Capturing, Processing and Rendering Real-World Scenes

Lars S. Nyland, Anselmo Lastra, David K. McAllister, Voicu Popescu, Chris Mccue, and Henry Fuchs

University of North Carolina

Abstract

While photographs vividly capture a scene from a single viewpoint, it is our goal to capture a scene in such a way that a viewer can freely move to any viewpoint, just as he or she would in an actual scene. We have built a prototype system to quickly digitize a scene using a laser rangefinder and a high-resolution digital camera that accurately captures a panorama of high-resolution range and color information. With real-world scenes, we have provided data to fuel research in many areas, including representation, registration, data fusion, polygonization, rendering, simplification, and reillumination. The real-world scene data can be used for many purposes, including immersive environments, immersive training, re-engineering and engineering verification, renovation, crime-scene and accident capture and reconstruction, archaeology and historic preservation, sports and entertainment, surveillance, remote tourism and remote sales.

We will describe our acquisition system, the necessary processing to merge data from the multiple input devices and positions. We will also describe (and show) high quality rendering using the data we have collected. Issues about specific rendering accelerators and algorithms will also be presented. We will conclude by describing future uses and methods of collection for real-world scene data.

Figure 1. Examples of our data. The scene is Thomas Jefferson's library at Monticello, VA. The top image is a range map, where lighter is further (black is no data), the center image is the reflected laser light image, and the bottom image is the color panorama taken from the same point.

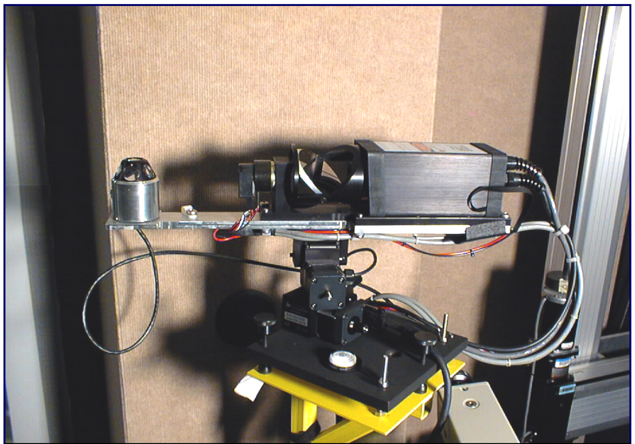

Figure 2. The laser rangefinder and scanning mirror mounted on an early panning system. Also included is the hi-ball tracker on the left end.

1 Introduction

While there are many scanning devices that capture the shape and color of objects [1-3], ours is the only portable system currently available that acquires the same type of data for entire scenes. The scanners from Cyra [1], Quantapoint [4], and Riegl [5] each acquire range data, but to date, they do not capture the color data. It has been our goal to explore the digitization of panoramic scenes, acquiring both shape and color information. Our original motivation was to support *image-based rendering* research at UNC [6, 7], but has grown to support many more research projects.

2 Data Acquisition System

In 1997, we decided to build our own system after failing to find any high-resolution, commercially available scene digitizers (we also considered many vision-based methods). Our prototype data collection system and its follow-on are described.

2.1 Prototype Data Acquisition System

Our range data is acquired using a line-scanning rangefinder from Acuity Research [8]. The rangefinder has suitable parameters to support our work: its effective range is 0 to 50 feet, its error is less than 1/2 inch, it takes up to 50,000 samples per second, and it has a high-speed interface for a legacy PC's (ISA bus slot required).

The device performs better (smaller error) if the sampling rate is reduced, the same is true if the maximum range is limited. We therefore typically run it no faster than 25,000 samples per second, and when possible, we limit the maximum distance to 24 -- 30 feet (providing a diameter of 48 -- 60 feet). This provides much better error characteristics than using the longest and fastest settings.

Coupled with the rangefinder is a scanning mirror whose drive motor has a 4096-position shaft encoder attached. Since our goal was to match the resolution of a high-quality digital camera, we required more than 4096 positions. Several interpolation schemes were explored to accurately estimate the position of the mirror resulting in much greater precision.

