International Conference on Networking, Sensing and Control, Ft. Lauderdale, FL, April 23-25, 2006

Applications of a 3D Range Camera Towards Healthcare Mobility Aids

R. Bostelman, P. Russo, J. Albus, T. Hong, and R. Madhavan

Abstract— The National Institute of Standards and Technology (NIST) has recently studied a new 3D range camera for use on mobile robots. These robots have potential applications in manufacturing, healthcare and perhaps several other service related areas beyond the scope of this paper. In manufacturing, the 3D range camera shows promise for standard size obstacle detection possibly augmenting existing safety systems on automated guided vehicles. We studied the use of this new 3D range imaging camera for advancing safety standards for automated guided vehicles. In healthcare, these cameras show promise for guiding the blind and assisting the disabled who are wheelchair dependent. Further development beyond standards efforts allowed NIST to combine the 3D camera with stereo audio feedback to help the blind or visually impaired to stereophonically hear where a clear path is from room to room as objects were detected with the camera. This paper describes the 3D range camera and the control algorithm that combines the camera with stereo audio to help guide people around objects, including the detection of low hanging objects typically undetected by a white cane.

I. INTRODUCTION

As part of the National Institute of Standards and Technology (NIST) Industrial Autonomous Vehicles (IAV) Project, NIST has been studying the potentially useful CSEM SwissRanger2 (SR2) 3D flash LIDAR (light detection and ranging) camera[1] for mobile robots and other intelligent systems [9,10,11]. This device can provide object detection and range to objects that can be critical for robot traversability and other intelligence characteristics.

The IAV project studies standards, measurements and advanced technology for industrial mobile robots, such as automated guided vehicles (AGVs). AGVs are typically used in warehouse and factory applications where people and other vehicles can also be. Sensors for detecting these and other objects within the vehicle path are critical to the safety of workers and other vehicles within this environment. The SR2 has recently been used to detect standard size objects detailed in British and American standards. The standard and the results of these experiments are detailed further in this paper. As the SR2 can detect these standard objects, the authors believe it can also be used in other environments such as for object and travel path detection applications for the blind or visually impaired.

Audio feedback for the blind has been studied for many years and most recently for graphical environments to hear information about virtual objects including their actual position within this environment [1]. In real environments, the "Victor Trekker" system uses global satellite positioning to tell blind and partially sighted users which road they are walking down, which shops and buildings are near them and when they are coming up to a junction [2]. Even audio maps are being researched where, for example, when the user moves a curser over a water element, a particular sound is played, when crossing a border another audio clip is heard, etc. Voice annotation is also used to speak the features labeled on the map [3].

The NIST Healthcare Mobility Project has recently been studying audio feedback for the blind or visually impaired when sound is combined with an SR2. The combined LADAR with stereo audio feedback can guide a person between objects, through a doorway and even detect low hanging objects that are typically difficult for the blind to detect when using a white cane for guidance. This in fact was also a motivator for the work in [4] on JiL, which used ultrasonic mobility aids and minimal audible warnings when an obstacle was first detected at specified ranges. Other research, although not exhaustive, in the area of audio feedback coupled with ultrasonic sensors are the SonicGuide [5], the SonicPathfinder [6], the NavBelt, the GuideCane [7], and the Robotic Guide [8]. All of these devices use ultrasonic sensors to provide more information about the environment including relative obstacle locations and/or travel path. The Robotic Guide also uses a laser line scanner to augment the ultrasonic sensors to better map the robots environment and includes a radio frequency identification (RFID) detector to locate products on shelves in stores. The position of the ultrasonic in the SonicPathfinder (chest), the Navbelt (waist) and the GuideCane (floor) cannot provide sufficient overhead obstacle information without further modification. However, the SonicGuide is worn by the user as a pair of spectacle frames and therefore, can detect overhead

---

All authors are with the National Institute of Standards and Technology (NIST), Gaithersburg, MD 20899 U.S.A. Corresponding author: E-mail: roger.bostelman@nist.gov, Tel: 301-975-3426.

P. Russo is a student at the University of Maryland and interns with NIST.

[1] Commercial equipment and materials are identified in this paper in order to adequately specify certain procedures. Such identification does no imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

obstacles based on acoustic sensor cone size and user head movement.

