Adaptive Direct Illumination Sampling

Petr Vévoda*
Charles University, Prague
Render Legion, s.r.o.

Jaroslav Křivánek†
Charles University, Prague
Render Legion, s.r.o.

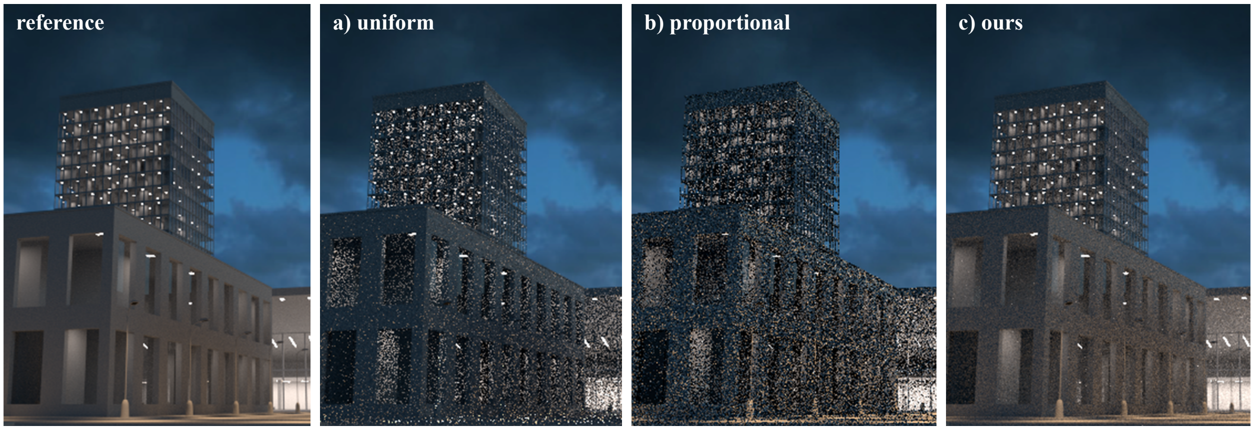

Figure 1: *Same time (20 s) comparison of our method and two other sampling techniques in a scene with 4000 lights: (a) uniform sampling (fast but poor), (b) sampling proportional to unoccluded light contributions (good but slow), (c) our method. All images are rendered using path tracing with next event estimation, direct illumination sampling is therefore carried out at each path vertex.*

Abstract

We present a new method for sampling direct illumination in scenes containing a large number of light sources. We divide a scene into cells and find a suitable clustering of the light sources for each of the cells. The clusters are then importance 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.

Keywords: sampling, direct illumination, clustering

Concepts: **•Computing methodologies → Rendering;**

1 Introduction

Computing direct illumination in a scene involves integrating contributions from all light sources in the scene at every point of the scene. However, evaluating contribution of each light source and summing them together is infeasible for more than trivial number of light sources. Ward [1994] saved some computations by performing visibility tests only for the first few light sources with the highest unoccluded contribution. For greater numbers of light sources this still takes too long so new methods based on clustering of light sources followed [Paquette et al. 1998; Walter et al. 2005]. These methods hierarchically cluster light sources into a tree and then use contribution of adaptively found tree cuts as estimates of direct illumination. They scale well with the number of light sources at the price of introducing bias.

Unbiased solutions usually rely on Monte Carlo estimation of the direct illumination integral, i.e. on computing contribution of only one light source picked randomly (sampled) according to some probability distribution. Variance in the resulting image then highly depends on the choice of the sampling distribution. Ideally, one would like to sample according to a distribution directly proportional to contributions of light sources. Such approach would yield zero variance but requires explicit knowledge of all the contributions. Simplified distributions considering only some part of the light contribution are therefore used. However, as long as the sampling distribution is not constant, its construction for greater numbers of light sources (hundreds and more) soon becomes prohibitive. On the other hand, constant distributions are fast to construct and sample, but they perform poorly (see Figure 1 and 2).

