High-performance laser range scanner

John Hancock”, Eric Hoffman’, Ryan Sullivan[b], Darin Ingimarson[b], Dirk Langer”, Martial Hebert[a]

“The Robotics Institute, Carnegie Mellon Univ., Pittsburgh, PA 15213

[b]K^{2}T, Inc. One South Linden St., Duquesne, PA 15110

ABSTRACT

Laser scanners, or ladars, have been used for a number of years for mobile robot navigation. Although previous scanners were sufficient for low-speed navigation, they often did not have the range or angular resolution necessary for mapping at the long distances required by high-speed navigation. Many also did not provide an ample field of view.

In this paper we will present the development of state-of-the-art, high speed, high accuracy, laser range scanner technology This work has been a joint effort between CMU (project lead) and K^{2}T (scanning mechanism) in Pittsburgh and Zoller + Fröhlich (laser) in Wangen, Germany.

The scanner mechanism provides an unobstructed 360^o horizontal field of view, and a 30^o vertical field of view. Resolution of the scanner is variable with a maximum resolution of approximately 0.06 degrees per pixel in both azimuth and elevation. The laser is amplitude-modulated, continuous-wave with an ambiguity interval of 52 meters, a range resolution of 1.6 mm, and a maximum pixel rate of *500* kHz.

This paper will focus on the design and performance of the scanner mechanism and will discuss several potential applications for the technology. One application, obstacle detection for Automated Highway applications will be discussed in more detail. Example data will bc shown and current mechanism improvements from the CMU prototype will also be discussed.

Keywords: lidar, lascr radar, ladar, laser scanner, scanner calibration, obstacle detection

1. INTRODUCTION

The ability to measure surfaces and objects in **3-D** is becoming increasingly important for many fields such as autonomous vehicle navigation and obstacle detection, quarry mapping, landfill surveying, and hazardous environment surveying. The current state of the art for scanning mechanisms is unable to meet the demand of many of these applications. Typical scanners arc slow, unable to measure with an unobstructed view and inflexible in their ability to use different types of range sensors. We have created a system that overcomes these limitations and provides a system that will meet the existing demand for more advanced scanning mechanisms.

A laser radar operates by measuring the time, either directly or indirectly, between transmitting a laser beam and receiving its reflection off a target object. The range to the target is proportional to that time. The laser can also provide a measure of the power of the reflected beam. This is the reflectance or intensity. Our imaging laser scanner can thus provide two images, a range image and a reflectance image.

This paper focuses on some of the design and performance issues of the laser scanner. In particular, the typical laser radar problems of crosstalk and mixed pixels are discussed as are results from scanner precision tests. The paper considers potential applications, and discusses one of them, that of reflectance-based obstacle detection for automated highways, in detail. Preliminary results arc given.

2. SCANNER DESIGN

The scanner device consists of a two-axis mechanical scanning mechanism which collects three dimensional information by steering laser energy through a 360 degree azimuth by up to 70 degree elevation field of view. Individual scanners can bc manufactured with different vertical fields of view (the CMU scanner has a **30** degree vertical field of view). The sensor produces a stream of 3D range points and intensity values arranged in a helical scan pattern. By combining a high speed, highly accurate scanning mechanism with high speed single point laser range sensors, we have created a data collection device which can provide remarkably high resolution range images of the sensor's physical surroundings in short periods of time.

The scanning device is comprised of a gold-coated aluminum mirror, a yoke assembly which allows the mirror to pivot vertically, a spindle assembly which rotates the yoke assembly horizontally, electronic circuits to establish the position of the mirror in space. electronic circuits to collect and store range data, a mechanical housing used to mount a range detection sensor, and a commercially available range detection device (see Figure 1). The laser is located beneath the mirror assembly and points upwards. The scanner operates by continuously rotating the mirror in the two degrees of freedom. Data can be collected on the initial downward pivot of the mirror (downswing) and on the return sweep upwards (upswing). At each instant, the mirror deflects the laser beam at a known azimuth (determined by the rotation of the yoke assembly) and elevation (determined by the vertical position of the mirror). The laser travels outward and upon striking an object or reflective mass (such as dust or fog), some of the scattered laser energy returns to the detector along the same path. The range to the object is calculated by directly or indirectly measuring the time it takes for the laser signal to return.

The current maximum azimuth rotational speed of the scanner is 2400 rpm. Typical mirror nodding rates are 0.02 to 4.0 Hz. The frame rate is double the mirror nodding rate since a frame can be collected on each upswing and each downswing of the mirror. Depending upon the laser type used, data rates can reach 500,000 data points per second with an absolute accuracy of 1 to 2 cm with 2 to 8 mm resolution. The data can be formatted into either Cartesian or spherical coordinates in space as measured from the center of the scanning mechanism. Range can vary from 0.10m to 400m, depending on the type and power of the laser.

