

ADAPTIVE DIRECT ILLUMINATION SAMPLING

Petr Vévoda and Jaroslav Krivánek

Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic
Render Legion, s.r.o.

ABSTRACT

We present a new method for sampling direct illumination in scenes containing a large number of light sources. We divide the scene into cells and find a suitable clustering of the light sources for each of the cells. The clusters are then sampled based on their estimated contribution. Our solution is unbiased, scales well with the number of light sources and can be easily incorporated into any existing path tracing rendering system.

MOTIVATION

Direct illumination is often the most dominant component of scene illumination. However, its effective computation in scenes containing large number of light sources (such as the city of Prague at night, shown below) is still an open problem as straightforward evaluation of each light contribution and summing them together is feasible only for trivial number of lights.



Clustering

Ward [1994] suggested reducing the number of visibility tests performed. It still takes too long so methods based on clustering of lights followed [Paquette et al. 1998; Walter et al. 2005]. They scale well with the number of lights at the price of introducing bias.

Sampling



Unbiased solutions usually rely on Monte Carlo estimation, i.e. on random sampling of the lights. However, for a greater number of lights (hundreds and more) even a construction of any non-trivial sampling distribution at each sampling point becomes prohibitive. Shirley et al. [1996] suggested to properly sample only a preselected group of important lights but that can significantly increase variance if any important light is missed.

Combination

Wang and Akerlund [2009] combined sampling of the lights with clustering. They clustered lights similarly to Walter et al. [2005] but then used the clusters for importance-sampling instead of direct illumination estimation. However, their solution is limited to point lights. Donikian et al. [2006] also proposed sampling of clustered lights but their method operates in image space only.

RESULTS

The figure below shows a scene with 4000 light sources rendered in the equal time (20 s) using three different sampling techniques: (a) uniform sampling, (b) sampling proportional to unoccluded light contributions, (c) our method. The uniform sampling is very fast but every iteration yields only a small improvement. On the other hand, the proportional sampling has the highest improvement per iteration, but is slow. Our method has good improvement per iteration and is fast at the same time. As a result, our method significantly outperforms the other two.



The left figure presents the convergence plot of the three sampling techniques in the same scene. Our method clearly converges the fastest. Comparison of the techniques with respect to scalability is shown in the right plot. As expected, the uniform sampling needs almost constant time per pass no matter the number of light sources while the proportional sampling scales linearly. Speed of our method is also close to constant.


CONCLUSION AND FUTURE WORK

The presented method is fast, unbiased and scales well with the number of light sources. However, there is still space for improvements. We are working on removing the main limitation of our method – sampling the light sources irrespective of their visibility. We are experimenting with updating the cached cuts by visibility information obtained during their sampling. We would also like to incorporate surface BRDF in the estimation of cluster contributions and improve the cut search algorithm based on a rigorous analysis of how variance in the resulting image changes during the tree descent.

OUR APPROACH

Similarly to Wang and Akerlund [2009] we also combine sampling of lights with clustering *but we cluster the scene as well*. This allows us to cache clustering of lights for each scene part along with partially estimated contributions. Our solution is fast, unbiased, capable of handling large number of light sources of different kinds and can be used at any point in the scene (e.g. at every vertex of a traced light path so both the direct and indirect illumination is improved).

In a preprocess step, we first hierarchically cluster the light sources into a binary tree (similarly to Walter et al. [2005]) and divide the scene into cells using a regular grid. Then during rendering, we lazily compute one tree cut per scene cell and use these cuts to importance-sample the light sources.



Finding a cut

For a given point we look up the scene cell containing this point and if it does not contain any tree cut yet, a new one is found. Our cut search algorithm is based on estimated radiance with which a light cluster can contribute to any point in the cell. It starts with the tree root and repeatedly replaces the cluster with the highest estimate by its children. The descent stops when estimated contributions of all clusters in the cut fall below a certain threshold.

Sampling the cut

Once the cell contains a cut, the contributions of its clusters to the shading point are estimated. These contributions are then normalized to produce a discrete sampling distribution. The distribution is used to select a cluster, a light source within the cluster is then picked proportionally to its intensity. This way even a very large number of light sources can be sampled fast and adaptively with respect to the current shading point.

ACKNOWLEDGEMENTS

The work was supported by the Charles University Grant Agency, project GA UK 1172416, by the grant SVV-2016-260332, and by the Czech Science Foundation grant 16-18964S. The scene used in this poster was modelled by Ludvík Koutný.

REFERENCES

- DONIKIAN, M., WALTER, B., BALA, K., FERNANDEZ, S., AND GREENBERG, D. P. 2006. Accurate direct illumination using iterative adaptive sampling. *IEEE Transactions on Visualization and Computer Graphics* 12, 3, 353–363.
- PAQUETTE, E., POULIN, P., AND DRETTAKIS, G. 1998. A light hierarchy for fast rendering of scenes with many lights. *Computer Graphics Forum* 17, 3, 63–74.
- SHIRLEY, P., WANG, C., AND ZIMMERMAN, K. 1996. Monte Carlo techniques for direct lighting calculations. *ACM Transactions on Graphics* 15, 1, 1–36.
- WALTER, B., FERNANDEZ, S., ARBREE, A., BALA, K., DONIKIAN, M., AND GREENBERG, D. P. 2005. Lightcuts: a scalable approach to illumination. *ACM Transactions on Graphics* 24, 3, 1098–1107.
- WANG, R., AND AKERLUND, O. 2009. Bidirectional Importance Sampling for Unstructured Direct Illumination. *Computer Graphics Forum* 28, 2, 269–278.
- WARD, G. J. 1994. Adaptive Shadow Testing for Ray Tracing. *Proceedings of the Second Eurographics Workshop on Rendering*, 11–20.