Supporting the rangefinder and scanning mirror is a panning motor. We've employed several, always looking for precision and strength. Our current system from Compumotor is more than satisfactory.

A single PC controls all of these devices. The rangefinder is controlled by a serial line and has a high-speed interface for receiving range and scanning motor information, while the panning motor is controlled using a simple serial line.

For color data, we use a Kodak/Canon DCS-520 35mm camera. It has a high quality CCD, can acquire color data up to 12 bits per channel, and has high-capacity disk drives that hold up to 120 images on a disk. It also has full control over the camera settings, allowing us to fix the focus and exposure settings. We take multiple exposures at different settings to produce higher dynamic range images through methods such as those described in [9].

We use a 14mm flat-projection lens. We felt this lens was necessary to acquire a wide field-of-view on the smaller-than-film CCD. The field-of-view is nearly the same as a 24mm lens on 35mm film, and is approximately 77 by 55 degrees. With a resolution of 1768 by 1152, this is approximately 1 milliradian pixel spacing in each direction.

2.2 Commercialization

The success of our system had led to the development of a commercial system that is much smaller, lighter, and more robust. 3rdTech (see www.3rdTech.com for more details) builds the DeltaSphere-3000, a system based on ours, and we currently use their data collection software. All of the complexity of power supplies, motor controllers, data cables, etc., are wrapped up in a small, portable package that weighs 22 lbs., and only requires external connections for 12 volts of power and an ethernet connection. Cooperative agreements allow us to benefit from their improvements (and vice-versa).

3 Data Collected

Color and range data are acquired sequentially from the same position by physically replacing the rangefinder with the camera on the panning motor. The details of the data are described in this section.

3.1 Rangefinder Data

The laser rangefinder produces samples that are not on a regular grid, since the motor that spins the scanning mirror varies slightly in speed. In addition, manufacturing errors distort the data slightly, so each sample must be undistorted, moving it to a more accurate position.

A single range reading includes not only the range, but also the strength of the returned laser light. The signal strength allows us to produce an image that is lit by the laser, and has no shadows, since the sensor is along illumination path. The top two images in figure 1 show panoramas of the range (as a gray map, where dark is closer) and the signal strength.

3.2 Camera Data

After the range data is acquired (in about 30 minutes), the camera replaces the rangefinder on the panning motor, and a panoramic set of photographs is taken. We typically take photos at 12 positions, spaced 30 degrees apart. In choosing this number, most portions of the scene are in two photographs, making subsequent registration simpler. We usually take several sets of bracketed exposures to obtain enough data to produce high-quality images.

3.3 Undistortion

The camera is calibrated once using public-domain software [10] to determine the intrinsic parameters. The parameters are used to undistort the images, using bicubic resampling, to create images that match a pinhole model. This is the only time the color data is resampled, in an effort to keep it as close to the original as possible.

The rangefinder data must also be undistorted. The inaccuracy stems from (at least) two sources: the panning motor and the 45 degree scanning mirror. All of the stepper motors we have used for panning do not match their published specifications--- most turn slightly less than one full revolution when asked to do so (even when backlash is accounted for). If the panning motor can move more than one full revolution, the laser can be used to measure this inaccuracy. Some panning systems are limited to less than one full revolution, so an alternate scheme was developed that simultaneously measured the inaccuracy of the panning motor and the mirror angle error simultaneously.

The experiment involves taking four measurements, as shown in figure 3. The panning motor is set at –90 degrees and the laser beam is aimed horizontally to the right. Its spot is marked on the wall and the distance is recorded. Then rotating the scanning mirror over the top, through straight up, and pointing to the left on the opposite wall, another spot is marked and its distance is taken. The panning motor is then set to +90 degrees, and the laser spot should coincide with the first spot. Its position is marked, and then the scanning mirror is again spun to point to the left, where it should coincide with the second spot. The distances between spots from spots 1 to 3 and 2 to 4 are measured, providing enough information to compute both the panning motor error and the mirror angle error as detailed in figure 3. Our typical values are 179.44 degrees panned with a mirror at 44.88 degrees. Appropriate corrections can be applied to the data after it has been taken, as the angular information is part of the dataset. Of course, once the errors are accurately determined, they can be applied during collection.