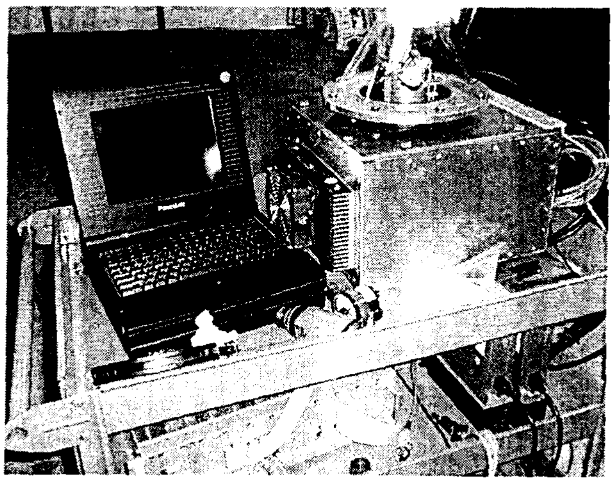

Figure 1. The laser scanner is on the right. A Sparc laptop, on the left, is used to interface with the scanner.

Currently Integrated Range Measurement Systems

The scanning mechanism has been tested with several range detection devices, including two models of a pulsed laser spot sensor developed by Riegl and an Amplitude Modulated Continuous Wave (AMCW) laser range finder developed by Zoller and Friihlich (Z+F) of Wangen, Germany.

The high accuracy pulsed laser spot sensor operates by emitting a concentrated laser energy pulse at each pixel. The pulse travels away from the sensor, strikes a surface, and returns. By measuring the time elapsed between the beginning of the pulse and the leading edge of the return pulse from the receiver, the distance to the object can be calculated by using the speed of light through air as a constant. The range is $d = 0.5ct$ where t is the time-of-flight of the laser pulse, and c is the speed of light (we multiply by a half since we are measuring the time for both the outgoing and return trip of the pulse). This laser device has an accuracy of +/- 2 cm and a resolution of 4 mm with a maximum data rate of 12,000 points per second. The laser also provides 8-bit intensity for each pixel.

The second pulsed laser system was created to see through dust clouds or fog. When operating in such environments, it is pos sible for some of the outgoing laser energy to be reflected by the dust or fog while some of the energy travels until it reaches a solid surface and returns to the detector. In this case, the receiver will see multiple return pulses. The laser measures the time elapsed between the beginning of the pulse and the *last* pulse returned to the receiver. Provided the dust or fog is not too dense, this can find the range to the first solid surface. A last pulse detection method was implemented in the laser system to create the high penetration system. This high-penetration version of the sensor has an accuracy of +/- 10 cm and a resolution of 10 cm and provides 8-bit reflectance data.

The current CMU version of the scanner uses the Z+F AMCW laser. The Z+F laser uses a dual frequency amplitude modulated signal and can operate at pixel rates as high as 500,000 per second. The receiver measures the phase difference between the original and returned laser signal at both modulation frequencies. The distance is proportional to the phase, up to an ambiguity at a 360 degree phase difference. 12-bit accuracy in the phase shift measurement of the high frequency modulation (23 MHz) provides a range resolution of 1.6 mm, while the phase shift in the low frequency modulation (2.875 MHz) is used to disambiguate between multiple cycles in the high frequency, and allows for an overall ambiguity interval of 52 meters. The laser has an overall accuracy of +/- 2 cm and can see up to 100 meters in the dark and provides a 16-bit reflectance value. Using the Z+F laser, the scanner can provide images of up to 6000 columns by several hundred rows over the 360 by 30 degree field of view with an angular resolution of 0.06° per pixel (see Figure 2). All of the experimental results described later were obtained using the Z+F laser

