Direct Illumination from Dynamic Area Lights

Greg Nichols*
University of Iowa

Chris Wyman†
University of Iowa

Figure 1: *An indoor garden illuminated solely by a large area light source displaying a video. Different frames from the video produce illumination of varying intensity and colors. The above figure displays three examples.*

1 Introduction

Area light sources are common in the real world, and thus important in realistic images. However, interactive rendering with area light sources is challenging, as each surface in a scene can receive light from every point in the area light. This problem is similar in nature to the rendering of single-bounce indirect illumination, and can be addressed with similar techniques.

One previous method of rendering indirect illumination [Dachsbacher and Stamminger 2006] computes a reflective shadow map (RSM), chooses VPLs from the RSM, and splats each VPL's contribution onto the scene. Prior work [Nichols and Wyman 2009] describes a multiresolution variant of reflective shadow maps. Each VPL's splat is adaptively refined into a set of multiresolution "subsplats" at the appropriate eye-space resolution, greatly reducing the fill-rate requirement of reflective shadow maps.

These techniques extend naturally to allow rendering of direct illumination from area light sources: we simply use a different set of VPLs. Our method requires no precomputation, enabling lights with textures and even video sources, as depicted in Figure 1. We also describe a light-clustering method that allows hierarchical selection of "patches" in both eye-space and light-space; this approaches the idea of hierarchical radiosity, but works in image-space instead of object-space.

2 Our Approach

Like multiresolution splatting [2009], we adaptively subdivide the scene into eye-space patches, each of which is rendered with illumination from a set of VPLs. Instead of choosing these VPLs by sampling the RSM, however, we create them by sampling the surface of the area light source.

To reduce artifacts arising from VPL undersampling while still maintaining interactive performance, we place VPLs finely around luminance discontinuities in the area light's texture or video source, and more coarsely in areas without them. A single geometry shader pass processes initial, densely-sampled VPLs and discards those without nearby discontinuities, weighting the remaining VPLs appropriately. Discontinuity detection relies on a min-max mipmap similar to that used during eye-space refinement. Figure 2 illustrates an example of this process.

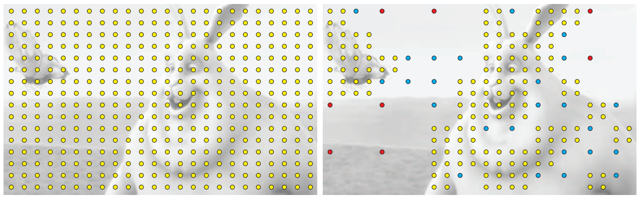

Figure 2: *Beginning with an initial dense VPL sampling (left), we cluster VPLs in areas without luminance discontinuities (right).*

During rendering, each eye-space fragment gathers illumination from the appropriate set of VPLs. Currently, each VPL contributes light to every eye-space fragment; we are exploring a variation in which the best set of VPLs is selected on a per-fragment basis, chosen to reduce fragment error below a user-specified threshold.

Our method renders the scene in Figure 1 at 25-30fps, with full-motion video on the light source. Like other RSM approaches, our method currently ignores visibility. In future work, we plan to address this, enabling direct shadows and improving our approximation.

Acknowledgements: All video stills are from "Big Buck Bunny", copyright © Blender Foundation (www.blender.org).

References

Dachsbacher, C., and Stamminger, M. 2006. Splatting indirect illumination. In *Proceedings of the Symposium on Interactive 3D Graphics and Games*, 93–100.

Nichols, G., and Wyman, C. 2009. Multiresolution splatting for indirect illumination. In *Proceedings of the ACM Symposium on Interactive 3D Graphics and Games*, 83–90.

---

[*] e-mail: gbnichol@cs.uiowa.edu
[†] e-mail: cwyman@cs.uiowa.edu