

Fast Optical Hazard Detection for Planetary Rovers using Multiple Spot Laser Triangulation.*

Larry Matthies, Tucker Balch & Brian Wilcox
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109

Abstract—A new laser-based optical sensor system that provides hazard detection for planetary rovers is presented. The sensor can support safe travel at speeds up to 12cm/second for large (1m) rovers in full sunlight on Earth or Mars. This is at least a 5 times improvement over the sensor aboard NASA’s Mars Pathfinder rover. The system overcomes limitations in the older design that require image differencing to detect a laser stripe in full sun. The new system ensures the projected laser light is detectable in a single image, eliminating the requirement for additional difference images. The improvement is significant since any reduction in image gathering or processing time provides for faster rover motion. The savings are even more important in the case of a Mars rover since power and radiation-hardening requirements lead to severely constrained computational resources. The paper includes a thorough discussion of design details and tradeoffs for optical hazard sensing that will benefit future efforts in this area.

I. INTRODUCTION

This article presents a new laser-based optical system for mapping the terrain before a planetary rover. Although some components of the design have appeared in an earlier system, this is the first in-depth discussion of many of the factors involved in the design of both.

Planetary rovers must be able to accurately, quickly, and reliably assess the terrain before them. Several active and passive technologies are available for the task, including fixed light stripe systems [8], mechanically-scanned light stripe systems [4] and stereo vision [7]. In the case of rovers embarking to other planets, the rigors of space travel and the target planet’s environment

must be considered. Reliability is crucial to mission success since repair of failed components is unlikely after launch. For this reason, solid-state solutions are preferred over mechanically scanned systems that might not survive the vibration and G-forces of launch, atmospheric entry and landing on Mars. Also, in the case of Mars, computing power is an important consideration. Radiation hardening and flight certification for instance, have restricted the Mars Pathfinder rover to an 8085-based CPU. This impacts the processing power available for all tasks including vision. Stereo vision approaches aren’t practical for such computationally restricted rovers. This work began with the realization that a few simple changes to the existing Pathfinder sensor design would greatly improve performance while maintaining flight qualification. Some specific constraints and goals for the hazard sensor system include:

- **Flight-qualifiable components:** This is always a requirement for planetary rovers, but the acceleration of Mars exploration efforts underscores the importance of using *already* qualified components and materials to the maximum extent possible.
- **Mass:** All sensor systems should be low-mass to provide for larger science payloads and the potential of sample return missions.
- **Power:** Most rover designs utilize solar or battery power for all on-board systems. Mission success and duration depend heavily on low power consumption.
- **Mechanical simplicity:** Reliability in space requires mechanically simple systems with few or no moving parts.
- **Computational simplicity:** image processing requirements should be minimized.

*ICRA-97: *IEEE International Conference on Robotics and Automation*




Fig. 1. Laser and camera configuration aboard the LSR-1 rover. The shaded areas represent distributed laser beams.

II. BACKGROUND AND RELATED WORK

The sensor is a derivative of the hazard detection system now flying aboard the Mars Pathfinder rover [8]. Pathfinder uses a laser stripe-based structured light system. Five lasers project stripes onto the terrain imaged by the rover's cameras. Two images are taken for each stripe: one with the laser turned off and another with it on. The difference between the two images reveals the stripe on the terrain. The lasers and cameras are arranged so that the height of the terrain relative to the rover can be determined using simple geometry once the stripes are detected.

The Pathfinder system is reliable and accurate. It is also fast compared to stereo vision approaches since the computational requirements are significantly lower. There are, however, several limitations. It takes approximately 20 seconds to complete a hazard scan which returns only 20 terrain data points. The sensor's performance is due to several interacting factors. Pathfinder's power budget allows only one of the five laser stripes to be active at a time. Since laser light intensity is reduced as it is spread over a stripe, image differencing is required to find the stripe in an image. These factors combine to require 10 separate exposures for each hazard scan. The process is especially time consuming on the 8085-based system since image acquisition is slow. The new approach addresses the first two of these limitations by improving the optical system so that only two lasers and two exposures are required to compute 30 data points. Technologies




Fig. 2. The Rocky III rover. This front view of the robot shows how the laser spots are arranged as a "push-broom" ahead of it.

for reducing image acquisition time are not addressed here, but are expected to provide a significant performance improvement as well.