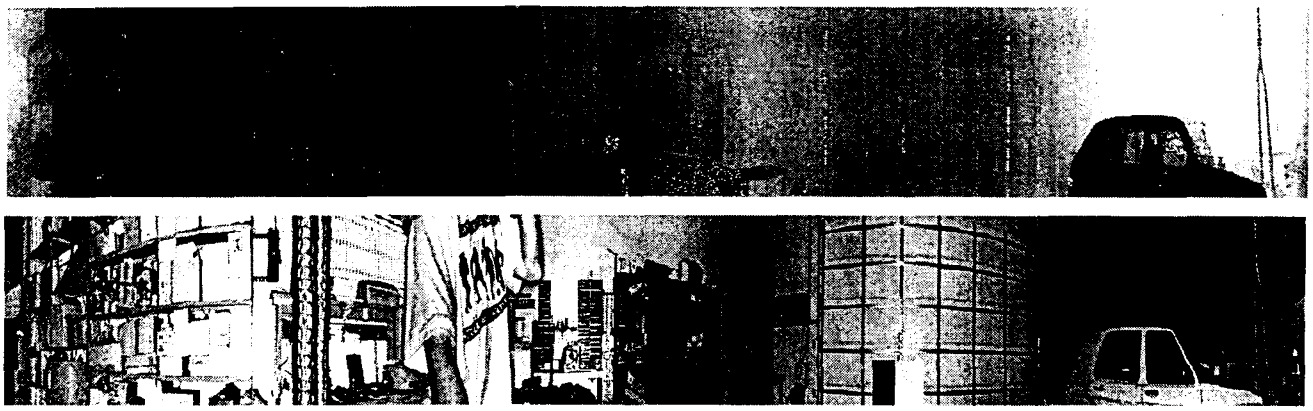

Figure 2. Range and reflectance images taken in our lab with the scanner using the Z+F laser. We have subsampled each of the images in the horizontal direction by a factor of 3 to make them fit on to the page. In the range image, top, darker pixels are closer and brighter pixels are farther. The range values have been scaled for printing *so* that a range of 14 meters or more appears as white. In the reflectance image, bottom, brighter pixels correspond to points of higher reflectance. The reflectance values have been scaled for better printing as well.

Important Characteristics

Unlike other traditional means for obtaining depth information such as stereo vision or depth from focus, the laser scanner works in the dark as well as in ambient light. Performance is enhanced with little or no ambient light. Some versions of the scanner can be programmed on-the-fly to return data from user selected regions of the scan volume. Scan speeds and fire rates can be varied to change spatial resolution and scan image sizes to better adapt to specific applications. A principal method of improving data quality is slowing the fire rate of the laser, allowing for a longer integration time and effectively improving the signal to noise ratio. The highest data rates are suitable for real time monitoring of dynamic environments for security, tele-operation or navigation. Unlike camera-based systems, **all** data returned is fully three dimensional so structural, CAD, or solid models can be built from the raw data with post processing. The laser data is clean and dense enough to use as is for virtual walk-through once transformed to Cartesian coordinates. This allows real-time modeling and VR applications to be satisfied with the option of storing the data for later post processing. Intensity information is also available and may potentially be used to recover some information about surface material types.

Environmental Issues/Covering the Scanner

A feature of the Z+F laser is its coaxial design. A coaxial laser has the benefit of having the transmitter and receiver fields of view overlap at all distances. This makes for more efficient laser detection at all distances. The coaxial design, however, presented a challenge when it came time to sealing the laser from the outside environment. The sensor ncedcd a transparent cover to keep dust and dirt from collecting on optical and mechanical surfaces.

The first cover designed for the sensor was a cylindrical glass with an anti-reflective optical coating. The optical coating was designed to reduce laser reflections by 98%. We discovered, however, that the glass caused problems near the center of the vertical field-of-view. Because of the coaxial design, light reflected from the glass at normal incidences was directed back at the receiver. Although these reflections do not pose a problem for the high penetration pulsed laser system, the reflections did cause problems with the Z+F continuous wave system. At near-normal incidences (-5 to 5 degrees), the laser light reflected from the inside of the glass cover overpowered the signal returned from secondary surfaces resulting in poor range values and bright reflectance values towards the center of our laser images. We surmised that although the glass resulted in laser reflection at other incidences, these reflections did not present a problem since the reflected light was not directed back at the receiver aperture.

Our second cover was designed to avoid near-normal incidence problems. We opted for an uncoated, molded optical-grade plexiglass (type P-84) design in the shape of a truncated cone. The slope of the cone was chosen to avoid normal incidences over the entire vertical field-of-view of the laser. Since our vertical field of view could be adjusted to provide laser elevation from $+15^{\circ}$ to -25°, we chose the slope of the cone to be 30°. This shape avoided normal laser incidences over the entire field-of-view, with points above the horizon (the least important areas for obstacle detection for terrain-navigating autonomous robots) having incidences closest to normal with a maximum incidence angle of 75° (15° from normal).

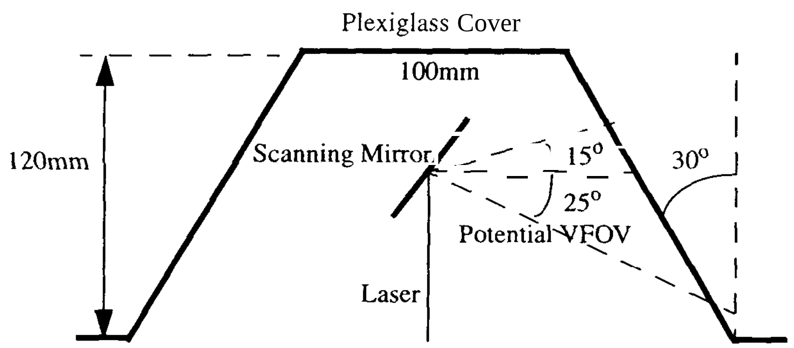

Figure 3. The plexiglass cover was designed to protect optical and mechanical surfaces from dust without producing poor data from normal incidence reflections.

