Interactive Photorealistic Inside-Looking-Out Automated 3-D Modeling

Voicu Popescu
Purdue University
West-Lafayette, IN
popescu@cs.purdue.edu

Mihai Mudure
Google Inc.
Mountain View, CA
mmudure@gmail.com

ABSTRACT

Digital 3-D models of real world scenes are important in many applications in defense and beyond. Constructing 3-D models that faithfully capture the complexity of the real world is a difficult problem. A promising approach is automated 3-D modeling based on directly recording the geometry and color of the scene. Current systems suffer of important disadvantages such as inadequate scene coverage due to slow acquisition, lack of immediate feedback on the quality of the acquired 3-D model, restriction to small scenes, and unintuitive operation.

In this paper we give an overview of the ModelCamera automated indoor 3-D modeling project at Purdue University, and we present novel research results and their implementation. The ModelCamera is designed to handle the challenging "inside-looking-out" 3-D modeling case of large, room-sized scenes. The system is interactive, it provides immediate feedback to the operator, it is efficient, and the resulting 3-D model supports rendering the 3-D scene photorealistically at interactive rates, as needed by applications such as virtual training.

ABOUT THE AUTHORS

Voicu Popescu is an associate professor of computer science at Purdue University. His research interests span the fields of computer graphics, visualization, and computer vision. Current projects focus on large scale automated 3-D modeling, on high-quality visualization of finite element analysis simulation data, and on effective distance learning. He received a computer science Ph.D. degree from the University of North Carolina at Chapel Hill in 2001. He is the author of over 50 peer reviewed publications.

Mihai Mudure is a research scientist with Google Inc. His research interests include automated 3-D modeling of real-world scenes for computer graphics applications. He received a computer science Ph.D. degree from Purdue University in 2008.

Interactive Photorealistic Inside-Looking-Out Automated 3-D Modeling

Voicu Popescu
Purdue University
West-Lafayette, IN
popescu@cs.purdue.edu

Mihai Mudure
Google Inc.
Mountain View, CA
mmudure@gmail.com

INTRODUCTION

Digital 3-D models of real world scenes are essential infrastructure for applications in a variety of domains, including science, engineering, defense, law enforcement, education, and entertainment.

Constructing such 3-D models remains challenging. The approach of manual modeling using CAD-like software is time consuming and fails to capture the full complexity of real world scenes. For many applications, including virtual training for defense, the realism of the digital environment is essential. Effective training requires that the real-world scene be captured with high fidelity. The experience of exploring the digital environment has to be believable in order to place the trainee in conditions that closely match the real world. A promising alternative to manual modeling is automated modeling based on recording depth and color in the scene. An operator scans the scene with an acquisition device and the depth and color data are combined into a photorealistic 3-D model.

However, current automated modeling approaches suffer from important disadvantages. Depth acquisition is slow, which limits the number of acquisition viewpoints, resulting in inadequate scene coverage. The problem is particularly severe for large scenes such as rooms which exhibit complex occlusion patterns. Most current automated modeling systems are limited to small one or a few object scenes. We call this scenario outside-looking-in modeling, since the acquisition device and the operator are outside the scene to be modeled. Inside-looking-out modeling, where a large scene is acquired from within, is considerably more challenging due to the sheer size of the scene and difficult accessibility. Model construction is performed off-line and off-site, so it is difficult to address any problems with data acquisition. Finally, modeling systems are complex, implying great hardware expenses and the need for expert operators.

In this paper we give an overview of the ModelCamera project at Purdue University, we present recent research results and their implementation, and we discuss the suitability of ModelCamera technology in support of defense applications.

The ModelCamera is an inside-looking-out automated 3-D modeling system for building interiors. The system is interactive—the operator sweeps the scene with an acquisition device and the depth and color data acquired is integrated in real time into a quality 3-D model constructed incrementally. The system provides immediate feedback to the operator—the evolving model is visualized continually on a nearby monitor so the operator can guide acquisition ensuring a high-quality model. The system is efficient—a team of two operators with a single acquisition device can acquire a building with 100 rooms in a few days. The resulting 3-D model is compact yet it allows rendering the 3-D scene photorealistically at interactive rates, as needed by applications such as virtual training.

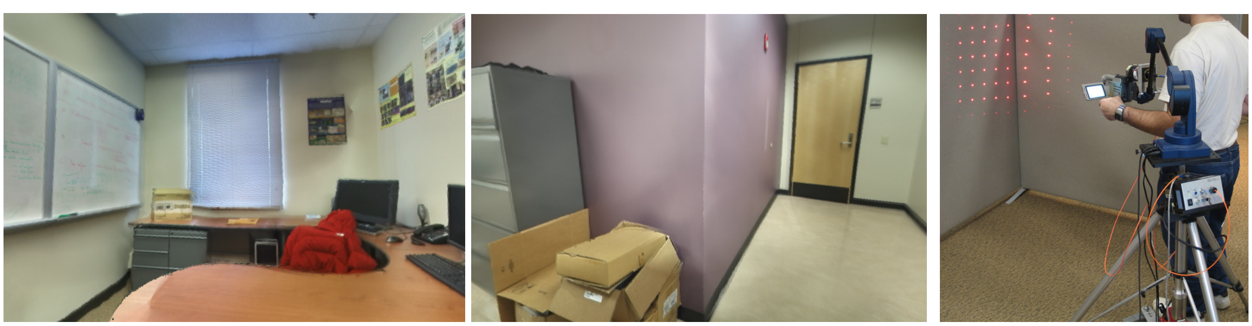

Figure 1. Images Rendered from Models Acquired with our System and Acquisition Device.

