Three Dimensional Imaging Utilizing Structured Light

Carl Christian Liebe, Curtis Padgett, Johnny Chang
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-354-7837
carl.c.liebe@jpl.nasa.gov

Abstract—This paper describes a method of remote sensing 3-dimensional structure of the proximity utilizing a laser, a holographic grating, and a single regular CCD camera. Basically, the laser beam is split by a holographic grating to form a regular spaced grid of laser beams that are projected into the field of view of a CCD camera. The laser source and the CCD camera are physically separated forming the base of a triangle. The exit angle of the laser beam and the angle measured by the camera where the beam intersects a surface is a function of the distance. These two angles and the distance between the source and the camera allow calculation of the range to the projected spot using triangulation. This paper describes an experimental proof of concept and an empirical calibration of the system. Encouraging results are achieved and presented. Also, an application in which this system potentially would be used for a Mars landing is discussed.

TABLE OF CONTENTS

1. INTRODUCTION 1
2. DESCRIPTION AND THEORY 2
3. EXPERIMENTS 6
4. CASE STUDY: STRUCTURED LIGHT UTILIZED FOR MARS LANDING SYSTEM 7
5. SUMMARY 9
REFERENCES 9

1. INTRODUCTION

For a number of different space based applications, a rover or a spacecraft requires a 3 dimensional map of the proximity to do its assignment. For autonomous mobility applications, range maps allow surface vehicles to plan paths to avoid obstacles and reach desired locations. For landing operations, range maps allow a descent vehicle to avoid hazardous conditions (rocks, craters, slopes) that are detrimental to subsequent operations (communication, power, deployment, etc.). Range maps also provide essential navigation information for other applications such as rendezvous and docking and formation flying.

The primary range map technologies employed by NASA in previous missions were based on visible stereo cameras. Future missions have considered alternative technologies such as laser radar [1] [2], microware radar [3] and structure from motion (single camera system), but all these systems are currently only in the developmental phase [4].

Radar has been used extensively on Earth for a variety of different applications. Radar (phased array, SAR) sends out energy and is able to construct a range map of the surface. Its major drawbacks for space applications have been its cost, mass and power consumption. A radar system is currently base-lined to provide a surface map for the 2009 Mars mission called Mars Science Laboratory (MSL) [3].

For small Mars rovers, the technology of choice has been a stereo camera system. This approach determines range by matching a feature in one camera frame to its position in the opposite camera image. The difference in the angular position of the feature in the two camera frames is related to range using triangulation [5]. The stereo system (as compared with radar) has excellent resolution at short ranges (range can be recovered for the great majority of shared pixels) and consumes only a few watts of power. With no active light source, a stereo system is limited to daytime operations (though a flash lamp could be added to provide a night capability).

A structured light system consists of an active light emitter that projects a pattern of light beams onto a surface in front of a camera that images the pattern [6]. The camera is offset from the light emitter. Each beam of light exits the light source grading at a fixed angle. Using that angle and the separation of the camera from the optics, range can be recovered from the image using triangulation. One way to generate a bundle of light beams is to pass a laser beam through a holographic diffraction grating. This will split the original laser beam up into many individual laser beams that compose the pattern.

The different technologies that can be used to generate a range map are summarized in Table 1.

[1] 0-7803-8155-6/04/$17.00 © 2004 IEEE
[2] IEEEAC paper#1488, Version 1, Updated December 3, 2003

The approach we describe in this paper explores how to utilize a grid of laser beams to generate a range map for hazard detection and navigation. The actual size and dispersion of the grid depends on the application. For instance, descent and landing operations typically require hazard information about the landing site at a very high altitude in order to make the necessary trajectory adjustments while the spacecraft still has control authority. The spread of the grid would need to cover the landing site and be of sufficient density to detect potential hazards (rocks). For rover surface operations, a wide field of view would be required for path planning and hazard avoidance (> 90°).

Table 1. Different technologies for generating range maps

	Structured light system	Stereo Vision	Laser Radar	Microwave radar
Mass	<1 kg	<1 kg	6 kg	~30 kg
Power	A few watts	A few watts	<40 W	~200W
Max operating distance	Tens of meters	Tens of meters	2.5 km	Kilometer range
Night time operation	Yes	No	Yes	Yes
Computational Demand	< 10 MIPS	> 100 MIPS	< 1 MIPS	>> 100 MIPS

The next section provides a description of the technology and theory. The section after that describes a system we have developed for proof of concept and we will discuss the results of our experiments. Finally, we provide an example of a system that addresses safe landing on Mars.

2. DESCRIPTION AND THEORY

The structured light system is based on the same principles as stereovision – it determines the distance to the laser spots utilizing triangulation. In stereovision, two cameras are located a distance apart and observe the same point simultaneously. The same image feature is found and correlated in the two images, and the 3D position is found by triangulation [7]. One of the difficulties in stereovision is to figure out which image features in the first image correspond to which image features in the second image. Also, stereovision will not operate in darkness without active illumination.

An isolated bright spot, for example, is very easy to identify in an image, and it can be accurately located (using a centroid). However, only one distance for each laser beam is found. Stereovision, on the other hand, finds the distance to a high number of pixels.