Unfortunately, this design also has a problem with poor data caused by laser reflections off of the plexiglass cover. In this case, the poor data occurs at the bottom of the image where there are some imperfections caused by the plastic molding process. Making a mechanical modification to the housing so that we can lower the cover further should fix the problem.

3. PERFORMANCE

From a cursory visual inspection, the new scanner clearly provides better images than past scanners. However, it is instructive to reexamine some of the problems from past laser radar scanners and see how the new scanner fares. The first three subsections on mixed pixels, crosstalk, and range and reflectance precision are particular to the Z+F laser. The last subsection on angular precision is particular to the scanner mechanism itself.

Mixed Pixels

Previous AMCW laser systems had significant problems with mixed pixels. Mixed pixels are those that receive reflected energy from two surfaces separated by a large distance. Mixed pixels can result in reported ranges that are on neither surface, but somewhere between the two ranges, or even worse, either behind or in front of both surfaces. This is an inherent problem with AMCW laser radars and cannot be completely eliminated [4]. However, the Z+F laser reduces the problem by having a significantly smaller spot size than previous laser systems. The ERIM, for example, had a laser divergence, or instantaneous field-of-view (IFOV) of 0.5 degrees[2]. The Z+F, on the other hand, has a beam divergence of only 0.5 mrads (0.03°). Since the beam divergence is smaller than the scanner resolution, a pixel is less likely to fall on the edge of an object (see Figure 4) and thus cause an erroneous range value. When mixed pixels do occur, however, they are generally isolated pixels and may be removed by a median filter. The higher resolutions provided by the scanner make it easier to ignore or filter away isolated pixels without throwing out large amounts of the data.

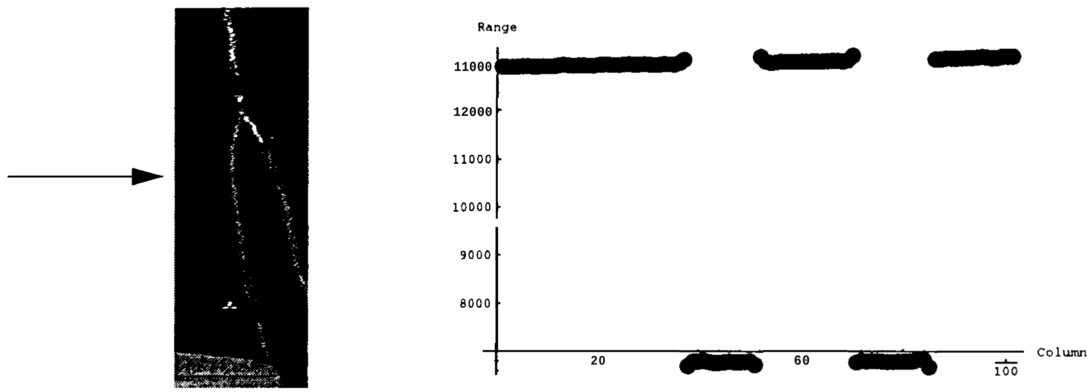

Figure 4. On the left is a portion of a reflectance image of an exhaust port. On the right is a graph of the corresponding range values of the indicated row in the reflectance image. Note that the range values immediately transition between the exhaust port and the far wall behind it. In this case, there are no mixed pixels resulting in erroneous ranges at the object edge.

Range/Reflectance Crosstalk

Ideally, a laser should measure the same median range to two targets which are equally distant but have different reflectance. In practice, however, this is not always the case. Crosstalk is a phenomenon in which reflectance or intensity values affect the measured range. To some degree, intensity always affects the range values. Lower intensity values decrease the signal-to-noise ratio and therefore increase the variance in measured range, although this should not affect the average range. However, as described in Hebert and Krotkov, the implementation of the electronics can also play a factor in crosstalk[4]. This is especially noticeable at edges between darker and lighter surfaces. In the range and reflectance images below (see Figure 5), the effect of crosstalk can clearly be seen. Abnormally high range values are reported for quite a few pixels at edge boundaries. Some of these may be explained by multipath reflections off of specular surfaces such as the bad data on the metal struts on the garage door. However, the pixels on the edge between the black and white posterboard in the foreground cannot be explained by specular multipath reflections, but are instead attributed to a combination of mixed pixel and crosstalk problems. Crosstalk can probably not be eliminated completely from any AMCW laser system.

