

Experimental Characterization of Commercial Flash Ladar Devices

Dean Anderson, Herman Herman, and Alonzo Kelly
The Robotics Institute – School of Computer Science,
Carnegie Mellon University, Pittsburgh, PA, USA
dranders@cmu.edu, herman@rec.ri.cmu.edu, and alonzo@ri.cmu.edu

Abstract

Flash ladar is a new class of range imaging sensors. Unlike traditional ladar devices that scan a collimated laser beam over the scene, flash ladar illuminates the entire scene with diffuse laser light. Recently, several companies have begun offering demonstration flash ladar units commercially. In this work, we seek to characterize the performance of two such devices, examining the effects of target range, reflectance and angle of incidence, as well as mixed pixel effects.

Keywords: flash ladar, sensor characterization, Canesta, CSEM, 3D imaging

1 Introduction

For some time now, laser radar (ladar) has served as the mainstay sensor for mobile robot applications. Ladar's accuracy, range and robustness make it ideal for gathering terrain data at relatively high speed. However, almost all ladar devices use a single collimated laser beam. In order to acquire full 3D data, this single beam must be scanned, usually in raster-fashion. This scanning process increases significantly the amount of time it takes to capture a full frame of range data. Furthermore, scanning from a moving platform introduces motion artifacts, which requires precise timing and localization to correct.

To remove these difficulties, a new-generation of scannerless ladar devices are currently being developed by several different groups[1][2][3][4]. These devices are commonly referred to as "flash ladar" as they illuminate the entire scene with diffuse laser light at once. Range measurements in these devices are computed using a 2D sensor array, usually either a modified CCD or custom CMOS array. In current flash ladar devices, this allows frame rates greater than 30 fps.

Recently, several companies have begun selling flash ladar development kits[5][6][7]. In this paper, we seek to characterize two such devices: the SwissRanger2, developed by the Centre Suisse d'Electronique et de Microtechnique SA (CSEM) and the CanestaVision DP205 by Canesta, Inc. Specifically, we will examine the effects range, reflectivity and angle of incidence to the target have on range measurements.

2 Device Descriptions

2.1 Principle of Operation

Both the CanestaVision and SwissRanger2 are amplitude-modulated continuous-wave (AMCW) measurement devices. This class of device measures the phase shift between the transmitted beam and its reflection off a target. Measuring this phase requires four samples of the returning waveform, separated by 90° . However, phase is only uniquely determined up to half a wave-length, which is referred to as the *ambiguity distance*.

2.2 SwissRanger2

The CSEM SwissRanger2 is built around a custom CMOS/CCD sensor array[5][8]. In particular, the device uses a "2-tap" pixel architecture which is capable of capturing and storing two samples at each pixel-site. Over each integration period, the sensor is able to sample the received light twice, 180° apart. A second integration period is required to fully disambiguate the phase, but requiring only two integration periods significantly decreases the effects of motion blur.

Illumination comes from a bank of 48 infrared LEDs. These are modulated at 20 Mhz, which corresponds to an ambiguity distance of 7.5 m. Further details about this sensor can be found in table 1.

2.3 CanestaVision

The CanestaVision DP205 uses a CMOS device for its sensing array[6]. Each pixel consists of two modulated



Figure 1: The CSEM SwissRanger2 flash lidar device



Figure 2: The CanestaVision DP205 flash lidar device

gate elements, where the modulation signal is synchronized with the light source. The modulation between the first gate and second gate is phase shifted by 180° , so when one is high, the other is low.

This gate modulation increases or decreases the gate's response to received light, causing the output of each gate to be a mixture of both the intensity of the received light and the gate modulation. This mixture causes the differential voltage output of both gates to be directly proportional to the phase difference. A second sample, shifted by 90° is required to fully disambiguate the phase.

This device uses a small laser diode array through a translucent diffuser for illumination. Furthermore, the modulation frequency is user selectable to 13, 26, 52, or 104 Mhz, corresponding to ambiguity distances of 11.5, 5.8, 2.9 and 1.4 m respectively. This can allow the sensor to operate in a dual-frequency mode: using a low frequency to increase the ambiguity distance while using a high frequency to increase resolution. This would require two full range images to be capturing, increasing susceptibility to motion effects. Further details on this device can be found in table 1.

