Interactive Relighting of Arbitrary Rough Surfaces

Tobias Alexander Franke · **TU** Darmstadt · **Fraunhofer** IGD · tobias.franke@igd.fraunhofer.de

(a) (b) (c) (d)

Figure 1: Results. All images rendered on a NVIDIA GTX 780 at 20ms to 35ms per frame. (a) A voxelized synthetic object in a real scene. (b) The augmenting Buddha receives a red indirect bounce from the real wall and is reflected on a rough metal surface. (c) Indirect reflections of synthetic objects, especially noticeable on glossy or specular surfaces, miss secondary indirect bounces. Bounces from far-field geometry do not reach the synthetic object because of the near-field extents of the volume (notice how both the reflection below the chin and the back of the head are too dark). (d) Diffuse indirect bounces are easily simulated with the same method.

The Relighting Problem

Introduction & Related Work

When presenting synthetic objects in a real environment – for instance for pre-visualization in advertisements – special attention needs to be directed at the mutual interaction of light reflecting off synthetic and real surfaces to form a coherent appearance. If the user is to be convinced that the synthetic object is part of the real scene, a relighting method has to handle shadowing and reflecting illumination between both synthetic and real surfaces. Even though a range of relighting methods are available for static scenes such as photographs, this aspect has been traditionally ignored in real-time augmented reality (AR) systems. A method often employed to merge synthetic light and shadows with a real background is Differential Rendering, leaving out indirect illumination. Attempts have been made to resolve this issue with Differential Instant Radiosity [4], which however requires many VPLs to suppress flickering. Delta Light Propagation Volumes [2] cluster many VPLs in a small volume, but suffer from heavy light bleeding artifacts. A GPU raytracer in [3] supports diffuse bounces with Differential Irradiance Caching, albeit at much higher cost than rasterizer based counterparts.

I present a novel relighting solution called **Delta Voxel Cone Tracing** (DVCT) to enable mutual diffuse, glossy and specular indirect bounces between real and synthetic geometries. The method is temporally coherent and to my knowledge the first real-time solution to support arbitrary glossy reflections in AR.

Theoretical Framework

Given a reconstructed real scene with light sources, geometry and an additional object O to insert, we can imagine a snapshot of the radiance field of the real scene L_μ. We can furthermore think of another radiance field L_ρ of the same scene containing the additional object O. The difference between both - the *Delta Radiance Field* - contains the residual light scattered by the additional object as well as *antiradiance*, compensating for the light which is blocked by it. The expansion of both radiance fields into a Neumann series yields a new linear transport operator T_Δ, which can be used to relight an existing radiance field.

$$L_\Delta = L_\rho - L_\mu = \sum_{i=0}^{\infty} \mathbf{T}_\rho^i L_e - \sum_{i=0}^{\infty} \mathbf{T}_\mu^i L_e = \sum_{i=0}^{\infty} \mathbf{T}_\Delta^i L_e$$

Prerequisites

I start by importance sampling one or more real light sources from a fish-eye lens camera (Figure **2**). A marker is used to track a model of the reconstructed real scene alongside an augmenting object O. A mask identifies real and synthetic pixels in the final image.

References

1. Crassin, C., Neyret, F., Sainz, M., Green, S., and Eisemann, E. 2011. ***Interactive indirect illumination using voxel cone tracing.*** Computer Graphics Forum.
2. Franke, T. 2013. ***Delta light propagation volumes for mixed reality.*** In 2013 IEEE International Symposium on Mixed and Augmented Reality.
3. Kan, P., and Kaufmann, H. 2013. ***Differential irradiance caching for fast high-quality light transport between virtual and real worlds.*** In 2013 IEEE International Symposium on Mixed and Augmented Reality.
4. Lensing, P., and Broll, W. 2012. ***Instant indirect illumination for dynamic mixed reality scenes.*** In 2012 IEEE International Symposium on Mixed and Augmented Reality.