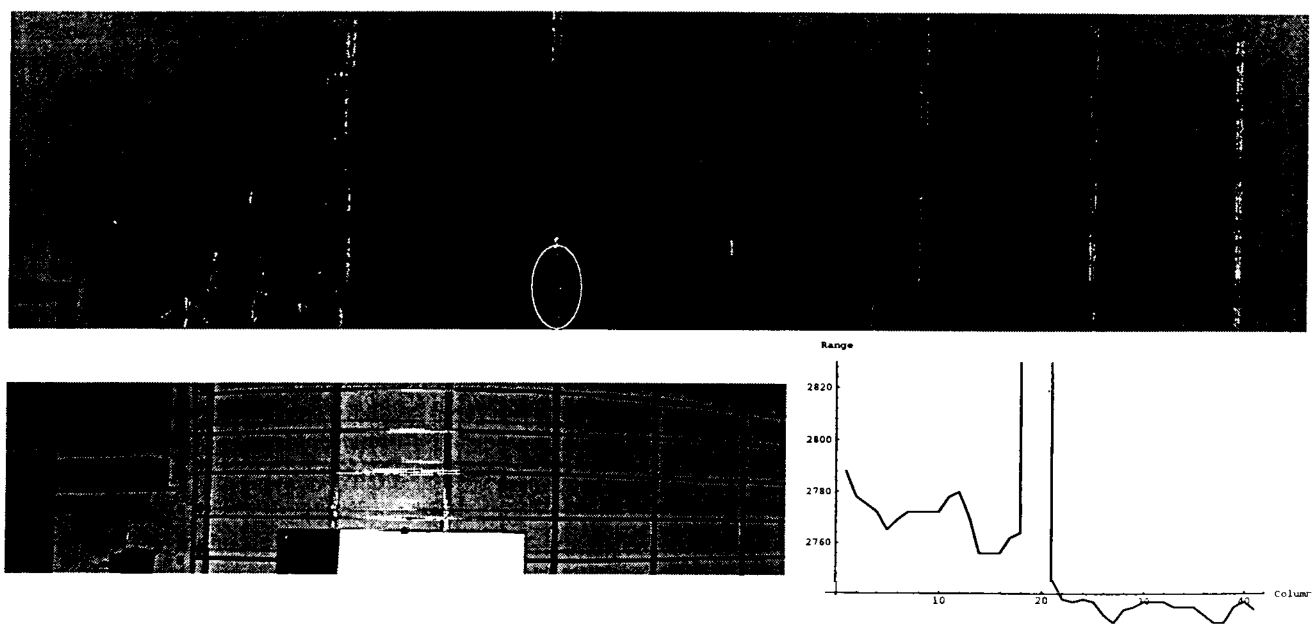

Figure 5. The range image, top, shows erroneous range values (the white pixels) corresponding to a number of edges in

the reflectance image, bottom. In particular, notice the white pixels circled in the range image at the boundary between the black and white posterboard in the foreground of the reflectance image (bottom left). The posterboard pieces are placed at the same distance from the scanner. On the bottom right is a graph of the range values for a section of one row of the image centered on the black/white posterboard boundary.

Range and Reflectance Precision

To test the precision of the range and reflectance channels, data was taken with two different targets at increasing distances from the scanner. The scanning motors were turned off for this experiment. At each location a series of 900,000 measurements was taken from which the mean and standard deviation of the range and reflectance was recorded (see Table 1).

Table 1: Precision of Range and Reflectance

White Target				Black Target			
Mean Range (mm)	**Mean Reflectance**	**Range Std. Dev. (mm)**	**Reflectance Std. Dev.**	**Mean Range (mm)**	**Mean Reflectance**	**Range Std. Dev. (mm)**	**Reflectance Std. Dev.**
3603	31518	3.48	167.2	3622	3471	10.76	46.3
8100	13741	5.34	71.1	8087	1165	19.2	21.4
13783	5470	8.61	45.6	13737	537		14.5
23445	1958	14.78	27.7	23506	144	501.0	10.0
32398	1099	19.76	20.72	34715	122	3121	8.2
42103	600	28.9	16.3	47481	16.4	1115	6.0

Although the black target is of much lower reflectance (approximately only 10% as reflective), it provides similar range readings within several centimeters (within the possible errors of target placement) until beyond 20 meters. At that point, reflectance is low enough that the range values begin to degrade, as illustrated by the rapidly climbing standard deviation of the range measurement. Mean range values for the black target are within one standard deviation of the white target range value except for the last measurement.