Figure 1 shows sample photorealistic images rendered at interactive rates based on 3-D models of office and hallway acquired with our system (*left and middle*) and our acquisition device mounted on a tripod for inside-looking-out acquisition (*right*). The office model was acquired in 31 minutes and consists of 342,000 depth samples (i.e. 3-D points) and 380 reference color frames. The hallway model comprises 260,000 depth samples and 446 color frames, and was acquired in 21 minutes.

The remainder of this paper is organized as follows. The next section discusses prior work. Section 3 gives an overview of the interactive modeling pipeline implemented by the ModelCamera. Section 4 reviews prior ModelCamera research results and their implementation in earlier systems upon which the current work builds. Section 5 discusses a novel ModelCamera system for inside-looking-out dense-viewpoint automated 3-D modeling. Section 6 presents results. Section 7 concludes the paper by sketching directions for future work including possible applications in defense.

PRIOR WORK

We give a brief discussion of prior work structured according to the method employed for recovering the geometry of the 3-D scene through depth acquisition.

Color Only Methods

Some automated modeling techniques rely entirely on color and bypass the depth acquisition or geometric modeling step altogether. Color panoramas (Chen, 1995) are 2D ray databases constructed from images that share a center of projection. Panoramas have the advantage of fast and inexpensive acquisition and support photorealistic rendering, but they restrict the user to a single viewpoint.

Light fields and their variants (Levoy 1996, Gortler 1996) are 4D ray databases that allow a scene to be viewed from anywhere within the scene field. Light fields support view-dependent effect and remain the only technique for rendering extremely complex materials such as feathers or fur (Matusik et al, 2002). However capturing and constructing light fields is challenging and the database is impractically large for sizeable scenes.

Manual Methods

User guided techniques rely on the user to manually specify the geometric models, either directly by using a 3-D modeling software tool such as AutoCad, 3dsmax or Maya, or indirectly by leveraging image or geometric constraints (Debevec et al., 1996, Gibson et al., 2002, Hidalgo and Hibbold, 2002). The user's knowledge scene is used to maximize the expressivity of the. However these methods oversimplify complex geometry, giving the resulting models an artificially clean appearance and failing to capture the details of a real scene.

Dense Depth Methods

Systems that rely on dense depth acquisition can be used to produce high quality models of complex scenes. Such examples include the digitization of Michelangelo's statues (Levoy et al., 2000), the digitization of a large model of imperial Rome (Guidi, 2005), Jefferson's Monticello (Williams, 2003), the Parthenon (Stumpfel, 2003), cultural treasures of ancient Egypt (Farouk, 2003), or the ancient city of Sagalassos (Pollefeys, 2001).

However dense depth systems have the disadvantage of long per view acquisition times which limits the number of acquisition viewpoints. The operator has little or no control over the acquisition process. Due to the delay between acquisition and the time when the model can be inspected, addressing problems is usually impractical due to the cost of returning to the scene.

Interactive Depth Methods

We are not the only researchers to notice that many if not all of the challenges of dense-depth acquisition can be overcome by employing an interactive modeling pipeline. Several interactive 3-D modeling systems have been developed. One such system employs a laser stripe triangulation rangefinder coupled with an electromagnetic (Fisher et al., 1996) or mechanical tracker (Hilton et al., 2000). However, the laser stripe degrades quickly with the distance to the acquired surface, which requires that the acquisition device be kept close to the scanned surface or to use laser power levels that are not eye safe. Complex geometry is also a challenge for laser stripe approaches due to laser scattering and secondary reflections.

A different approach for achieving interactive acquisition from many viewpoints is to move the targeted scene in the field of view of a fixed scanner using a rotating platter (Terauchi et al., 2005), or by relying on the operator to rotate the object (Rusinkiewicz et al., 2002). Such systems are object scanners and scale poorly with the size of the scene. Our ModelCamera project has researched and implemented several interactive 3-D modeling techniques, which we describe in detail in Section 4.

SPARSE-DEPTH DENSE-COLOR INTERACTIVE 3-D MODELING—PIPELINE OVERVIEW

The most challenging task in automated modeling is recording the scene geometry through depth acquisition. As discussed, current dense depth methods are based on acquiring detailed depth maps of the scene one viewpoint at the time. Acquiring dense depth is a lengthy process, which requires searching for correspondences in depth from stereo, or sequential scanning in laser rangefinding. The long per viewpoint acquisition time limits the number of viewpoints, leading to poor scene coverage, and prevents an interactive 3-D modeling pipeline.

Instead of this conventional dense-depth sparse-viewpoint (DDSV) scanning, our ModelCamera automated modeling approach relies on sparse depth acquisition from thousands of viewpoints. Sparse depth can be acquired quickly and robustly for each viewpoint, which enables an interactive pipeline. In such sparse-depth dense-viewpoint (SDDV) automated modeling, depth samples quickly accumulate over the many viewpoints yielding great modeling power. The operator benefits of immediate visual feedback which allows him to avoid over-scanning scene sections with simple geometry and to concentrate the depth acquisition effort on complex sections.

