Evaluation of Canesta's Range Sensor Technology for Urban Search and Rescue and Robot Navigation

Jeff Craighead
Institute for Safety, Security and Rescue Technology
University of South Florida
4202 E. Fowler Ave.
Tampa, FL, USA
craighea@cse.usf.edu

Brian Day
Institute for Safety, Security and Rescue Technology
University of South Florida
4202 E. Fowler Ave.
Tampa, FL, USA
bjday@cse.usf.edu

Robin Murphy
Institute for Safety, Security and Rescue Technology
University of South Florida
4202 E. Fowler Ave.
Tampa, FL, USA
murphy@cse.usf.edu

***Abstract*— This paper evaluates a new time of flight range sensor developed by Canesta, Inc for use with urban search and rescue (USAR) robots for mapping and navigation. Two field tests were conducted using the sensor; one to evaluate the sensor for robot navigation by a human operator, and the second to evaluate the sensor for use with autonomous navigation behaviors. The results of the tests show that the Canesta sensor provides enough information to navigate a robot without the use of other sensors, both for human operators and autonomous behaviors.**

I Introduction

Range sensing in man-packable robots has been limited by the size of available sensors and their cost. Man-packable robots are systems which fit into one or two backpacks that can be carried by the operators. The Inuktun micro-VGTV series, where the entire system (operator control unit, batteries, robot, etc.) weighs about 27kg, is an example of a man-packable robot. The iRobot Packbot series borders on the upper extreme of a man-packable robot.

The Canesta EP200 series of range cameras, shown in Figure 1, which obtains depth using the time of flight of reflected infrared laser pulses, appears to be an alternative to the Sick laser scanner. The camera provides a 64-by-64 pixel range image and corresponding grayscale image in real time. Example output can be seen in Figure 2. The field of view varies by model between 30, 55, or 80 degrees. The nominal size of the sensor housing is 12.7cm wide x 5.08cm tall x 5.08cm deep. It weighs approximately 544.3 grams. The small size and low weight allow this sensor to be used on a variety of small robots that are used for USAR.

Fig. 1. The Canesta EP200 Series range camera

The maximum unambiguous range (accurately resolvable distance) is 11.5m. The minimum unambiguous range was tested to be 25cm. This limit is a result of the pulse frequency. This is an adjustable setting which also affects the depth resolution of the sensor. The unambiguous range and depth resolution are inversely related; the maximum depth resolution, tested to be 5mm on average, is achieved with the minimum unambiguous range, 1.44m. This tradeoff between resolution and unambiguous range is present in all TOF sensors that use phase shift for range detection. In [1] they note: "An important issue for TOF sensors is the aliasing effect

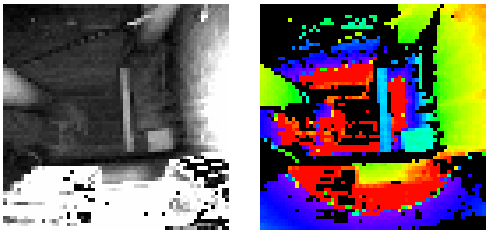

Fig. 2. Example output from the EP208 sensor. Left: Active Brightness Image (grayscale). Right: Range Image.

arising from the periodicity of the modulated signal whereby the distances to objects differing in phase by 360 degrees of phase shift are not distinguishable."

This paper discusses a preliminary field evaluation of the Canesta range sensor for robot navigation in confined spaces such as those encountered in urban search and rescue, military, and utility maintenance environments. A second test was conducted in a rubble cluttered outdoor area. The tests described in the following sections were conducted using the Canesta sensor mounted on an Inuktun Extreme VGTV. Section II identifies some recent work on the use of 3D range finders on robots. Section III describes the two test environments and the data collected. Section IV summarizes the results and concludes with a discussion of problems due to bright sunlight and non-reflective materials. Section V discusses noise reduction algorithms as a potential solution to the issues discussed in Section IV. Section VI concludes that the sensor has great potential, but more work is needed to address the problems with daylight operation.

II Related Work

Robots must use some type of range sensor to detect distances to objects in their operating environment. In the past, inaccurate sonar sensors have been used extensively for obstacle detection. Optical range finders, which are far more accurate than sonar and have become relatively cheap, are now standard equipment on many robots. Some range systems, such as computationally intensive stereo vision can generate a high resolution three-dimensional map of the environment. Other systems such as the SICK line of planar range finders produce only a two-dimensional scan but require very little computational power to identify the distance to objects. The goal of many robotics researchers is to produce a system that requires little processing power and is capable of creating a 3D environment map.