In all of these electronic travel aid systems, ultrasonic sensors: provided a range of about 4.6 m (15 ft); detected objects of approximately 30 mm (1.2 in) or larger; required continuous scanning of the environment where some scanning included user movements; may require additional measurements of obstacles detected to determine the dimensions and shape of the object; and have acoustic interference (called masking), with sound cues from the environment, reducing the blind person's ability to hear these essential cues [7].

NIST considered the above technology factors and incorporated the new SR2 3D range camera instead of ultrasonic sensors. The SR2 provides faster, higher resolution, and longer range data and it is also a compact, lightweight, self-contained sensor system (sensor, light source and processor) as compared to acoustic sensors. The SR2 may not currently be cost effective for most potential blind users but, has reduced in prototype cost fourfold from its SR1 predecessor while increasing in resolution by four. Also, the manufacturers have predicted that the cost will decrease to cost-effectiveness through mass-production.

Sections II and III of this paper describe the SR2 and the tests towards non-contact safety sensors implemented on mobile robots, providing a sound basis for SR2 obstacle and travel path detection. Section IV describes the SR2 incorporated into a second application of an audio feedback system for the blind or visually impaired when used as an electronic travel aid. In section V, we discuss deficiencies and planned experiments with the SR2. And in sections VI and VII we provide a conclusion and propose future research followed by a list of references, respectively.

II. CSEM SwissRanger2 3D Range Camera

The SR2 3D range camera [12], shown in Figure 1, is based on the Time-Of-Flight (TOF) principle and is capable of simultaneously producing intensity images and range information of targets in indoor environments.

Figure 1 – CSEM SwissRanger2 3D Range Camera

This range camera is extremely appealing for obstacle detection in industrial applications as it will be relatively inexpensive as compared to similar sensors. It can deliver range and intensity images at a rate of 30 Hz with an active range of 7.5 m while incorporating no moving parts, such as a spinning mirror found in many off-the-shelf laser sensors.

The SR2 camera is a compact, robust and cost effective solid-state device capable of producing 3D images in real-time. The camera measures 14.5 cm x 4 cm x 3 cm (5.7 in x 1.6 in x 1.2 in), has a field-of-view of 0.7 radians (42 degrees) horizontal x 0.8 radians (46 degrees) vertical, and is capable of producing range images of 160 pixels x 124 pixels.

The camera is equipped with a light source that is ideal for use in dark areas such as within collapsed structures and caves, although indoor lighting had little affect on the SR2 data. The camera can provide 3D distance to objects up to 7.5 m (25 ft.) including stairs and measurement of step height/depth to within 2.5 cm (1 in) and even detect such negative objects such as stairwells and holes if the camera is angled properly. It can also provide distance to piles of objects such as cinder blocks and boxes, and can measure the shapes and sizes of the piles.

III. SR2 use toward Advancing Industrial Vehicles Safety Standards

A recent change to the U.S. American Society of Mechanical Engineers (ASME) B56.5 standard [13] allows non-contact safety sensors as opposed to contact sensors such as contact bumpers on AGVs. Prior to this change, the B56.5 standard defined an AGV bumper as a "mechanically actuated device, which when depressed, causes the vehicle to stop." The recent change allows proven non-contact sensing devices to replace or be used along with mechanical AGV bumpers. The allowance of non-contact safety sensors on AGVs opens new areas of vehicle control generally supporting the safety notion of not only stopping the vehicle but, also slowing the vehicle around high-risk objects, such as humans. For example, should the vehicle sensors detect walls, shelving, posts, etc., the vehicle may continue to move at higher speeds than if the sensors detect what could be human arms or legs.

Figure 2 – Test Apparatus

Now, the American ASME B56.5 and British EN1525 safety standards [14] both consider non-contact safety sensors and specify that horizontal test pieces used to test sensors shall be 200 mm (7.9 in) diameter x 600 mm (23.6 in) long lying perpendicular to the vehicle path. Vertical test pieces shall be 70 mm (2.8 in) diameter and 400 mm

(15.7 in) tall completely within the vehicle path. Figure 2 shows a photo of the vertical apparatus used at NIST to test the SR2 toward meeting the US and British standards. The horizontal test apparatus and associated SR2 processed data can be seen in [10].

The center post is closest to the standard size suggested by the safety standards. However, it's interesting to note that the objects to the right of center are successively smaller down to 19 mm (0.75 in) wide and are still detected by the SR2 at ranges up to approximately 5 m (16 ft).