Assuming that the camera and laser/grating system is aligned perfectly (i.e. the x axis of the focal plane is parallel with the baseline), then the y component of the spot is observed at the same y value on the camera for all distances. This is true for a co-boresighted stereo camera system as well – a feature found in the left image will appear at the same y value in the right image. The distance from the camera to the feature will determine the horizontal displacement of the feature in the images.

The idealized scenario for a single laser beam triangulation is shown in Figure 2. In this scenario a coordinate system is defined with the origin at the exit aperture of the laser. The x-axis of the coordinate system is oriented towards the equivalent pinhole of the camera. The laser and the equivalent pinhole are separated by a baseline, B. The distance from the laser to the target is Z. It is observed in Figure 2 that the distance to the target can be determined from the equation:

$$\frac{B}{Z} = \tan\alpha \Leftrightarrow Z = B\cot(\alpha) \qquad (1)$$

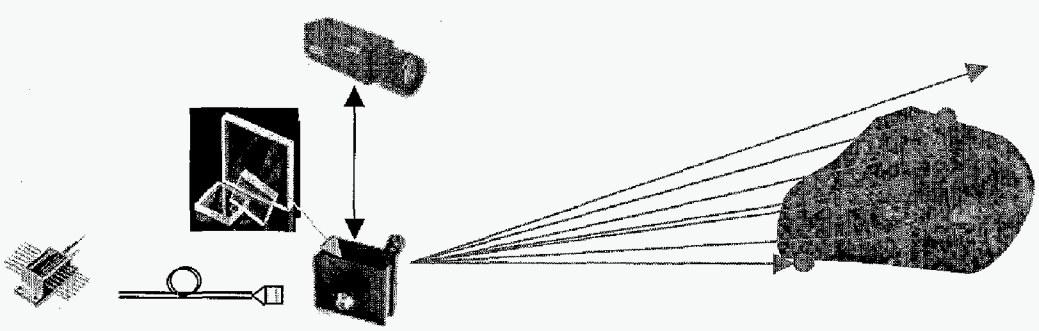

Figure 1 - Components of a patterned light system. From left to right: a laser and optional fiber coupling; a holographic diffraction grating, mechanical housing for the grating; and a camera, offset from the grating. The illustration depicts a laser beam being divided into a pattern of spots that are projected into the camera field of view.

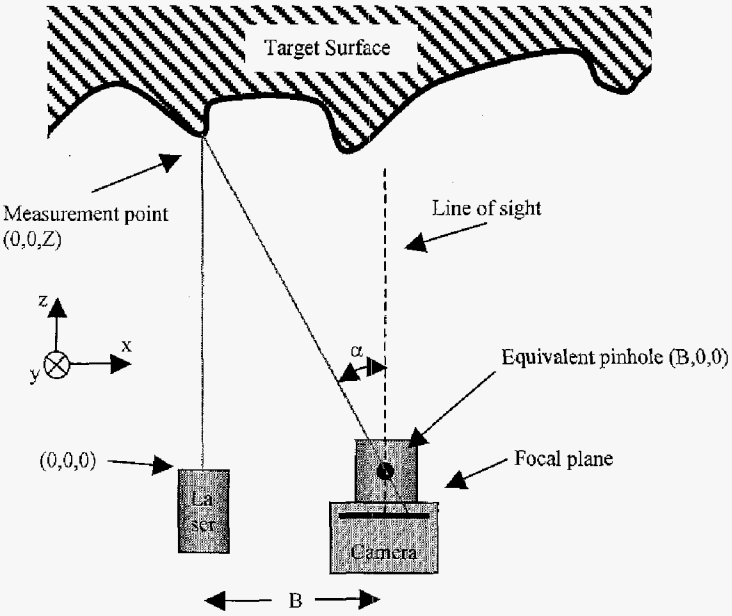

Figure 2 - Sketch of an idealized system for laser triangulation

The uncertainty on Z can be found by a sensitivity analysis with the appropriate uncertainties on B and α. The uncertainty of α depends primarily on 3 quantities:

- field of view of the camera (FOV),
- number of pixels
- the accuracy of centroiding the spot.

This is shown in Equation 2.

$$\varepsilon_\alpha = \frac{FOV \cdot \varepsilon_{subpixel}}{N_{Pixels}} \qquad (2)$$

As an example, a camera has a FOV of 10 degrees, 512 pixels and is able to determine the centroid of spots to an accuracy of a 1/10 of a pixel. This means that the average angular accuracy is 0.002 degrees. Assuming a baseline of 20 cm, the depth accuracy at 10 meters is then ~8 cm.

The discussion up to this point has assumed an idealized system. However, the components of the system will be misaligned. The way that we have defined the coordinate system (the x-axis starting at the laser exit aperture and going towards the equivalent pinhole of the lens) prohibits unwanted translations in the system. All errors show up as rotational errors. A non-idealized system is sketched in Figure 3.

In Figure 3, it is observed that there are 6 parameters that describes the geometry: 1) The baseline (B). 2) The elevation (El) and azimuth (Az) angle that the laser beam is offset relative to pointing along the z axis and 3) a rotation matrix A (3 degrees of freedom) that describes the rotation of the camera relative to the ideal situation. This is a total of 6 degrees of freedom. Also, internal parameters of the camera (e.g. focal length, intersection of line of sight and the focal plane and the optical distortion) are additional degrees of freedom. For a conventional calibration, taking a large number of measurements and solving an over-determined set of equations would derive the unknown constants.

