Efficient Complex Shadows from Environment Maps

Aner Ben-Artzi
Columbia University

Ravi Ramamoorthi
Columbia University

Maneesh Agrawala
Microsoft Research

1. Introduction

Complex illumination adds realism to computer-generated scenes. We focus on realistic lighting from environment maps, as described in [Debevec 1998]. Correct shadowing requires solving for the visibility of each pixel with respect to each light direction. This is extremely expensive, since the illumination can come from any direction in the environment map. We leverage recent sampling methods like [Agarwal et al. 2003], which reduce the environment to a few hundred directional lights. Shadow testing for all sources at each pixel is still the bottleneck in a conventional ray-tracer or other renderer.

We show that the visibility function can be efficiently calculated by exploiting coherence in both the angular and spatial dimensions. Existing work, such as [Guo 1998] and [Hart et al. 1999], makes use of coherence for scenes illuminated by point or area lights. Most of these techniques find discontinuities in image-space, and share visibility within regions bounded by discontinuities, thus reducing the number of primary rays needed. Such methods are difficult to use for sampling environment map illumination because hard discontinuities are not discernible in the image. Our work can be seen as an extension of [Agrawla et al. 2000], who exploited coherence in the visibility of an area light source to reduce the number of shadow rays cast. We present a simple method for efficient coherence-based evaluation of visibility for environment maps.

We identify two important components of an algorithm — reusing visibility calculations to predict the results of tracing shadow rays, and explicitly tracing new shadow rays in regions of uncertainty. The first component involves reusing and possibly warping geometry, and is well-studied in image-based rendering (IBR). However, the reuse of sampled geometry alone does not always provide a sufficiently accurate notion of visibility. Hence, the second component is essential—determining regions or light directions prone to errors in the reconstruction.

2 Exposition

One approach is to evaluate the image in scanline order. For each pixel, we use blockers from the previously evaluated pixels above and to the left. We can then warp those blockers to reconstruct an approximation to the geometry seen at the current pixel. If any blocker warps onto the cell corresponding to some light, we predict that this light is blocked. If no blockers land in the cell of a light it is predicted visible.

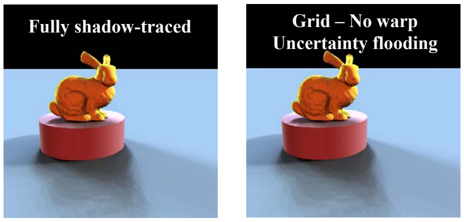

Figure 1. *The bunny scene - 200 light samples*

	shapes				bunny			
number of lights:	50	100	200	400	50	100	200	400
true	100	100	100	100	100	100	100	100
scan, warp, bound.	40.5	34.3	26.3	22.4	52.4	42.6	33.6	25.9
grid, warp, bound.	22.0	17.9	15.1	12.4	21.5	16.6	13.5	11.6
grid, no warp, bound.	21.7	17.1	13.8	10.6	21.2	15.8	12.0	9.3
grid, warp, uncert.	10.4	10.7	11.0	11.3	9.1	9.4	10.2	11.5
grid, no warp, uncert.	9.4	8.8	8.3	7.5	8.1	7.7	7.2	6.7

Table 1. ***Percentage of Shadow-Rays Traced.*** *The entries for the table indicate the percentage of shadow rays where $N \bullet L > 0$ that were traced. Various sampling rates of the env-map were used.*

An alternative approach is to evaluate the image in a coarse-to-fine, grid-based order. The blockers from each of the 4 neighbors (not immediate) are warped into the local light-space for the current pixel. Cells that contain many (four or more) warped blockers are marked as blocked. Cells that contain no warped blockers are marked visible. If a cell contains few (one, two, or three) warped blockers, it is marked as uncertain, and a shadow-ray is explicitly traced to determine visibility.

We discovered that the method of [Agrwala et al. 2000] for tracing lights that are on the boundary of visible/blocked, is not very efficient for environment maps. A new measure, uncertainty flooding, works better. When predicting form multiple previously evaluated pixels, we explicitly trace any light for which the other pixels cannot agree. This tracing is flooded to nearby lights until all shadow-traces agree with the predictions.

We show results of possible combinations for evaluation order and flooding in Table 1. Timings appear in [Ben-Artzi et al. 2004]

3 Conclusions and Future Work

When strong coherence exists in the visibility function, as in most scenes, efficient evaluation can be achieved using very simple coherence-based methods. It is important to couple predictions based on coherence with an appropriate measure for discovering areas of uncertainty. IBR uses the techniques of warping and splatting to recreate geometry. In a ray-tracing environment, we can benefit from the added ability to introduce new samples wherever we believe they are needed. A good measure of uncertainty will guarantee that even scenes with weak coherence, will be rendered accurately, though at a lower prediction rate. Our method can typically reduce the number of shadow tests needed by an order of magnitude, with essentially no loss in quality, thus allowing scenes to be efficiently traced under environment map illumination. Animations are available at www.cs.columbia.edu/cg/. See [Ben-Artzi et al. 2004] for details.

4 References

AGARWAL, S., RAMAMOORTHI, R., BELONGIE, S., AND JENSEN, H. W. Structured Importance Sampling of Environment Maps. *ACM Transactions on Graphics*, 22, 3, 605–612.

AGRAWALA, M., RAMAMOORTHI, R., HEIRICH, A., AND MOLL, L. Efficient Image-Based Methods for Rendering Soft Shadows. *ACM SIGGRAPH 2000*, 375–384.

BEN-ARTZI, A., RAMAMOORTHI, R., AGRAWALA, M. Efficient Shadows from Sampled Environment Maps. *Columbia University Tech Report CUCS-025-04 2004.*

DEBEVEC, P. Rendering synthetic objects into real scenes, *ACM SIGGRAPH 1998*, 189-198

GUO, B. Progressive radiance evaluation using directional coherence maps. *ACM SIGGRAPH '98*, 255-266.

HART, D., DUTRE, P., GREENBERG, D. Direct illumination with lazy visibility evaluation. *ACM SIGGRAPH '99*, 147-154.