Surmann, et al. in [2] demonstrate the use of a 2D planar range finder rotating on the horizontal axis to generate a 3D environment map. The system is capable of mapping a 150 degree by 90 degree area in 4 seconds with a point cloud size of 22500 points. The system will image the same area in 12 seconds and generate a point cloud with 115000 points. While this is useful for generating a very detailed map of the environment, a 4-12 second refresh rate is simply not fast enough for a highly mobile robot to avoid obstacles in the environment. Assuming a useful range of 4 meters the robot would be limited to traveling at 1 meter per second or less. In addition the system is bulky which limits its use to robots that are larger than those typically used for urban search and rescue.

Davidson, et al. in [3] demonstrate the use of a single wide angle camera 3D SLAM system. The paper states that the system is capable of processing 30 frames per seconds when running on a desktop computer. While the wide angle camera system itself is compact and could probably be replaced with a robots built in camera the computational requirements are beyond what is available. Typically robots used for USAR have little or no onboard processing capabilities which necessitates a 3D sensor system with low computational complexity.

Hoey, et al. demonstrate the use of Canesta sensors for collision avoidance on an electronic wheel chair in [4]. The system described uses two Canesta range sensors to fill an occupancy grid in front of the robot. When the wheel chair is guided too close to an object, the chair stops and audibly warns the user and suggests a direction to turn. The high speed and resolution of the Canesta sensors allow the chair to detect thin objects that some planar laser systems would miss. From the photos in the paper it appears that this system was only used indoors.

For a review of current off-the-shelf 3D range technologies available, see Blais' article in [5].

III Approach

Two field tests were conducted to evaluate the Canesta sensors. The first test was conducted in Richmond, Missouri during an urban search and rescue training event. The second test was conducted at the University of South Florida's outdoor robot test bed. The sensor was mounted on an Inuktun Extreme VGTV using zip-ties and industrial hook-and-loop fasteners, Figure 3. The sensor was connected to a laptop running Windows XP and the Canesta EPToolkit software for the first test. Custom navigation software built using the Canesta API was used for the second test.

III-A Indoor Confined Space

The indoor field test was intended to validate the usefulness of the sensor in USAR conditions. The test site was a crawl space located in the basement of a collapsed building. This area was strewn with rubble,

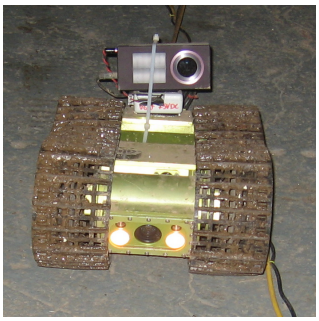

Fig. 3. Canesta Sensor on an Extreme VGTV with hook-and-loop fasteners and zip-ties.

steel pipes, hanging wires and air ducts. See Figure 4 for an image of the crawl space where the test was conducted. The environment was completely free of light. Figure 5 is a picture of the basement taken without a flash, the only light available is coming from the robot's headlights. To conduct the tests, the robots were placed in maintenance crawl spaces and run down the length of the space, turned around, and then run back to the entrance. This type of test was repeated with various frame rate and frequency settings in order to determine which settings provided the best trade off between frame rate and object visibility.

Fig. 4. A crawl space in the collapsed building. Photo taken with a Canon Powershot camera.

Fig. 5. Operating area in Richmond, MO. Notice the only light source is the robot's headlights. Photo taken with a Canon Powershot camera.

III-B Outdoor Test Area

The navigation test was conducted at the University of South Florida's robot test bed. This site is composed of sand and gravel with a pile of concrete rubble and sewer pipes. This test was divided into two phases, an obstacle avoidance task and a corridor following task. The obstacle avoidance task was set up to test the sensor's ability to work in bright, indirect sunlight. Various obstacles were placed in the testbed. The robot with the sensor was aimed directly at the first object and the test program was activated. For the corridor following task the robot with the sensor was placed near a sewer pipe angled into the wall of the pipe before the test program was activated.

IV Results

The results of the navigation tests provide insight into the sensitivity of the sensor in various lighting conditions, notably the presence of noise when used in sunlight. The sensor performs best in complete darkness, producing clear images even at long distances. When exposed to indirect sunlight or shade, moderate image corruption is present, however the result is still usable as noted in Section IV-B. When exposed to direct sunlight the image becomes unusable.

The accuracy and quality of the range image produced by the Canesta depends on several internal settings. The shutter speed and CMR (external light compensation) settings play the most significant role in producing a smooth image. In addition the intensity of the laser can be adjusted, however adjusting the laser produces less noticeable results. Generally the laser is used at full power and the shutter speed adjusted to compensate for over or under saturation.