An alternative to the conventional calibration method is to employ an empirical method to calibrate the system. The reason for choosing an empirical calibration is that it requires less effort to implement.

A setup is made where the non-idealized system is placed in front of a flat wall at a minimum distance (e.g. 1-meter). An image is acquired and the centroid of the spot is calculated. The system is then moved to a new distance (e.g. 1 meter and 10 centimeters) and a new image is acquired and the centroid is calculated. This procedure is repeated up to the maximum distance that the system is required to work. The scenario is shown in Figure 4.

The measured centroids during the calibration sequence are superimposed on an artificial image as shown in Figure 5.

After the calibration, the system is pointed at a surface of unknown distance and an image is acquired and the centroid is calculated. An example centroid of unknown distance is shown in red in Figure 6.

It is observed that the measured centroid lies between 1.4 meters and 1.5 meters The unknown distance is therefore in that range. A more accurate estimate is made the following way: A straight line is projected through all of the calibration points in an RSS sense. This is shown as the black dotted line in Figure 6. The red dot is then projected down onto the dotted line. Now all calibration points are plotted in a new coordinate system. The distance (in pixels) from the first calibration point is the unit on the x-axis and the distance to the target is the unit on the y-axis. This is shown in Figure 7.

A polynomial is fitted though the points and it is possible to estimate the Z distance based on the projection onto the line between the calibration points. In the example from before, the red dot is estimated to be at a distance of 1.44 meters.

In the discussions up to this point, we have discussed the theory for a single spot. However, suppose that two laser beams illuminated the scene. It would then be possible to determine the 3D position of the two spots independently utilizing the same theory on each spot. We can continue to add spots as long as we can uniquely associate each spot with its individual calibration curve.