Angular Precision

To measure the angular precision of the scanner, a set of *6* images was taken with a stationary scanner and scene, all on the downswing of the mirror. Six feature points were selected in the first image and matched in the remaining images. The reported azimuth and elevation angles were compared for each point and averaged over the *6* images. **A** more sophisticated and complete approach for measuring angular precision is described by Hebert and Krotkov[4]. Unfortunately, we could not use this approach since the scanner provides no frame synchronization and thus image row can not be used to match features or as a measure of angular precision.

For each feature in all images, the azimuth varied by less than +/- 2 pixels (at maximum scanner resolution) or 0.12^{o} from the average value. The elevation angle varied by less than +/-0.6^{o} from the average. Reported elevation angles for the same feature can vary even more if the downswing elevations are compared to upswing elevation angles. **This** imprecision in the elevation is caused by backlash, or slop, on the vertical pivot gear for the mirror. During a single-direction scan this results in fairly small errors, but when the gear changes directions to drive the mirror in the reverse direction, the difference in swing directions results in an error of approximately 3^{o} between calculated elevation. Factory calibration could reduce this effect for the user, but has not been performed on this scanner. For this reason, we currently only use data from the downswing of the mirror. Current versions of the scanner address this problem and improve angular resolution but have not yet been tested.

We have delayed angular accuracy tests until the elevation gear mechanism is replaced.

4. APPLICATIONS

A scanner which can run at variable scan rates and provide very high resolution images (up to 6000 columns by 500 rows) in seconds could be used for a number of tasks. A few potential applications are briefly discussed below.

Building Modeling

Full interior and exterior structure and content information can be quickly generated to model or map buildings. Low resolution scans can give room shapes and obstacles in less than one second. High resolution scans can yield highly detailed models in which people, light fixtures, and even door knobs can be imaged and located (by human eyes) with ease and with no post processing. Additional processing using either custom or existing software packages can yield surface or solid models which can be rendered for visualization purposes. Floor plans can be generated as well.

Surveying

Contour maps can be generated at a rate **far** exceeding manual survey techniques. The scanner can be applied to quarry mapping, site mapping, stock pile estimation, cut and fill estimation, blast preparation, and a host of other standard civil engineering applications in which data must be collected.

Security

Video cameras and near IR devices can be fooled with camouflage. Even thermal imaging devices can be deceived in some cases. However, it is not possible to mask a physical volume moving through space. By scanning an area in 3D at high speed and high resolution, it is next to impossible to move across the scanned area without being detected. Fairly straightforward algorithms can be used to determine what has moved between scanned images and the relative volume of the mass which is moving can be calculated. This has great applications in situations where a high value asset must be guarded or a site must be monitored for intruders.

Tele-operation **or** Dynamic Situation Monitoring

High speed range images can be supplied in real time to an equipment operator to take the place of, or augment, 2D cameras for tele-operation applications. High resolution point clouds can show the world to the operator in 3D with fairly low computational overhead. This is most useful in environments in which video is not suitable such as night operations where lights are not possible, in smoke, and in fog (one version of the Franklin Scanner can penetrate smoke and fog).

5. LASER REFLECTANCE FOR OBSTACLE DETECTION

One potential application for laser scanner technology is obstacle detection for automated highway applications. Highway obstacle detection is a challenging problem. Highways present an unknown and dynamic environment with real-time constraints. In addition, the high speeds of travel force a system to detect objects at long ranges. Systems such as the Automated Highway System (AHS) which demand high levels of safety are not feasible unless these critical problems are addressed[1].

Traditional methods of laser processing for off-road obstacle detection involve transforming the range image into an elevation map. Since the elevation map is generally sparse and noisy, smoothing is generally necessary. Then a terrain-typing algorithm is run to classify areas of the elevation map **as** navigable or unnavigable. Besides being a fairly complicated process, this method has another drawback in that elevation maps may not represent obstacles well, since they cannot represent vertical planes because of the discrete grid size. On the other hand, laser reflectance ought to provide **us** with a more direct means of finding obstacles or vertical surfaces. At the long lookahead distances and grazing angles typical of high-speed travel, horizontal surfaces should provide very weak (or nonexistent) laser returns. Vertical surfaces, however, should result in stronger signals[3].

Preliminary tests with the Franklin scanner have shown that laser reflectance can detect medium-sized obstacles at distances over 50 meters using simple processing. Detection of a pallet at 54 meters was possible by constructing a reflectance model from the reflectance values and then finding the residual between the predicted reflectance based on the model and the actual reflectance at each pixel (see Figure 6). The model used was a quadratic as a function of image row. In this example, masking of non-road image pixels was performed by hand. However, automatic masking should be possible using **a** road follower such as RALPH, or by performing road-edge detection in the reflectance image.

