Capturing Dense Environmental Range Information with a Panning, Scanning Laser Rangefinder

Lars S. Nyland

Abstract— **We use a laser rangefinder with a scanning mirror on a panning platform to acquire range data in a local environment (range < 15m). In this paper, we discuss the hardware components, the software, show several examples of the range data, and discuss problematic issues such as material properties. We conclude with our future plans for improving the data and integrating it into our research program.**

Keywords— **laser rangefinder, scanning, range acquisition**

I. Introduction

THE driving goal of our research group is to render 3D scenes using images, rather than polygons, as the source data. This requires full-color images augmented with range data on a per-pixel basis. Each of these images is used as an input primitive to a rendering technique called *image-based rendering* [2].

In this paper, we review the technology we have developed for capturing depth of an environment. The hardware consists of a time-of-flight laser rangefinder, a line-sweeping mirror, and a panning unit all attached to a conventional PC. Custom software is required to control this hardware and is described. After a brief description of calibration and error compensation, we present several examples, showing both the quality of the range readings and the problems encountered in real-world environments.

II. Image-Based Rendering

Image-based rendering can be viewed as the fusion of computer vision (creating models from images) and 3D computer graphics (rendering images from models), where we render images from images without explicitly creating a geometric model. To render from images, we require not only the photographic images of the scene, but range data on a pixel-by-pixel basis. Given the digital image and a corresponding range image, we can render the scene over a wide range, providing realistic renderings. Multiple images of the same scene from different locations allow renderings from virtually any position in the environment without "holes" created by occlusions.

Acquiring digital photographs is commonplace with today's technology. Our photograph acquisition relies on digital cameras, such as the Canon EOS D2000 or Olympus DL300 camera, or scanned photographs. Acquisition of range data is not nearly as simple. The main point of this paper is to describe our active range acquisition system, contrasting it with other active and passive systems.

Department of Computer Science, Univ. of N. Carolina, Chapel Hill, NC 27599, `http://www.cs.unc.edu/~nyland`

This work funded by in part by NSF and DARPA.

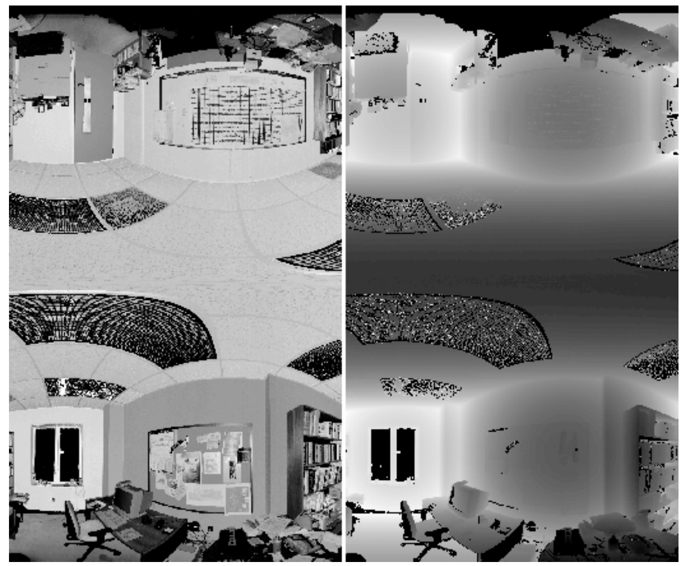

Fig. 1. The spherical image of reflected laser light and range data. Each column starts at 60° below the horizon going up, over the North Pole (middle of the image) continuing down the other side, until 60° below the horizon is again reached for a total of 300° per column. The horizontal field-of-view is 180°, providing a full panorama of the environment that is only missing a 60° cone towards the floor.

III. An Example of Range Data

Figure 1 shows a typical dataset of reflectance and range data, resampled onto a grid). The projection is spherical (each column is a longitude where the samples are spaced at equal angles), as it most naturally matches the rangefinder's rotating mirror and panning unit.

The data is shown in subsequent figures (2, 4, and 3) visualized as point-clouds, using the reflected intensity of the laser light to color each point. The images are captured from a visualization tool as we move through the data. We use this tool to explore the data, modifying it with a variety of techniques (described later in the paper) to eliminate errors in the data. Figures 2 and 3 are rendered from a single set of range data collected in figure 1. Figure 4 is from an earlier set acquired in the same room during a calibration experiment (described fully in section VI-A).