4 Merging and Simplifying the Data

The original goal of our system was to augment the photographs with range data. To meet this goal, we fuse the range and color data together. Additionally, panoramas taken from different locations must be registered. Simplification is also required, since the number of samples is too large to render interactively.

4.1 Registration of multiple panoramas.

When multiple scans are made from different locations (of a single scene), there is the problem of aligning one scan to the next. To do so, either the scanning positions can be very accurately measured, fiducials can be placed in the scene, or post-processing alignment can be done. We've chosen the last, as the first requires significant hardware that is not portable, while the second alters the scene.

We currently used the simplest of all methods for registration, a software tool that allows the user to move and rotate the scenes with respect to one another, showing the view in 3D. Thus far, this produces the best results in the shortest amount of time.

Still, we are always considering automated methods. We have developed an asymetric search strategy called the "Empty Space Registration Method" that aligns one set against another. It uses the knowledge that all the samples in the data represent the 1st surface struck by the rangefinder, so all the space along the ray is known to be empty between the rangefinder and the sample. What is beyond the sample is not known. The error metric then becomes the squared distance of a

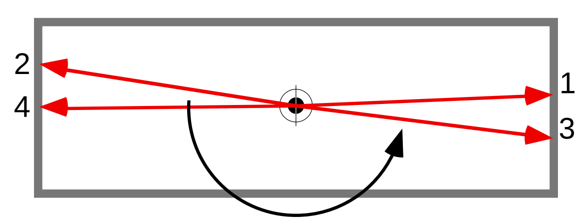

Figure 3. A diagram showing how to calibrate both the rotational error of the panning motor and the error of the 45-degree mirror. Readings are taken at positions 1—4, rotating the mirror first, then the panning motor, then the mirror again. Enough data is available to determine both values.

Figure 4. The ceiling over the workbench is highlighted with a pattern, indicating that the ceiling was the "strongest" planar area in a 3D-Hough transform of the range data. Finding and matching planar areas in different scans is one way of determining the rigid-body transformation between two scans.

sample in one set that is in space that is known to be empty in the other to the nearest unknown area. This search can work quite well, but often gets stuck (so it must be started from many locations), or leads to the inappropriate solution of completely separating the two sets.

Another method looks for large planar areas, and then trys to establish a correspondence between the planar areas in the different scans. The planes are found using a 3D Hough transform, which is a straightforward calculation since the rangefinder essentially does all the edge-finding for the second part of the transformation. Figure 4 shows an example of plane-finding by highlighting those areas in the photo that were on or very close to the most highly ranked planar area in the Hough tranform of the data. The walls and floor are reconized by the subsequent peaks in the Hough transform.

4.2 *Fusion of Color Images with Range Data*

Fusion is the process of using the range data to determine the depth of every pixel in the color images. The input to this process is the panoramic set of color images, the range data (both range and signal strength), and intrinsic camera parameters. A simple method for performing data fusion requires several correspondence points to be marked, and then applying a least-squares method to overlay one set of data on the other (which also involves reprojection). The labor required led us to an alternate, more automated version that searches for edge alignment between the two sets of data sets.

4.2.1 Edge Detection

Our automated method of fusing the data relies on the automatically detected edges in the different images. One difficulty is that CCD images are typically noisy, especially in the blue channel, making edge detection difficult. There are 2 approaches for alleviating this problem, and we have used both. One is to take multiple color images at the same exposure and then blend them together. With a digital camera, alignment is simple, and a simple average can be taken for each pixel.

The other alternative for dealing with the noise is blurring. To keep the salient edges but remove those caused by noise, we use a blurring method called *variable-conductance diffusion* [11] that treats smooth areas as conductors (using an electrical metaphor) and edges as insulators. The result of VCD provides edge detection algorithms with suitable input for finding actual edges in the scene, much better than standard blurring techniques.