3 Experimental Setup

The objective of our experiment was to characterize the performance of both sensors. In particular, we were interested in how the range, reflectance and angle of incidence of a target affect range accuracy and precision. We also sought to examine common artifacts of lidar systems, such as mixed pixels.

For both cameras, an integration time of 50 ms was used. Because the SwissRanger2's modulation frequency is fixed at 20 Mhz, we set the CanestaVision's modulation frequency at 26 Mhz so their maximum ranges would be comparable. Additional processing such as amplitude filters and common mode rejection was disabled in both cameras.

3.1 Performance Model

Variance in range measurements is a valuable metric for predicting the ability of a sensor to discriminate between objects at different ranges. In general, the variance of range measurements of an AMCW range finder has an inverse dependence on the irradiance of the illumination received[9]. This can be written as

$$\sigma_r \propto \frac{\lambda R^2}{\rho \cos \alpha} \quad (1)$$

with σ_r the standard deviation of range, R the range, λ the wavelength, ρ the reflectance of the target, and α the angle of incidence.

We designed our experiments to examine each of the three variables in equation 1, following descriptions of previous lidar characterization found in [10].

3.2 Range, Reflectivity and Angle of Incidence

We performed all experiments indoors, in a large high-bay. Ambient illumination consisted primarily of lamps, but sunlight diffused through translucent ceiling panels was also present.

Both devices were mounted on a stationary tripod about 120cm above the floor. For targets, we used three colored, foam board cards measuring 50cm by 76cm: one black, one gray, and one white, to exhibit different reflectances. These targets were also mounted on a stand about 120cm above the floor. That stand allowed us to adjust the height and rotate the target about a horizontal axis (for changing the angle of incidence).

To measure the effect of range and reflectance, the target stand was placed along the sensors line of sight at distances between 0.5m and 8.0m, at intervals of 0.5m. At each distance, at least 200 range images were captured of each of the three different targets. This setup can be seen in figure 3.

Table 1: Specifications of the CanestaVision DP205 and the CSEM SwissRanger2

	CanestaVision	SwissRanger2
Pixels	64(H) x 6 4(V)	160(H) x 124(V)
Field of View	55°	40°
Maximum Range	~ 10 m	7.5 m
Max. Frame Rate	50 fps	30 fps
Illumination	Laser Diode Array	LED Array
Power and Wavelength	1 W @785nm	0.8 W @870 nm
Modulation Frequency	13, 26, 52, or 104 Mhz	20 Mhz
Sensor Type	CMOS	CMOS / BCCD
Dimensions	12.5cm x 6cm x 6.3cm	13cm x 4cm x 3cm
Power	10W (2A @5V)	20W (1.5A @12V)
Interface	USB1.1	USB2.0



Figure 3: Experimental setup for characterizing the effects of range, reflectance, and angle of incidence have on range measurements of a flash lidar device.

To measure the effects of angle of incidence, tests were conducted with the white cards at distances of 1.5m, 3.0m, 4.5m and 6.0m. At each of these distances, the target stand was rotated in 10° increments, from vertical to fully horizontal. Again, at least 200 range images were captured of each scene.

3.3 Mixed Pixels

In lidar devices, it is generally assumed that the laser spot size is infinitesimal. However, in reality this spot has some finite size and it is fairly common for this spot to lie on two objects simultaneously. For the two sensors tested, the projected pixel size at 5m is on the order of 3cm by 3cm. This leads to a range artifact known as “mixed pixels”, where the range measurement is some mixture of the range to the object in the foreground and that of the background.

In this test, we use two foam board cards (both white). One measuring 10cm by 76 cm in the foreground, and the other measuring 50cm by 76cm in the background. The background card was placed at a range of 5.0m from the camera, and experiments were performed with the foreground target at ranges of both 3.0 and 4.0m. (see figure 4)




Figure 4: Experimental setup to determine the effect of mixed pixels on flash lidar devices

From the foreground object’s initial position, it was translated horizontally in increments of 2mm to a maximum displacement of 6.0cm. This caused the foreground object to fully traverse several pixels, partially occluding them in the process.


Due to device availability, this test was only performed using the SwissRanger2

4 Results

To characterize the performance, we computed the mean value and standard deviation of each pixel across the entire 200 image set for each scene. We then isolated a window of pixels that lay on the target in each scene, and computed the average of the mean and standard deviation for these pixels. For the SwissRanger2, this window was sized 10 by 10; for the lower resolution CanestaVision, a window size of 5 by 5 was used. Typical range images are shown in figure 5.