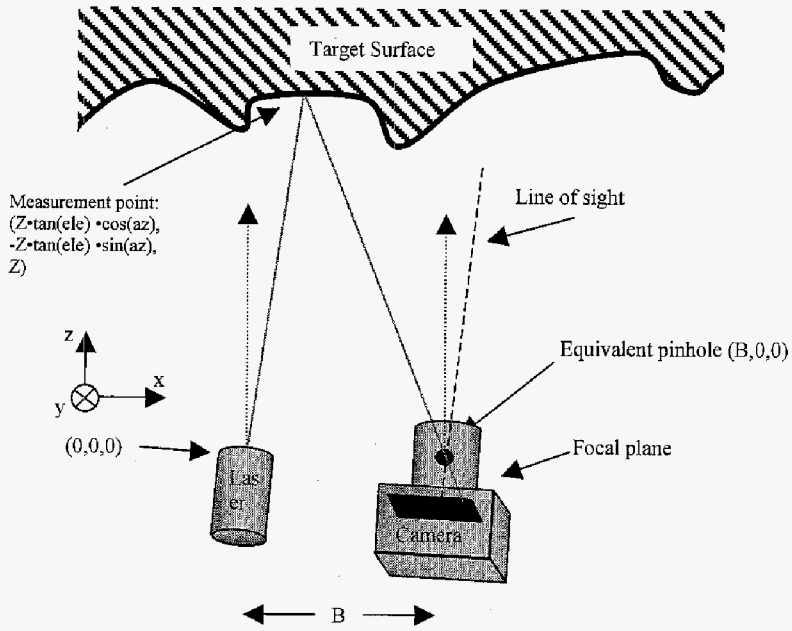

Figure 3 - Sketch of the non-idealized system

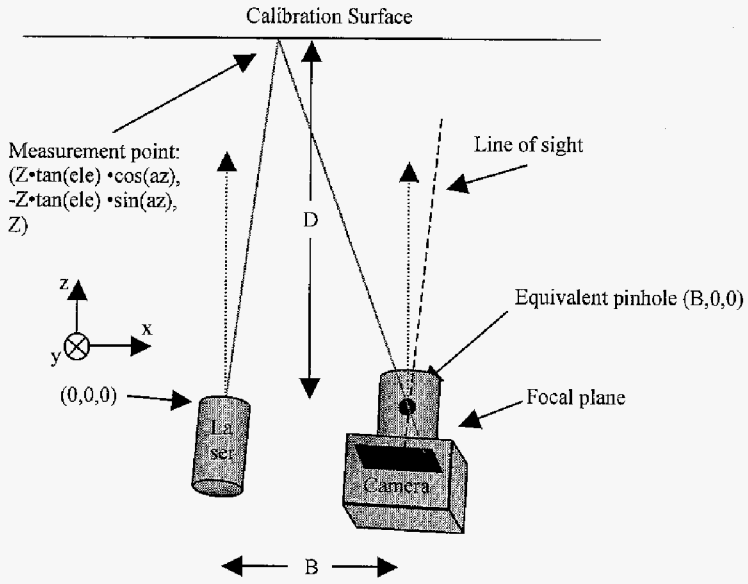

Figure 4 – The calibration setup

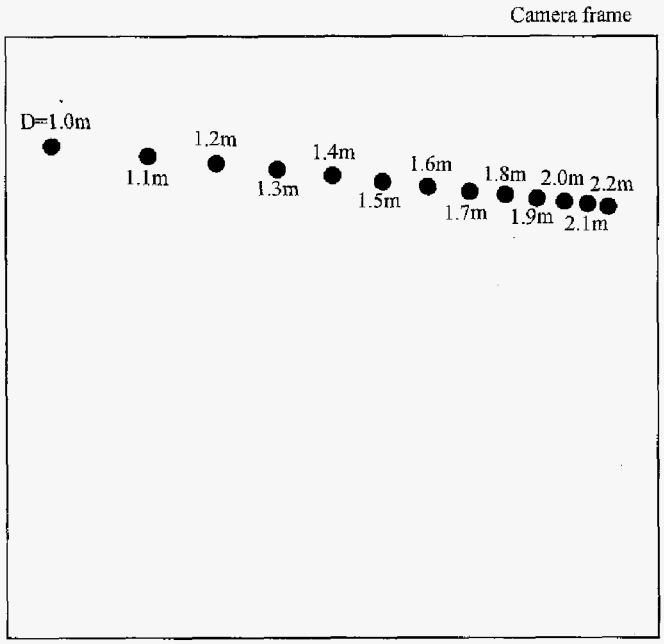

Figure 5 - An example of a set of centroids from the series of calibration images

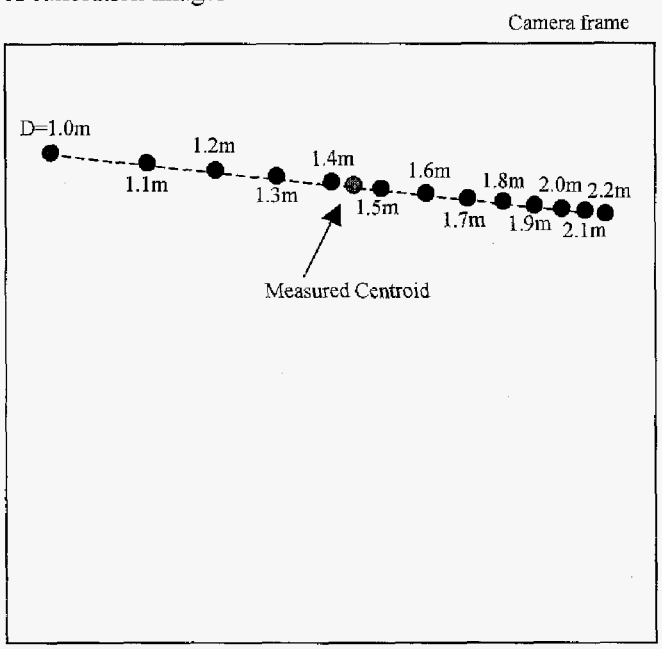

Figure 6 - Measurement of an unknown distance (red dot) utilizing a calibration curve

In practice, it is impractical to add 1000 individual lasers to illuminate the scene in front of the camera. Therefore, the laser beam is passed through a holographic diffraction grating. The interference pattern induced by the grating can split the outgoing laser beam into m x n beams. Each beam exits the optic at a fixed angle (α, γ) from the incident beam. The number and angle of a spot is determined by the properties of the grating and the laser wavelength [8].

If a row of spots all fall on the same image row, different Z distances will tend to displace the image locations (x's) closer together, possibly interfering with each other. It is also possible that some portions of the surface might shadow the camera from imaging some of the spots (spot dropout). In either case, the problem of identifying the spot and its associated angle is made difficult without taking some precautions to guard against aliasing or dropout.

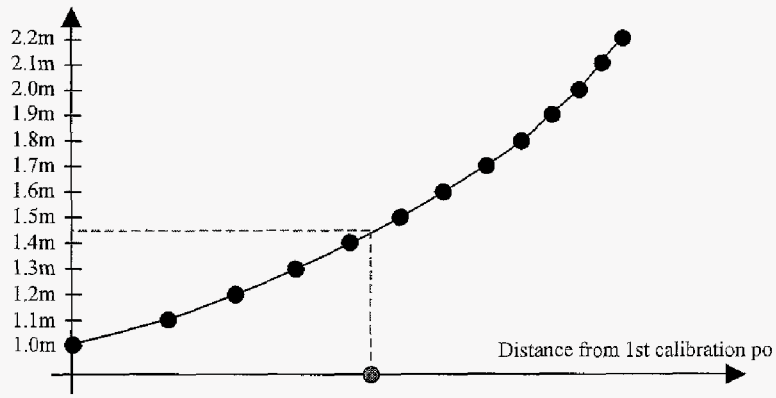

Figure 7 - The calibration curve. The red line represents the centroid from the unknown distance

The approach we've taken to uniquely identify each spot involves rotating the grid produced by the laser/grating system by an angle r around the direction of the laser beam. Also, the baseline, B and the geometry of the setup is constrained to limit the amount of dispersion (=movement on the focal plane), d, experienced by a spot over the operating ranges of the system. To insure that neighboring spots do not overlap, r and B are selected to provide each spot a unique area on the image plane. The area needs to be of sufficient width to provide reasonable range resolution and of sufficient height to insure that neighboring spots do not overlap. This is sketched in Figure 8.