IV. Range Acquisition Hardware

The hardware used for range acquisition consists of a laser rangefinder, a rotating mirror to scan a vertical plane, a panning unit, a conventional PC, and a cart to move all the equipment from location to location. We have also included a 12 volt, deep-cycle battery with an inverter,

Fig. 2. A reprojection looking toward the windows using the data from figure 1. The scenery outside the windows is too far away to reflect enough light and is shown as black.

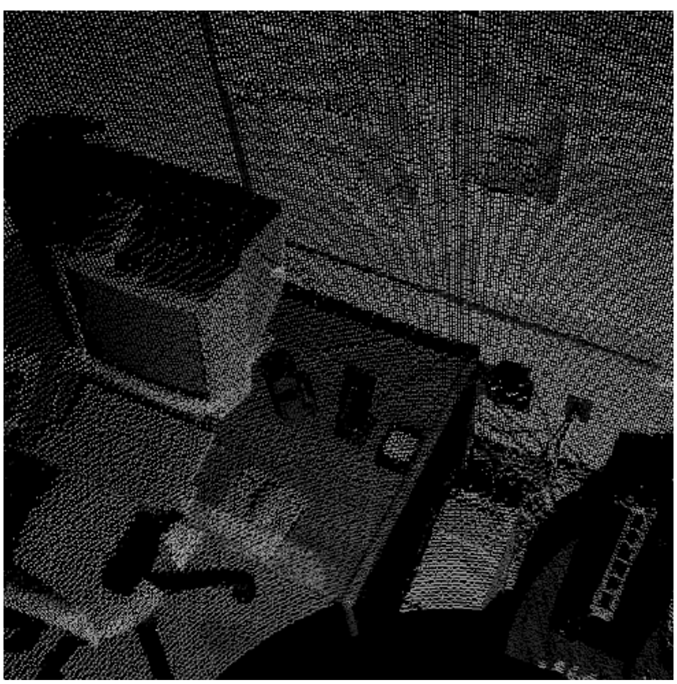

Fig. 3. A view from above a computer table. A chair, keyboard and monitor are on the left. Next to the table, on the right, is a computer with all of its wires shown behind it.

Fig. 4. A view showing the ceiling tile pushed upward, allowing us to calibrate the azimuth.

allowing us to keep power cords out of the images. This section describes each in more detail.

A. Laser Rangefinder

At the heart of our range acquisition system is a commercial laser rangefinder, the AccuRange 4000-LIR from Acuity Research, Inc. in Palo Alto, CA. It modulates the laser for time-of-flight measurements to accurately measure distances from less than 1 inch to 50 feet (15m). It weighs less than 4 lbs, and is reasonably compact. It can take measurements at speeds up to 50 kHz, but we usually run it between 8-25 kHz to improve quality. We purchased the infrared model (the 4000-LIR) because of higher powered laser (8mW vs. 5mW) that gives better readings at long range and in sunlight. The high-speed interface card is required to obtain data at rates higher than 1 kHz, which requires a computer with an available ISA slot. More information is available at http://www.GoodEnuf.com/client/Acuity/.

The high-speed interface returns an 8-byte sample for each range reading. The information contained includes 19 bits of range information, 8 bits of signal strength (laser light reflected), 8 bits of ambient light, 8 bits indicating motor position, and a few flag bits indicating such things as buffer overflow and scanning motor zero.

The laser is not eye-safe, due to its power, invisibility, and lack of hardware interlocks. When in operation, the nominal hazard zone is only a few meters in diameter. Precautions are taken when the laser is in operation, we currently use it in closed environments where all occupants are wearing eye protection (we are hoping to relax this restriction with our safety officers).

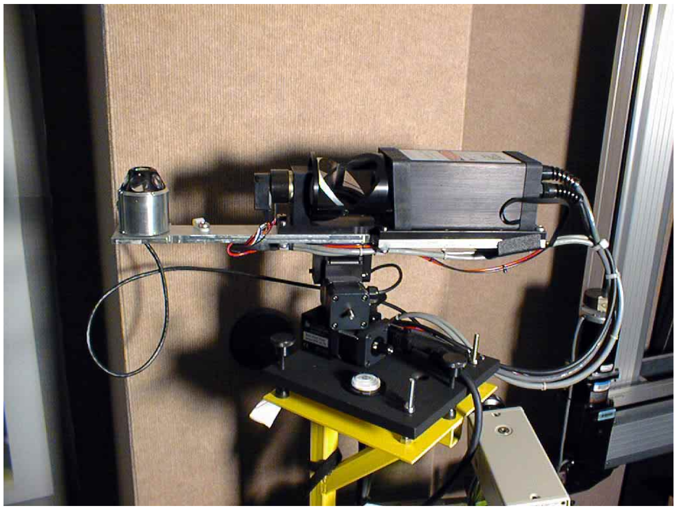

Fig. 5. The rangefinder, scanning mirror, and pan-tilt unit. The hiball tracking unit is on top, with 6 lenses looking at the ceiling.