This article covers design of the system, and considerations for choices made in selection of its optical components.

III. APPROACH

The new system has been integrated into two rovers: the Lightweight Survivable Rover (LSR-1) and Rocky III (Figures 1 and 2). Both robots utilize a six-wheeled rocker-bogey suspension like that of the Mars Pathfinder rover. The LSR-1 rover is equipped with one of the 8085-based computers developed for Pathfinder, while Rocky III uses a 486-based computer.

The rovers are equipped with two laser/camera pairs for hazard detection. Each laser/camera pair operates as an independent hazard sensor. For now, consider one of the pairs by itself. First, laser light is split into 15 co-planar beams using a commercially available diffraction grating. The central beam is aimed at the ground in front of the rover (50cm ahead for LSR-1). The rest of the beams fan out to the left and right. On flat terrain, the beams form a straight line of spots. If an obstruction is present, they follow its contour instead (Figure 3). The camera is offset horizontally from the laser at the same height above the ground. On-board software finds the spots in the image which are in turn used to determine coordinates




Fig. 3. Laser and camera geometry, left; example image acquired by the sensor, right. The bright spots are laser beams.

of the terrain. From the camera’s point of view, each spot shifts left or right depending on the height of the terrain the corresponding beam strikes.


Pathfinder’s laser energy is spread across a stripe designed to cover 300 pixels in the imaged terrain versus only 15 (one pixel per laser spot) for this system. The increased signal-to-noise ratio eliminates the need for image differencing. More details concerning the signal-to-noise ratio are covered in Section IV.

Disparity is computed between where a laser spot should appear if the rover were on flat ground and where it is actually detected. The approach is similar to other structured light systems [1]. Equations for each beam, and a model of the camera’s optics are computed ahead of time in a calibration phase. These equations are sufficient for computing the locations of the laser spots on the terrain.

One issue remains: how can the 15 beams be differentiated? Fortunately, the geometry of the system provides a convenient solution. The laser and camera are positioned so that a line drawn between their optical centers is parallel to scanlines in the image plane. This arrangement ensures each beam will always appear on the same scanline. If the plane of the beams is tilted slightly, each beam will appear on a separate scanline. The beams may then be found unambiguously by scanline. Furthermore, the knowledge that each beam will only ever appear on a specific scanline may be used to reduce the image processing necessary to find them.

IV. OPTICAL DESIGN

The primary optical components of the system include the laser, the laser beam splitter (diffraction grating),



camera lens, camera filters and the CCD. In this analysis we consider each of these and show why we expect the system to work well in conditions on Earth and on Mars.

A. Laser versus Solar Brightness

To meet performance goals for hazard detection, the system must detect laser spots at a nominal range of 65cm in full sun. A key factor regarding spot detectability is the ratio of spot brightness to reflected sunlight. The required laser brightness depends on the type of processing used to find spots, the range of albedo in terrain features and, depending on the processing technique, the spatial frequencies of albedo (texture). If spots are to be detected by thresholding alone for example, the following relation must hold:

$$\rho_{\text{low}}\Phi_{\text{solar}} + \rho_{\text{low}}\Phi_{\text{laser}} > \rho_{\text{high}}\Phi_{\text{solar}}$$

where ρ_{low} and ρ_{high} are the highest and lowest albedos of objects in the scene, and Φ_{solar} and Φ_{laser} are the power of illumination by the Sun and laser. If this relation is maintained, even when the laser strikes the darkest part of the terrain the reflected light is still brighter than the sun on the brightest surface in the scene. A 5 to 1 laser to solar illumination ratio, for instance, is the minimum required for terrain with albedos from 0.05 to 0.3 (6 to 1).