An algorithm then processes the SR2 data and determines not only the objects range, but also segments the objects from their environment and places them in a world model. Generally, the obstacle detection and segmentation algorithms combine intensity and range images from the range camera to detect the obstacles and estimate the distance to the obstacles.

We first calibrate the camera with respect to the AGV so that we can convert the range values to 3D point clouds in the AGV coordinate frame. Next, we segment the objects which have high intensity and whose elevation values are above the floor of the operating environment on the AGV path. The segmented 3D points of the obstacles are then projected and accumulated into the floor surface-plane. The algorithm utilizes the intensity and 3D structure of range data from the camera and does not rely on the texture of the environment. The segmented (mapped) obstacles are verified using absolute measurements obtained using a 2D scanning laser rangefinder with a range uncertainty of 3 cm (1.2 in). Figure 3 shows the resultant intensity, range, and segmented objects images.

Potential obstacles in the world model can be accumulated as the AGV moves. Figure 4 shows an obstacle map representation that is part of the world model – overhead map of obstacles. The obstacles map is shown at 10 cm grid resolution. Further details of the algorithm and SR2 experiments can be seen in [9, 10, 11] including: outdoor experiments with the SR2 in cloudy, shaded conditions and combining two SR2's to make a single, wider FOV image.

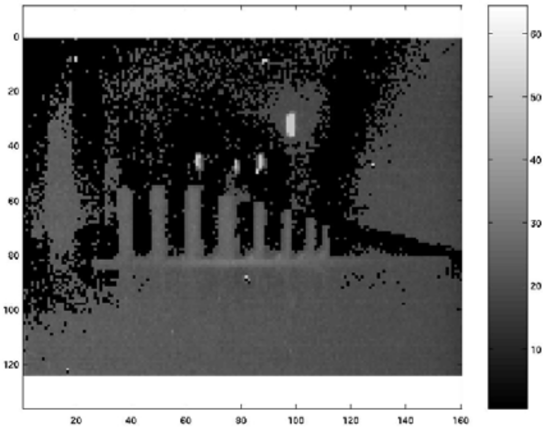

(a)

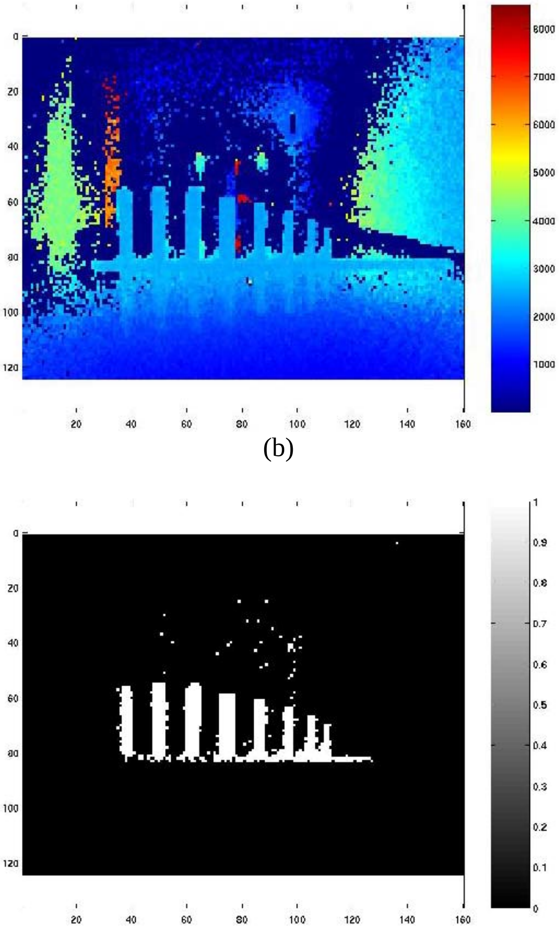

(b)

(c)

Figure 3 – Results of the obstacle detection and segmentation algorithm for the experimental setup shown in Figure 2. The resultant intensity, range, and segmented object images are shown in (a), (b), and (c), respectively.