The edges in the range images are far easier to detect, as there is little noise, and edges clearly stand out. Edges in the rangefinder reflectance images are also used, as these have more in common with the color images (but have different lighting).

4.2.2 Alignment Search

A search determines the three rotational angles needed to align the color and the range data. The other three degrees of freedom are already known, since the camera's nodal point is aligned with that of the rangefinder. The error metric is computed by aligning the edge images. The edge images are convolved with a kernel that provides a sharp peak at the edge centers, but has wide support so that moderately close solutions provide some reinforcement. The edge images are projected onto one another, and the edges from one image that overlap edges in the other are used to compute the goodness of the match.

Since there are many local minima, a grid in the search space is used as starting locations for a series of searches. The best at one resolution of the grid is kept, and then finer searches are done around that location. Figure 5 shows the results of a search by superimosing the edges on a color image.

Figure 5. Registration of depth with color. The edges extracted in the depth image, shown in yellow, are superimposed on the color image.

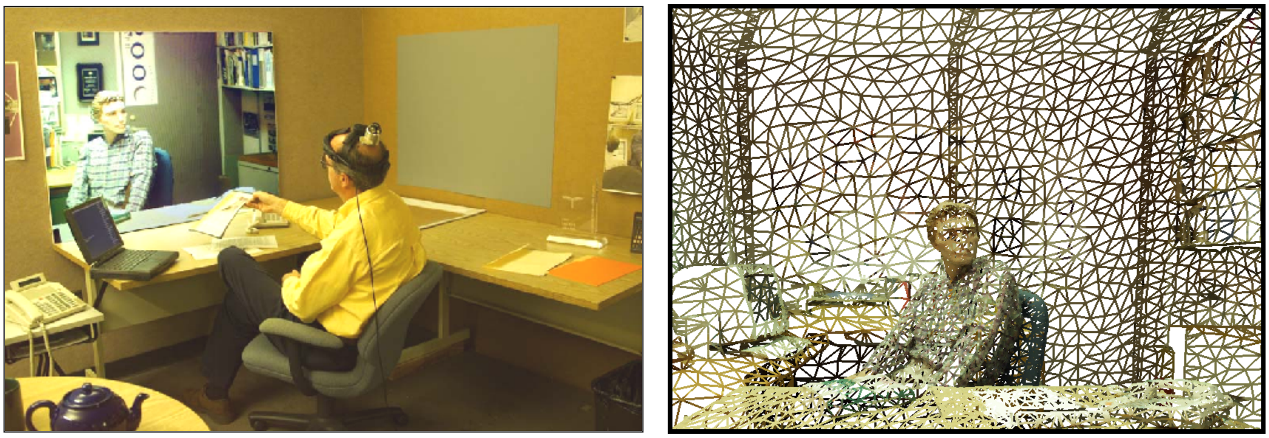

Figure 6. On the left is a view of the immersive tellecolaboration setup. The person in the office is viewing a stereo image of one of our 3D models of an office in our department. The person is tracked (device on head) to accurately update the display as he moves. The scene "through the window" is static in this image, but there are results with real-time reconstruction of participants over the network.

The image on the right shows the simplified model created from our data. The triangles are normally texture-mapped to provide realistic detail.

4.2.3 Discontinuity determination

Several methods for determining discontinuities have been tried. The simple methods yield comparable quality to more complicated methods.

The simplest triangulates adjacent pixels, creating breaks where the normal of the triangle is nearly perpendicular to the viewing ray. The cutoff angle is about 72 degrees, which, even if the surface is broken incorrectly, a surface at that angle is very poorly sampled, so using the data from this point of view is probably not a good idea anyway.

An alternate method yields the convex hull of the objects. It takes pairs of views, moving one to the viewpoint of the other. The rays in the stationary image can be used to determine inappropriately stretched surfaces.

4.2.4 Data Produced

One output of the fusion process is a set of IBR-Tiff images [12]. These are multi-layer tiff images that have the undistorted color as the first layer, the range (usually the disparity) as the second layer, followed by a 4x3 matrix as the third "layer" that relays viewpoint and orientation information.