(a)



(b)


Figure 5: Typical range image of the white test target from (a) the CanestaVision and (b) the SwissRanger2 at a range of 4.0m

4.1 Range Accuracy

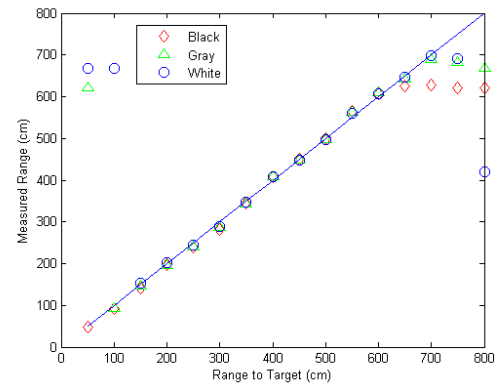
A plot of the average measured range for both cameras can be found in figure 6. In general, range accuracy is within 5cm for both devices and independent of range.

In the CanestaVision device it is apparent that the performance of the sensor against a low reflectivity target is especially poor. Range accuracy for the black target is an order of magnitude worse than that of the higher reflectivity targets, and no range values are returned after a mere 3.0m. Increasing the integration time would significantly improve the performance of the device with low reflectivity targets, but at the expense of frame rate and possibly motion blur.

The primary artifact of the SwissRanger2 occurred at extremely short ranges. At ranges below a meter, the received reflected energy actually saturated the sensor, making it impossible to measure range. This could be addressed by shortening the integration time. Otherwise, accuracy is excellent up to the ambiguity distance.



(a)



(b)

Figure 6: Average measured range as a function of range for (a) the CanestaVision and (b) the SwissRanger 2 flash ladars.

4.2 Range Variance

4.2.1 Variance with Range and Reflectivity


A plot of variance as a function of the range and the target can be found in figure 7. Again, the data indicates that the CanestaVision camera has significant problems with low reflectivity targets.

Otherwise, the data agrees with the performance model discussed above. The standard deviation increases with R^2 for the same target with a constant of proportionality dependent on the reflectivity of the target.


4.2.2 Variance with Angle of Incidence

The variance of range measurements as a function of angle of incidence is shown in figure 8. Again, the data agrees with the performance model. For a constant range, the variance increases with the secant of the angle of incidence as predicted by equation 1.


Performance for glancing angles of incidence was generally much poorer for both sensors. At ranges above 1.5m and angle of incidences greater than 20 degrees, both devices were unable to measure ranges to a black target (hence, only the white target was fully tested.)




(a)



(a)



(b)



(b)

Figure 7: Standard deviation of range measurements of (a) the CanestaVision and (b) the SwissRanger2 as a function of range and target reflectance. The fit lines are of the form $\sigma = kR^2$. The angle of incidence is fixed at 0.

Figure 8: Standard deviation of range measurements of (a) the CanestaVision and (b) the SwissRanger2 as a function of range and angle of incidence. The fit lines are of the form $\sigma = k \sec(\phi)$. All data is for the white target at the range indicated.

4.3 Mixed Pixels

For analyzing mixed pixels, the mean and standard deviation of the range value for each pixel was again computed for each set of 200 images was computed. A column of five pixels on the background target was then chosen. As the foreground target was translated, this column of pixels gradually became fully occluded.

Given the projected area of a pixel on the target, we then back-computed what fraction of the pixel was occluded as a function of the translation of the foreground target. The average range value as a function of the fraction of pixel area occluded by the foreground target is shown in figure 9.

5 Conclusions

We have examined the performance of two commercially available flash lidar devices. Although they may never fully replace scanned lidar, as flash lidar devices are further developed and refined they should prove ideal for use in short-range mobile robot applications. Our future work will evaluate the suitability of such sensors for obstacle detection and avoidance. We are




Figure 9: Average range measurements of mixed pixels as a function of the fraction of pixel area occluded by the foreground object.

currently working to mount two SwissRanger2 flash lidar on a small outdoor mobile robot for such purposes.