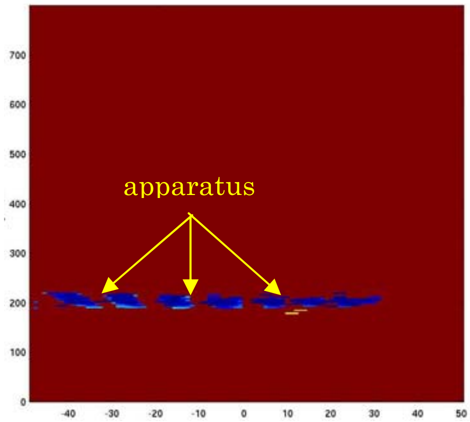

Figure 4 – Segmented objects rotated to show a top-down view for placement in a world model. Bottom axis shows width, left axis shows distance from camera.

IV. SR2 USE TOWARD GUIDING BLIND PEOPLE

A. GUIDANCE METHODOLOGY

The realization of a tool for blind guidance faces several challenges. The guidance device must either convey object data accurately enough to the user in order to allow him/her to navigate on his/her own, or the device must choose a perceived destination and convey to the user the path to the chosen target.

Converting complex visual data into non-visual data is a challenge. Sound is a tempting substitute for sight, as an individual can typically hear and locate the source of many sounds at the same time. Fortunately, there is an array of tools for synthetically creating three-dimensional sound effects, including the Open Audio Library (OpenAL) and Microsoft's DirectSound. Convincing three-dimensional sound effects can be heard with even an inexpensive pair of headphones – an array of speakers is not necessary and in fact is not as effective.

Given the large number of 3D data points generated by one or more three-dimensional imagers (in this case the SR2), the data is too complex for direct conversion to audio data. Our approach is to extract objects, and find a safe path, and then guide the user to his or her destination.

The first step of our approach is to remove the floor data to not mistake it as an obstacle. With the SR2 physically on a fixed mount, removing the floor from the image data is easy. Since the height and vertical angle of the camera are known, the data can be rotated and a simple threshold can be used to extract the floor. On an unfixed mount, the problem becomes much more complicated. The surface of the floor has to be detected first [13].

The second step of the approach is to segment objects, which are tracked. This is performed using the 3D Connected Components Algorithm, an efficient adaptation of the Connected Components algorithm for grouping connected points together.

This approach was chosen because it is independent of the SR2 from which the data originated as long as the data is registered and placed into the same coordinate system. Grouping data points into objects (obstacles) is a more manageable way to convert sight into sound cues and provides the option of playing a sound at the most hazardous position on each object as a warning. The user can then use the warnings to navigate around the obstacles.

The obstacles were then integrated into a local map centered on the user [15]. Each obstacle also included the space the user takes up to prevent the sound guidance from attempting to steer the user through a space that is too small or clip the corner of objects as he moves past them. With the search graph built, the A* search [16] is used to plot a path from the current location to the intended destination of the user. The Manhattan distance (distance between two points measured along axes at right angles) is used as the heuristic function. The algorithm tries to suggest a path for the user to get from where he is to 7.5 m directly in front of him.

Although the algorithm can find a full path of how to get the user from where he is to where he would/might like to be, it is hard to convey this data without visual cues. The method chosen for this algorithm is to only inform the user about what direction he should take. To accomplish this, a sound is played in each ear that seems to be coming from the suggested direction of travel. The object data and path are updated every frame, and the location of the sound is adjusted accordingly. For example, if there is an object in front of the user, a sound will be played in either the right or left ear depending upon how the path planner calculates a clear path to send the user.

B. EXPERIMENTAL SETUP AND DATA COLLECTION

The SR2 could have been worn on a person's chest, belt, head or elsewhere to study it's use for guiding the blind or disabled. However, as the person walks they move horizontally and vertically complicating the experiment. For this paper, NIST instead considered a preliminary case where the SR2 is fixed mounted to a stationary support providing the basis for blind person guidance as well as, applying 3D range imaging to yet another area – powered wheelchair guidance. The SR2 was mounted on the top of a vehicle currently being researched at NIST in the Healthcare Mobility Project. The vehicle is a NIST-designed, powered lift-wheelchair. The lift- wheelchair, called the NIST RoboChair and shown in Figure 5, incorporates not only powered horizontal mobility but, vertical mobility as well.

The RoboChair, was designed to allow lift for wheelchair dependents to reach high objects at least a meter above what typical wheelchair-users can reach. A follow-on version will also have the capability to rotate and seat wheelchair users on chairs, beds, toilets, and other seats. However, should the vehicle be kept at the high position while moving laterally, the RoboChair may hit high objects, such as door frames, requiring a need for detection of these obstacles too (see Figure 6). The first version of RoboChair, incorporating powered mobility and lift as well as 180 deg chair rotation, also served as a testbed for the SR2 audio feedback system for the blind. And existing IAV Project safety standard data collection code was also adapted for the blind guidance system.