Ideally, the rotation should be sufficient to insure that the divergence of a spot d should not interfere with its neighbors that have a separation of s (also in pixels).

$$r = \arcsin(s/d) \tag{3}$$

In the example in Figure 8 the grid is rotated so that the next spot in the grid is 4 scan lines displaced. With a rotation, spots on other rows in the grid are rotated onto a given spot's scan line. Even though off-row spots are separated further in angular space, spot disparity could still lead to spot aliasing. To eliminate this possibility, the horizontal angular separation between spots can be increased, or the baseline between the camera and the laser/grating can be shortened. Given a maximum allowed divergence d (in angle), and a minimum range (Z_{min}), the maximum baseline can be approximated to be:

$$B = Z_{min} * \tan(d) \tag{4}$$

Decreasing the dynamic range or increasing the grid spacing provides additional flexibility in designing the structured light system.

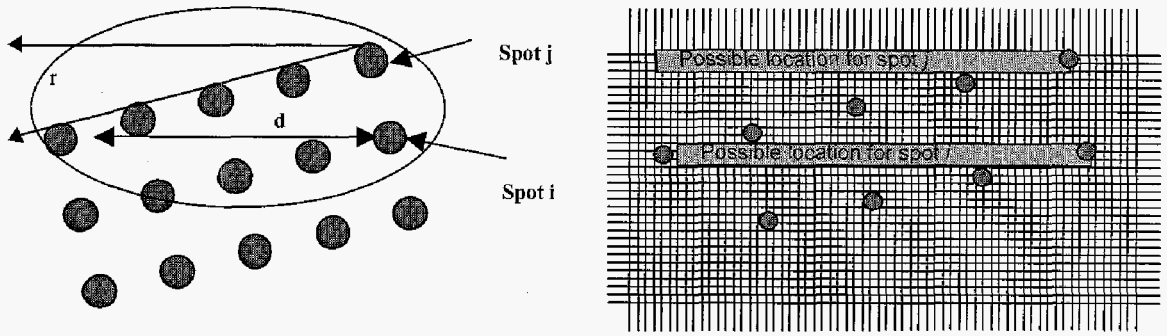

Figure 8 - The left sketch shows the grid rotated by an angle *r* while the right sketch depicts the motion induced on a spot by allowed ranges. The leftmost spot divergence is at the minimum range while the red dots show the spot at maximum distance.

For baselines less than B in Equation 4, a spot is constrained to move in an area that is unique to the point. Spots found in the area bounded by their maximum divergence and their location at Z_{min} will have valid ranges. Aliasing is not a problem as each spot's area is unique and spot loss simply means that the range map has one less entry.

3. EXPERIMENTS

An experiment was performed to evaluate the structured light method. The laser source for the experiment was a regular laser pointer. It was mounted behind a holographic grating that generated 31 x 31 beams in a rectangular spaced pattern. The spot generator was placed on an optical table. The baseline to the CCD camera was ~20 cm. The camera was a commercial SBIG CCD camera [9]. The camera utilized a 25-mm lens, resulting in a 14 degree FOV. The lens f/# was set at maximum (small aperture) so the camera could focus over a range of distances without having to refocus the lens. In front of the camera, lens was mounted a narrow bandpass filter of 632 nm to increase the signal to noise ratio. The setup is shown in Figure 9.

The setup was placed on a sliding rail in front of a white wall. Images were acquired at distances from 1 meter to 2.2 meters. An example of a calibration images acquired at 1.2 meters and 2.2 meters is shown at Figure 10.

The method was tested on a number of different test objects. In Figure 11 (top left) is shown an image of a stone. The image was acquired with no bandpass filter in front of the lens and no laser illumination (for illustration purposes). It was then reacquired utilizing laser illumination and a bandpass filter. The image is also shown in Figure 11 (top right). The described distance calibration was applied to the spots and the estimated 3D shape of the stone is also shown in Figure 11 (bottom row).

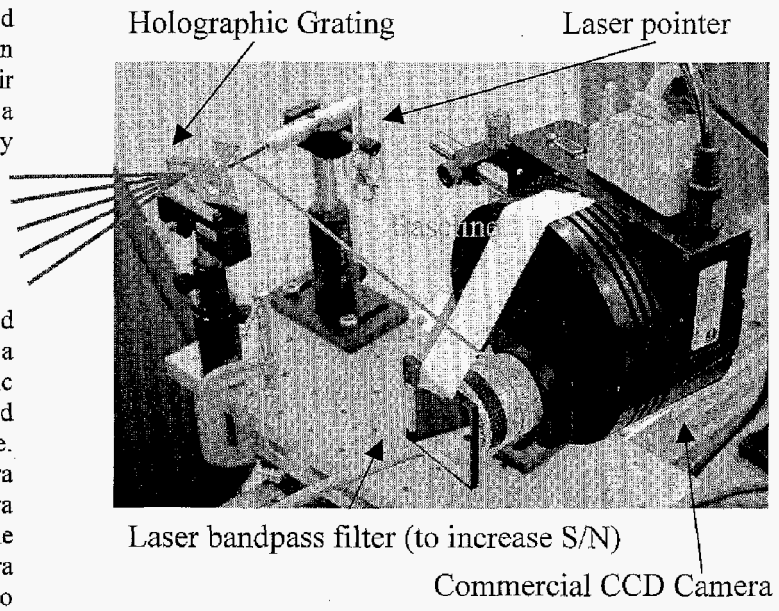

Figure 9 - Image of the experimental setup

4. CASE STUDY: STRUCTURED LIGHT UTILIZED FOR MARS LANDING SYSTEM

The flexibility of a structured light system allows it to be configured to support many different types of scenarios. In this section, we examine an application where the proposed system is utilized as a safe landing system.

