Eurographics Symposium on Rendering 2010
Jason Lawrence and Marc Stamminger
(Guest Editors)

Volume 29 (2010), Number 4

Multiresolution Image-Space Rendering for Area Lighting (Additional Discussion)

4. Additional Results and Discussion

Unfortunately, length restrictions on the paper submission limited our ability to fully explore the behavior of our multiresolution rendering algorithm. This supplementary material explores our results in somewhat greater detail.

If not clear elsewhere, unless stated differently all results in the paper and this supplement use 256 VPLs sampled uniformly on the light using a regular, 16×16 grid and use a $1024^2 \times 128$ voxel grid for visibility.

4.4. Varying Light Size

While we designed our algorithm with large, dynamic area lights in mind, it can be applied to smaller sources, as demonstrated in Figure 13. As light size decreases, at some point the rationale for using our work disappears. Our main advantage over more traditional "soft shadow" work (e.g., [SS07]) is twofold: our avoidance of artifacts from the single silhouette assumption and our capture of lighting variations in unshadowed regions that arise from radiance variations across the light surface. With small area lights, the single silhouette assumption is usually acceptable and a single color light source closely approximates the results for most pixels. In such a scenario, traditional soft shadow algorithms may be faster.

Furthermore, as the light size decreases one would expect the voxelized visibility to become more apparent. This is, in fact, visible for the smaller lights in Figure 13, though the 128-bit voxel depth exacerbates the problem. For smaller light sizes, a 256, 512, or 1024-bit voxel depth could easily be added (as in [ED08]).

4.5. Comparison to Light Propagation Volumes

Recently published work introduced the idea of Cascaded Light Propagation Volumes [KD10]. This work was published after the initial submission of our research, and its still-recent publication date makes it difficult to provide comparison images. Below, we instead provide a brief qualitative comparison.

Basically, Kaplanyan introduces a multiresolution implementation of irradiance volumes, using spherical harmonics to represent illumination and visibility samples at the grid lattice points. These SH coefficients are populated by sampling VPLs from a reflective shadow map [DS05] and projecting them to a SH basis. Because they compute illumination and visibility only at these lattice points, instead of the hundreds of thousands of fragments we use, their work runs significantly faster than ours. They also handle participating media, which we do not address.

However, their coarse sampling makes rendering of high frequency illumination and visibility quite difficult. They claim to sample on 32^3 grids, which suggests their aliasing from low volumetric resolution (e.g., poor contact shadows, difficulty with sharp illumination boundaries) will be significantly worse than ours. Our use of a 2D image-space structure suggests we might scale to higher sampling rates than their 3D structure, due to the curse of dimensionality, though we must reconstruct our structure each frame.

Finally, while reliance on a SH-basis enables propagation volumes to easily shade normal mapped surface, they are unable to render high frequency BRDFs or sharp shadows, such as our Phong images from Figure 8.

It may be possible to use a hybrid between our work and theirs, sampling light and visibility to a SH basis in a multiresolution screen space fashion, though this is future work.

4.6. Comparison to Imperfect Shadow Maps

Ritschel et al. [RGKS08] proposed imperfect shadow maps (ISMs), the other work closely comparable to ours. Figure 7 compares our work with a rendering using many perfect 64^2 shadow maps. These shadow maps were generated using the traditional rasterization pipeline instead of using the precomputed, uniform sampled points and the subsequent hole filling algorithm proposed for ISMs. This leads to a comparison in Figure 7 with significantly slower performance than ISMs, but higher quality, due to the elimination of aliasing and shadow map holes arising from coarsely sampled points.

Our rendering times are faster than those reported by Ritschel et al. even though we output higher resolution images. However, their speeds may be comparable to ours after accounting for GPU improvements. We likely run slower on high polygon models but faster on lower polygon models, due to ISMs' use of user-specified numbers of point samples.

As far as quality, the coarse resolution of ISMs makes hard contact shadows impossible. While our voxel buffer also leads to errors for contact shadows, we can capture higher frequencies in these regions and a finer voxelization reduces our errors for a modest cost. The banding visible in