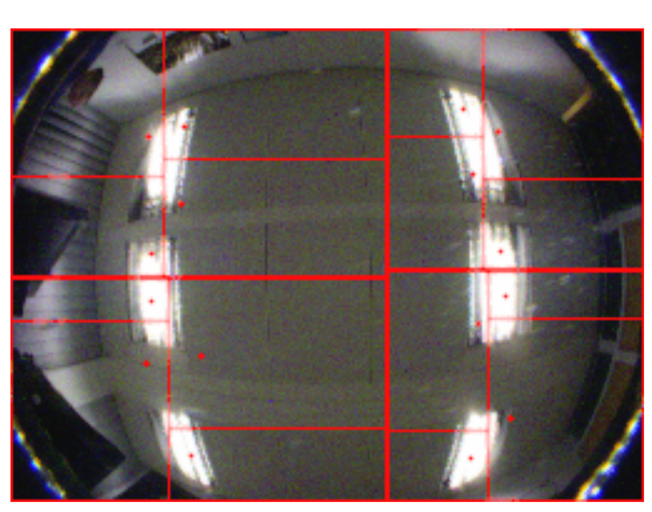

Figure 2

Delta Voxel Cone Tracing

Creating the Delta Volume

First, the reconstructed real scene and the synthetic object O are voxelized into a small volume V_η around O with 256^3 voxels or more (Figure **3a**). It contains surface normals and an occlusion marker which - when filtered into higher levels of a hierarchy - represents the average occlusion in each voxel.

For each light source, I render two Reflective Shadow Maps (RSM): R_μ of the reconstructed scene, and R_ρ of the same scene where O is included (Figure **3b-c**). In the next step called *split-injection*, all VPLs created from R_ρ are injected into a volume V_Δ first, and VPLs created from R_μ are then injected negatively (Figure **3d**). The latter VPLs emit antiradiance, correcting excess light in the real background image.

Augmenting Reality with a Synthetic Object

I use a deferred renderer to create the final image. A geometry buffer is populated with albedo, normals and depth information. For each pixel belonging to a real reconstructed surface, the radiance from the background image is augmented by the cone-traced radiance from the Delta Volume V_Δ.

In Figure **1b-d** indirect contribution on each real pixel is sampled from its hemisphere with 9 wide-angle cones to gather diffuse bounces, and one narrow cone - where the solid angle depends on the surface roughness - to gather specular contribution. Shadows are evaluated by casting *shadow cones* (which only sample occlusion from V_η) into the direction of each light source. The occlusion is multiplied with the lit reconstructed real surface and added as antiradiance onto the background image (Figure **4**).

To relight O I also create another volume V_ρ (Figure **3c**): because this volume includes indirect light from both the real and the synthetic scene, O receives indirect bounces from itself and its surrounding included within the extents of V_ρ. Any other relighting method for the synthetic object can be chosen however, such as precomputed radiance transfer combined with an environment map captured from the 180 degree fish-eye lens in Figure **2**.

Figure 4

Conclusion

DVCT is a real-time method to relight real images for the change in illumination caused by the introduction of a new synthetic object. It drastically minimizes the error compared to a previous volumetric solution [2], rivals path-tracing based approaches [3] and allows for a wider spectrum of effects, such as glossy reflections. In the future, I want to explore synthetic emissive surfaces and multiple indirect bounces.

A limitation of the current approach is that the spatial extents of the volume around O define the maximum range for real bounces that can reach the synthetic object. It is therefore only suitable for near-field relighting. Real surfaces however gather light from V_Δ (which fully contains O) and therefore indirect bounces from the synthetic object can be handled on the entire reconstructed real scene.

Because DVCT operates on separately filtered volumes V_Δ and V_ρ, the corrective term introduces bias, which increases with higher glossiness (Figure **5**). Glossy and perfect specular reflections also show structural artifacts bounded by the volumetric resolution.