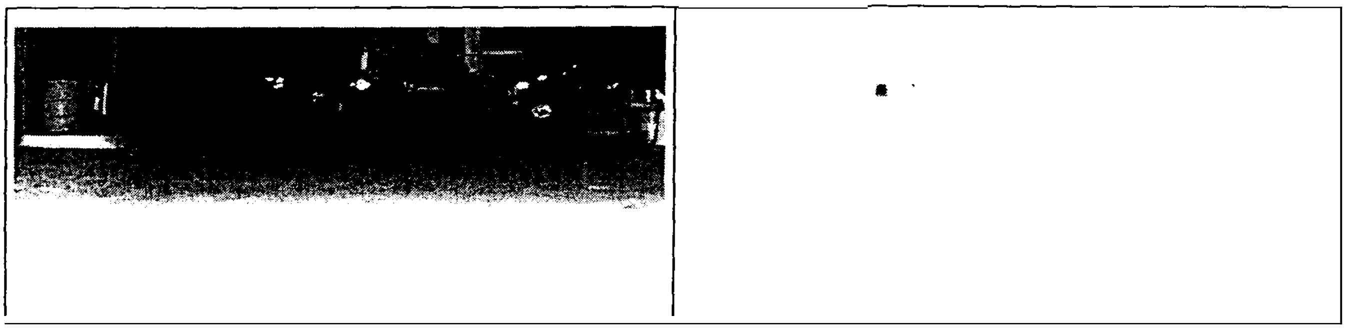

Figure 6. A laser reflectance image (left) and a processed version of the image (right). Pixels not belonging to the ground or target object were first masked out. Then a quadratic model as a function of image row was fit to the reflectance data. The processed image shows the residual between the modeled reflectance value and the actual reflectance, where white pixels have a small residual error, and dark pixels have larger errors. The target object, a pallet at 54 meters away, shows up clearly in the processed image as a dark rectangle.

Although this result is encouraging since it used only implicit models built from the data, it will be necessary to model laser reflectance explicitly to improve detection accuracy and eliminate false positives. More robust results should be possible with explicit modeling of the laser sensor. Analysis of laser reflectance is complicated by its dependence on other factors such as range and material type of the sensed object. Assuming the diffuse component of the laser reflection is Lambertian, we can model the laser reflectance signal with the following relation:

$$P_{return} \propto \frac{\rho \cos\theta}{z^2}$$

where P_{return} is the power in the returning laser pulse, ρ is the actual surface reflectance ($0 \le \rho \le 1$), θ is the angle of incidence of the beam with the surface and z is the depth[6]. Ideally, given a flat road model or other estimated road model, we could invert the previous equation and predict the surface reflectance for an image point:

$$\rho(x, y) = \frac{z^2 P(x, y)}{\cos\theta}$$

Note that for vertical obstacles, z is smaller than predicted and $\cos\theta$ is much larger than predicted, so pixels corresponding to vertical obstacles will be given a very high ρ value. Any points with unusually high ρ values should correspond to vertical obstacles. Currently, we are working on verifying this model. Although this reflectance model is a good representation at the microlevel, it is not clear how strongly the returned power depends on the angle of incidence at the macrolevel. Microfacet structure and orientation may play a larger role in determining the effective angle of incidence than the macro-orientation of the surface. Some empirical evidence seems to suggest this is true.

Assuming a flat ground plane, we can relate the angle of incidence to the range:

$$y = \alpha + ky$$

$$z = \frac{h}{\sin y} = \frac{h}{\cos\theta}$$

where α is the inclination of the sensor, γ is the inclination of an individual laser ray, k is a row factor, and h is the height of the sensor. Combining our equations, we see:

$$P_{return} \propto \frac{\rho \cos\theta}{z^2} = \frac{ph}{z^3}$$

This would indicate that to obtain a consistent (normalized) reflectance for pixels on a flat road, obtained at a constant height, we would need to multiply the intensity by the range cubed. However, as the graphs below illustrate (*see* Figure 7). multiplying the reflectance by the range squared seems to provide consistent values. This indicates that the macro angle of incidence does not have as large an effect as was initially predicted. Further experimentation will be necessary to characterize the effect of surface orientation on the reflectance. The oscillations in the second graph are as yet unexplained.