Descent and landing are a part of many Mars missions. Numerous strategies can be employed to insure that the spacecraft touches down in an area free of hazards. Future Mars missions currently baseline a strategy of collecting range maps of the likely touchdown site and evaluating the site for hazards. The potential hazards in a landing site are many. Large rocks can topple the spacecraft, large slopes may interfere with vehicle/instrument deployment or communication, and craters might limit surface exploration.

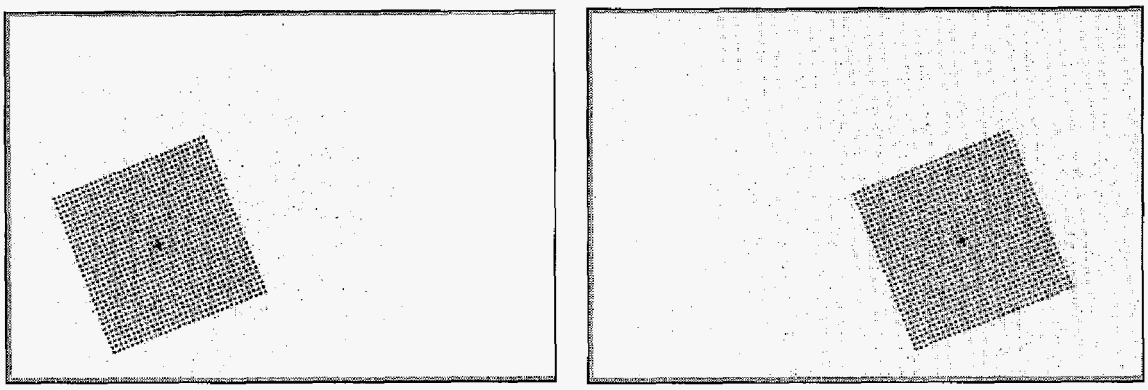

Figure 10 - Two calibration images acquired at 1.2 meters (left) and 2.2 meters (right). Shown in negative

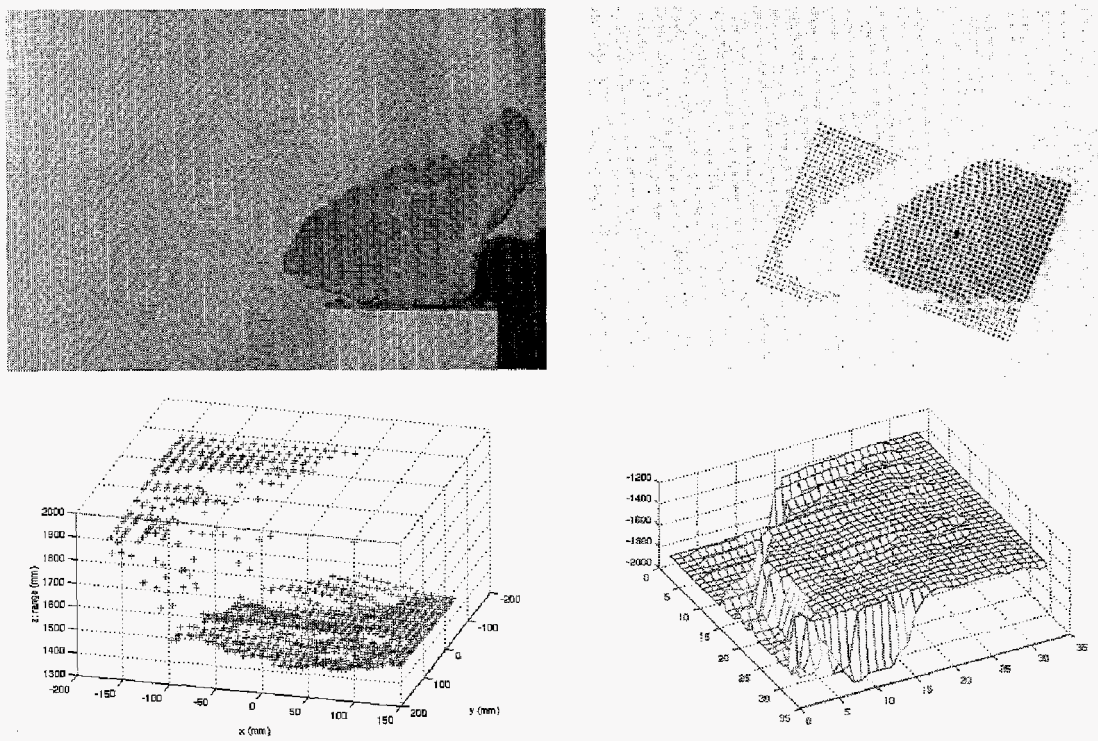

Figure 11 - Visual image of the test object (top-left), the test object with laser illumination and bandpass filter (top-right), and the estimated 3D structure (bottom row)

Large-scale hazards (> 30 meters from expected landing site) such as craters would be hard to detect with a structured light system. The distance from the surface limits the accuracy of the range maps, and it is unlikely that a good map could be generated at the altitude needed to detect such hazards. However, local hazards, such as rocks or slopes, could be detected by a properly configured structured light system.