Image processing using a matched filter or other peak-detection logic can reduce the required laser brightness significantly. Although a complete analysis of required laser brightness has not been completed, our experience over many experiments in JPL’s “Mars Yard” (a Mars analogue test area) indicates a 5 to 1 ratio of laser to solar illumination is sufficient to allow

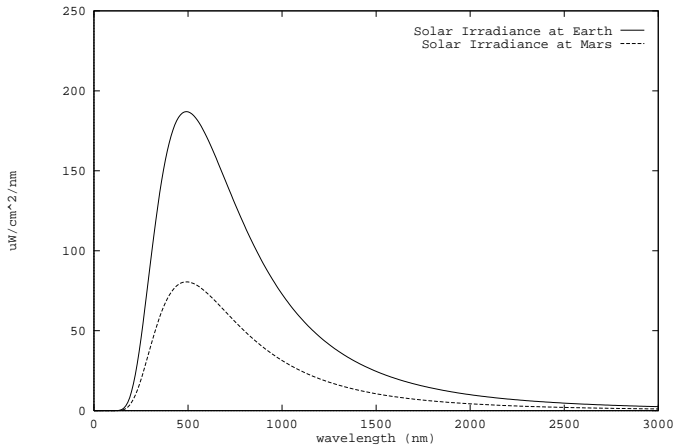


Fig. 4. Comparison of solar irradiance at Earth and Mars, with the Sun modeled as a 5900 degree K black body [3].

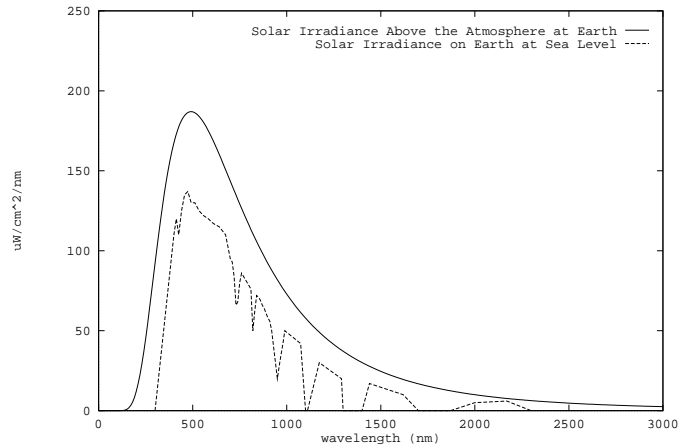


Fig. 5. Solar irradiance on the surface of Earth showing absorption of some frequencies by the atmosphere [3].

detection of laser light in nearly all terrain configurations.

This design seeks to maximize the ratio in two ways: first by using a laser whose wavelength (860nm) is in a region where solar power is low (see Figure 4); second, by filtering the reflected light to a narrow band around that wavelength before the scene is imaged. As Figure 4 illustrates, solar irradiation at the distance of Mars from the Sun is 43% of that at Earth. The system was designed for the brighter conditions of Earth for convenient testing, so performance will be even better on Mars.

The power and efficiency of the laser, the number of beams the laser light is split into and the efficiency of the beam splitter all contribute to the final brightness of the spots. 30% efficient 860nm space-flight qualified lasers are employed on Pathfinder and this derivative system. When powered they emit a 100mW beam. Laboratory measurements of the laser optics show that the diffraction grating transmits 76% of the laser energy into the 15 primary beams. Each of the 15 spots is more than 5 times as bright as the Sun on Earth out to 1 meter from the robot (in the 800-900nm band).

B. Camera Optics

Before introducing the equations governing performance of the camera system, it is useful to consider the path light follows from the Sun to a pixel in the robot's camera. Upon reaching the Earth or Mars, collimated sunlight passes through the atmosphere, where several frequencies are fully or partially absorbed before striking the terrain (see Figure 5). The terrain absorbs some more of the light, depending on its albedo,

and reflects the remainder in all directions across a hemisphere of 2π steradians (sr), assuming a planar surface. The intensity of light that bounces in the direction of the camera depends on the angle at which it struck the surface (incidence) and the angle from which the surface is being viewed (reflectance). Finally, the light is gathered by the lens and focussed on a pixel. The larger the diameter of the lens, the more light it can gather and the brighter the apparent image. Laser light projected by the sensor follows a similar path, but atmospheric attenuation is negligible for such a short trip. We consider the overall irradiance at a patch of terrain to be the sum of direct solar and laser light at that point.