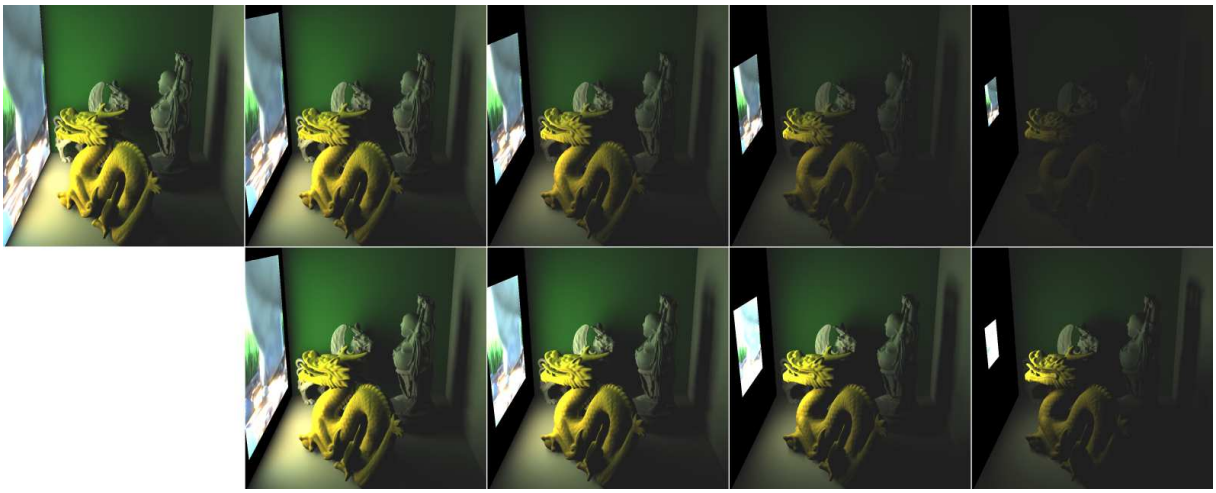

Figure 13: *Results of our technique as the light size varies. In the top row, the watts$/m^2$ is kept constant, yielding less illumination as the light size decreases; in the bottom row, the overall wattage of the light source remains constant.*

the shadow map comparison in Figure 7 comes from the low resolution shadow maps, which our per-pixel visibility sampling avoids. As the BRDF becomes more specular, shadow map texel aliasing becomes more objectionable for ISMs as fewer shadow maps are used to average visibility. Conversely, ISMs sample the light surface more densely, giving better quality when high VPL sampling is needed (see Section 4.7 below).

One key difference: ISMs require a preprocess to point sample the scene geometry. This somewhat limits geometric changes. Additionally, due to fixed sampling, the samples used to create individual imperfect shadow maps may be outside the light frustum or too sparse on nearby geometry, potentially leaving visible parts of the scene undersampled.

In general, we view ISMs as somewhat orthogonal to our work. As stated in our conclusion, we envision using ISMs in our incremental rendering process instead of ray-marched voxel visibility, combining the strengths of these different techniques.

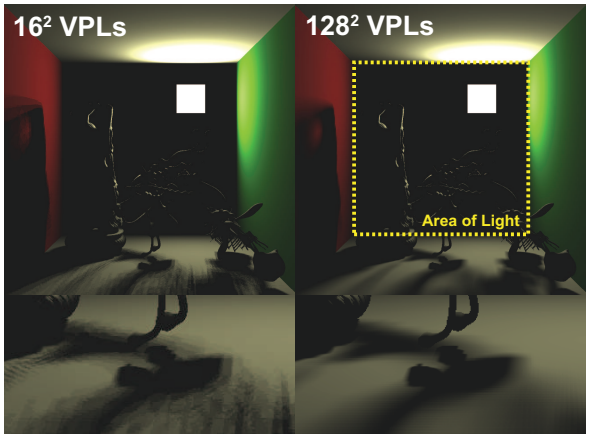

Figure 15: *A pathological scene demonstrating artifacts from a fixed* 16×16 *VPL sampling. Here, a video on the back wall has a white square bouncing around a black screen, so most VPLs uselessly represent large black regions on the light.*

4.7. VPL Sampling

For diffuse and slightly glossy surfaces, our implementation uses a fixed set of 256 VPLs. We found this coarse sampling sufficient; Figure 14 compares a scene with varying VPLs counts. There is no perceptible difference between the images (even in a difference image), despite the factor of 64 change in VPL count. While the larger VPL count does not often yield increased quality, it certainly increases cost. A naive implementation scales linearly with VPLs, though the incremental visibility detailed in Section 3.3.2 may allow sub-linear scaling. We did not explore scaling issues, as we found 256 VPLs a good quality-performance tradeoff; lower sampling does start introducing banding.

While 256 VPLs suffice for many diffuse scenes, pathological scenarios will show artifacts. Figure 15 compares 256 and 16,384 VPLs in such a scene; the geometry is identical to Figure 14, but the video on the back wall contains a white square bouncing around a mostly black screen. In this case our uniform VPL sampling wastes most samples on black regions of the video and undersamples the square.

