 0.75 s is set by the 667 ns pixel transfer time. This speed limitation is a result of both the maximum digitization rate of the ADC and the maximum frequency of the analog clocks. The latter is, in turn, limited by the low-current voltage sources for the clock rails. Although higher rates could be achieved with this basic architecture by using higher-current (and hence higher-power) voltage sources and a faster (and hence higher-power) ADC, resource limitations dictated the present values.

4.2.6. Digital electronics. The digital electronics are responsible for clock pattern generation, sampling of the CCD signal, conversion of the 12-bit samples to 8-bit encoded pixels, storage of the pixels, and, finally, readout of the pixels to the spacecraft. The DSP in the DAS permits full speed software emulation of much of the usual analog processing, including




Figure 6. Mars Surveyor 1998 MARDI flight unit.

CDS. Using software emulation, the zero reference (“reset”) level for each pixel is digitized and stored in a register. The sum of the video plus zero reference (“video”) level is then digitized, and an arithmetic subtraction is performed to produce the final result. The CCD output requires only scaling to the ADC range; no analog sampling, delay, or differencing is required.

MARDI uses a Motorola DSP56166, selected because it can process 3 Mpixel/s (30 million instructions per second (MIPS)), it incorporates $4K \times 16$ -bit data and $2K \times 16$ -bit program memory on-chip, it has two 15 Mb/s serial ports that can be used simultaneously for a throughput of up to 30 Mb/s, and its firmware can be booted over these serial lines. For the nominal frame period of 0.75 s, the 56166 executes up to 20 arithmetic instructions per pixel. This performance enables the DSP to generate the CCD clocks, read the reset and video levels from the ADC, perform the correlated double sampling, convert the pixel from 12 to 8 bits using a nonlinear transfer function, apply lossless ($\sim 1.8:1$) first-difference Huffman compression, and transmit the compressed image data digitally with handshaking over the serial communication interface to the spacecraft CPU. The MARDI uses only the 8 K data static random access memory (SRAM) internal to the DSP.

The DSP clock rate at power-up is 10 MHz, supplied by the off-chip oscillator. This permits fixed-rate 2 MHz serial data from the spacecraft to be loaded. Once active, the DSP runs at 60 MHz (30 MIPS) using its internal phase-locked loop. Lower-speed operation can be used to save power in standby and idle periods.

The DAS also includes a 1 MByte SRAM buffer. This memory acts as a buffer between the roughly 0.8 s it takes to read out a full MARDI image and the 5.3 s it takes to transfer that image to the spacecraft C&DH through its Payload Activity Command Interface (PACI) card.

4.2.7. Power supply and heater. MARDI uses the Modular Devices MDI2690R-T15 power converter, which produces isolated $+5$ and ± 15 V from a nominal 28 V input. The ripple on the ± 15 V is excessive for the front-end video circuitry, so it is downregulated with series-pass regulators (LP2951 for $+12$ V and LM2990 for -12 V) to remove ripple. The clock rail and DC bias voltages required by the CCD are generated by source-follower opamps and zener diodes and are appropriate to the current required.

Owing to a limitation on power switching imposed by the spacecraft, two Hewlett-Packard HSSR-7111 optocoupled sol-

