Near-Field Illumination for Mixed Reality with Delta Radiance Fields

Tobias Alexander Franke, Fraunhofer IGD

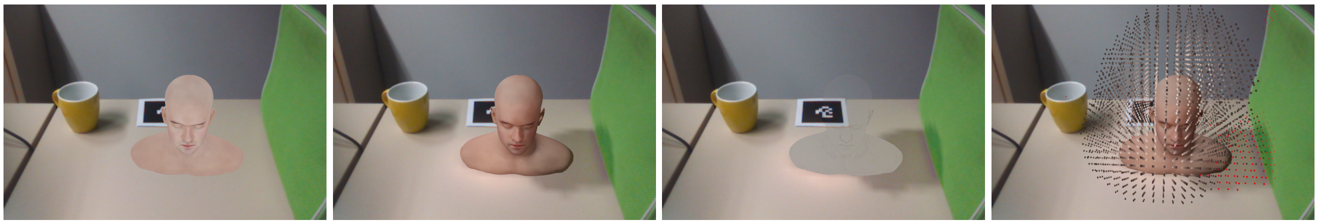

(a) No illumination (b) Full DLPV rendering (c) DLPV effects only (d) Volume visualization

Figure 1: *Infinite Head model inserted into a real scene with one reconstructed light source: (a) a synthetic object is inserted without illumination, (b) visible first bounce around the base as well as low resolution shadow (32 propagations,* 512^2 *VPLs, 11ms per frame), (d) indirect effects without synthetic object for better visualization, (d) visualization of the DLPV (red dots indicate negative values).*

1 Introduction

Apart from geometric registration, fusing synthetic objects with a real context requires believable interaction of real and virtual light. Where shadows provide the visual cues to locate a synthetic object in a real scene and transferred light from real light sources let it appear in harmony with its surroundings, the mutual indirect interaction of illumination is a necessary detail to convince an observer that the rendered result is not merely augmented but part of the scene. Such a *mixed reality* (MR) system has applications in movie production, advertisement of unfinished products or cultural heritage visualization. While there are a range of relighting tools available for offline renderers or static scenes (e.g. photos), interactive MR systems usually disregard proper lighting entirely or greatly simplify shading. A method often employed is to merge virtual light and shadows cast from synthetic objects with a real background with Differential Rendering, leaving out any other light interaction such as indirect light bounces. Attempts have been made to resolve this issue with Instant Radiosity [Knecht et al. 2010; Lensing and Broll 2012]. To suppress flickering, a large number of virtual point lights (VPL) is necessary, drastically taxing execution speed. We propose to model all light, virtual and real, as a unified radiance field instead to avoid performance issues from oversampling and to maintain temporal coherence.

2 Technical Approach

A manually reconstructed model of the real scene is geometrically registered with a marker and real light sources are importance sampled from a hemispherical image provided by a fish-eye lens camera. Consider an existing radiance field L^μ. When we insert another non-emissive object O into the scene covered by L^μ, the properties of the radiance field change: light is either blocked or scattered by the newly inserted geometry. To account for this *change* in the radiance field, consider another field L^ρ which has the same light configuration as L^μ and additionally contains O. The change in a radiance field L^μ by introducing O is therefore $L^\Delta = L^\rho - L^\mu$. To generate L^Δ – dubbed the *Delta Radiance Field* – we calculate two Reflective Shadow Maps (RSM) for each reconstructed real light source: one contains the reconstructed scene geometry with the introduced object O, and one without O. The difference between both is used to initialize the residual direct light and indirect bounces into a volume representing the radiance field. To model this radiance field, we make use of small volumetric textures (32^3 pixels) with spherical harmonic encoded directional light similar to Light Propagation Volumes [Kaplanyan and Dachsbacher 2010] called Delta-LPV (DLPV). We use a two step injection phase, where first light from the RSM containing O is injected into the volume and afterwards light from the second RSM without O is injected negatively. Indirect light is approximated with VPLs created from each RSM. For direct light, we use the flux and direction of the reconstructed real light source, but determine the position inside the volume with the help of the RSMs. Because of the differential the volume will contain negative values (i.e. Antiradiance). When superimposed on a background image, the DLPV introduces indirect bounces and shadows, correcting the real radiance field for the introduction of O. An example can bee seen in Figure 1. The evaluation cost of a DLPV differs from an LPV only in the double injection for direct and indirect light.

By modeling a radiance field at a low resolution the final image suffers from artifacts such as light and shadows bleeding through thin geometry, self-illumination and self-shadowing. Considering the low-frequency nature of indirect diffuse light however, these artifacts are not as apparent as the heavy aliasing along shadow borders from the direct injection. A solution is to increase the spatial resolution to model a more accurate Delta Radiance Field. Furthermore, we want to include indirect specular bounces from synthetic objects on real surfaces. With Sparse Voxel Octrees [Crassin et al. 2011] all of the problems mentioned above may be resolved.

References

CRASSIN, C., NEYRET, F., SAINZ, M., GREEN, S., AND EISEMANN, E. 2011. Interactive indirect illumination using voxel cone tracing. *Computer Graphics Forum*.

KAPLANYAN, A., AND DACHSBACHER, C. 2010. Cascaded light propagation volumes for real-time indirect illumination. In *Proc. of the 2010 ACM SIGGRAPH symposium on Interactive 3D Graphics and Games*, ACM, I3D '10, 99–107.

KNECHT, M., TRAXLER, C., MATTAUSCH, O., PURGATHOFER, W., AND WIMMER, M. 2010. Differential Instant Radiosity for Mixed Reality. In *2010 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 99–107.

LENSING, P., AND BROLL, W. 2012. Instant indirect illumination for dynamic mixed reality scenes. *2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 109–118.

SIGGRAPH 2013, July 21 – 25, 2013, Anaheim, California.

ACM 978-1-4503-2261-4/13/07