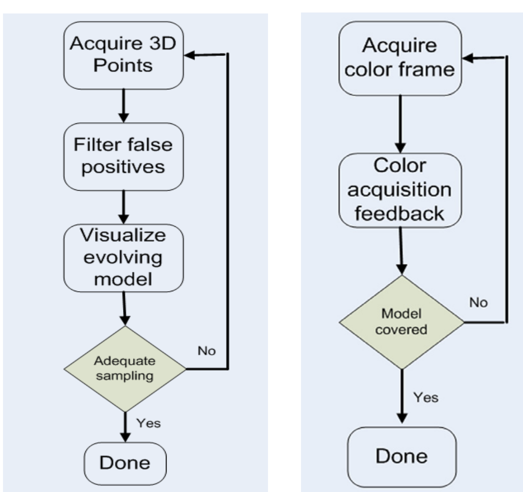

Figure 2. Interactive Depth (*left*) and Color (*right*) Acquisition Pipelines.

The SDDV automated 3-D modeling approach implemented by the ModelCamera system separates depth from color acquisition. The ModelCamera acquisition device acquires depth based on structured light, which modifies the scene lighting and cannot proceed simultaneously with high-quality color acquisition. The depth and color acquisition pipelines are shown in Figure 2. The operator relies on immediate visual feedback to decide whether the scene geometry was captured sufficiently well and whether there are enough color frames to color (i.e. cover) the scene geometry.

PRIOR MODELCAMERA RESEARCH RESULTS AND THEIR IMPLEMENTATION

The ModelCamera automated indoor 3-D modeling project started at the Computer Science Department of Purdue University in 2001. Four important research and implementation milestones have been reached:

1. Development of device that acquires quality color and sparse depth robustly, at interactive rates, and that can be easily maneuvered by the operator in real time (Popescu et al., 2003, Bahmutov et al., 2004).
2. Interactive single-viewpoint inside-looking-out 3-D modeling (Bahmutov et al., 2004, Bahmutov et al., 2006).
3. Interactive SDDV outside-looking-in 3-D modeling (Mudure and Popescu, 2008).
4. Interactive SDDV inside-looking-out 3-D modeling (this paper).

The following three subsections briefly describe the first three milestones, for completeness. The next section describes the fourth.

Acquisition Device

We developed an acquisition device that consists of a video camera and a laser system attached to the camera

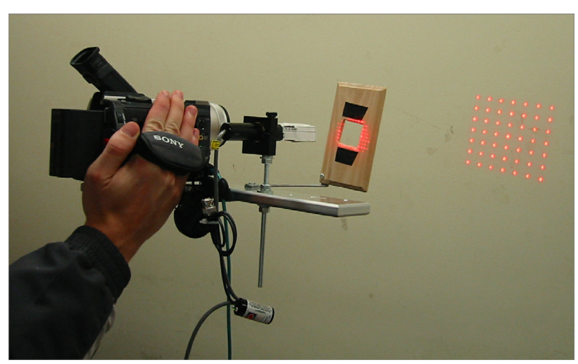

Figure 3. ModelCamera Acquisition Device.

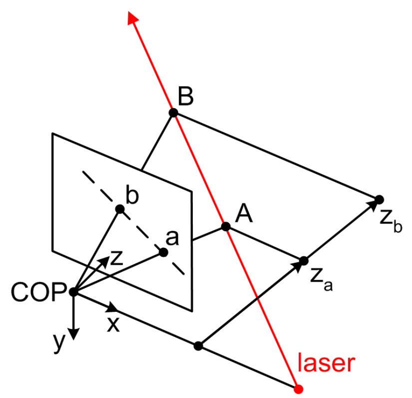

Figure 4. Depth Acquisition by Triangulation.

(Figure 3). The laser system projects a 7x7 pattern of laser dots into the field of view of the camera. A laser beam creates a bright dot into the video frame which is detected and triangulated into a depth sample. In Figure 4, if the dot is detected at image point *a*, the 3-D point *A* is computed as the intersection of the camera ray through *a* and the known (i.e. calibrated) laser beam. The baseline of 20 cm between provides a useful range of 50 cm to 400 cm and an average accuracy of 2mm for a target at 100 cm. When the lasers are turned off, the video camera acquires quality color frames (3 CCD's, progressive scan).

By design the laser dots move on disjoint epipolar segments into the field of view of the camera which makes depth detection efficient and robust. Figure 5 shows a fragment of a video frame acquired in a dark room with the lasers turned on. Each laser beam projects to a disjoint video frame segment, illustrated by the colored "epipolar" lines. Dots are located along epipolar lines by searching for an intensity peak in 1-D, followed by additional 2-D circularity tests. The resulting device acquires robustly 720x480 color and 7x7 depth samples per frame, at 15Hz.

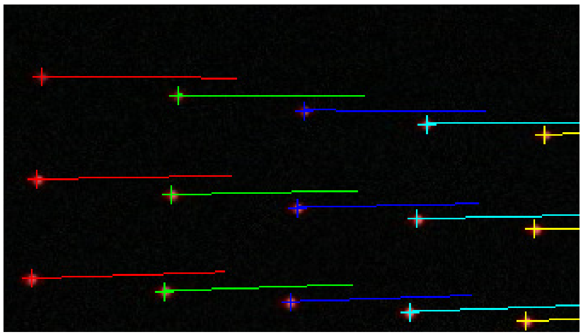

Figure 5. Laser Beams with Non-Intersecting Video Frame Projections.