At first glance, the image produced by the Canesta seems to provide relatively low resolution when compared to some state-of-the-art stereo vision or laser based range devices. However the size and power requirements, as well as the use of mechanical pan/tilt mounts on some systems, make them impractical for use on a small robot. When compared to a typical stereo vision based system, the Canesta provides a low resolution image, but does so without the need for well defined lighting conditions or a power hungry computer. When compared to a SICK laser based system the lack of resolution is not so apparent.

The Canesta provides a fixed 64 pixels of horizontal and vertical resolution, and variable fields of view depending on then lens, which provide 30, 55, or 80 degrees FOV. A SICK laser has a fixed planar field of view, 180 degrees, and variable resolutions of 180, 360 or 720 pixels. These resolutions provided by the SICK

correspond to 1 degree, 0.5 degrees and 0.25 degrees per pixel. The Canesta comparatively provides 1.25 degrees, 0.86 degrees and 0.47 degrees per pixel. Thus it can be seen that the absolute spatial resolution of the Canesta sensor falls within that of the SICK laser. In addition it provides the vertical resolution without the need for a pan/tilt mechanism.

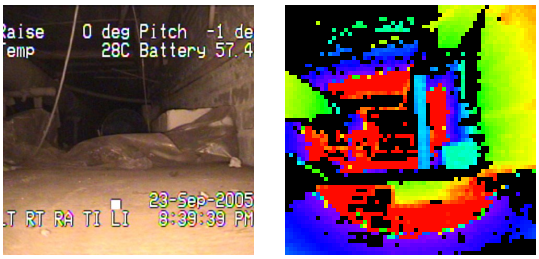

Fig. 6. Crawlspace at the Richmond test site. Left: Video from robot camera. Right: Range image from EP208

IV-A Indoor Navigation

In Figure 6, the walls, pipes, and wires are clearly visible in the depth image while they are less visible in the video image from the robot camera. The robot is visible in the bottom portion of the image, with the track clearly identifiable on the left and right. This extra information allowed the operator to better understand the layout of the environment while driving the robot through the crawl space. Also visible in the image is the aliasing effect caused by the unambiguous range limit.

Notice that the bottom half of the robot is a different color, blue, than the rest of the robot, red. Blue is typically mapped to the farthest unambiguous distance while red is mapped to the closest. A blue area is visible again next to a red area in the center of the image, past the vertical green pipe. This is the limit of the unambiguous range and shows an inherent problem with the Canesta sensor in that there is no method to automatically identify the false depth readings. This is not as much of a problem for a human operator, who will be able to identify the spatial relationships and segment out the false readings, but for an automated routine this could be disastrous, causing the robot to avoid an obstacle that isn't actually there.

IV-B Outdoor Navigation

Compare the two images in Figure 7. Notice that the right image has more random noise in the detected objects. This is noise due to the non-reflectance of certain materials, in this case wood, and interference from sunlight. This is an example of a particularly bad

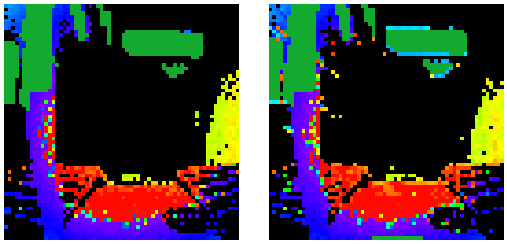

Fig. 7. Example of noise in the depth image. Left: A noise free frame. Right: A frame with random noise.

frame; generally there are enough less noisy frames for an autonomous navigation system to work correctly.

The object avoidance test program was able to successfully navigate the robot around several obstacles in its path. The course was constructed from plywood sheets that are approximately 45cm tall x 45cm wide arranged as shown in the figure. The wood obstacle was a stack of wood approximately 30cm wide x 30cm deep x 60cm tall. The robot was initially placed directly facing the wood obstacle. It was able to detect and move around this obstacle and continue on the path. The robot encountered outcroppings in the course and successfully avoided these as well. The application for this test based the obstacle detection on a 22 pixel wide x 3 pixel high area in the center of the depth image. The goal of the application was to avoid hitting anything in front of the robot while constantly moving forward. The 22 x 3 window size was determined experimentally to provide the best results.

The corridor following test program successfully navigated the robot into and through a 1.52m long concrete pipe. The application used for this test used a 1 pixel wide strip on either side of the depth image to identify the corridor walls. This application attempted to keep the robot centered in the corridor while moving forward.

IV-C Sensitivity to Lighting

Figure 7 is an example of a particularly bad frame produced by the sensor in indirect sunlight. This frame shows that there are several types of image corruption that occur when the sensor is used in sunlight. The solid green patches are caused by over saturation of the CCD.