Our analysis includes the following factors and assumptions:

- **Solar illumination** with a spectrum as depicted in Figure 5.
- **Lambertian reflectance:** the surface reflects light uniformly in all directions. Such a surface appears equally bright from all viewing directions.
- **Angle of incidence for solar illumination:** assumed to be normal, to simplify analysis.
- **Angle of incidence for laser illumination:** assumed to be normal, to simplify analysis.¹

¹ A potential concern with this assumption is that it may tend to exaggerate the brightness of a laser spot in the image. However, research into the reflectance properties of diffuse surfaces ([5, 10]) indicates that no matter what the angle of incidence, if the viewer and illumination source are co-located, the surface appears even brighter than predicted by Lambert's Law for normally incident illumination. In the case of the LSR-1 rover, the

- **Angle of reflectance:** is assumed to be normal for laser and solar illumination.
- **Albedo:** Viking orbiter measurements show Mars surface albedo to vary from 0.10 to 0.36 [2]. For this analysis we selected $\rho = 0.3$.
- **Optical filters:** to admit a 100nm band of light centered at 850nm.
- **Aperture:** the diameter of the camera lens' opening.
- **Focal length** of the lens.
- **Quantum efficiency** of the CCD imaging device.
- **Size of pixels** in the CCD imaging device.
- **Well depth** of pixels in the CCD imaging device.
- **Exposure time:** 0.1sec . The shortest reliable exposure for the camera hardware is 0.01sec. Using 0.1sec provides room for adjustment later.

We do not consider:

- **Focus**
- **Diffraction effects**
- **Phase angle**

Assuming laser light incident on the terrain is approximately 5 times as bright as sunlight in the filter pass band, camera design can proceed with the goal of exposing pixels to $\frac{1}{6}$ of their capacity for nominal images of terrain. This exposure will provide the greatest dynamic range in the image and facilitate differentiating the laser spots from other image features.

The analysis will proceed from each “end” of the problem. First, the power necessary to expose one image pixel to $\frac{1}{6}$ capacity is considered. Next, filter and lens components to meet that requirement are selected.

C. Pixel Exposure

In this section the energy required to expose one pixel to $\frac{1}{6}$ (16.7%) of its capacity is calculated. The computation is based on manufacturer’s specifications for the Kodak KAI-0370NI CCD, but the approach is appropriate for other imaging devices as well.

Important factors in CCD pixel performance are **size, well depth, quantum efficiency, and fill factor**. Size refers to on-center spacing between pixels.

laser and camera are approximately 11 degrees apart from the point of view of the nominal terrain.

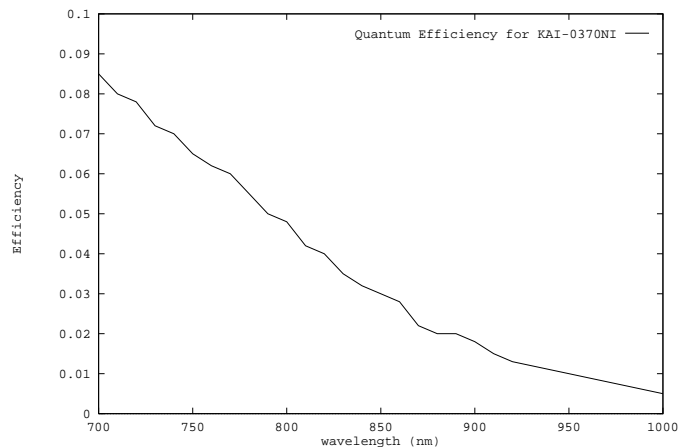


Fig. 6. Quantum efficiency of the KAI-0370NI CCD.

Spacing is most often rectangular, but not strictly square. Incident photon energy at a pixel is measured as the charge on a capacitor, with well depth being the fully saturated charge (in electrons). Quantum efficiency is the rate at which the device converts photons to electrons at the capacitor. Figure 6 shows how quantum efficiency varies with wavelength in the region of the spectrum around the laser’s wavelength. Finally, fill factor is the percentage of pixel area able to gather light. The KAI-0370NI specifications incorporate fill factor into the quantum efficiency chart.