Subsequent processing is usually performed to eliminate areas that are covered by more than one view. Our current choice is to keep the data from the view that is closest to orthogonal from the sampled area. We do this on a tile-by-tile basis, where a tile is a 32x32 pixel subregion of a larger image. Future plans include using all of the data to partially reconstruct view-dependent lighting, much like that described in [13]. The result is an *environment* file suitable for rendering. The process described in this section is fully detailed in [14].

5 Rendering

Current rendering hardware is optimized for texture-mapped triangles, not 3D point-samples in space. In this section, we describe rendering on conventional graphics accelerators, and our proposed architecture for rendering samples.

5.1 Rendering on currently available graphics accelerators

On computer systems that have OpenGL support, we use conventional triangle-strips to render our environments. The samples can be rendered as triangle meshes, colored points, or simplified texture-mapped triangles. One approach is to triangulate the image tiles, connecting samples in adjacent rows and columns with a mesh. The mesh is broken at large depth discontinuities to avoid drawing artificial surfaces that are almost aligned with the viewing rays. This is typically the fastest method of rendering, as it is the fastest path through the graphics accelerator. Unfortunately, it often has poor quality, due to improper triangulation and artifacts introduced by the triangulation (the choice of diagonal is often incorrect, yielding saw-tooth edges). Points often look much better, and are a better representation of the data, but GL points are fixed in size on the screen, so as the viewer moves closer to a surface, the surface breaks apart when the sampling distance exceeds the pixel spacing. We have explored alternatives on our PixelFlow system, using Voronoi primitives, with much better reconstruction. See [14] for details on customized rendering primitives.