However, significant barriers for the use of flash lidar still exist. The experiments described in this work were performed entirely indoors; when taken outdoors, both the SwissRanger2 and CanestaVision

are unable to measure range well in the presence of sun-light while moving. The CanestaVision sensor possesses background-subtraction processing that makes it somewhat less sensitive to ambient light, but this feature is only useful in static scenes.

Interference from ambient lighting could be mitigating by narrower band filters on the lens, although this may require additional optics. Eye-safety concerns currently limit the feasible output power of any flash lidar transmitter. The ideal transmitter is infinitely small; any distance between the light source and the receiver can cause phase error, limiting the feasible size of illumination arrays.

Furthermore, the relatively short modulation wavelengths used and the short maximum range also limit the utility of these sensors on outdoor applications. This can be addressed by using lower modulation frequencies, especially when coupled with dual-frequency operation.

As neither of the two devices examined were not “flash lidar” in a pure sense (they require two full integrations to measure range), motion blur may still be an issue. This in particular limits their use in high-speed mobile robots. We plan on further studying motion effects once the sensors are mounted on a mobile robot.

With improved robustness to ambient illumination, increased maximum range, and increased spatial resolution, these low cost flash lidars are well positioned to overcome the limitations of scanned lidar and ultimately become the mainstay sensor of the mobile robot community.

References

- [1] P. Garcia, J. P. Anthes, J. T. Pierce, P. V. Dressendorfer, I. K. Evans, B. D. Bradley, J. T. Sackos, and M. M. LeCavalier, “Nonscanned lidar imaging and applications,” in *Proceedings of SPIE*, vol. 1936, pp. 11–22, October 1993.
- [2] B. L. Stann, A. Abou-Auf, S. Frankel, M. Giza, W. Potter, W. C. Ruff, P. H. Shen, D. R. Simon, M. R. Stead, Z. G. Sztankay, and L. F. Lester, “Research progress on scannerless lidar systems using a laser diode transmitter and fm/cw radar principles,” in *Proceedings. SPIE Laser Radar Technology and Applications VI*, vol. 4377, pp. 12–22, September 2001.
- [3] M. J. Halmos, M. D. Jack, J. F. Asbrock, C. Anderson, S. L. Bailey, G. Chapman, E. Gordon, P. E. Herning, M. H. Kalisher, L. F. Klaras, K. Kosai, V. Liquori, M. Pines, V. Randall, R. Reeder, J. P. Rosbeck, S. Sen, P. A. Trotta, P. Wetzel, A. T. Hunter, J. E. Jensen, T. J. DeLyon, C. W. Trussell, J. A. Hutchinson, Raymond, and S. Balcerak, “3d flash lidar at raytheon,” in *Proceedings of SPIE* (G. W. Kamerman, ed.), vol. 4377 of *Laser Radar Technology and Applications*, September 2001.
- [4] R. Stettner, H. Bailey, and R. D. Richmond, “Eye-safe laser radar 3d imaging,” in *Proceedings of SPIE*, vol. 4377 of *Laser Radar Technology and Applications*, September 2001.
- [5] T. Oggier, M. Lehmann, R. Kaufmann, M. Schweizer, M. Richter, P. Metzler, G. Lang, F. Lustenberger, and N. Blanc, “An all-solid-state optical range camera for 3d real-time imaging with sub-centimeter depth resolution (swissranger),” in *SPIE conference on optical system design*, (St. Etienne), September 2003.
- [6] S. B. Gokturk, H. Yalcin, and C. Bamji, “A time-of-flight depth sensor system description, issues and solutions.” URL: <http://www.canesta.com/technicalpapers.htm>.
- [7] “3DV Systems Company Website.” Online, July 2005. <http://www.3dvsystems.com/>.
- [8] R. Lange and P. Seitz, “Solid-state time-of-flight range camera,” *IEEE Journal of Quantum Electronics*, vol. 37, pp. 390–397, March 2001.
- [9] M. Hebert and E. Krotkov, “3-d measurements from imaging laser radars: How good are they?,” *International Journal of Image and Vision Computing*, vol. 10, pp. 170–178, April 1992.
- [10] I. S. Kweon, R. Hoffman, and E. Krotkov, “Experimental characterization of the perceptron laser rangehder,” Tech. Rep. CMU-RI-TR-91-1, The Robotics Institute, Carnegie Mellon University, January 1991.