A structured light system consisting of a large pulsed laser and a 30x30 grid of points over a 12x12 meter-landing site could effectively identify 0.75-meter radius rocks and 20-degree slopes (deemed hazardous by future missions). The proposed laser for this system is a 25 mJ, 532 nm (doubled) actively Q-switched system pulsing at 1 Hz. Each beam would have a 1.5 mrad divergence (full angle) that delivers 22 μJ per pulse per beam after the holographic grading. This assumes a holographic grating efficiency of 0.80. Such lasers and grating systems are available from commercial vendors and the laser has been flown in previous missions (Mola, Clementine). To minimize reflectance from the Martian surface, a bandpass filter would be employed to limit wavelengths outside of the laser band. Figure 12 shows a flight-qualified laser for the proposed structured light system.

The detector for this system would consist of a single camera with a 3.0-meter baseline from the structured light generator. Flight cameras are available having a pixel resolution of 1024x1024 [5]. A fast exposure time ~ one ms can be used to limit the amount of sun illuminated surface reflectance.

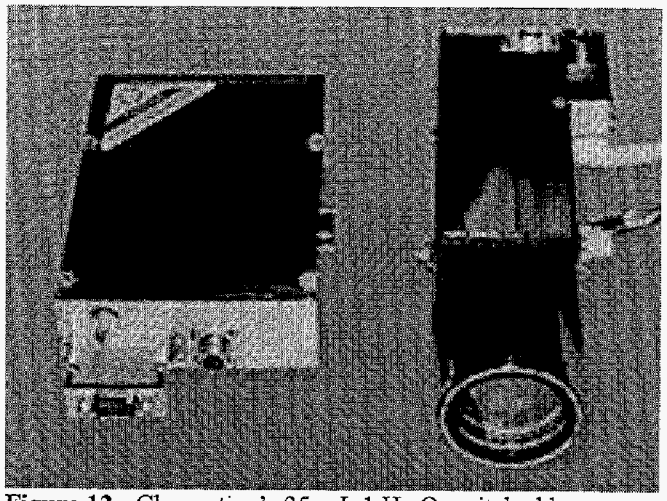

Figure 12 - Clementine's 25 mJ, 1 Hz Q-switched laser

In most Mars spacecraft landing scenarios, the spacecraft jettisons its chute during entry, and the expected touchdown site is evaluated for local hazards (rocks or slopes) while the spacecraft hovers. The structured light system will collect 2 images synchronized with the structured light generator at 37 and 35 meters (1.0 s between images). The first image will be used for hazard evaluation and the second image for verification. Hazard evaluation will be based on the range measurements obtained from the 30x30 sample points in the 1024x1024 image.

There are three substantial issues that need to be addressed for a structured light system to meet the requirements for local, hazard detection and avoidance:

1. Determining if such a system can be configured to provide sufficient range accuracy to separate hazards from acceptable terrain features.
2. Providing sufficient ground resolution to insure that hazards contained in the landing site are actually sampled.
3. Determining that the ground spots have enough energy to be detected accurately

The range accuracy for a structured light system can be calculated using the theory previously described. Since the structured light system employs spots and centroiding, typical pixel accuracy is likely to be higher than correlation based matching techniques—on the order of $1/10^{th}$ of a pixel, one sigma. Figure 13 shows the range resolution derived with this pixel accuracy. The horizontal line reflects the accuracy needed to separate 0.56 meter rocks that are acceptable to the mission from 0.75 meter rocks the a hazardous to the mission with the required detection/false positive rates (the most difficult hazard for this sensor).

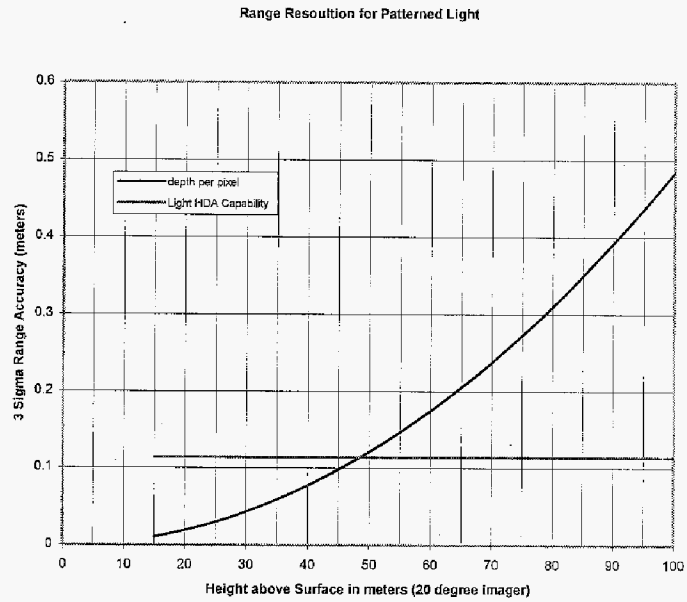

Figure 13 - Range resolution using a stereo system with $1/10^{th}$ pixel accuracy, a 20-degree camera with 1024x1024 pixels. The red line indicates the needed accuracy to separate out hazardous rocks from non-hazardous ones.