The equation governing intensity registered at a pixel is:

$$\text{intensity} = \frac{\Phi_i \times \text{exposure_time} \times \text{quantum_efficiency}}{E_p \times \text{well_depth}}$$

Where intensity refers to the level of exposure from 0.0 to 1.0 (fully saturated). Φ_i is the power incident on one pixel, E_p is energy per photon. E_p is computed for light in the middle of the filter bandpass region (850nm). Since (16.7%) exposure is desired, intensity = 0.167. Other parameters are easily computable, dictated by hardware constraints, or are available from the manufacturer’s specification sheet. For this system:

Pixel Size	158 μm^2
Exposure Time	0.100 sec
Quantum Efficiency	0.0300 (at 850nm)
Well Depth	60,000 electrons
Photon Energy	2.34×10^{-19} J (at 850nm)

Now, solving Equation 1 for power, and substituting

in the known quantities:

$$\begin{aligned}\Phi_i &= \frac{\text{intensity} \times E_p \times \text{well_depth}}{\text{quantum_efficiency} \times \text{exposure_time}} \\ &= \frac{0.167 \times 2.34 \times 10^{-19} \text{J/photon} \times 60,000 \text{electrons}}{0.0300 \text{electrons/photon} \times 0.100 \text{sec}} \\ &= 7.82 \times 10^{-7} \mu \text{W}\end{aligned}$$

$7.82 \times 10^{-7} \mu \text{W}$ is the power of light at 850nm required to expose one pixel 16.7% in 0.100 seconds.

D. Lens and Filter Parameters

This section considers lens and filter components that will deliver $7.82 \times 10^{-7} \mu \text{W}$ to each pixel for an image of nominal terrain. Since only the the light incident on a single pixel is of interest, it is assumed for now that the surface being imaged is a small patch of just the right size that its focused image exactly covers one pixel. The result generalizes across the entire image. The power of light at a pixel is given by [6]:

$$\Phi_i = \omega L \cos \theta_e A_s T_f$$

where

- $\omega = \pi d^2/4r^2$ is the solid angle subtended by the lens, as viewed from the surface patch [sr]. d is the diameter of the lens [m] and r is the distance from the lens to the surface patch [m].
- L is the radiance of the surface [$\text{W m}^{-2} \text{sr}^{-1}$].
- $\cos \theta_e A_s$ is the fore-shortened area of the surface patch [m^2], where θ_e is the emittance angle between the surface patch and the receiver line of sight.
- T_f is the transmittance of the filter.

For now, the lens parameters are fixed and an appropriate filter transmittance (T_f) is computed. To proceed, A_s , area of surface patch, d , the diameter of the lens, r , the range to the surface patch, L , radiance of the surface, and θ_e , and the emittance angle must be found.

The lens has a focal length of 8.00mm and an f-number of 4.00. f-number is defined as $\frac{\text{focal_length}}{d}$, so $d = 2.00 \times 10^{-3} \text{m}$. The camera is (on average) 0.650m from the laser spots on the surface. Since the angle of reflectance is normal to the surface, $\theta_e = 0$. With these parameters, ω is computed as:

$$\begin{aligned}\omega &= \pi d^2/4r^2 \\ &= \pi(2.00 \times 10^{-3} \text{m})^2/4(0.650 \text{m})^2 \\ &= 7.44 \times 10^{-6} \text{sr}\end{aligned}$$

Next, A_s is calculated. The area of the surface patch depends on the area A_i of an image pixel, the focal length, f , of the lens, the range, r , and the viewing angle θ_e as

$$\begin{aligned}A_s &= \frac{A_i r^2}{f^2 \cos \theta_e} \\ &= \frac{158 \mu \text{m}^2 (0.650 \text{m})^2}{(8.00 \text{mm})^2 \cos 0} \\ &= 0.0104 \text{cm}^2\end{aligned}$$

