Exploring Global Illumination for Virtual Reality

Roger Hoang, Steve Koepnick, Joseph D. Mahsman, Matthew Sgambati, Cody J. White, and Daniel S. Coming
Desert Research Institute and University of Nevada, Reno*†

1 Introduction

Real time global illumination (GI) is a difficult and steadily researched area. Advances in the field could potentially benefit virtual reality applications by increasing users' sense of presence. In immersive virtual environments (IVE) like CAVEs, applications must support perspective-corrected stereoscopic rendering. Dmitriev et al. [2004] performed GI in a CAVE using Precomputed Radiance Transfer, which requires a static scene. Mortensen et al. [2007] also performed GI in a CAVE using Virtual Light Fields which did not allow moving lights or geometry. We present our attempt to find GI techniques that support dynamic lights and scene geometry in our 6-sided CAVE-like IVE (*DRIVE6*). We implemented two separate illumination techniques: GPU photon mapping (GPM) and multiresolution splatting for indirect illumination (MSII). Each technique makes trade-offs between image quality and speed, and appropriate use of each depends on the needs of the application. Anecdotal evidence suggests that these techniques increase the sense of presence, warranting formal study. A user study is planned.

2 Our Methods

Image space photon mapping [McGuire and Luebke 2009] rasterizes a scene from the light's perspective to generate initial photon bounces, traces all further bounces on the CPU, and scatters the effects of the resulting photons by rendering them as ellipsoids with a special shader. This technique does not support area lights or accurate reflection or refraction. Another method implements irradiance caching and adaptive photon mapping and gathering on the GPU but suffers from low frame rates [Wang et al. 2009].

Using OptiX (www.nvidia.com/object/optix.html), we combine these methods into our GPM technique that supports reflections, refractions, and area lights. After rasterizing the scene and photon ellipsoids, we ray cast on reflective and refractive surfaces. We perform photon gathering by casting a ray from the hit point along the surface's normal to find which photon ellipsoids the ray intersects.

We also tested MSII [Nichols and Wyman 2009], which creates virtual point lights where the scene receives direct illumination, and their contributions are splatted into a multiresolution buffer, using min-max mipmaps of normals and depth to select regions for higher resolution rendering. We ported a publicly available implementation of this technique (www.cs.uiowa.edu/˜cwyman) including improvements using a stencil buffer technique [Nichols et al. 2009].

3 Results and Conclusion

Each technique was tested on a cluster of 12 nodes, each with 24 GB of RAM, an Intel Xeon 3.2 GHz processor, and an Nvidia Quadro FX 5800 connected to a 1920x1080 projector. Rendered at 800x800 and upscaled, GPM achieved frame rates of 11.37, 11.29 and 3.84 fps using scenes containing 70k, 100k, and 346k triangles, respectively, without reflective and refractive surfaces. With these enabled, frame rates dropped to 2.15, 1.78, and 1.33 fps. MSII reached 16.87, 15.82, 14.72, and 11.00 fps using scenes containing 116k, 186k, 436k, and 782k triangles, respectively.

Figure 1: *(Left) GPU Photon Mapping and (Right) Multiresolution Splatting running in a six sided CAVE-like environment.*

MSII achieved higher frame rates than GPM, but the visual quality of the latter is significantly higher, since it can support reflections, refractions, and multiple indirect bounces. MSII also proved to be significantly more scalable than GPM, because it is an image-space technique and its speed is independent of geometric complexity. We suspect that the poor GPM frame rates may be partially explained by OptiX constructing the photon acceleration structure on the CPU rather than the GPU, requiring transfers from and to GPU memory.

We made some interesting observation while testing these methods. The virtual light drawn on the IVE walls would often illuminate the user, which we felt increased the feeling of presence. We also noticed that low frame rates were alleviated by the stereo rendering. With both eyes open, slow frame rates were less apparent than with one eye closed. We suspect that since DRIVE6 uses active stereo, the user's eyes receive twice the update frequency when combined.

This work was completed as part of a course at UNR, using equipment at DRI supported by the U.S. Army's RDECOM-STTC under Contract No. N61339-04-C-0072.

References

DMITRIEV, K., ANNEN, T., KRAWCZYK, G., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2004. A CAVE system for interactive modeling of global illumination in car interior. In *Proc. ACM Symp. on VR Software and Technology*, ACM, 137–145.

MCGUIRE, M., AND LUEBKE, D. 2009. Hardware-accelerated global illumination by image space photon mapping. In *Proc. 1st ACM Conf. on High Performance Graphics*, ACM, 77–89.

MORTENSEN, J., KHANNA, P., YU, I., AND SLATER, M. 2007. Real-time global illumination in the CAVE. In *Proc. ACM Symp. on VR Software and Technology*, ACM, 145–148.

NICHOLS, G., AND WYMAN, C. 2009. Multiresolution splatting for indirect illumination. In *Proc. 2009 Symp. on Interactive 3D Graphics and Games*, ACM, 83–90.

NICHOLS, G., SHOPF, J., AND WYMAN, C. 2009. Hierarchical image-space radiosity for interactive global illumination. *Computer Graphics Forum 28*, 4, 1141–1149.

WANG, R., WANG, R., ZHOU, K., PAN, M., AND BAO, H. 2009. An efficient GPU-based approach for interactive global illumination. *ACM SIGGRAPH 2009 28*, 3, 1–8.

*{skoepnick, matthew.sgambati, dan.coming}@dri.edu
†{hoang, mahsmanj, cjwhite}@cse.unr.edu

SIGGRAPH 2010, Los Angeles, California, July 25 – 29, 2010.
ISBN 978-1-4503-0210-4/10/0007