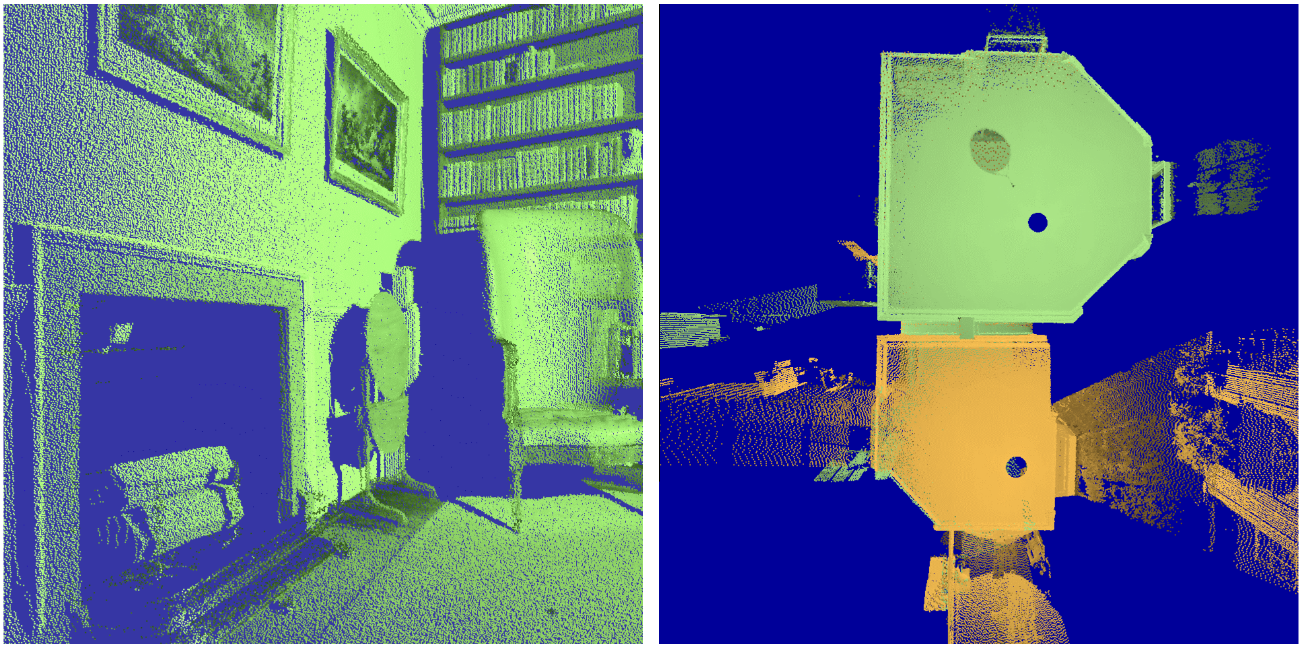

Figure 7. Jefferson's Library. The view on the left shows his reading chair and the nearby fireplace. The view on the right is from high above, showing 2 scans (the device is centered on the small holes). The room at the top of the image is the library and the room below it is the library annex. These 2 scans represent approxiamately 30 million range samples.

5.2 Our design for future graphics accelerators

An alternate approach to rendering IBR-tiff images is currently under development at UNC [15]. The WarpEngine is a hardware architecture specifically designed for rendering natural scenes from arbitrary viewpoints. Significant work is done to provide high-quality reconstruction to eliminate the problems typically found on conventional rendering hardware. The design is scalable, with a small system capable of driving a VGA-sized display, and a larger system, though still PC-sized, will display HDTV resolution images at 50 Hz. The WarpEngine is a single ASIC design, configured in such a way that several can be put together into a larger, more powerful system. It is also a very programmable SIMD system, allowing for experimentation in algorithms for displaying data from the real-world.

6 Example Data

As a demonstration of our work, several of our data sets and their uses are shown.

6.1 The Library at Monticello

In August, 2000, we had the opportunity to scan some rooms at Monticello, Jefferson's historic house. As a pilot project, we took nine scans in the library and the library annex. This project will allow us to build models of a premier example of American architecture, and will allow the historians to evaluate how this dense data will affect their goals in historic preservation. For instance, they suggested scanning the grounds, looking for subtle depressions in the ground that could be slave gravesites. They also suggested using the model to compute the illumination from candles and lamps, an operation they dare not perform due to the danger involved. Figure 1 shows the panoramic data from scan #2. Figure 7 shows the range data from two of the nine scans, with one view of Jefferson's reading chair, and the other view showing a plan view from high above. The different colors are from the two different views. This project shows the applicability of our work in historic preservation and archaeology.

6.2 A (Simulated) Automobile Accident

The next set of data comes from a simulated automobile accident, with a view shown in figure 8. The goal of this project was to demonstrate the ability to capture a temporally volatile scene of keen interest. Auto accidents are only one application of this sort; other applications are crime scenes and other accidents. These all have the characteristic that they need to be examined in detail, but they also need to be cleaned up quickly, so normal activities can resume.

6.3 Immersive Collaboration

One application that uses real-world data is immersive collaboration. There is a large research emphasis on this project at UNC in the "Office of the Future" project, with the larger goal of creating an immersive environment that will allow several remote users to work together. This is far grander than shared multimedia on a workstation in that life-sized images of people appear in front of each participant, with fully detailed environments around them. Additionally, computer-generated imagery, such as CAD models or MRI data, can be superimposed for collaborative work.

Currently, the real-world data collection system described here is used to acquire the static scenery, while other, more aggressive methods are used to capture the active participants.

Some images from this system are shown in figure 6.

6.4 *As-built Verification*

Manufacturing and processing plants are built from many complex systems that must all fit and function together. The scene shown in figure 9 is from an emissions control system being built at NIST in Gaithersburg, MD [16]. Their use of this kind of 3D data is to ensure, perhaps on a daily basis, that what was designed to be built is, and what can't to be detected as early as possible so that other physical systems can be adjusted, or that planning can be accommodated. Path-planning is also important in this area in that it is important to know that a new device can actually be moved into its destination location without impediments along the way.