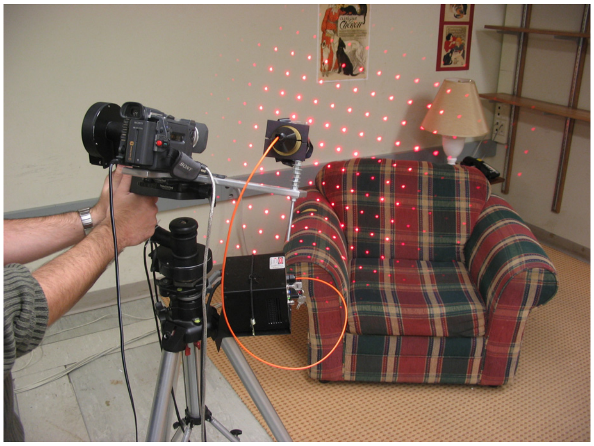

Figure 6. Parallax-Free Pan-Tilt Acquisition.

Interactive Single-Viewpoint Inside-Looking-Out 3-D Modeling

The first step towards the goal of efficient photorealistic 3-D modeling of large-scale indoor scenes was to restrict the motion of the acquisition device to two rotational degrees of freedom. A parallax-free pan-tilt bracket mounted on a tripod allows the operator to freely acquire depth and color from the center of a room. Figure 6 shows an acquisition device with a more powerful laser system consisting of a laser source (black box) and laser head connected with fiber optics (orange wire). The acquisition device is panned and tilted by the operator to sample the room geometry.

The resulting depth and color data are combined at interactive rates into a photorealistic model as follows:

- the color frames are registered and stitched together into a color panorama implemented as a cube map;
- the depth samples are Delaunay triangulated in 2-D on the faces of the cube maps;
- the connectivity so inferred is applied to the 3-D points producing a 3-D triangle mesh;
- the 3-D triangle mesh texture mapped with the color cube map, called a depth enhanced panorama or DEP, provides a photorealistic model of a room.

The DEP shown in Figure 7 was acquired in 30 minutes and comprises 100,000 triangles. Several DEPs can be combined leveraging same-plane constraints (i.e. same floor, ceiling, side wall) to model large-scale indoor scenes. Figure 8 shows a building model (6 floors, 20 rooms, 1,400 m^2) that was acquired in 40 hours by a two-member team using a single acquisition device.

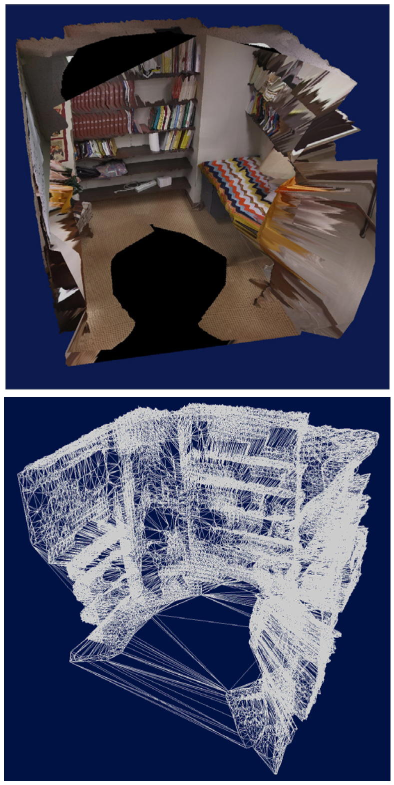

Figure 7. Depth Enhanced Panorama Texture-Mapped (*top*) and Wire-Frame (*bottom*).

A DEP samples the scene from a single viewpoint, which simplifies registration—frames are stitched together using only the color data—and model construction—depth samples are triangulated in 2-D in the domain defined by the pan and tilt angles,

Figure 8. Multistory Building Modeled with Multiple DEPs.

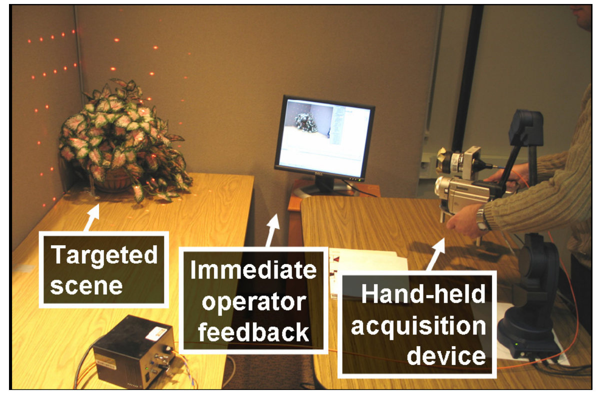

Figure 9. Outside-Looking-In SDDV Interactive 3-D Modeling.

bypassing the problematic triangulation in 3-D. However, the single acquisition viewpoint also implies the shortcoming of only storing color and depth for surfaces visible from the center of the panorama. A DEP provides a quality visual depiction of the scene it models while the desired viewpoint remains close to the acquisition viewpoint. Large viewpoint translations reveal surfaces not visible from the acquisition viewpoint, which produces the "stretching" artifacts visible on the coat hanger in the lower right corner of the top image in Figure 7. This fundamental limitation can only be addressed by dense viewpoint acquisition.

Interactive SDDV Outside-Looking-In 3-D Modeling

The first step towards dense viewpoint acquisition was to limit the scene size to a 0.5m x 0.5m x 0.5m cube, which can be modeled in the outside-looking-in configuration (Figure 9). Registration is provided by a 6 degree-of-freedom mechanical arm.