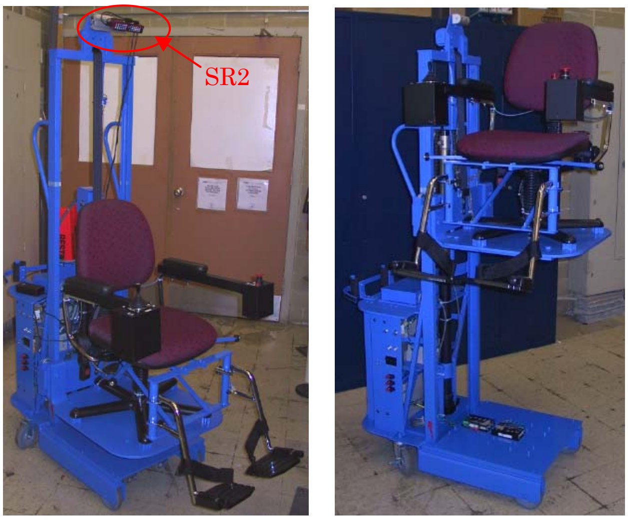

Figure 5 – Photos of the NIST RoboChair (left) in a low, high-speed horizontal mobility configuration and (right) in a raised vertical mobility configuration.

A typical task a blind person must accomplish in any room is to find the way to the next room. Therefore, such a task was considered a good test of the functionality of the blind guidance system.

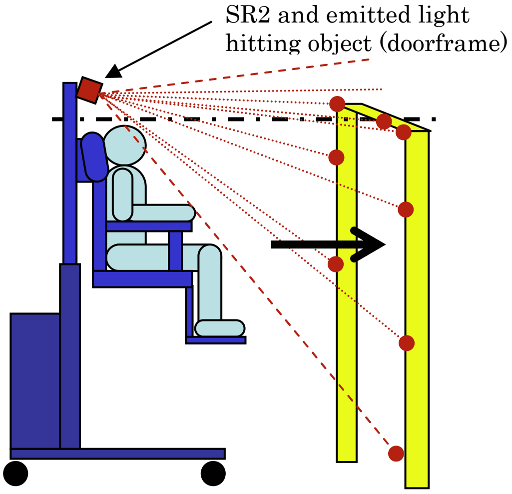

Figure 6 – Graphic showing a potential hazard (lift-wheelchair passing through a doorway) that can be detected by a 3D range imaging camera.

The SR2 was set up in front of a doorway with obstacles blocking the path. Leading up to the door is a narrow hallway formed by a wall on the left and cabinet backs on the right. The doorway has two doors, one of which is closed and the other is open. Beyond the door is a hallway that runs perpendicular to the one in which the test was conducted. Figure 7 shows a picture of the scene and data obtained during the test. The SR2 was mounted 1.9 m (74 in) above the ground and angled 0.45 rad (26 deg) downwards so that the top of the field of view of the camera was slightly above horizontal.

The algorithm was very successful at completing the planning task. The algorithm generated a path around obstacles placed in front of the 3D range imager. When the objects or camera were moved, the path replanned around the new object positions. Audio played to each ear corresponding to how the path changed. For example, as the path shifted to go left around an object instead of previously to the right, the audio signal to the ears changed accordingly. The process only failed when the camera came too close to an object and the high-amplitude return signal saturated the demodulation pixel array.

V. SR2 DEFICIENCIES AND PLANNED EXPERIMENTS

The SR2 proves to be a useful tool in both manufacturing and blind guidance. Its real-time frame rate, small size, and low relative cost make the SR2 well suited for many mobility applications.

However, the SR2 has a number of shortcomings that may be overcome in applications where the user remains the main decision maker. For example:

- The camera accuracy and precision are degraded under conditions where a weak signal is returned. Weak signals can be caused by long-distances and by scanning dark-colored objects. In most cases, close range detection has proven fairly reliable.
- The SR2 does not perform well in areas with a lot of near-infrared background noise, such as occur in outdoor environments. A potential solution here is to replace the illumination system of light-emitting diodes with perhaps laser diodes.
- The SR2 also has difficulty with very high amplitude signal returns. High amplitude returns may be caused by either a reflective object in the camera's field of view, or by an object that is too close to the camera. High returns cause spill effects that cause blur into neighboring pixels. This problem of high amplitude renders the SR2 inoperable at any range less than approximately 1 m (39 in). It may be possible to adjust integration settings to improve both high amplitude and weak signal results.