6.5 *Relief Texture Maps*

An alternate use for rendering this kind of data is less radical than that proposed in section 5, and relies upon what current graphics accelerators already do well. The idea of *relief texture maps* [17] is that a standard texture map is augmented with a per-pixel depth, represented as an 8-bit value (perhaps replacing the alpha-channel). At rendering time, the texture can be warped appropriate to the desired viewpoint. A full warp can be decomposed into a standard texture-mapping operation, preceded by a pre-warp step. The pre-warp step is simple compared to the computation required for texture mapping with mip-maps, so it could easily be added to a conventional graphics pipeline. The additional benefit of relief textures is that the number of triangles is drastically reduced for a model. Figure 10 shows 2 views of a model with only 10 triangles rendered with relief textures. To achieve the same level of detail with conventional triangles, thousands would be required.

7 Conclusions

It has been exciting to produce realistic data that has been enthusiastically used by several of our research groups. While we have successfully captured many scenes, there are several research and engineering issues still to be solved.

7.1 *Solutions demonstrated*

Our biggest success is in building a system that captures scenes of the real world in high detail. Every book on the shelf, every paper on the desk, even the phone and its cord are part of our models. The color is as good as electronic devices can capture, and the scenes are far more compelling than hand-built models.

A key to this success is the data fusion operation. Placing accurate range data on color images in the proper location is difficult, but when done properly, yields excellent results. Examples beyond those shown

Figure 8. A view from the side of our simulated car wreck data. The acquisition device was centered in the circular area with no data on the ground.

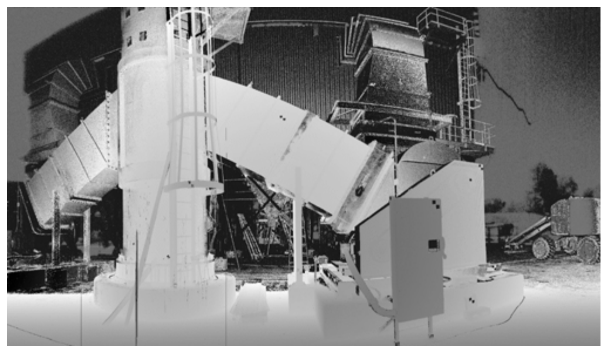

Figure 9. A view of the air scrubber under construction at NIST. Data was taken during the day in bright daylight (data courtesy of 3rdTech).

Figure 10. Two views of a library represented by 10 triangles with relief texture maps. The per-pixel relief allows proper motion and occlusion of the scene.

in this paper are available at our IBR web site, www.cs.unc.edu/~ibr (as is a color, electronic version of this paper).

Substantial work has been done calibrating the rangefinder and the associated positioning system. This includes building models to explain deviations from a perfectly built system, and the mathematics to correct for the errors. It also includes the interpolation method for the shaft encoder where we require resolution that is at least 10x the number of positions on the shaft.

7.2 Just the beginning

There are many problems yet to be solved, and which we believe will yield useful research results.

7.2.1 Simple, Geometric Models

There are very few people interested in the point samples that are produced by our system. What most people would like is a simple geometric model that not only has low polygon count, but might also be properly segmented to allow, say, the separation of one car from the other in an automobile accident. Many research groups and companies are exploring this problem, and our data has added a new facet to the problem by providing inside-looking-out data.

7.2.2 Real-time

Our scans take less than an hour to acquire. During that time, we try to ensure that nothing moves and the lighting stays the same. These restrictions are difficult to enforce.

As a research goal, we are seeking faster methods of acquiring range (and coupled color) data. Promising methods have been developed at Sandia National Labs, and our future work may be an outgrowth of their work [18].

7.2.3 Range-to-Range Registration

Aligning multiple scans is a difficult problem. To date, we have tried four methods, and find it dissatisfying that the simplest method of pushing and rotating the model seems to work the best. Among the other methods we have explored are a corresponding plane method, a method called the "Empty Space Registration Method" and some methods based on a Hough transform of the data to find the large planar areas. Iterated Closest Point has also been explored, but this method tends to get stuck (it is best used as a refinement tool, not an overall solution). This problem seems to have all that is necessary to find the answer, yet a practical solution still eludes us.

7.2.4 View-dependent Effects

Since we capture the appearance of objects from several locations, we have an opportunity to use the multiple views to build view-dependent renderings. Simple examples are reflective highlights, such as those created from overhead lights on reflective furniture. Reflections that are "glued down" instantly break the immersive feeling, and are noticed as soon as the viewer moves.