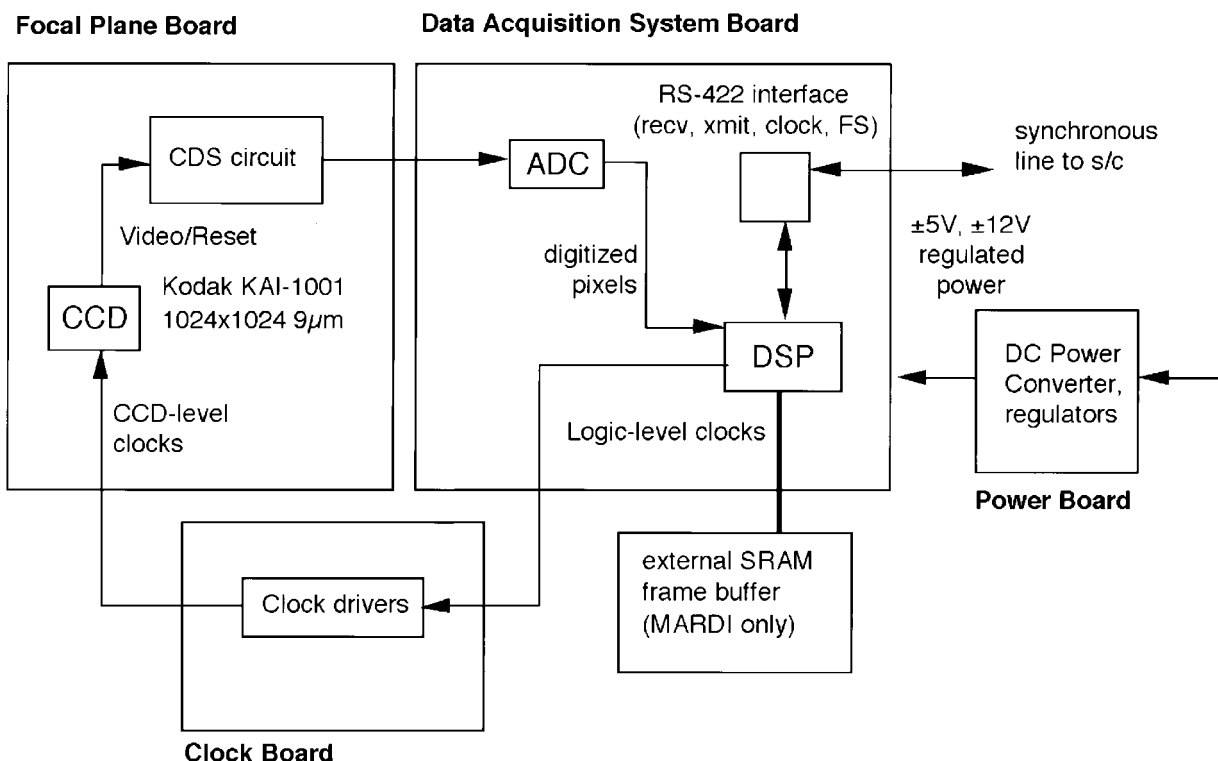


Figure 7. MARDI system block diagram.

id-state relays are used internally to switch the +28 V input power between the electronics and the survival heater. The survival heater is a 2 inch square, thin-film resistive heater fabricated by Tayco Engineering. It is mounted to the base of the camera and dissipates 3 W at the minimum spacecraft bus voltage of 22 V.

4.2.8. Latchup protection. The DSP is a commercial CMOS component with a low latchup threshold typical of such devices. To protect the DSP from damage caused by radiation-induced latchup, a current-sense circuit on the power board senses increased latchup current and opens a crowbar transistor, shorting the +5 V output of the DC-DC converter to ground. This triggers the internal short-circuit protection of the converter, which removes power from the system for roughly 10 ms. This latchup reset was originally intended to be sensed by monitoring software in the spacecraft, which would then have rebooted the DSP. However, the MSP Project decided not to implement this form of latchup protection. Instead, two power resets are directly commanded by the spacecraft: one ~30 s into the parachute descent and another just prior to the last nonpowered descent image. All other components used in the MARDI are either adequately radiation hardened or draw sufficiently low currents to be protected by current-limiting resistors.

4.2.9. Health monitoring. No dedicated internal instrument health monitoring is provided. Two temperature sensors (Analog Devices AD590) are provided, one on the FPA and one on the DAS board. These are monitored by the spacecraft computer. Instrument health diagnostics, including coarse measures of CCD voltage and temperature, are extracted from dark-current imaging.

4.2.10. Mechanical form. The MARDI electronics (Figure 8) consist of four printed circuit boards: the Focal Plane

Assembly, the clock generation board (CLK), the Data Acquisition System, and the power regulation and switching board (PWR). The FPA is a five-layer rigid-flex polyimide-glass 1.3 inch by 1.3 inch board whose rigid part is entirely contained within the instrument housing. The FPA flex cable passes through a slot in the housing and connects to the top of the CLK board.