Shirley et al. [1996] suggested to properly sample only a preselected group of important light sources but that can lead to significant increase of variance if any important light source is missed. Wang and Akerlund [2009] combined sampling of the light sources with clustering. Similarly to Walter et al. [2005] they also find cuts in the light tree but they use the cuts for importance sampling the light sources instead of taking them directly as illumination estimates. However, their solution is limited to point light sources. Sampling of clustered light sources was proposed also by Donikian et al. [2006] but their method operates in image space only.

We also combine sampling of the light sources with their clustering but we cluster the scene as well. This allows us to cache clustering of light sources 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).

---

*e-mail:petrvevoda@seznam.cz

†e-mail: krivanek@cgg.mff.cuni.cz

SA '16, December 05-08, 2016, , Macao
ISBN: 978-1-4503-4540-8/16/12
DOI: http://dx.doi.org/10.1145/3005274.3005283

2 Our approach

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.

The lazy initialization works as follows. For a given shading point we look up the enclosing scene cell and if the cell does not contain any tree cut yet, a new one is found. Our cut search algorithm is based on estimated radiance each light cluster can contribute to any point in the cell. We developed radiance estimates for light clusters with different kinds of light sources including point, area, directional or environmental. The cut search algorithm starts with the root of the light tree and repeatedly replaces the cluster with the highest estimate by its two children. The descent stops when estimated contributions of all clusters in the cut fall below a certain threshold. The resulting cut is then stored in the cell along with parts of the cluster estimates that can be reused for all points in the cell. Since often only small percentage of all scene cells is actually used (e.g. only 6% in the scene in Figure 1), lazily finding cut only for cells that are really needed saves a lot of work and memory.

Once the cell contains a tree cut, the contributions of its clusters to the shading point are estimated (with utilization of the cached partial estimates). 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. Note that final sampling of a point on a light source is not subject to our work.

Since we do not take the estimated contributions directly as illumination approximation but use it for importance sampling of the light sources, our method is unbiased. It also allows us to use less conservative radiance estimates as their accuracy influences only variance in the resulting image. This is necessary for the cut search algorithm since conservatively bounding contribution of a light cluster to a scene cell often results in exaggerated estimates.

3 Results

Figure 1 shows a comparison of a scene rendered with our method and with two other sampling techniques, in particular with uniform sampling and with sampling proportional to unoccluded contribution of light clusters. Figure 2 then presents the convergence plot of these three techniques. We can see that our method significantly outperforms the other two. Comparison of the techniques with respect to scalability is shown in Figure 3. 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.

4 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 the surface BRDF in the estimation of cluster contributions and improve the cut search algorithm based on a rigorous analysis of how variance in the image changes during the tree descent.

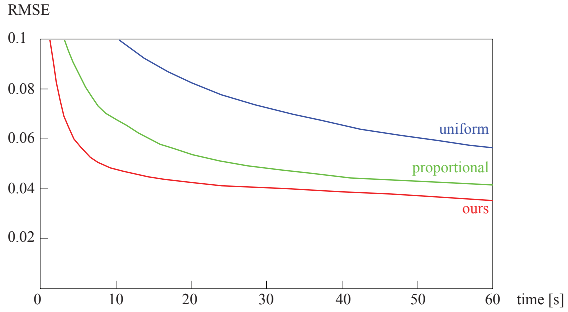

Figure 2: *Plot of convergence in a scene with 4000 lights for our method (red), uniform sampling (blue) and sampling proportional to unoccluded light contributions (green).*

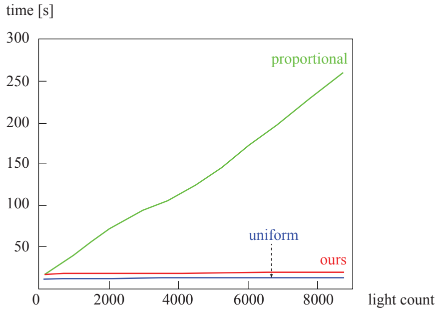

Figure 3: *Plot of rendering time versus the number of light sources for our method (red), uniform sampling (blue) and sampling proportional to unoccluded light contributions (green).*

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 shown in Figure 1 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.