Depth is acquired by freely moving the acquisition around the scene. 3-D points quickly accumulate to overcome complex occlusions such as those created by the many small leafs of the plant in Figure 9. The points are visualized on a nearby monitor, so the operator can easily ensure adequate scene coverage. In Figure 10 the acquisition path is visualized with a color changing from red to yellow, revealing that the operator decided to take a second pass at the scene, increasing the depth sampling rate. This possibility of adding geometric detail at linear cost is an important advantage of SDDV versus DDSV acquisition. In the case of dense depth maps (DDSV), additional viewpoints bring diminishing returns at the same high

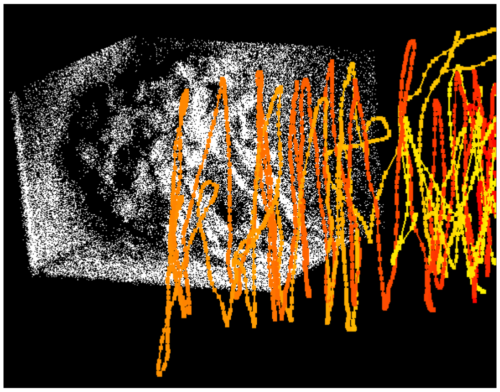

Figure 10. Visualization of 3-D Points and Dense Viewpoint Acquisition Path.

initial cost. The acquisition path in Figure 10 comprises 3,684 viewpoints.

Color frames are acquired in a second pass. All color frames see the entire scene volume, which is an advantage of outside-looking-in scanning of small scenes. To ensure adequate view-direction sampling, the operator is guided by visual feedback (Figure 11). The resulting model enables photorealistic rendering of complex scenes (Figure 12), which are difficult to model and render with prior approaches.

INTERACTIVE SDDV INSIDE-LOOKING-OUT 3-D MODELING

Acquisition

In order to move to inside-looking-out modeling suitable for large scenes, the first problem that needed to be addressed was to place the acquisition device—with the mechanical tracking arm—on a tripod, which

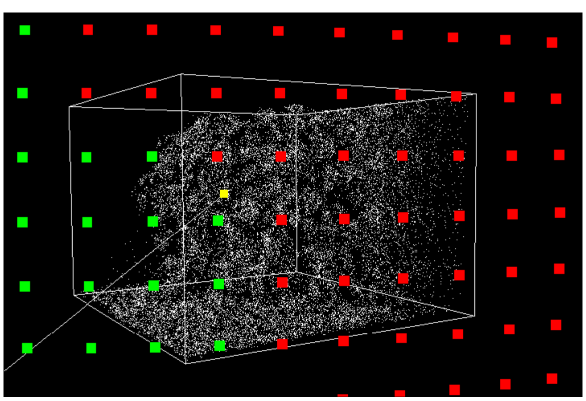

Figure 11. Feedback during Outside-Looking-In Color Acquisition.

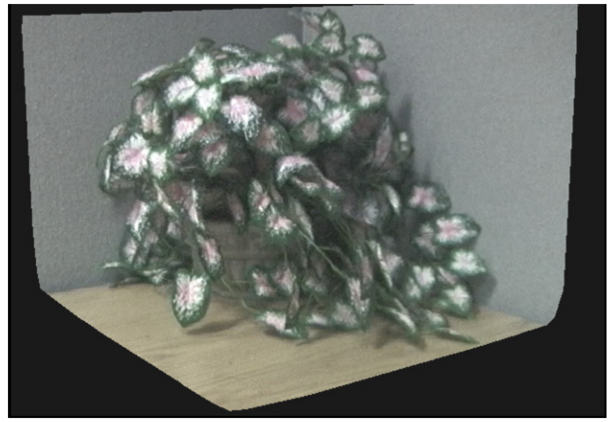

Figure 12. Image Rendered from Plant Model.

allows scanning a room from its center, in inside-looking-out fashion. A special bracket holds the mechanical arm and also provides a convenient way of parking the acquisition device for transfer from one room to the next (Figure 13).

The second problem introduced by the inside-looking-out operation was view selection for the visual feedback during depth acquisition. Whereas before the view was fixed on the 0.5m cube containing the scene, a dynamic viewpoint is required now. We found through experiments that operators preferred the feedback viewpoint to be placed behind the video camera and slightly higher. Figure 14 shows two visual feedback modes available to the operator during depth acquisition.

A third problem pertains to color acquisition. Whereas before all frames saw the entire scene, and one could focus on sampling the space of view directions, inside-looking-out modeling requires tiling the scene with

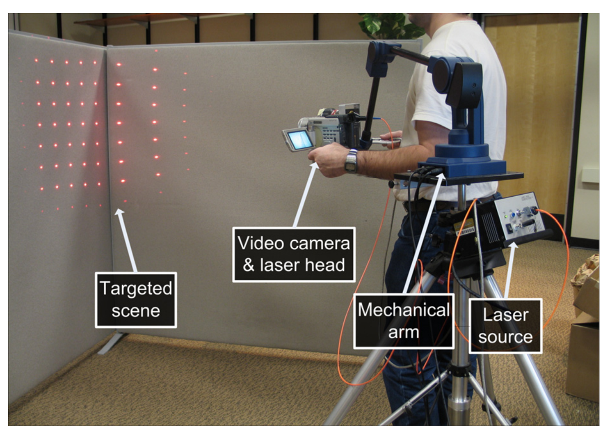

Figure 13. Inside-Looking-Out SDDV System.