Finally, L , radiance of the surface, is given by:

$$L = \frac{\rho E}{\pi}$$

Where E is the surface irradiance in W/m^2 and $\rho = 0.300$ is the surface albedo. E can be approximated by inspecting Figure 5. Irradiance near the center of the 800-900nm band is $70 \mu \text{W}/\text{cm}^2 \text{nm}$, or approximately $7000 \mu \text{W}/\text{cm}^2$ across the whole band. So $L = 668 \mu \text{W}/\text{cm}^2 \text{sr}$. Therefore, the power at a pixel (no filter) is:

$$\begin{aligned}\Phi_i &= \omega L \cos \theta_e A_s \\ &= 7.44 \times 10^{-6} \text{sr} 668 \mu \text{W}/\text{cm}^2 \cos(0) 0.0104 \text{cm}^2 \\ &= 5.17 \times 10^{-5} \mu \text{W}\end{aligned}$$

Based on this, and the requirement for $7.44 \times 10^{-7} \mu \text{W}$ at the pixel, the desired transmittance for the filter is $T_f = \frac{7.82 \times 10^{-7} \mu \text{W}}{5.17 \times 10^{-5} \mu \text{W}} = 0.0151$ or 1.51%.

Reviewing the calculations up to this point: a filter system that transmits 1.51% of the light uniformly between 800 and 900nm, and 0% elsewhere will expose one pixel 16.7% in 0.1sec. A pixel will only be fully saturated when the corresponding surface patch is illuminated with laser light 5 times as bright as the Sun.

The 100nm pass-band filter is created by combining two colored glass filters: one cuts light below 800nm (RG-830), and another cuts light above 900nm (KG-1). The combination of the two results in 4% transmission between 800 and 900nm (Figure 7) The difference between the desired 1.51% and 4% can be made up by adjusting exposure time of the CCD appropriately.

V. RESULTS

The primary goals of this work are to provide fast, reliable and accurate information about the terrain for rover navigation. Each of these will be examined in turn.

The overall time required for a hazard scan depends primarily on time required to expose and collect images of the projected laser spots. Since the system



terrain reconstruction




Fig. 8. Sample terrain (left), and a reconstruction based on sensor data (right).




Fig. 7. Two filters are used in combination to restrict light to an 800 to 900nm band. About 4% of the light in that band is transmitted.

uses simple logic for spot detection, image processing is negligible. On the 1486-equipped rover (Rocky III) exposure time is 100ms per camera, acquisition takes 600ms, and an additional 30ms are required to turn the lasers on and off, find the spots and compute hazard heights. As a result, Rocky III can complete a hazard scan in 830ms. This is over 20 times the rate for Mars Pathfinder hazard scans. We estimate the 8085-based LSR system will require 4 seconds per scan due to the slower image acquisition time. This is about 5 times as fast as Mars Pathfinder using the same computer and camera technologies.

To avoid getting their wheels stuck in gaps between rocks rovers must sample the terrain for hazards at

least every $\frac{1}{2}$ wheel diameter of forward motion. At .83 sec per scan, the new sensor provides for speeds up to 12cm/sec on an LSR-sized rover (20cm wheel diameter) using a 486-class computer, or 2.5cm/sec using an 8085. This is 5 to 20 times as fast as could be maintained using the older sensor aboard Mars Pathfinder. Rocky III (12cm wheel diameter) can travel safely at 7.2cm/sec.

The speed of the system arises from the fact that simple peak-detection logic is used to find the laser spots. The reliability of this approach was evaluated by measuring the rate of spot detection in a series of 24 test images acquired over an area of simulated Mars terrain in full sun (Figure 8). In the 24 images, 339 spots were visible. Of these, 336, or 99% were detected. Also, no image features were mistaken for laser spots.

The sensor's accuracy in reporting the coordinates of obstacles was evaluated using the same test images. The reported height for each rock was found to be within 1cm of the true height. Accuracy in lateral and fore-aft dimensions was evaluated qualitatively by inspecting a reconstruction of the scene generated using the sensor data (Figure 8). Every hazard was detected.