B. Supporting Hardware

B.1 Scanning Mirror

Acuity Research sells a line-scanning attachment which is a rotating mirror tilted at 45°. The mirror rotates at speeds from 2 Hz to 40 Hz, allowing range measurements in a 300° sweep (the supporting frame obscures 60°). We typically run the scanning mirror at 2-10 Hz (allows acquisition of 5–20 samples per degree). The motor has a shaft encoder with 2000 positions per revolution (however, the high-speed interface card can only read the lower 8 bits, leading to some interesting estimation software) with a pulse generated at an unknown "zero-position." We have mounted the scanner to measure vertical planes.

B.2 Pan-Tilt Unit

The scanning rangefinder must be rotated (or translated) to acquire a 3D data set. To rotate the device, we attached it to a pan-tilt unit (model PTU-46-70) from Directed Perception, Inc (see http://www.DPerception.com for more information). We do not require the tilt operation for our task, and fortunately it can be disabled. The panning unit attaches to a computer with a standard serial cable. The panning unit has a claimed position resolution of 0.771 arc minutes, allowing us to make nearly 80 scans per degree. We chose this model over the 46-17.5 for the higher positional resolution. It also has the advantage of more power and slower velocity since lower gearing is used. The pan-tilt unit is very easy to interface to. The rangefinder with the scanning mirror and pan-tilt unit are all shown in figure 5.

C. Complete Rangefinder System

The complete system consists of the line-scanning laser rangefinder with the high-speed interface card, the pan-tilt unit, a Dell PC with a 200 MHz Pentium Pro. In addition to the 2 special ports on the high-speed interface card, the rangefinder also requires a serial port, as does the pan-tilt unit. A power supply (12V) is also required to power the scanning mirror motor.

Fig. 6. The movable cart containing the necessary hardware to acquire range data (including power). The rangefinder and panning motor are in the upper right, there is a computer display and B/W monitor (for showing the laser position) in the center, underneath is a Dell PC, and at the bottom is a deep-cycle marine battery and power inverter.

All of these components have been mounted on a cart, shown in figure 6. The PC has been outfitted with a flat LCD display for two reasons: the first is size, we want to minimize the appearance of the system in the rangefinder's field-of-view; the second is power consumption, since we often rely on battery power.

To allow the operator to "see" the laser, we have placed an infrared-sensitive video camera with a wide-angle lens and a small video monitor on the cart. And to be mobile, we have added a deep-cycle marine battery (12V) with an inverter to get 120 VAC to eliminate external power needs. The system can run for several hours on a single charge.

To complete the system, we have built a bracket to hold our digital camera with its center-of-projection aligned with the center-of-projection of the panning, scanning rangefinder (see figure 7). After an environment is scanned, the laser is removed from the panning unit and replaced with the camera, whereupon we take 12 images of the environment.

D. Errors and Compensation

D.1 Sources of Errors

The precision of the measurements made by this system is limited by each of the components. First and foremost is the rangefinder device itself. It appears to have range

Fig. 7. The Canon EOS D2000 with mounting bracket to position the center-of-projection of the lens at the same location as the center-of-projection of the rangefinder. Since the light-sensor in this camera is smaller than film, it is necessary to use an extremely wide-angle lens for image collection. Here we show the 14mm lens, which acquires an image similar to a 24mm lens on 35mm film.

accuracy better than 1 part in 500. Since data is collected in a spherical manner, the precision of theta and phi (rotation around the z-axis and angle from z-axis) also play an important role. Each is addressed separately.

The rotational angle of the scanning mirror is reported by a shaft encoder with 2000 positions. While this sounds plentiful, consider that in 90°, there are only 500 positions, thus a 1024×1024 image with a 90° field-of-view would have more than 2 pixels per shaft position.

The motor that drives the mirror has limited power, requiring some time to accelerate the motor to the desired operating speed. This turns out to be a benefit as there are no sudden accelerations, so we can safely assume that the velocity is constant during the time required to acquire 1 buffer (1024 samples requires 1/16 second at 16k samples per second). A least-squares linear fit of the data is performed to map the integral data into real values. The least-squares fit adequately determines *relative* changes in phi.