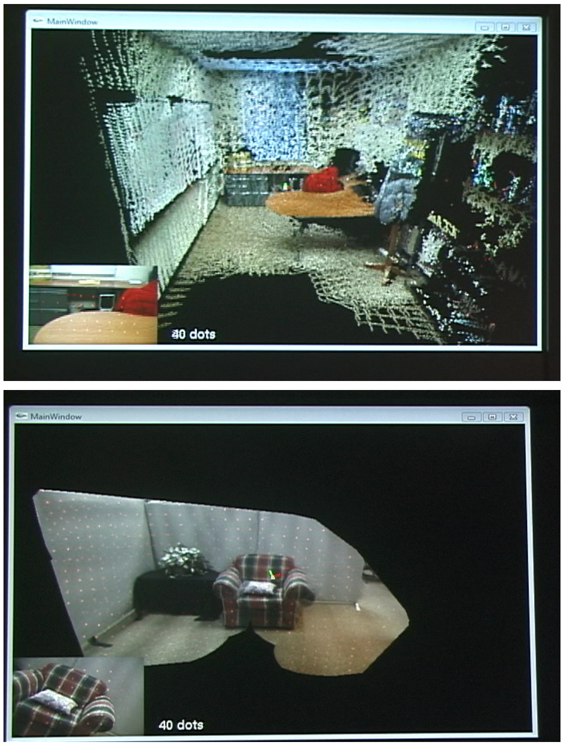

Figure 14. Colored Point Cloud (*top*) and Texture-Mapped Mesh (*bottom*) Visual Feedback during Depth Acquisition.

color frames. Figure 15 shows how the visual feedback highlighting in green the points that are colored by the frames acquired so far—the operator intuitively aims the acquisition device at the remaining points.

Rendering

Existing algorithms such as splatting (e.g. surfels (Pfister et al., 2000), q-splat (Rusinkiewicz et al., 2000)) or 3D triangulation (Hoppe et al., 1992, Solem and Heyden, 2004) could be used to render the acquired depth and color data. However, while these algorithms would produce good results for scenes where the geometry is densely sampled, they would fail in the case of scenes with undersampled geometry.

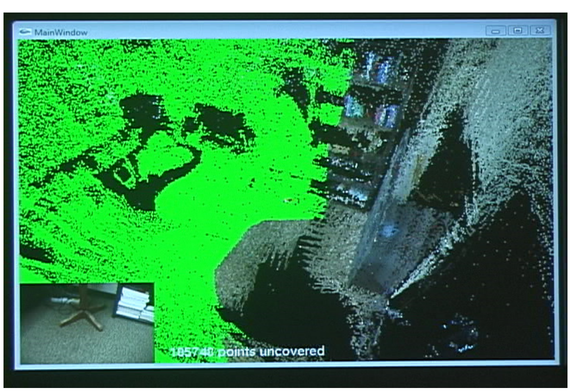

Figure 15. Feedback during Inside-Looking-Out Color Acquisition.

We have developed a reconstruction algorithm that performs well even in the case of undersampled geometry. Our algorithm is view-dependent: it uses the acquired point cloud and reference color frames to compute a reconstruction for each desired view. The algorithm uses the depth samples to automatically define a morph between the reference color frames, thus providing the flexibility to compensate the undersampled geometry by increasing the density of the color acquisition viewpoints.

The algorithm performs the following steps for each frame:

A. Visibility
The set of visible points is computed by splatting all the points and keeping those whose splats remain visible in at least one pixel. Two splatting strategies are used. During live acquisition when the model is rendered for operator feedback we employ square screen aligned splats. The model-space size of the splats is estimated on the fly. When rendering the model after the acquisition phase, we use model-space oriented splats. The splat normals are computed in a pre-rendering step.

B. 2D triangulation
The projections of the visible points are Delaunay triangulated in the image plane and the connectivity information is applied to the corresponding 3D points in order to derive a view-dependent mesh.

C. 3D mesh coloring
During the coloring phase we provide color for each pixel frame. A GPU shader projects each fragment into the set of color frames that is used for the current view. Visibility in each color frame is assessed by using a z-buffer pre-computed as shown in steps A & B.

In order to efficiently compute the set of color frames to be used for coloring the current view, we pre-compute for each 3D point the "best" color frame: from the set of frames where the point is visible we pick the one where the point is closest to the center. This ensures maximum coverage by the color frame in the scene and has the added advantage of consistency by assigning the same color frame to neighboring points, thus minimizing the size of the set.