The other three boards have an area for parts roughly 2.3 inches square. They are sandwiched between housing sections and are connected by flex cables on overhang regions outside of the housings. The DAS and CLK boards are 10-layer and 6-layer, respectively, polyimide-aramid boards for thermal stability. The PWR board is a four-layer polyimide-glass board. All boards have parts on both sides.

In order to maintain a complete Faraday cage for interboard signals, the overhang region of each board is plated with a copper shield, except for a cutout for the flex cable pads. The flex cables are custom fabricated with copper-mesh-plated Kapton layers on top and bottom and have shield traces on their outer edges conductively coupled to plated-through mounting holes. Once installed, copper-plated Kapton shields are placed over the connector pads on the board exteriors and conductively coupled to the board copper shields. The circuit boards are in direct conductive contact with the housing sections. Hence signals conducted between the boards via the flex cables are completely within the instrument's Faraday cage. The flex portion of the FPA is similarly shielded and connected.

This scheme was somewhat compromised for the spacecraft connectors, which are surface-mount types soldered directly to pads on the overhang regions of two of the circuit boards (DAS and PWR) on the side opposite to the interboard flex cables (conventional connectors with shells were not available in form factors compatible with physical space constraints). These con-




Figure 8. Expanded view of mechanical elements of the MARDI design.

nectors are 50-mil-pitch strip types whose pins are embedded in a Ryton matrix without metal shields. Metal tape is used to shield the connectors.

Radiation shielding is provided by aluminum housings with wall thicknesses of ~ 150 mils. This reduces total dose to the 5–10 Krad(Si) level for the CLK, DAS, and PWR boards. Additional shielding has been added to keep the FPA under 1 Krad(Si).

Electronics mass, including housings, is ~ 350 gm.

4.3. Software

MARDI software runs on two processors: the main spacecraft C&DH CPU and the DAS DSP. The CPU is responsible for instrument operational commands and downlink processing. The DSP is responsible for generating the CCD clocks, performing signal processing on the digitized pixels, and transmitting the data output to the C&DH.

4.3.1. DSP software. The DSP software is written entirely in assembly language. As indicated previously, it performs all

clock generation and pixel processing for the camera. Mission-specific versions of the software have been written to tailor the camera to the capabilities needed. The DSP code captures the image from the CCD and immediately converts it from 12-bit to 8-bit per pixel form via a table lookup, storing the pixels in the external SRAM buffer. Image acquisition into SRAM is completed before readout to the spacecraft begins. During readout the DSP compresses the pixels using a fixed Huffman encoding of the first-order difference of each pixel value. The compressed data can be resynchronized on a line-by-line basis in the event of data loss or corruption. Each image is compressed to a fixed size assuming a nominal 1.7:1 compression factor. If the compressed image is larger than this, it is cut off on the bottom, and if smaller, fill data are sent.

4.3.2. C&DH software. MARDI performs most of its command and image processing functions internally. However, image acquisitions are commanded from the C&DH CPU, using the serial communication line. The C&DH is based on a

20 MHz RAD6000 (a rad-hard derivative of the IBM/Motorola PowerPC architecture) processor board built by Lockheed Martin Federal Systems Division. High-level commanding, downlink formatting, and minimal fault protection are provided by instrument-supplied software running in the C&DH CPU.

The C&DH flight software is written in ANSI C for the VxWorks environment provided for the C&DH by Lockheed Martin Astronautics (LMA), the spacecraft manufacturer. It relies on LMA system services for basic functions such as uplink commanding, downlink data transfer, and communication with instrument hardware. The MARDI version of this code provides internal command sequencing for timed image acquisition and housekeeping functions, image memory allocation, communication with the camera hardware, and image processing.

4.4. Spacecraft Interface

The Mars Polar Lander (MPL) hardware interface is quite simple. The instrument is powered from an unregulated +28 V source. A logic-level discrete line is provided to switch the power internally to the electronics and to the survival heater. Two temperature sensors are provided for external monitoring of the electronics and focal plane. The camera typically draws 3 W while imaging and <2 W idle.