Although the SR2s resolution (160x124 pixels) is suitable, a higher resolution would be beneficial. In future blind guidance systems, it will not be sufficient to simply detect objects. It will also be necessary to classify certain key objects such as doors, cars, people, roads, and crosswalks. The SR2 can currently detect that an object exists with good reliability. However, classifying the object is more difficult.

In our experiments, we have demonstrated that the SR2 can be used as a non-contact cane for detecting obstacles. NIST envisions using the SR2 as a non-contact cane or wearable obstacle detection device. We are investigating the possibility of conveying to the user the position of all objects in the scene via the three-dimensional sound synthesizer that the planner uses. Current problems with this approach include the difficulty of measuring and conveying complex geometric information without confusion due to multiple sound sources.

NIST plans to test the system with up to four SR2s to provide a wider field of view.

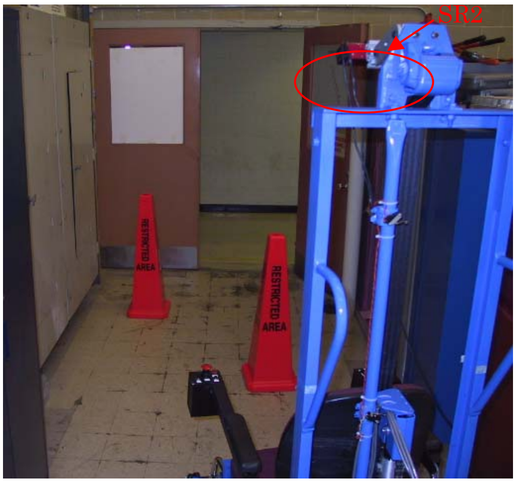

(a)

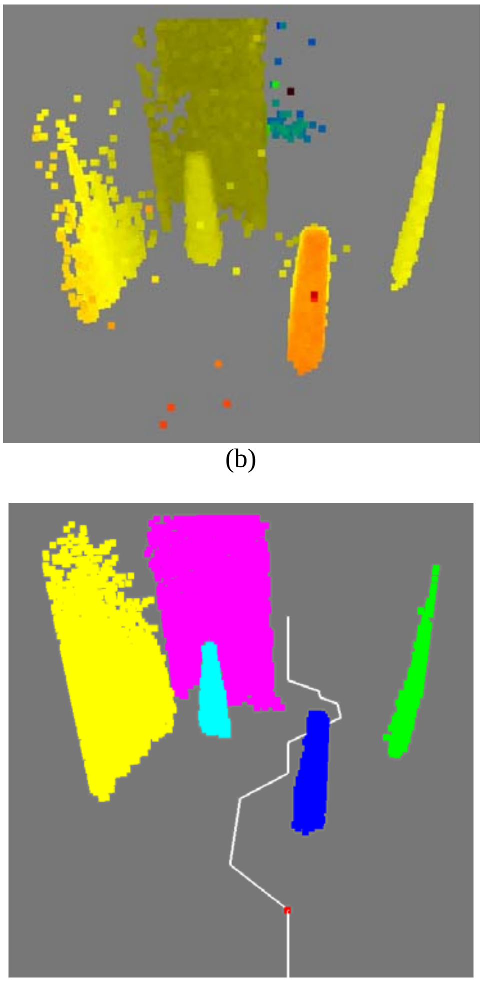

(b)

(c)

Figure 7 – (a) Photo of the test scene, (b) Input to the planner algorithm, (c) Range data, objects and path planned.

VI. CONCLUSIONS AND FUTURE RESEARCH

The SR2 is sufficient for detecting standard size objects in a robot path [17]. The obstacle segmentation algorithm was very successful in isolating detected objects placed in front of the camera. The path planning algorithm was also successful at planning a path and for directing sound to the ears laying groundwork toward a guidance system for the blind. The NIST RoboChair was used as a testbed demonstrating that the 3D imager can also be applied to lift devices for detecting high objects. The SR2 obstacle detection process only failed when the camera came too close to an object and the high-amplitude return signal saturated the demodulation pixel array.