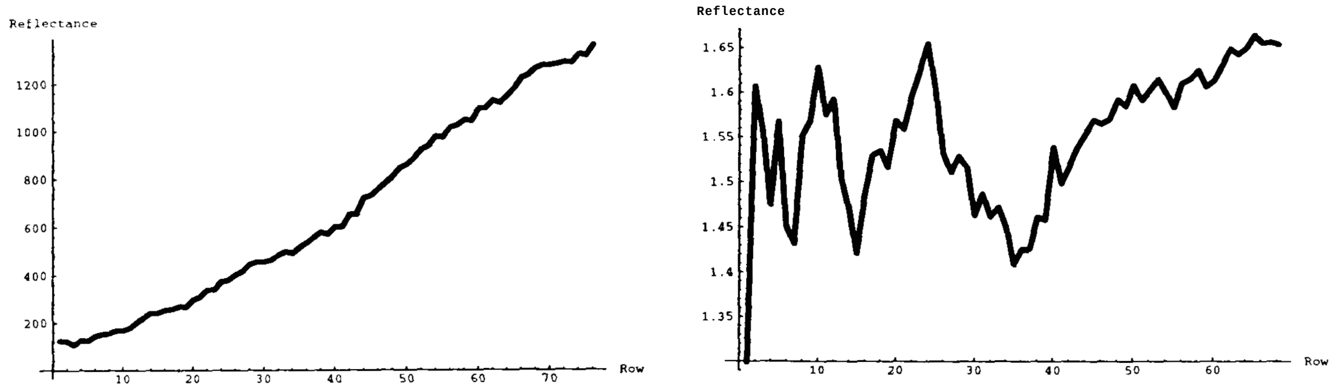

Figure 7. The left graph shows the average reflectance of road pixels vs. image row for an image similar to Figure 6. Lower image rows are at the top of the image and are farther away so they have lower reflectance. The right graph shows the average normalized reflectance, i.e. the reflectance values once they have been multiplied by the square of the corresponding range values. Notice that normalizing by the range has resulted in reasonably consistent values (within +/- 10%) across all image rows except the first row.

With better sensor reflectance models, we hope to demonstrate that laser reflectance can provide a viable obstacle detection method for onroad navigation or other mild terrain settings. A significant benefit of using laser reflectance as opposed to range is that an optical sensor that provided only reflectance would be much simpler and cheaper to build than current laser scanners. Such a device might be affordable enough to place them on commercial vehicles.

6. FUTURE WORK

There are a variety of areas which demand future work. First, some refinements need to be made to the scanner design. The scanner needs to be faster to cope with the real-time obstacle detection requirements **for** autonomous vehicles. We expect that we can achieve rotational speeds of 5000 rpm. Another mechanical problem is the elevation gear backlash or slop. Gear backlash causes problems in elevation angle accuracy, especially at the transition point between the upswing and downswing of the mirror where the gear changes directions. Changing the mechanism to use a gear with a finer pitch should reduce or eliminate the problem.

Although the scanner provides visually convincing **3D** data, we have not been able to compare the **3D** data with ground truth information. More evaluation and calibration of the scanner mechanism should be performed. Although angular and range precision tests have been performed, we have not performed any angular accuracy tests. In particular, elevation accuracy needs to be tested once the gear backlash problem is solved. It **is** also unknown how much vibration affects the angular accuracy of the system. Although some range accuracy tests were performed by Z+F, additional tests should be performed in the context of the scanner.

Finally, many experiments remain to be done to better characterize the laser reflectance data.

ACKNOWLEDGEMENTS

This research was partly sponsored by: the Ben Franklin Technology Center of Southwestern Pennsylvania, the USDOT under Cooperative Agreement Number DTFH61-94-X-00001 as part of the National Automated Highway System Consortium, and a DARPA/TACOM grant, “Technology Enhancements for UGVs”, DAAE07-96-C-X075.

REFERENCES

[1] Bishop, J. R. “Intelligent Travel: The Automated Highway System.” ***Proceedings of the International*** *Conference on Intelligent Autonomous* ***Systems*** *(IAS-3)*, 1993.

[2] Eberle, K. Reflectance Processing. "Erim Staff Report." *Range and Reflectance Processing Workshop Proceedings,* Warren, MI, December 1987.

[3] Hancock, J. *High-speed Obstacle Detection for Automated Highway Applications.* Thesis proposal. Carnegie Mellon Technical Report, CMU-RI-TR-97-17, May 1997.

[4] Hebert, M. and E. Krotkov. "3D Measurements from Imaging Laser Radars: How Good Are They?" *Image and Vision Computing,* vol. 10, no. 3, April 1992.

[5] Kweon, I., R. Hoffman, and E. Krotkov. *Experimental Characterization of the Perceptron Laser Rangefinder.* Carnegie Mellon Technical Report, CMU-RI-TR-91-1, Jan. 1991.

[6] Nitzan, D., A. Brain, and R. Duda. "The Measurement and Use of Registered Reflectance and Range Data in Scene Analysis." *IEEE Proceedings*, Vol. 65, No. 2, February 1977.