Progressive Lightcuts for GPU

Tomáš Davidovič
Saarland University
Intel VCI Saarbrücken

Iliyan Georgiev
Saarland University
Intel VCI Saarbrücken

Philipp Slusallek
DFKI, Saarland University
Intel VCI Saarbrücken

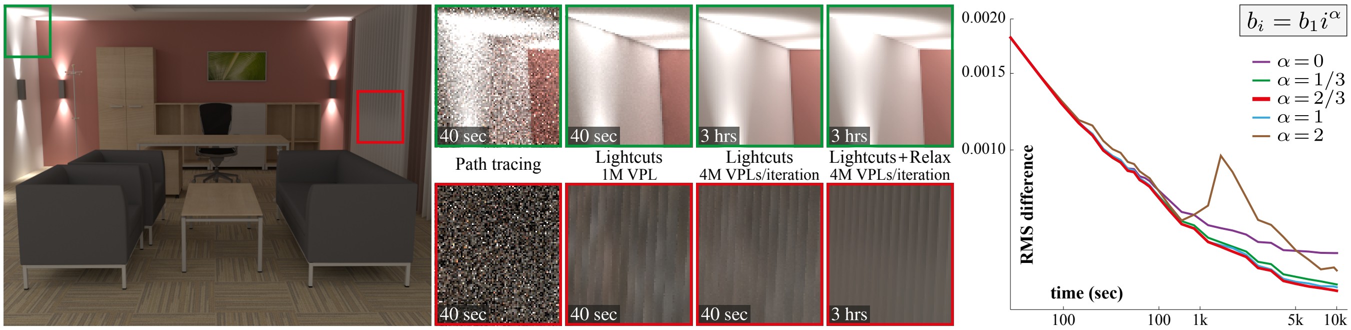

Figure 1: *We test our progressive GPU Lightcuts algorithm on a scene with strong secondary lighting and detailed geometry. In the same time, Lightcuts clearly remains superior to path tracing. However, one million virtual point lights (VPLs) cannot capture the illumination on the blinds in full detail. Increasing the number of VPLs improves quality but consumes more memory; further improvements require accumulating multiple rendering iterations. In order to remove the VPL clamping bias, we additionally apply our progressive clamping relaxation scheme. The plot on the right compares the convergence of different clamping relaxation rates, measured against the reference image on the left.*

Abstract

Lightcuts [Walter et al. 2005] is an attractive rendering algorithm that efficiently handles complex lighting by adaptively approximating the illumination at a point using clusters of virtual point light (VPL) sources. Two of its limitations are the infinite amount of memory required for variance convergence and the bias introduced by VPL contribution clamping. We present a consistent progressive Lightcuts variant, which converges to the correct solution with a bounded memory footprint. This is essential for high quality rendering, especially considering the tight memory budget on the GPU.

Overview

Figure 1 demonstrates that on the GPU Lightcuts remains superior to brute-force path tracing on scenes with non-uniform illumination. However, it also shows that the limited amount of GPU memory is often insufficient to store enough VPLs, e.g. more than 1 million in 1.5GB, to capture the illumination on complex geometry in full detail. We address this problem by performing rendering iteratively, and obtain a high fidelity result via progressive averaging of Lightcuts images, each generated using an new independent set of VPLs. This is made possible by noting that Lightcuts is essentially an unbiased adaptive VPL stratification algorithm [Walter et al. 2005].

Since VPL contributions can be infinite due to a weak singularity at corners, a common approach is to clamp each contribution to some maximum value b. This, however, introduces bias that Kollig and Keller [2004] compensate for using solid angle integration, which avoids the weak singularity. This path tracing based solution, however, in turn again increases variance, which is particularly noticeable at corners in the otherwise smooth Lightcuts images.

Our approach to making Lightcuts consistent, while maintaining its low variance, is based on the observation that the amount of necessary clamping is closely related to the number of samples, i.e. the total number of VPLs used. We develop a progressive clamping relaxation scheme, which reduces clamping as more VPLs are accumulated to the image over multiple independent rendering iterations.

The clamping parameter b trades off variance and bias. An asymptotically unbiased image can be obtained if both the variance and bias of the accumulated result diminish over time. It can be shown that if the variance of the individual rendering iterations increases sub-linearly, consistency is still achieved. For an arbitrary signal bounded by b, its variance is bounded by $b^2/4$. We show that for the specific case of light transport around corners, where the weak singularity occurs, pixel variance increases only at a rate of $O(\sqrt{b})$. At each iteration $i > 1$, we use a clamping constant $b_i = b_1 i^{2/3}$, where b_1 is the initial bound. We show that with this scheme both the cumulative pixel variance and bias diminish at optimal rates.

Results

We numerically compare different clamping relaxation rates, with a clamping constant computed as $b_i = b_1 i^\alpha$, for different values of α. We measure the root mean squared (RMS) difference against a path traced reference solution over time, plotted in Figure 1 right. Without clamping relaxation, i.e. with $\alpha = 0$, the image soon converges to a wrong solution. For any $\alpha \in (0; 1)$ the algorithm is consistent, and on this scene values of $1/3$ and 1 perform similarly to the optimal $\alpha = 2/3$. Relaxation rates $\alpha \geq 1$ are too aggressive.

While we present clamping relaxation is an alternative to clamping compensation, the two can in fact complement each other. If used together, clamping relaxation progressively increases efficiency by reducing the amount of compensation necessary.

Acknowledgements. The office scene was modeled by Jiří Friml.

References

KOLLIG, T., AND KELLER, A. 2004. Illumination in the presence of weak singularities. In *MCQMC Methods*.

WALTER, B., FERNANDEZ, S., ARBREE, A., BALA, K., DONIKIAN, M., AND GREENBERG, D. P. 2005. Lightcuts: a scalable approach to illumination. In *ACM SIGGRAPH 2005 Papers*, ACM, New York, NY, USA, 1098–1107.