The *absolute* position of the scanning mirror has not been specified by the manufacturer, and we have been in an ongoing process to calibrate it. The value available from the shaft encoder is a counter that starts at 0 when the device is powered on. In the data packet returned from the rangefinder, there is a bit that signifies that the shaft encoder has passed a preset "zero" position. It is not specified at what angle this occurs, and this is the value we have been seeking. To find it, we created a range event near the azimuth (a ceiling tile was raised, shown in figure 4). A scan was taken in one direction, followed by a 180° horizontal rotation where another scan was taken. The data from second scan was reversed ($\phi' = 180 - \phi$), and the two compared. If the horizon is well-known, then the two events marking the missing tile should coincide. A process of refining our estimate of the zero position of the shaft encoder has lead us to a value refined to the nearest 0.01°. We have concluded that the mirror's zero position occurs 131.77° prior to the beam being horizontal.

The panning-tilting unit is another source of inaccuracy. The manufacturer claims it can be positioned every 0.771 minutes of arc, or 14,000 positions in 180°. However, when a pan of 180° is requested, the position is off by several tenths of a degree.

Another source of imprecision is the mirror that performs the line-scanning operation. It should be angled at 45° from the laser beam, but the accuracy here is suspect. The plane scanned out is not truly a plane, but a nearly-flat cone. Once the angle of the cone is known, compensation of the error can be made in software.

D.2 Calibrating and Correcting Errors from the 45° Mirror and Panning Motor

We have performed a simple experiment to estimate the error of both the panning motor and the 45° mirror. It consists of taking 4 measurements in a large room as follows.

1. Place the rangefinder in the approximate center of the room.
2. Aim the laser horizontally at the center of one wall, and mark its position (called A).
3. Rotate the mirror 180° vertically, and mark its position on the opposite wall (called B).
4. Rotate the panning unit 180° horizontally, repeating the above markings (called C and D).

The measurement results are shown (exaggerated for effect) in figure 8. The calibration result will define values for x and y, and are determined as follows.

If the mirror is at 45°, then x will be 180°. The value of x is determined by $360 = 2x + \alpha + \beta$. This yields $x =$ 179.54°. Thus the angle between the cone and the plane is $(180-x)/2$ since the error is equal on both sides of the cone. It is $\epsilon = 0.23^\circ$. Knowing this value, we can recalculate the actual position of the laser.

The rotation of the panning motor is captured by the value for y. It is computed by $y = x + \alpha = 179.7$, since α is covered twice (note the rotation arrow in figure 8). This error is assumed to be linearly distributed over the span of the panning motor and is large enough that compensation is required. Fortunately, the compensation involves changing a single constant, the ratio of positions per degree, thus the correction is trivial.

V. Range Acquisition Software

A Java application was developed to control and collect data from the laser rangefinder and associated hardware. Symantec's Visual Cafe software allows a developer to build graphical interfaces for Java programs quickly. The application's interface to control the rangefinding system is shown in figure 9. The software is multi-threaded, using the JVM default scheduler. The threads are:

1. A data collection thread. This thread continuously waits for the rangefinder's buffer to be more than half full at which point it reads half of the buffer from the laser to the PC. It uses 2 buffers in the PC, so that buffers are not overwritten while other threads are reading the data.
2. The control thread that repeatedly positions the panning unit and reads columns of data.