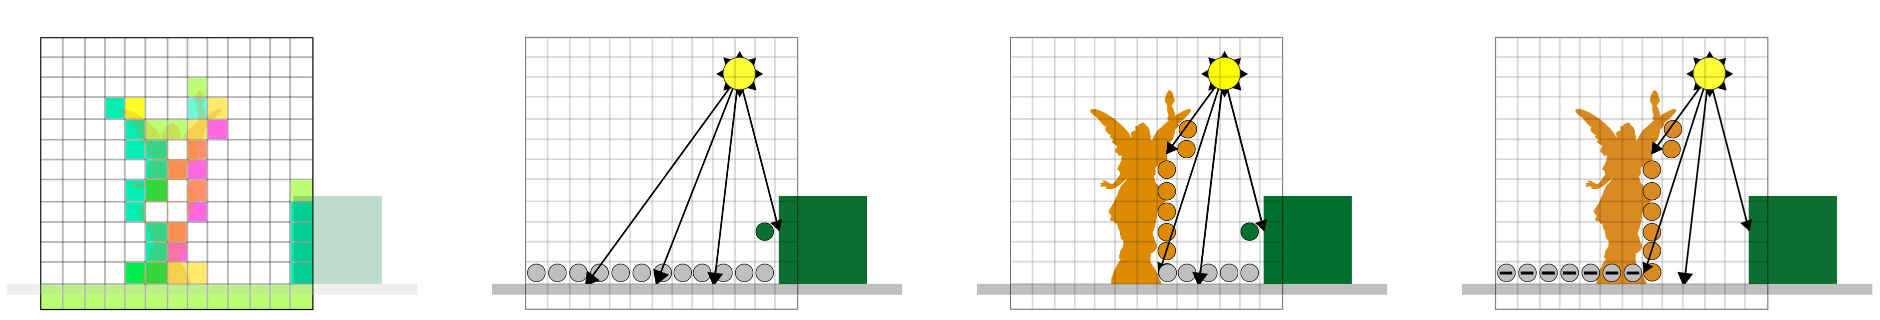

(a) (b) (c) (d)

Figure 3: Creating the Delta Volume V_Δ. (a) The geometries of both synthetic and real reconstructed parts of the scene are voxelized into a small volume around O. (b, c) Two RSMs of the scene are rendered: R_μ without and R_ρ with the synthetic object O. Injection of each RSM into a volume would yield V_μ and V_ρ respectively. (d) Injecting R_ρ and -R_μ consecutively into the *same* volume encodes the *change* in illumination V_Δ.

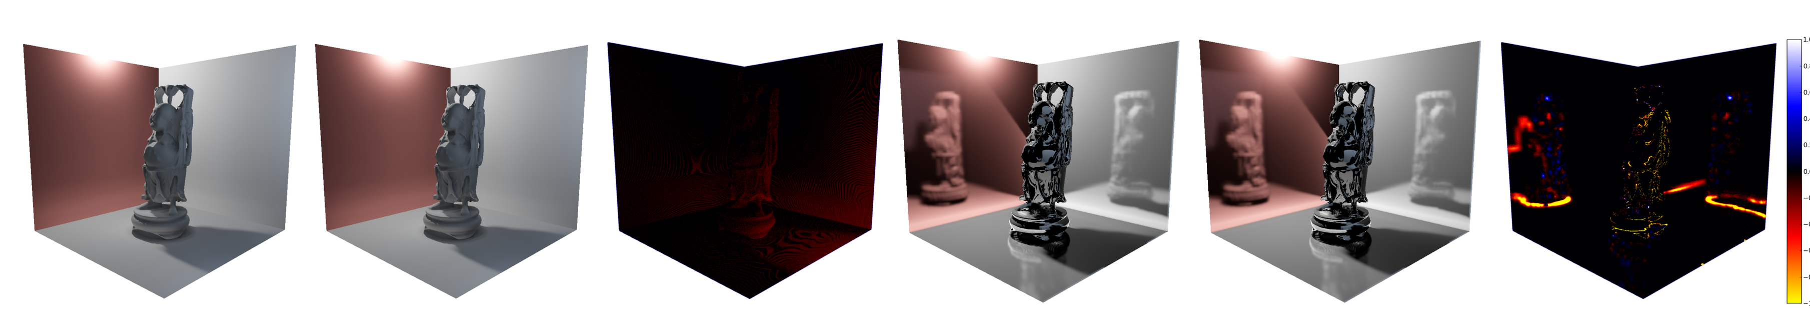

(a) (b)

Figure 5: Error Analysis. The impact of relighting an existing radiance field with DVCT in both image sequences is extracted with $E = [\ V_\rho - (V_\mu + V_\Delta)\ ]$. Red-yellow are negative, blue-white are positive differences (enhanced by 16x). (a) Diffuse scene with augmenting Buddha. The error is spread across the scene. (b) E is concentrated along edges in highly glossy reflections.

SIGGRAPH2014 · TECHNISCHE UNIVERSITÄT DARMSTADT · Fraunhofer IGD