To insure that the spots provided by the structured light system are easily detected (with simple thresholding) by the camera system, we need to demonstrate that the energy from the spot is significantly greater than that from the neighboring surface. We assume the irradiance of the Martian surface is 3500 μWatts/cm^2 for the 500-550 nm band and a surface albedo of 0.3 [6]. Calculating the radiance we have: L = 3500 μWatts/cm^2 * 0.3 / π = 334 μWatts/cm^2 sr. The amount of energy irradiated from a surface area covered by a 1.7 cm^2 ground pixel area works out to 0.3 μJ with a 0.5 ms exposure.

With a 1.5 mrad spot, the incident energy on the area viewed by this pixel from the laser pulse is:

22 μJ/spot * 1/4 spot/pixel * 0.3 = 1.65 μJ

This works out to a ratio of greater than 5:1 pulse to sunlight ratio that should be sufficient to detect the spot with a simple thresholding technique [6].

5. SUMMARY

A rectangular grid of spots produced by a structured light system provides a low power, night-capable, range map sensor suitable for automated navigation operations. Aliasing and dropout problems can be eliminated by the prudent selection of baseline and by rotation of the grid with respect to the camera. We have shown that this simple system can produce reliable range maps that are efficiently

processed and accurate for near range operations. The system components are highly reliable and available for immediate flight needs.

REFERENCES

[1] C.C.Liebe, A.Abramovici, R.Bartman, R.Bunker, J.Chapsky, C.Chu, D.Clouse, J.Dillon, B.Hausmann, H.Hemmati, R.Kornfeld, C.Kwa, S.Mobasser, M.Newell, C.Padgett, W.T.Roberts, G.Spiers, Z.Warfield, M.Wright: Laser Radar for Spacecraft Guidance Applications, In proceedings of the 2003 IEEE Aerospace conference, Montana, March 2003, in print.

[2] R.Kornfeld, R.Bunker, G.Cucullu, J.Essmiller, F.Y.Hadaegh, C.C.Liebe, C.Padgett, E.Wong: New Millennium ST6 Autonomous Rendezvous Experiment (ARX), In the proceedings of the 2003 IEEE Aerospace conference, Montana, March 2003, in print.

[3] Brian D. Pollard, Gregory Sadowy, Delwyn Moller, and Ernesto Rodríguez: A Millimeter-Wave Phased Array Radar for Hazard Detection and Avoidance on Planetary Landers, In proceedings of the 2003 IEEE Aerospace conference, Montana, March 2003, in print.

[4] A. Johnson, Y. Cheng, L. Matthies: Machine Vision for Autonomous Small Body Navigation, 2000 IEEE Aerospace conference, March 2000.

[5] A.Eisenman, C.C.Liebe, M.W. Maimone, M.A.Schwochert, and R.G. Willson: Mars Exploration Rover Engineering Cameras, Proceedings of the SPIE #4540: Sensors, Systems, and Next Generation Satellites V, Toulouse France, September 2001.

[6] Matthies, L. and Balch, T. and Wilcox, B. Fast Optical Hazard Detection for Planetary Rover Using Multiple Spot Laser Triangulation. ICRA 1997.

[7] Rafael C. Gonzalez, Richard E. Woods: Digital Image Processing, Addison-Wesley Pub Co; 2nd edition (January 15, 2002) ISBN: 0201180758.

[8] Optical Principles and Technology for Engineers. James E. Stewart. Marcel Dekker, Inc.1996.

[9] Santa Barbara Instrument Corporation, Santa Barbara, US, URL: http://www.sbig.com/sbwhtmls/online.htm , cited 11/12/03.

ACKNOWLEDGEMENT

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology and was sponsored by the National Aeronautics and Space Administration. References herein to any specific commercial product, process or service by trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

BIOGRAPHIES

Dr. Carl Christian Liebe received the M.S.E.E. degree in 1991 and the Ph.D. degree in 1994 from the Department of Electro physics, Technical University of Denmark. Since 1997, he has been an employee at the Jet Propulsion Laboratory, California Institute of Technology. Currently, he is a senior member of the technical staff in the Laser Remote Sensing Group in the Interferometry and Large Optical Systems Section. His research interests include new technologies and applications for avionics sensors.

Dr. Curtis Padgett received his Ph.D. from the Computer Science and Engineering Department at the University of California at San Diego (UCSD) in 1998. He received his M.S. degree from the same department in 1992. He has been employed at the Jet Propulsion Laboratory, California Institute of Technology, Pasadena since 1993. He is a Senior Member of the Technical Staff in the Machine Vision group of the Mobility Section where he works on remote sensing applications for space systems. His research interests include algorithm optimization, machine vision, and artificial intelligence applied in classification and pattern recognition tasks.

Johnny Chang is a software engineer and a member of the technical staff in the Simulation and Verification Group at Jet Propulsion Laboratory, California Institute of Technology. At JPL, he is currently involved in developing high-performance distributed and parallel simulation software. He received his B.S in 2000 from U.C. Berkeley in Computer Science and Electrical Engineering, and his M.S. in 2002 in Computer Science from University of Illinois at Urbana-Champaign, where he focused on surface reconstruction research.