The serial communications interface uses RS-422 differential signaling to carry four signals. These four signals consist of a synchronous serial data transfer scheme common to Motorola, Texas Instruments, and Analog Devices DSPs; they are clock, frame sync, receive data, and transmit data. The maximum clock rate is 2.5 MHz for boot and 15 MHz thereafter. The clock runs continuously and the frame sync is a 1-bit-wide pulse which signals the start of each 16-bit word or "frame." Although the interface allows full-duplex operation (data can flow both to and from the camera simultaneously), this capability is not typically used in practice, in part because of limitations in the C&DH serial interface implementation.

4.5. Test and Calibration

MARDI performance was characterized and calibrated prior to delivery to the spacecraft. Details of this calibration will be included in the postflight science papers, primarily because in-flight calibration will also be used as constraints. Among the measurements made in the laboratory and during flight are absolute photometric performance, system noise performance, optical distortions (geometric and photometric), system MTF over the full FOV, and stray light susceptibility. Instrument orientation with respect to spacecraft axes was measured by photographing fiducial marks on the spacecraft assembly stage that had been previously surveyed.

5. Concluding Remarks

Descent imaging has not to date been considered a primary mission requirement for Mars landers. Rather, it plays a secondary or supporting role in planetary lander missions. Both the Mars Surveyor 1998 and Mars Surveyor 2001 Projects have compromised the potential return from their respective descent imaging systems by locating their cameras where their FOVs are obscured by descent engine plumes, by accepting large, rapid, and uncontrolled attitude excursions during parachute descent, and by taking no steps to mitigate the substantial vibration induced by pulse-modulated descent rocket engines. Despite the formidable challenges of imaging under

these conditions, both descent imaging systems are likely to provide important and useful results.

Recent Mars Global Surveyor Orbiter Camera observations indicate that much of Mars may be inaccessible to soft landing systems that do not employ active terminal hazard avoidance. Thus descent imaging may play a much more important role in future planetary lander missions, one that indeed may be enabling. The performance of the Mars Polar Lander MARDI may be critical to future Mars landing missions.

6. Loss of Mars Polar Lander

The Mars Polar Lander was launched on January 3, 1999. As the Polar Lander is completely enclosed within the atmospheric entry protection system, no calibration observations were possible of astronomical objects. However, observations were made twice during the cruise to Mars to search for any changes in instrument performance. Surprisingly, these images showed that a substantial amount of light was making it into the aeroshell; the source of the light is believed to be illumination through a thin, light-transmitting multilayered insulation gasket around the cruise rocket engine mounting bracket. These images demonstrated that the instrument was working well following launch and were the only imaging data returned by the Mars Descent Imager. The spacecraft and its science instruments were lost during entry, descent, or landing on December 3, 1999. An extensive search for the lander or effects of its encounter with the Martian surface was made using the Mars Global Surveyor Orbiter Camera, but as of the middle of January 2000, no trace of the Polar Lander had been found. It is likely that the Mishap Review Board for the Mars Polar Lander will conclude that the loss of the spacecraft was at least in part the result of poor training, insufficient personnel, and poor management decision-making processes.

Acknowledgments. The authors express their appreciation to the MARCI/MARDI hardware development team, led by M. Caplinger: Paul Otjens, Jeffrey Lewis, Jose Tamayo, Charles Schmitz, and Michael Ravine. We are extremely indebted to Steve Bailey for his efforts on our behalf during spacecraft assembly, test, and launch operations. MARDI was first conceived and developed under Planetary Instrument Definition and Development Program funding; the authors are indebted to William Quaide (now retired) of NASA Headquarters for his confidence in our effort and to T. Soulanille, K. Andrews, and T. Cushing (deceased) for the initial design study that eventually led to MARDI. Juan Reyces designed the MARDI optics; Optical Corporation of America refined that design and fabricated the lens. This work supported under JPL contract 960480.