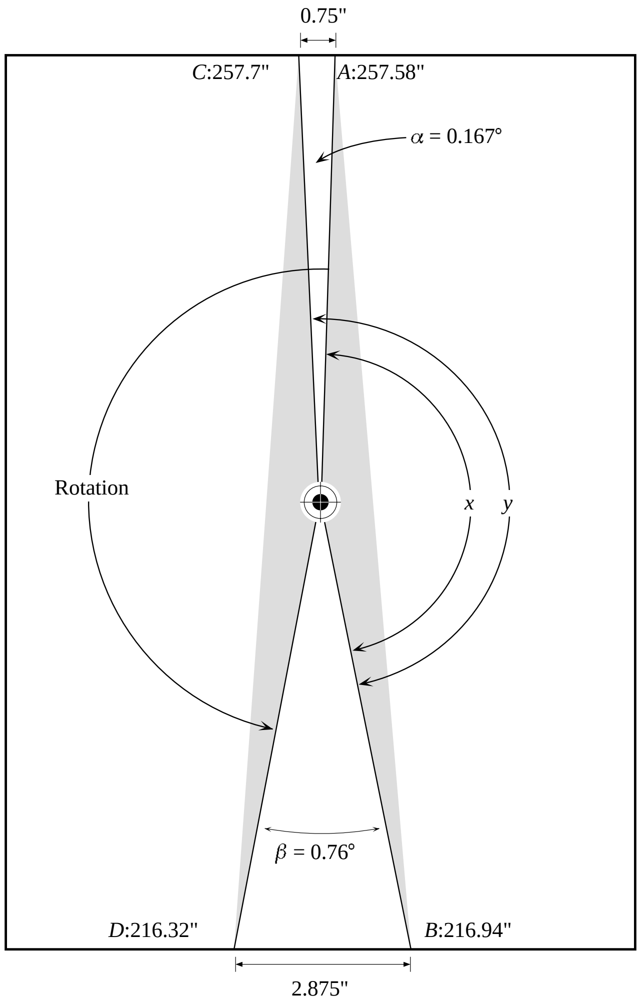

Fig. 8. The measurement setup for determining error in the angle of the 45° mirror and the panning motor. The angles are greatly exaggerated here. The gray areas designate the cone-shaped areas (rather than plane) swept by the laser as it reflects off the 45° mirror. Measurements at A, B, C, and D were taken from which the spread distances could be measured. The angles α and β were determined assuming right triangles.

3. A scanning mirror monitor. The job of this thread is to simply show a histogram of the scanning motor's speed so the user can adequately judge when the motor has achieved the desired speed.
4. The built-in Java threads that control and dispatch menu, slider, and keyboard events.

A. Java Considerations

The data collection thread is run at a higher priority than all the other threads. The reason for this is that we don't want to lose any data, since that might require repositioning the panning motor and waiting for the scanning mirror to return to the point where the last good data was received.

Since Java is object-oriented and manages all object-memory in the JVM, a strong effort was made to allocate all the necessary objects prior to data collection so that

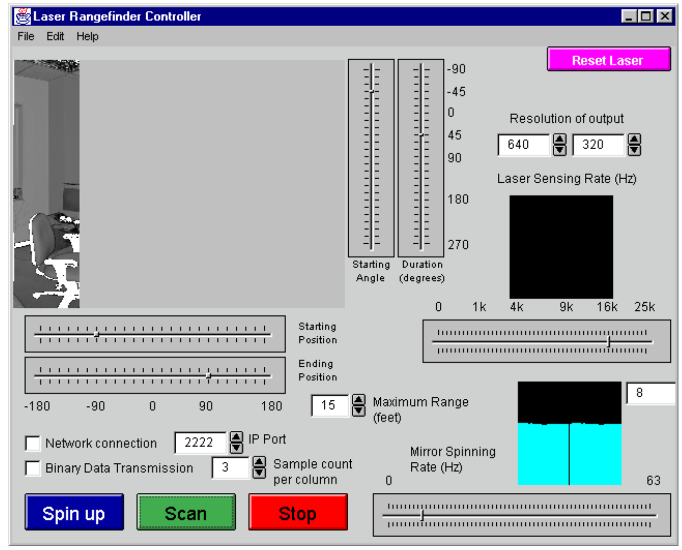

Fig. 9. The GUI for the rangefinder data-collection application. This clip shows a scan in progress. The data collected has a resolution of 640×320, covering $180^\circ \times 90^\circ$. The sample rate is 16k samples/second, and the scanning mirror is running at about 5 Hz. The partial scan is shown in the upper left.

the Java garbage collector would not run during data collection. We have had success here, it is noticeable when the collector runs, and we do not experience any associated delays.

With multiple threads that share data, it is important to ensure proper access. For example, the thread collecting data from the rangefinder should not overwrite data being read elsewhere. One way to do this is to use the predefined synchronization primitives defined in Java. Another way is to carefully design access to the shared data structures so that they are read only when they have been properly defined. In this application, it is possible to perform the latter, so the cost of the Java critical section implementations has not been an issue.

Unfortunately, the virtual memory size used by the JVM is neither self-adjusting nor controllable from within Cafe (under Windows), thus, for big data collections, the application must be run from an MS-DOS window specifying the memory size. Typically, we set the memory size to 100–200 MB for our collections.

B. Collecting Data

The "Spin up" button is pressed to start the line-scanning motor running. While it is accelerating, the controls can be set for the desired resolution and scan angles. Once everything is set, the "Scan" button is pressed. The panning motor moves to the beginning location, and data collection begins.

