Acquisition of Large-Scale Surface Light Fields

Wei-Chao Chen [*]
NVidia Corporation

Lars Nyland[†] Anselmo Lastra[†] Henry Fuchs[†]
University of North Carolina at Chapel Hill

Abstract: This sketch reports the development of an acquisition methodology to acquire high-resolution surface light fields of an office-size environment.

1 Introduction

In recent years we have witnessed exciting new research in image-based rendering and modeling that brings us closer to the realization of 3D photography, *e.g.*,[4; 3; 2]. Despite the tremendous progress, however, these techniques have not demonstrated the combination of real, high-resolution scenes with unconstrained camera locations, and these scenes are often limited to outside-looking-in configuration such as buildings and statues.

With recent techniques, we are capable of rendering much more complicated models than ever before. Lacking suitable datasets to stress our rendering algorithms, we set forth to develop a complete pipeline for acquiring surface light fields beyond the desktop scale. Our acquisition system uses off-the-shelf hardware components and is applicable to most static environments.

2 Methodology

Acquisition To date, large-scale fully-automatic scanning remains an active research topic. In our system, we perform acquisition planning and part of the registration with human assistance. It consists of a commercial laser rangefinder[1] (Figure 1,left) and an off-the-shelf digital camera. The two-step acquisition process starts with the acquisition and production of an unified geometric model. We then acquire color photographs and register the photographs to the geometric model.

The laser rangefinder acquires depth images on a spherical coordinate system. For the *Office* scene, we scan seven panoramic depth images, each containing approximately 8 million depth samples (Figure 1,center). The depth images are registered together using a commercial implementation of Iterative Closest Point (ICP) algorithm[2]. We then merge the geometry from individual depth scans into a unified geometric model and simplify it. For the color information, we take photographs at grid locations in the environment. Because the surfaces are quite diffuse, 100 wide-angle images cover most of the surfaces several times. We calibrate the camera and remove nonlinear distortions from the photographs (Figure 1, right). We then compute a local camera pose for each photograph by using manually-selected correspondences between the photograph and a depth image. The global camera pose is equivalent to transforming the local camera pose from depth image coordinate to that of the unified geometry. Please refer to [1] for details on the process.

Compression and Rendering The size of the raw dataset is on the order of several GBs. To render it efficiently, we choose our work[5] because of its compactness and applicability to hardware-accelerated rendering. The algorithm process uses multi-pass texture-mapping during decoding. To limit the size of light field texture maps, we partition the geometric model with breadth-first search and generate light field maps for each partition independently. We found this approach is applicable to geometric culling algorithms without introducing a texture thrashing problem. Figure 2 shows the synthesized *Office* scene. This model can be rendered at highly interactive rates on a PC with commodity graphics.

Figure 1: Left: Rangefinder. Center: Partial depth image (red: low-confidence samples). Right: Photograph of the *Office*.

Figure 2: *Office* rendered with 3-term PCA+Texture Compression. Its physical dimension is $15\,ft(W) \times 10\,ft(D) \times 8\,ft(H)$.

3 Discussion and Conclusion

This sketch reports our first step in in large-scale surface light field acquisition. Although the process still requires some human intervention, we believe this is one of the first successful experiments of similar kind. With a proper tracking device, autonomous acquisition can be realized in a very near future, and image-based models will be a viable alternative for applications such as interactive walkthroughs and virtual presence.

References

[1] W.-C. Chen. *Light Field Mapping: Efficient Representation of Surface Light Fields*. PhD thesis, Dept. of Computer Science, UNC-Chapel Hill, May 2002.

[2] C. Buehler *et al.* Unstructured Lumigraph Rendering. In *SIGGRAPH 2001*, pages 425–432, August 2001.

[3] D. N. Wood *et al.* Surface Light Fields for 3D Photography. In *SIGGRAPH 2000*, July 2000.

[4] P. E. Debevec *et al.* Efficient View-Dependent Image-Based Rendering with Projective Texture-Mapping. In *Eurographics Rendering Workshop 1998*, June 1998.

[5] W.-C. Chen *et al.* Light Field Mapping: Efficient Representation and Hardware Rendering of Surface Light Fields. In *SIGGRAPH 2002*, July 2002.

[*] e-mail: ciao@cs.unc.edu (Research done while at UNC-Chapel Hill)
[†] e-mail: {nyland,lastra,fuchs}@cs.unc.edu

[1] DeltaSphere 3000, 3rd Tech Corporation, North Carolina, USA.
[2] Polyworks 6.0, InnovMetrics Incorporated, Quebec, Canada