Similarly, caustic reflections in the lenses need to be removed from the scene. The interreflection of bright lights within the lens elements currently are glued onto whatever was behind them in the scene. For example, in the auto accident scene, several lights were placed around the autos, and the caustics produced ended up on the sides of the building behind. Viewers find this inconsistent with their notion of the real world.

7.3 Our goals for the future

In the short term, we are exploring drastic simplification of the models created, along with efficient representation. A single scene currently swamps most computing systems, not only in terms of rendering, but in memory and disk storage as well. The task of acquiring all of the rooms at Monticello is currently daunting, but an obvious and desirable goal. Thus the goal is in extracting the desired information from the data acquired, and reducing the computational and storage requirements.

Better automatic registration is also an important goal. Imagine our device mounted on a mobile platform, with sufficiently smart software (and battery power) to wander about a building acquiring data. The registration of the different panoramas could have a rough guess (the platform needs to navigate and have a rough idea of where it is), but would certainly have errors of several inches and several degrees of rotation.

Automatic acquisition demands better planning of acquisition sites. This is another area of research studied in object data collection [19], but has not been tackled for devices that look to the outside.

8 References

[1] Cyra, "The Cyrax 3D Laser Imaging System," http://www.cyra.com, 1999.

[2] Minolta, "Vivid 700 non-contact digitizer," http://www.minoltausa.com/vivid/specs.html, 2000.

[3] Cyberware, www.cyberware.com, 2000.

[4] "Quantapoint (formerly K2T Home Page)," http://www.quantapoint.com, 1999.

[5] Riegl, "Riegl USA," http://www.rieglusa.com/, 2000.

[6] "UNC IBR Home Page," http://www.cs.unc.edu/~ibr/, 2000.

[7] L. McMillan, "An Image-based Approach to Three-Dimensional Computer Graphics," : University of North Carolina at Chapel Hill, 1997.

[8] Acuity Research, "Acuity Research Scanners and Sensors," http://www.acuityresearch.com/, 1998.

[9] P. E. Debevec and J. Malik, "Recovering High Dynamic Range Radiance Maps From Photographs," presented at SIGGRAPH '97, Los Angeles, 1997.

[10] R. Y. Tsai, "An Efficient and Accurate Camera Calibration Technique for 3D Machine Vision," presented at Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, Miami Beach, FL, 1986.

[11] T. S. Yoo and J. M. Coggins, "Using Statistical Pattern Recognition Techniques to Control Variable Conductance Diffusion," UNC-CH Computer Science, Technical TR94-047, 1994.

[12] M. Oliveira, "IBR TIFF File Format," http://www.cs.unc.edu/~iflow/ibr_tifflib/README, 1997.

[13] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs," presented at Siggraph99, Los Angeles, 1999.

[14] D. K. McAllister, L. Nyland, V. Popescu, A. Lastra, and C. McCue, "Real-Time Rendering of Real-World Environments," presented at Rendering Techniques '99, Proceedings of the Eurographics Workshop on Rendering, Granada, Spain, 1999.

[15] V. Popescu, J. Eyles, A. Lastra, J. Steinhurst, N. England, and L. Nyland, "The WarpEngine: An Architecture for the post-polygonal age," presented at SIGGRAPH00, New Orleans, 2000.

[16] NIST, "NIST Fire Research Facility Emissions Control System (ECS)," http://cic.nist.gov/vrml/ecs.html, 2000.

[17] M. M. Oliveira, G. Bishop, and D. McAllister, "Relief Texture Mapping," presented at SIGGRAPH 2000, New Orleans, La, 2000.

[18] J. Sackos, B. Bradley, B. Nellums, and C. Diegert, "The Emerging Versiatility of a Scannerless Range Imager," presented at Laser Radar Technology and Applicatiosn, Orlando, FL, 1996.

[19] R. Pito, "A Solution to the Next Best View Problem for Automated Surface Acquisition," *IEEE PAMI*, vol. 21, pp. 1016-1030, 1999.