However, to develop a more robust system, more advanced 3D imaging devices will need to be developed to enable object classification and recognition. A better ground plane detection algorithm needs to be developed to extract the floor without knowing height and angle information in advance, such as when the user wears the sensor on his/her moving head. In future research, we plan to mount the sensors on other moving support structures, for example suspended, mechanical systems, that can turn a wheel-chair victim around and seat them on a bed, chair, or toilet. Also, we plan to test the performance of the system and incorporate 4-D/RCS [18] to the system to provide real-time modular software architecture for plug-and-play of test algorithms that can optimize path planning and world modeling.

VII. REFERENCES

[1] G. Romain, "Device Lets Blind Feel Graphics", Betterhumans Staff, September 16, 2003.

[2] news.bbc.co.uk/2/hi/health/2503149.stm, "Satellite Guidance Helps the Blind", November 22, 2002.

[3] B. Campin, W. McCurdy, L. Brunet, and E. Siekierska, "SVG Mapping for People with Visual Impairment", SVGOPEN Conference 2003, Vancouver, Canada.

[4] P. Blenkhorn, S. Pettit, and D. Evans, "An Ultrasonic Mobility Device with Minimal Audio Feedback", www.dinf.ne.jp/doc/english/Us_Eu/conf/csun97_019.htm, November 22, 1997.

[5] L. Kay, "Evaluation of the Ultrasonic Binaural Sensory Aid for the Blind", Evaluation of Sensory Aids for the Visually Handicapped, National Academy of Sciences, 1971.

[6] A.D. Heyes, "Sonic Pathfinder",Wireless World, April 1984.

[7] S.Shoval, I. Ulricch, and J. Borenstein, "Computerized Obstacle Avoidance Systems for the Blind, and Visually Impaired", Invited chapter in "Intelligent Systems and Technologies in Rehabilitation Engineering," CRC Press, pp. 414-448, December 26, 2000.

[8] V. Kulyukin, C. Gharpure, and J. Nicholson, "RoboCart: Toward Robot-Assisted Navigation of Grocery Stores by the Visually Impaired", Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems, August 2005.

[9] R. Bostelman, T. Hong, and R. Madhavan, "Obstacle Detection using a TOF Range Camera for AGV Safety and Navigation", Journal of Integrated Computer-Aided Engineering, Vol. 12, No. 3, pp. 237-249, IOS Press, 2004.

[10] R. Bostelman, T. Hong, and R. Madhavan, "Towards AGV Safety and Navigation Advancement – Obstacle Detection using a TOF Range Camera", Proceedings of the International Conf. On Advanced Robotics, pp. 460-467, Seattle, WA, July 2005.

[11] R. Bostelman, T. Hong, and R. Madhavan, "Experiments Toward Non-Contact Safety Standards for Automated Industrial Vehicles", Proceedings of the SPIE Defense and Security Symposium, Orlando, FL, March-April, 2005.

[12] T. Oggier, M. Lehmann, R. Kaufmann, M. Schweizer, M. Richter, P. Metzler, G. Lang, F. Lustenberger, and N. Blanc. "An All-Solidstate Optical Range Camera for 3D Real-time Imaging with Sub-centimeter Depth Resolution", Proceedings of the SPIE Conference On Optical System Design, September 2003.

[13] American Society of Mechanical Engineers' Safety Standard for Guided Industrial Vehicle and Automated Functions of Manned Industrial Vehicle, Technical Report ASME B56.5, 2004.

[14] British Standard Safety of Industrial Trucks – Driverless Trucks and their Systems, Technical Report BS EN 1525, 1998.

[15] T. Chang, T. Hong, S. Legowik, and M. Abrams, "Concealment and Obstacle Detection for Autonomous Driving", Proceedings of the IASTED Robotics and Applications Conference, Santa Barbara, October 1999.

[16] P.H. Winston, "Artificial Intelligence", 3rd Edition, Addison-Wesley, 1992.

[17] Y.K. Hwang and N. Ahuja, "Gross Motion Planning – A Survey", ACM Computer Surveys, 24(3), pp. 219-291, September 1992.

[18] J. Albus, "4-D/RCS Reference Model Architecture for Unmanned Ground Vehicles", Proceedings of the International Conference on Robotics and Automation, pp. 3260-3265, 2000.