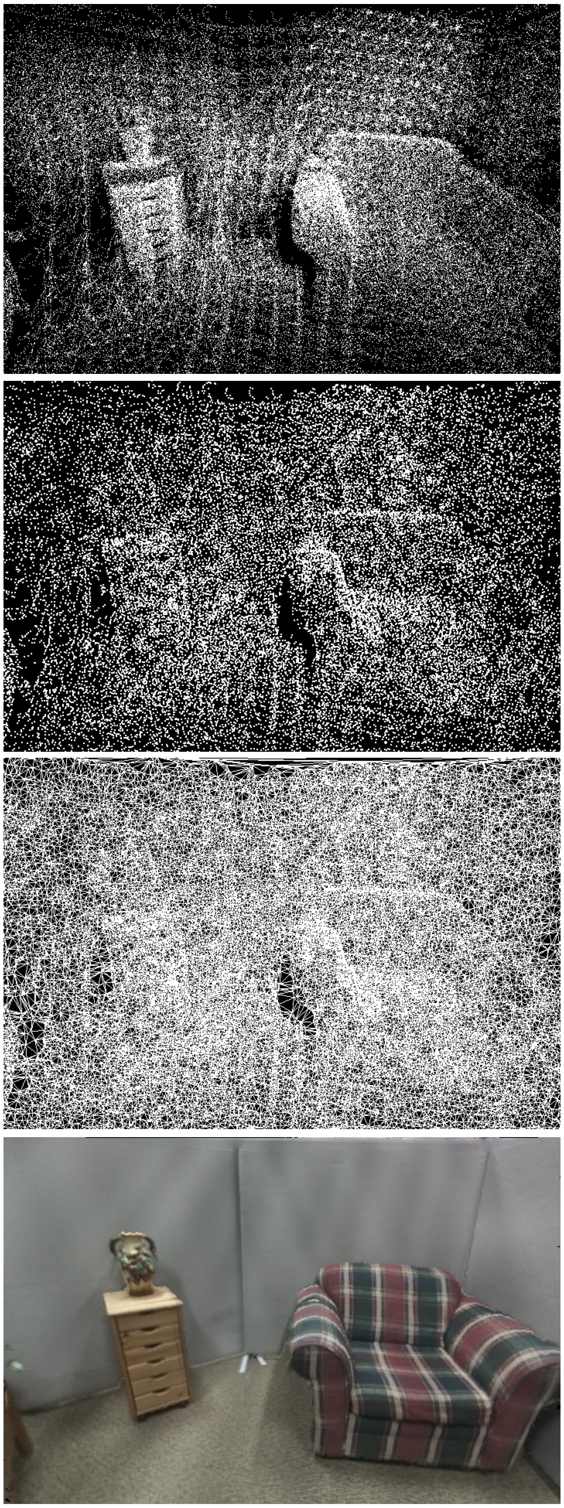

Figure 16. Stages of the Rendering Algorithm (top to bottom): Original Points, Visible Points, Triangle Mesh, Colored Mesh (Output).

The final color for the fragment is computed by blending the color contributions of the frames where

Figure 17. Scene Rendered w/o and w/ blending.

the fragment is visible. Blending is necessary because overlapping frames can have considerably different brightness as the camera adjusts exposure during acquisition (Figure 17). Our blending method minimizes these differences for smooth transitions.

Each color contribution from a reference frame will get a weight of 1.0 if the fragment projects at the center of the reference frame and 0.0 if the fragment projects at the edges. The weight contributions for each fragment are then normalized to [0, 1]. Because of shader limitations on the number of textures that can be used, blending is done in several passes. The color frame set is divided into N groups of k frames. At each step S the corresponding set of k color frames is used to color the mesh as described here. The result is then blended with the result from the previous N-1 steps. All the blending steps are performed on the GPU.

RESULTS

We tested our system with good results by acquiring several indoor scenes with simple and complex geometry (Figure 1 and 18). We placed the computer on a cart connected to an uninterruptible power supply to allow moving from one location to another without shutting down the system. A monitor was also placed on the cart to display the real time feedback for the operator. The acquisition device was placed on a sturdy wheeled tripod in order to facilitate transportation. Two persons were needed to move the setup: one to push the cart and another person to push the acquisition device.

	Office	**Hall**	**Couch**
Points [x 1000]	342	260	194
Color frames	380	446	48
Depth + color acquisition time [mins]	22 + 9	14 + 7	13 + 3

Table 1. Statistics for Acquired Models.

Figure 18. Images Rendered from Room, Hallway and Couch Models.

The system is efficient and robust: each model was acquired in a single take in approximately 15 to 30 minutes. The models have a compact memory footprint: they have fewer than 342,000 3D points and fewer than 446 color frames (Table 1).

CONCLUSIONS AND FUTURE WORK

We presented a system for large scale acquisition of room-sized interiors based on sparse depth acquisition from a dense set of view points. We implemented the sparse-depth/dense-viewpoint approach into an automated modeling system that handles a variety of scenes robustly. The system has ample modeling power, is fast and has good scene coverage. Our interactive modeling pipeline is more efficient than pipelines based on lengthy dense depth acquisition. The operator monitors and guides the acquisition process, receives immediate feedback and corrects the problems during acquisition. Quality models are obtained in a single scanning session. The acquired models are available for inspection throughout the scanning process. The depth and color data is combined into view-dependent model that allows photorealistic rendering at interactive rates. We validated our system by capturing several scenes.

We will continue to develop our system. One goal is to improve the maneuverability of the hardware. Simultaneous acquisition of color and depth would significantly reduce the acquisition time. This would require replacing the laser dot regions in a frame with color from nearby frames. Modeling specular effects is another improvement that would increase the quality of the walkthroughs. This would require detecting the light sources in the scene and replacing the currently static highlights with correct, dynamic ones. Also as future work we will investigate scanning with several acquisition devices in parallel and assembling the model on a server station.

At the current state, the ModelCamera system is already sufficiently mature for immediate deployment in virtual training applications in defense and beyond. The models constructed by the system capture the complexity of the real world in great detail and can provide safe environments where important skills are acquired effectively.

Secondly, the photorealistic 3-D models of building interiors are important infrastructure not only for training emergency response personnel, but also for virtual training civilians for improved preparedness should natural or man made disasters occur.

We are also interested in using our 3-D models as the basis for high-fidelity, large scale numerical simulations. We will leverage our experience with finite element analysis (FEA) simulations (Popescu et al., 2003, Hoffmann et al., 2004, Rosen et al., 2008). Detailed 3-D models of actual building interiors, including furnishings, are invaluable for increasing the confidence in FEA simulations.