The new system returns 50% more information about the terrain in each hazard scan than the older design (30 data points versus 20). When accurate odometry is available, the uniform arrangement of data points provides for 3cm grid terrain maps, similar to the one illustrated in Figure 8. These maps will be exploited in future work aimed at more efficient navigation. Conversely, the Mars Pathfinder sensor's 20 points are arranged in a non-uniform 4x5 grid that is not as appropriate for building terrain maps.

VI. CONCLUSION

We have presented the design and implementation of a hazard sensor that provides for safe travel at 5 to 20 times that of the system used on the Mars Pathfinder rover. Additionally, the sensor offers:

- **Reliability:** 99% of detectable spots are found using simple peak-detection logic with a very low false-alarm rate.
- **Accuracy:** The system has demonstrated better than 1cm accuracy in measuring obstacle height.
- **Flight-qualified components:** The laser, CCD and optical filters have already been qualified for flight aboard the Mars Pathfinder rover. The camera lens and laser beam-splitter should be easily qualified.
- **Low mass:** The sensor weighs approximately 200 grams, excluding cabling.
- **Low power:** Consumption peak is 1.1 Watt. 800mW is required for the cameras and 300mW powers each laser (turned on one at a time). The same sensor power budget as Mars Pathfinder provides hazard scans at 5 to 20 times the frequency, offering equivalent increases in safe traversal speed.
- **No moving parts.**

Several areas of future work will help further refine the sensor system. First, performance of the system will be more completely characterized in tests at the JPL "Mars Yard." A comprehensive data sample will be collected for quantitative evaluations. The design could be extended in several ways. At present, a primary bottleneck in processing is image acquisition. Future incorporation of an Active Pixel Sensor (APS) will reduce this overhead significantly [9]. An additional row of beams aimed down at a steeper angle will provide close-in hazard detection and allow for the case when the system must look over a hill as the rover approaches the top. Also the system will be integrated with stereo vision in a hybrid system. Finally, rover navigation strategies will be extended to take advantage of the higher resolution terrain data provided by the sensor.

VII. ACKNOWLEDGEMENTS

Todd Litwin developed the camera model and 3-dimensional line intersection software used in this system; he also provided several helpful suggestions. Hrand Aghazarian provided technical assistance on Rocky III and LSR-1.

REFERENCES

- [1] G.J. Agin and T.O. Binford. Computer description of curved objects. In *Proceedings of 3rd International Joint Conference on Artificial Intelligence*, pages 629–640, August 1973. Stanford, California.
- [2] P.R. Christensen and H. J. Moore. The martian surface layer. In Kieffer, Jakosky, Snyder, and Matthews, editors, *Mars*, pages 686–729. The University of Arizona Press, 1992.
- [3] C. Elachi. *Introduction to the Physics and Techniques of Remote Sensing*. John Wiley and Sons, 1987.
- [4] A. Gruss, S. Tada, and T. Kanade. A vlsi smart sensor for fast range imaging. In *Proceedings of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '92)*, pages 349–358. IEEE, July 1992. Raleigh, NC.
- [5] B. Hapke. Bidirectional reflectance spectroscopy 1. theory. *Journal of Geophysical Research*, 86(B4):3039–3054, 1981.
- [6] B. Horn. *Robot Vision*. MIT Press, 1986.
- [7] L. H. Matthies. Stereo vision for planetary rovers: stochastic modeling to near real-time implementation. *International Journal of Computer Vision*, 8(1):71-91, 1992.
- [8] L. Matthies, E. Gat, R. Harrison, B. Wilcox, R. Volpe, and T. Litwin. Mars microrover navigation: Performance evaluation and enhancement. *Autonomous Robots*, 2(4):291–312, 1995.
- [9] R Nixon, S. Kemeny, C. Staller, and E. Fossom. 128x128 cmos photodiode-type active pixel sensor with on-chip timing, control and signal chain electronics. In *Proceedings of the SPIE Conference on Charge-Coupled Devices and Solid-state Optical Sensors V*, volume 2415. SPIE, 1995.
- [10] M. Oren and S. Nayar. Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14(3):227–251, 1995.