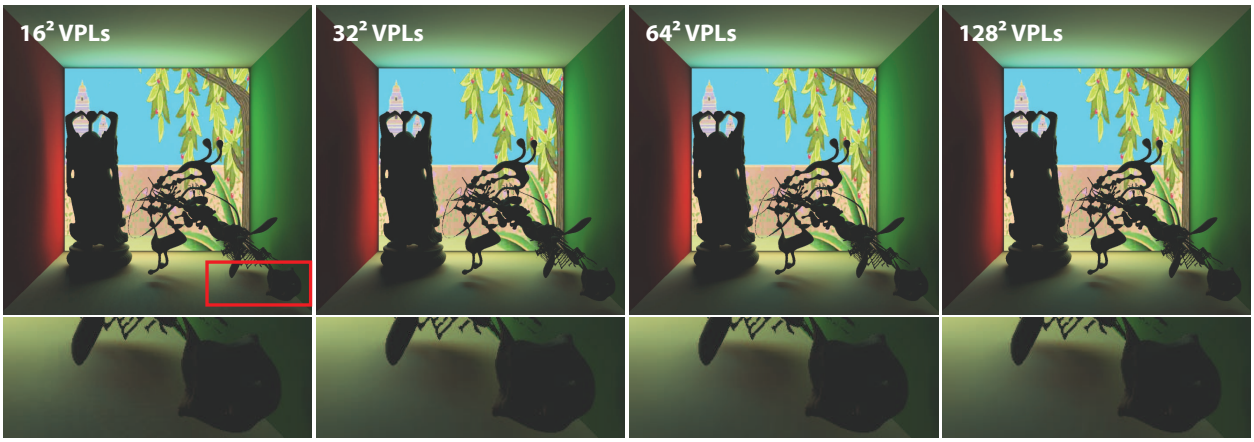

Figure 14: *A scene using* 16^2, 32^2, 64^2, *and* 128^2 *VPLs. Increasing VPL sampling generally yields little visible change for diffuse surfaces.*

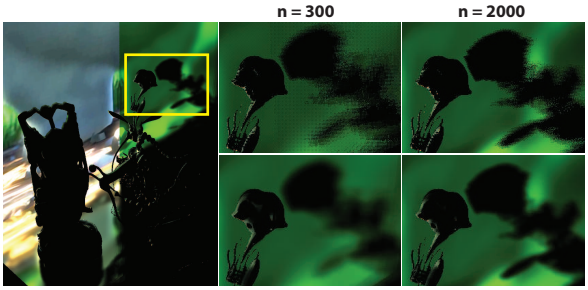

Figure 16: *With sharp BRDFs, simple binary visibility queries may require many VPLs to produce artifact-free results. Here, the limited set of VPLs produces banding and interleaving artifacts with binary visibility queries (top). Querying using filtered variance values yields much better results (bottom).*

An adaptive sampling would reduce artifacts, though a better option would simply treat the video as a dynamic square light.

Non-diffuse surfaces complicate matters. Light samples should be focused inside the material's reflection lobe. One approach would be to adaptively sample the light for each fragment, as proposed by Nichols et al. [NW09]. Instead, as outlined in Section 3.2.1, we use the regular sampling where sufficient and adapt to a per-lobe texture sampling scheme when fixed VPL samples become visible.

4.8. Variance Visibility Queries

Figure 16 illustrates banding and interleaving artifacts that may arise using binary visibility queries to a small number of VPLs. Applying filtering and using variance queries as described in Section 3.3.3 greatly reduces these artifacts.

References

[DS05] DACHSBACHER C., STAMMINGER M.: Reflective shadow maps. In *Proceedings of the Symposium on Interactive 3D Graphics and Games* (2005), pp. 203–231.

[ED08] EISEMANN E., DÉCORET X.: Single-pass gpu solid voxelization for real-time applications. In *Proceedings of graphics Interface* (2008), pp. 73–80.

[KD10] KAPLANYAN A., DACHSBACHER C.: Cascaded light propagation volumes for real-time indirect illumination. In *Proceedings of the Symposium on Interactive 3D Graphics and Games* (2010), pp. 99–107.

[NW09] NICHOLS G., WYMAN C.: Multiresolution splatting for indirect illumination. In *Proceedings of the Symposium on Interactive 3D Graphics and Games* (2009), pp. 83–90.

[RGKS08] RITSCHEL T., GROSCH T., KAUTZ J., SEIDEL H.-P.: Interactive global illumination based on coherent surface shadow maps. In *Proceedings of Graphics Interface* (2008), pp. 185–192.

[SS07] SCHWARZ M., STAMMINGER M.: Bitmask soft shadows. *Computer Graphics Forum 26*, 3 (2007), 515–524.