ACKNOWLEDGEMENTS

We thank Gleb Bahmutov and Elisha Sacks for their contribution to the ModelCamera project, Toby Lamb for his help devising and constructing brackets to hold the acquisition device, our graphics group for insightful suggestions and constructive criticism of this work, and Larry Rieger for patient and experienced guidance offered in his role of Bird Dog for this paper. This work was supported by Purdue University's Computer Science Department, by the National Science Foundation through grant SCI-0417458, and by the Tellabs Foundation.

REFERENCES

Bahmutov, G., Popescu, V., & Mudure, M. (2006). Efficient Large-Scale Acquisition of Building Interiors, *Computer Graphics Forum, volume 25, issue 3 (EG 2006)*, pp. 655-662.

Bahmutov, G., Popescu, V., Mudure, M., and Sacks, E. (2004). "Depth Enhanced Panoramas", *Proc. of Video Vision and Graphics 2005.*

Bernardini, F., et al. (2002). Building a Digital Model of Michelangelo's Florentine Pieta. *IEEE Computer Graphics & Applications, 22(1)*, pp. 59-67.

Chen S. (1995). Quicktime VR - An Image-Based Approach to Virtual Environment Navigation. *Proc. SIGGRAPH 95*, pp. 29-38.

Debevec, P., Taylor, C., and J. Malik (1996). Modeling and Rendering Architecture from Photographs: A Hybrid Geometry and Image Based Approach. *Proc. SIGGRAPH '96*, 11-20 (1996).

Farouk, M. et al. (2003). Scanning and Processing 3D Objects for Web Display. *In Proceedings of 4th International Conference on 3D Digital Imaging and Modeling (3DIM '03)*.

Fisher, R., et al. (1996). A Hand-Held Optical Surface Scanner for Environment Modeling and Virtual Reality. *In Proc. Virtual Reality World*, pp. 13-15.

Gibson, S., Cook, J., Howard, T., and Hubbold, R. (2002). ICARUS: Interactive Reconstruction from Uncalibrated Image Sequences, *In Proc of SIGGRAPH 2002 courses*.

Gortler S., Grzeszczuk, R., Szeliski R., Cohen M. (1996). The Lumigraph. *Proc. SIGGRAPH '96*.

Guidi G., Frischer B., Spinetti A. (2005). 3D digitization of a large model of imperial Rome. *Proc. of Fifth International Conference on 3-D Digital Imaging and Modeling* (3DIM'05).

Hidalgo, E. and R. J. Hubbold (2002). Hybrid geometric-image-based-rendering. *Proceedings of Eurographics 2002, Computer Graphics Forum, 21(3)*, pp. 471-482.

Hilton, A., Illingworth, I. (2000). Geometric Fusion for a Hand-Held 3D Sensor. *Machine Vision Applications 12(1),* pp. 44-51.

Hoffmann, C., Popescu, V. Kilic, S., Sozen, M. (2004). The Pentagon on September 11th. *IEEE Computing in Science and Engineering*, pp 52-60.

Hoppe H., et al. (1992). Surface reconstruction from unorganized points, *Computer Graphics*, 26:2.

M. Levoy et al (2000). The Digital Michelangelo: 3D Scanning of Large Statues, Proc. *ACM SIGGRAPH*.

LEVOY, M., Hanrahan, P. (1996). Light Field Rendering. *Proc. of SIGGRAPH 96*, pp. 31—42.

Matusik W., et al. (2002). Image-Based 3D Photography using Opacity Hulls. *Proc. of SIGGRAPH '02.*

Mudure, M., Popescu, V. (2008). 1001 Acquisition Viewpoints: Efficient and Versatile View-Dependent Modeling of Real-World Scenes. To appear in the *Visual Computer*.

Pfister, H., Zwicker, M., Baar, J. Van, and Gross, M. (2000). Surfels: Surface Elements as Rendering Primitives. *Proc. of SIGGRAPH '98*, pp. 335-342.

Pollefeys, M. et al. (2001). A Guided Tour to Virtual Sagalassos. *In Proceedings of VAST2001 (Virtual Reality, Archaeology, and Cultural Heritage)*.

Popescu, V., Sacks, E., and Bahmutov, G., (2003). The ModelCamera: A Hand-Held Device for Interactive Modeling. *Proceedings of International Conference on 3-D Imaging and Modeling*, pp 285-292.

Rosen, P., Popescu, V., Hoffmann, C., Irfanoglu, A. (2008) A High-Quality Physically-Accurate Visualization of the September 11 Attack on the World Trade Center. *IEEE Transactions on Visualization and Computer Graphics*, in press.

Rusinkiewicz, S., Levoy, M. (2000). QSplat: A Multiresolution Point Rendering System for Large Meshes. *In Proc. SIGGRAPH*.

Rusinkiewicz, S., Hall-Holt, O., and Levoy, M. (2002). Real-Time 3-D Model Acquisition. *In Proc. of SIGGRAPH 2002*.

Solem, H.E., Heyden, A. (2004). Reconstructing open surfaces from unorganized data points. *In Proc.* of *Conference on Computer Vision and Pattern Recog*.

Stumpfel J. et al. (2003). Digital Reunification of the Parthenon and its Sculptures**,** *4th Symp. on VR, Archaeology and Intelligent Cultural Heritage*.

Terauchi, T., Oue, Y., Fujimura, K. (2005). A Flexible 3D Modeling System Based on Combining Shapefrom Silhouette with Light-Sectioning Algorithm. *In Proceedings of the Fifth International Conference on 3-D Digital Imaging and Modeling*.

Williams, N. et al. (2003). Monticello Through the Window. *In Proceedings of the 4th International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST 2003*.