An effort was made to ensure that *every* rotation of the mirror provided useful data, in as much as possible. This has been successful in that the pan from one column to the next can usually complete prior to the scanning mirror reaching the beginning of the next column to scan. Thus, it is simple to compute the length of time required to scan an

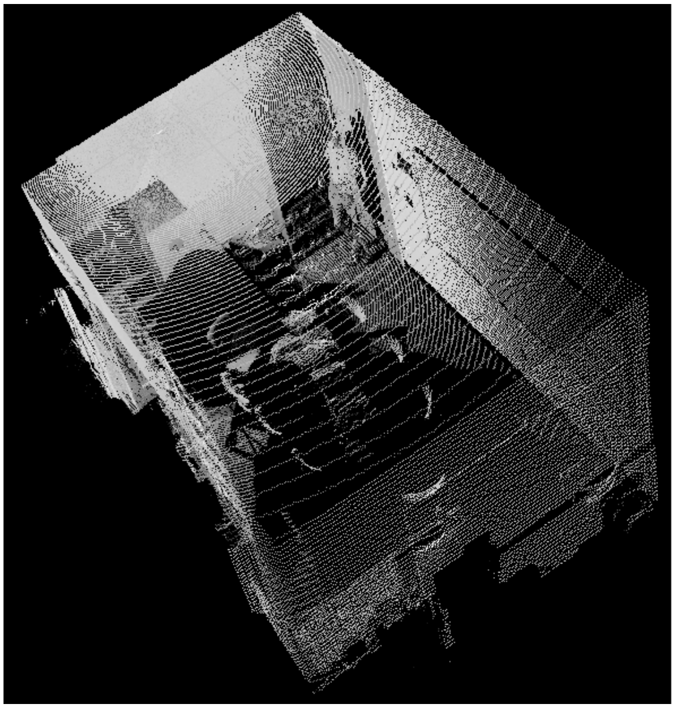

Fig. 10. View showing the entire MSL conference room. The scanning latitudes can be seen at the close end of the ceiling. Note the author standing in the far corner of the room.

Fig. 11. View showing a break in the ceiling where the counter-clockwise scans meet the clockwise scans, showing the seam where data is missing due to errors in the panning unit.

environment; it is the number of columns of data divided by the rotational speed of the scanner. For example, a scan with a resolution of 1000×1000 samples takes 125 seconds (2+ minutes) when the scanning motor is spinning at 8 Hz.

C. What Data Is Collected?

After an environment is measured according to the parameters, the user can select "Save" from the menu. Five files are written. They are:

1. An IVR file that contains some of the parameters of the scan (HFOV and VFOV). This file is in a format that allows panoramic plug-ins[1] to web-browsers[2] to immediately display (pan, zoom) the reflected-light data collected.
2. The reflected light image. This image is surprisingly good, simply showing the signal strength of the laser at each pixel. It has some strange characteristics since the projection is spherical, the light is infrared and the source of illumination is at the sensor (no shadows).
3. The ambient light image. Similar to above, but contains the level of light subtracted from the incoming signal to identify the AM laser signal. Since we have an IR filter, this signal is usually very dark. It is used to determine sensor-overload (if the sensor is aimed at the sun, for instance).
4. The distance image. This is written in 2 forms, one for visual inspection (scaled as a gray map from 0 to 255), and the other as the actual floating-point values.
5. The RTPI file. This file is the most useful, but most difficult to use. The above files all involve interpolation fitting the data on a regular grid, while this file contains the raw data. There is a header with the number of columns scanned, followed by an integer for each column indicating the number of valid readings in that column, followed by the range, theta, phi, and intensity values (all floats) for each reading. The range is in inches, theta (angle around azimuth) and phi (angle above horizon) are in degrees, and the intensity is from 0 to 255.

The RTPI data are not on a regular grid, their position is our best estimate of the actual location of the device when the sample was taken. This includes compensation for the panning motor (correcting theta), since it can only move to one of many discrete positions.

Samples of the reflected light and gray-map distance images are shown in figure 1. The majority of the other figures shown in this paper are from reprojections of the RTPI data converted to Cartesian coordinates.

VI. Sample Data Sets

In this section several data sets are shown reprojected into space and rendered as point clouds using the intensity of the reflected laser light to color the points. Shadows are very distinct since we render the background as black. There is a cone-like shape just below the laser where no data exists, this is where the data-collection platform is.