References

- Folkner, W. M., C. F. Yoder, D. N. Yuan, E. M. Standish, and R. A. Preston, Interior structure and seasonal mass redistribution of Mars from radio tracking of Mars Pathfinder, *Science*, 278, 1749–1752, 1997.
- Golombek, M. P., et al., Overview of the Mars Pathfinder Mission and assessment of landing site predictions, *Science*, 278, 1743–1748, 1997.
- Johnson, J. R., et al., Preliminary results on photometric properties of materials at the Sagan Memorial Station, Mars, *J. Geophys. Res.*, 104, 8809–8830, 1999.
- Maki, J. N., J. J. Lorre, P. H. Smith, R. D. Brandt, and D. J. Steinwand, The color of Mars: Spectrophotometric measurements at the Pathfinder landing site, *J. Geophys. Res.*, 104, 8781–8794, 1999.
- Mandelbrot, B. B., How long is the coast of Britain? Statistical self-similarity and fractional dimension, *Science*, 156, 636–638, 1967.
- Mandelbrot, B. B., *The Fractal Geometry of Nature*, 468 pp., W. H. Freeman, New York, 1982.

- Morris, E. C., Viking 1 Lander on the surface of Mars: Revised location, *Icarus*, **44**, 217–222, 1980.
- Morris, E. C., K. L. Jones, and J. P. Berger, Location of the Viking 1 lander on the surface of Mars, *Icarus*, **34**, 548–555, 1978.
- Oberst, J., R. Jaumann, W. Zeitler, E. Hauber, M. Juschel, T. Parker, M. Golombek, M. Malin, and L. Soderblom, Photogrammetric analysis of horizon panoramas: The Pathfinder landing site in Viking orbiter images, *J. Geophys. Res.*, **104**, 8927–8933, 1999.
- Parker, T. J., R. L. Kirk, and M. E. Davies, Location and geologic setting for the Viking 1 Lander, *Proc. Lunar Planet. Sci. Conf. 30th*, 2040, 1999.
- Scheidegger, A. E., Stochastic models in hydrology, *Water Resour. Res.*, **6**, 750–755, 1970.
- Stoker, C. R., E. Zbinden, T. T. Blackmon, B. Kanefsky, J. Hagan, C. Neveu, D. Rasmussen, K. Schwehr, and M. Sims, Analyzing Pathfinder data using virtual reality and superresolved imaging, *J. Geophys. Res.*, **104**, 8889–8906, 1999.
- Stooke, P. J., Horizon topography and the location of Viking Lander 2, *Earth Moon Planets*, **76**, 47–65, 1998.
- Stooke, P. J., Revised Viking landing site (abstract), *Proc. Lunar Planet. Sci. Conf. 30th*, 1020, 1999.
- Tarboton, D. G., R. L. Bras, and I. Rodriguez-Iturbe, The fractal nature of river networks, *Water Resour. Res.*, **24**, 1317–1322, 1988.
- Turcotte, D. L., Fractals in geology and geophysics, *Pure Appl. Geophys.*, **131**, 171–196, 1989.
- Turcotte, D. L., Fractals in geology: What are they and what are they good for?, *GSA Today*, **1**, 1–4, 1991.
- Xu, T., I. D. Moore, and J. C. Gallant, Fractals, fractal dimensions and landscapes—A review, *Geomorphology*, **8**, 245–262, 1993.
- Zeitler, W., and J. Oberst, The Mars Pathfinder landing site and the Viking control point network, *J. Geophys. Res.*, **104**, 8935–8941, 1999.

M. A. Caplinger and M. C. Malin, Malin Space Science Systems, P.O. Box 910148, San Diego, CA 92191-0148. (malin@msss.com)

M. H. Carr, U.S. Geological Survey, 345 Middlefield Road, Menlo Park, CA 94025.

S. Squyres, P. Thomas, and J. Veverka, Center for Radiophysics and Space Research, Cornell University, Ithaca, NY 14853.

(Received July 13, 1999; revised February 11, 2000; accepted February 29, 2000.)