Notice the green, nearly square blob at the top of Figure 7's images. This corresponds to a window on a building that is far outside of the sensors capable range. However, because of the reflected sunlight the window is visible, yet the depth reading is not valid. We will refer to this type of corruption as ghost images. In some cases when the depth image was obtained in direct sunlight the entire CCD is over saturated and the image has no useful

information.

The red pixels that appear near the left side of Figure 7's right image are an example of the second type of image corruption. This noise is caused by stray sunlight entering the lens and corrupting the depth reading at that pixel. This type of noise can also be seen on objects that are outside of the illumination range of the laser. Through our various runs we found that in all of the 10000 images collected, some form of noise was present. This noise has the potential to cause autonomous behaviors to operate in an unexpected manner when it occurs in large portions of an image over consecutive frames. A noise removal strategy that solves this problem is discussed in Section V.

It should be noted that an image in which an object is only partially saturated may be useful. In Figure 8 a paper soft drink cup was placed 30.5cm from the 80 degree FOV sensor. The laser was set to full power, the modulation frequency was set at 104MHz, and CMR was set to 30. The over saturated portions of the image are clearly visible as a circle on the table and an oval on the cup. The remaining portions of the cup appear correctly in the image and the depth values for those portions of the cup are correct, the values range between 27cm and 32cm. Notice in the active brightness image that the pixels next to over saturated pixels suffer from bleed over, which is a common problem with CCDs. Bleed over is caused by excess electrons, generated by an over saturated pixel, leaking to nearby pixels on a CCD. These pixels, which show up as dark gray pixels surrounding the over saturated region in the grayscale image, are also treated as invalid in the range image, which eliminates the possibility of an incorrect reading.

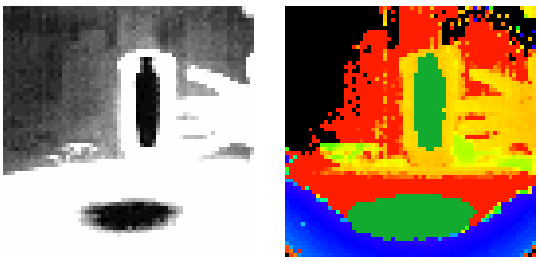

Fig. 8. Image of a paper soft drink cup on a desk. The cup was 30.5cm from the sensor. Notice the two over saturated areas in the image. One is a circle on the desk, the other is an oval on the middle of the cup.

IV-D Sensitivity to Material

Figure 9 is an example depth image, taken in our lab, of a 0.91m x 0.46m, black Hardigg case with a standard US letter size piece of white paper taped to it. The only object visible in the image is the paper. This is an example of the inability of the sensor to detect certain materials. The sensor emits laser pulses in the near infrared spectrum, if a material absorbs this energy instead of reflecting it the object is invisible to the sensor. Various materials in the lab and outdoors exhibited this behavior. Examples of materials invisible to the sensor are black plastic and water.

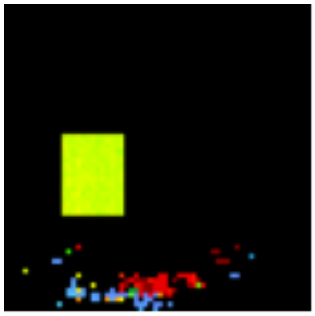

Fig. 9. Example of invisible materials. Image is of a large, black, Hardigg case with a white, US Letter size piece of paper taped on.

When running the outdoor tests we encountered the problem with detecting water, and even object that are wet. The wood objects used for the obstacle avoidance test were wet during several test runs. In these cases the sensor did not detect the wet wood at all. These results will have significant impact on the usability of the Canesta EP200 range sensors in certain environments. Given that the sensor simply cannot see certain materials, it is important to know about the IR absorption properties of the objects in the intended environment before using this sensor.

V Possible Mechanisms for Noise Reduction

Of the three problems encountered with the Canesta sensor, the complete over saturation of the CCD in direct sunlight and the inability to detect certain materials require changes to the hardware. However, the issues of sporadic noise and partial over saturation are solvable using software. The Canesta software attempts to remove the noise in the depth readings "through spatial and temporal averaging."[1] However it is apparent that this is not enough since the depth images still contain noise.

Using several simple techniques from computer vision, the ghost images and noise can be removed from a depth image. The noise reduction algorithm we use is a 4-step solution that operates on raw depth readings from the API. The four steps are: 1. remove known bad values, 2. fill holes, 3. erosion, and 4. dilation. The process is discussed in detail in the remainder of this section.