[1]Part of PhotoVista from LivePicture
[2]Netscape, Internet Explorer

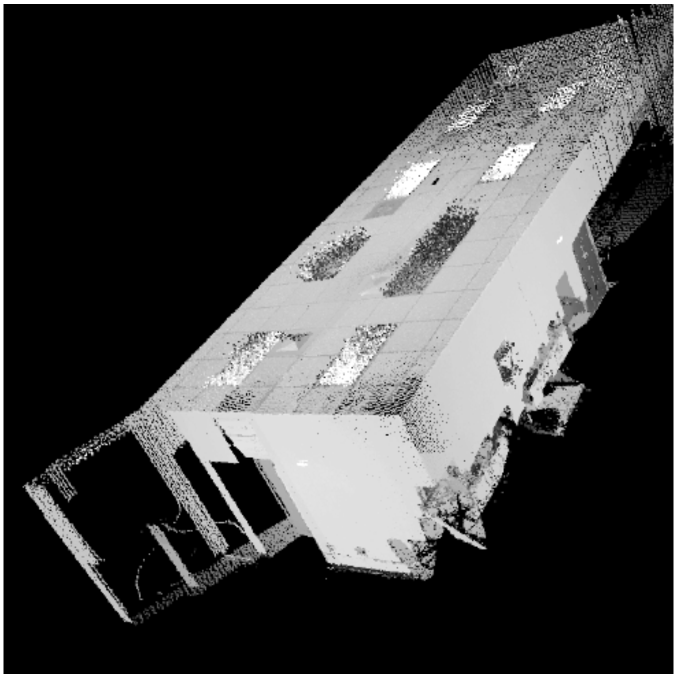

Fig. 12. A view showing the hardware teaching lab data from a far distance. The outcroppings in the lower left are places where the laser entered a closet.

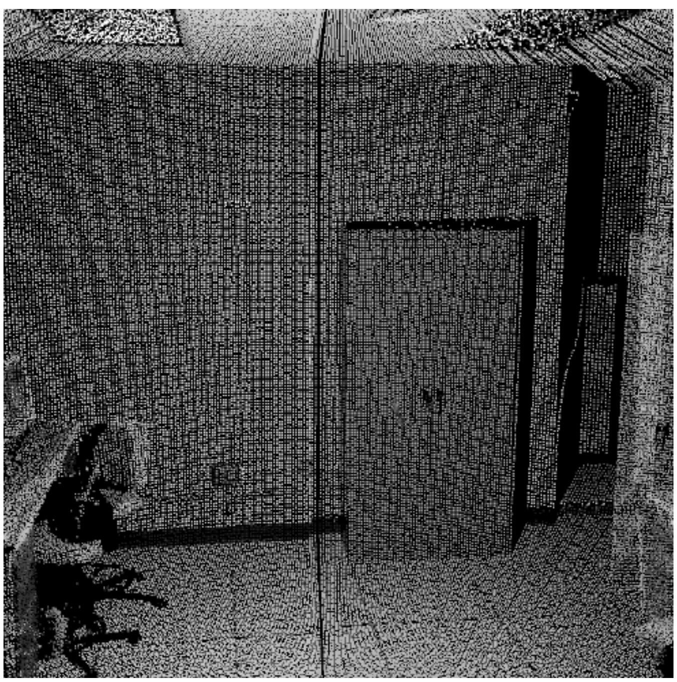

Fig. 13. This view shows a different view of the closets described in figure 12. This image is included to show how the rangefinder sometimes interpolates between distances as some portion of the beam falls on something near, while the remainder of the beam falls on a further surface. In the closet, on the right of the image, there is a trace that goes from a closer wall to an object inside the closet. Notice how it gradually moves from one to the other. This data is completely wrong, arising from interpolation by the rangefinding hardware. The same trace and another are visible in figure 12.

Fig. 14. A view towards the far wall. This image shows the maximum range of the rangefinder being set just a bit short, so the center of the far wall is within the maximum, but the corners are not. Also note the plant in the foreground.

Fig. 15. This view shows the electronics workbench area. The rangefinder's shadow is the large black area on the floor.

Fig. 16. Data collected while the scanning rangefinder was slowly pushed across the room. The position of the rangefinder is recorded when each column of data is collected. Note the irregular motion in the distorted ceiling tiles, walls and checkerboard patterns.

Fig. 17. Reprojection of the data collected in figure 16. Note how the ceiling tiles and checkerboard patterns are regular and the walls, tables and ceiling are flat.

A. Sitterson Hall, Room 239

In the first set of reprojected images (figures 1, 2, 4, and 3), we show the spherical panorama acquired and several reprojections. In Earth-like coordinates, the data was acquired at equiangular positions from 60° South, up across the North Pole and back down the other side to 60° South. Equiangular longitudes were taken, creating a spherical image.

This is the author's office which has been the primary data collection site during development. The walls have several protrusions, there are several tables, bookshelves, file cabinets and computer monitors in the room, adding to the complexity of the environment. The data from two panoramic sweeps (one is shown in figure 1) is reprojected and viewed from 2 different points of view to show the range data acquired in figures 2, and 3. Figure 4 is from the same environment, but from an earlier dataset.

B. MSL Conference Room

The next data set is from a small conference and coffee break room (shown in figures 10 and 11). It has several bookshelves, a white-board, a small cabinet with a sink and a meeting table in the center. The tabletop is dark plastic laminate, and returns very little data, thus only the viewable portions of the chairs around the table can be seen. The chromed light fixtures, despite their high reflectivity, seem to be visible to the rangefinder.

C. Hardware Teaching Lab

The largest room shown measured is our Hardware teaching lab, it is approximately $40 \times 20 \times 10\text{ft}^3$. Reprojections of two datasets are shown in figures 12, 14, 15, and 13. The lab consists of two areas, a lecture area with a table, chairs, and a white-board, along with a lab area where there are 4 electronic workbenches. These images show a plant, the effect of maximum range, and some erroneous data from the laser being split across two surfaces during acquisition (a close and far wall).

VII. Tracked Acquisition

An additional application for our rangefinder is data collection for multiple-center-of-projection (MCOP) images. Think of these as images where each column of data is acquired while the rangefinder is moving. If the position and orientation of the rangefinder is well known for each sample (or group of samples), then the scene can be properly reconstructed. This style of image allows close-up data collection where detail is necessary and quick data collection from afar when the data is not so important. See [3] for details on MCOP images.

To accommodate the MCOP project, we mount a hi-ball tracking device [1] on the rangefinder platform (shown in figure 5 on the left). The hi-ball is a device with 6 lenses looking towards a ceiling with sequentially flashing IR LEDs. The sensing of the flashes provides highly accurate position and orientation data.

The data collection proceeds by collecting each column with its position and orientation. Examples of the data collected and their reconstruction are shown in figures 16 and 17. This example shows that not all input must be in regular images, but can come from a variety of sampling techniques.

VIII. Coping with Real Data

A. Problematic Materials

Since the laser rangefinder works by reading its own reflected light, any material that does not reflect light back to its source is bound to be problematic. Some problematic materials are obvious: the metal legs on desks and chairs that are painted glossy black and shiny metal items that reflect light away, while ordinary walls reflect very well. Several surprising materials exist (black velvet reflects quite brightly back to the source). Plants, despite their green coloring, reflect IR light quite well.

B. Erroneous Data

The circuitry in this rangefinder reports "floaters" around specular items. The floating points tend to occur alone and can be removed with the following rule: a valid reading must be near at least one of its neighbors. This assumes that most of the data is valid, which it is. The floaters can be eliminated or replaced with a copy of a neighbor's position. We are also exploring estimations, since the environments scanned contain so many planar surfaces.

Another source of errors are specular objects, at the positions where the laser is reflected directly back on itself. The measurement for the highlights is closer than the surrounding area, thus the bright spots seem to hover off of the surface. While these errors have not been explored in great detail, it is our opinion that they stem from the strength of the signal returned, and can be compensated for.

IX. Conclusions and Future Directions

We have built a system that adequately measures environments. We can acquire precise, high-resolution range images consisting of 20–40 samples per degree. With these, we will be able to correspond range data to each pixel in color images of the same environment, allowing us to perform image-based rendering in these environments.

Several sources of error exist, but adequate methods of compensation have been developed. The result is that the range data thus far meets all of our expectations.

We are currently working on semi-automatic registration of range data taken from multiple locations in the same environment. The registered data sets will allow us to fill in shadows, providing a more complete set of data for an environment to support image-based rendering.

Acknowledgments

Special thanks to our image-based rendering research group, including John Thomas for building the cart, Paul Rademacher for his MCOP images, and everyone else for motivation and helping with error analysis.

We gratefully thank our funding sponsors at NSF and DARPA.

References

[1] UNC Tracker Group. Wide-area tracking navigation technology for head-mounted displays. http://www.cs.unc.edu/~tracker.

[2] Leonard McMillan and Gary Bishop. Plenoptic modeling: An image-based rendering system. In *Siggraph'95*, Los Angeles, CA, 1995.

[3] Paul Rademacher and Gary Bishop. Multiple-center-of-projection images. In *Siggraph'98*, Orlando, FL, July 1998.

Lars Nyland is a Research Associate Professor in the Computer Science Department at the University of North Carolina in Chapel Hill, NC. His current interests have to do with collecting and processing large sets of data with the highest possible performance. Additional interests all focus on high-performance computing, from building custom hardware to designing special-purpose programming languages (and compiling them). Nyland received his Ph.D. in 1991 at Duke University.