The first step is to remove over and under saturated pixels from the depth image. This can easily be done since these pixels show up as negative values in the depth image. The first step also eliminates pixels with untrusted values. As mentioned in Section I, 25cm is the minimum unambiguous range as measured in our lab. Any readings closer than this are not accurate and often, object physically closer than 25cm appear to be at the far end of the unambiguous range. Both of these corrections can be accomplished by setting to zero any pixel with a value less than 25cm as well as those within 10 per cent of the maximum unambiguous range.

The second step is to fill any holes in the image. Holes are portions of the image that are zero valued, but completely surrounded by good depth readings. The holes can be any size. They are filled with the average of the surrounding readings, producing smooth flat surfaces for navigation algorithms to use.

The third and fourth steps are complimentary operations. For the third step is an erosion using a 3x3 kernel with uniform weighting is performed on the depth image produced in step 2. This removes any noise that occurs from objects that are outside of the illumination range or stray light entering the lens. Finally the result from step 3 is dilated using the same 3x3 kernel to return detected objects to their correct sizes. Figure 10 shows the image of the cup on the desk (Figure 8) as it goes through the noise removal process. The images below were generated using MATLAB and have a different colormap than Figure 8, which was produced in Canesta's EPToolKit application.

VI Conclusion and Future Work

The tests described above have shown both the good and the bad performance of the Canesta EP200 series of range sensor. The fact that it is able to provide a real time depth image with virtually no computational overhead in a small package is a huge win and shows the potential utility of this sensor. The information provided could be used to construct detailed maps of a robots environment for real time or archival use. The performance problems presented are no small matter however. The fact that there is no way to identify which objects are past the unambiguous range could cause huge problems in an automated system. In addition the limited daytime outdoor capabilities are very limiting for general robotics applications. We feel that in general the benefits of this sensor outweigh the drawbacks, especially if it is used on a multi-sensor system. In such a case a robot would not have to rely solely on the reading from the Canesta sensor and thus will still be able to operate in conditions unfavorable to the Canesta.

This paper has shown that the Canesta EP200 series sensor can be used for autonomous robot navigation as well as an aide for human operators. Three problems have been identified, two of which will require some redesign of the sensor; one however, which is the most frequently occurring, can be solved with software. Future work will include integration of the behaviors used in the test programs into the Distributed Field Robotics Architecture to allow autonomous navigation and mapping using the Canesta sensor on many robot platforms. In addition these will incorporate the filtering techniques outlined in Section V to improve daylight operations.

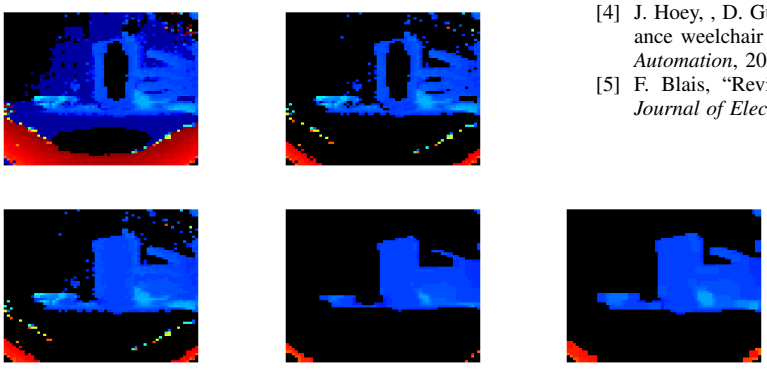

Fig. 10. The noise removal process. From top left: original, removed bad values, holes filled, erosion, and dilation.

References

[1] S. Grokturk, H. Yalcin, and C. Bamj, "A time-of-flight depth sensor – system description, issues, and solutions," in *2004 Conference on Computer Vision and Pattern Recognition Workshop*, June 2004, p. 35.

[2] H. Surmann, K. Lingemann, A. Nuchter, and J. Hertzberg, "A 3d laser range finder for autonomous mobile robots," in *32nd International Symposium on Robitics*, April 2001, pp. 153 – 158.

[3] A. J. Davidson, Y. G. Cid, and N. Kita, "Real-time 3d slam with wide-angle vision," in *Symposium on Intelligent Autonomous Vehicles*, July 2004.

[4] J. Hoey, , D. Gunn, A. Mihailidis, and P. Elinas, "Obstacle avoidance weelchair system," in *International Conference on Robotics Automation*, 2006.

[5] F. Blais, "Review of 20 years of range sensor development," *Journal of Electronic Imaging*, vol. 13, no. 1, pp. 231–243, 2004.