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

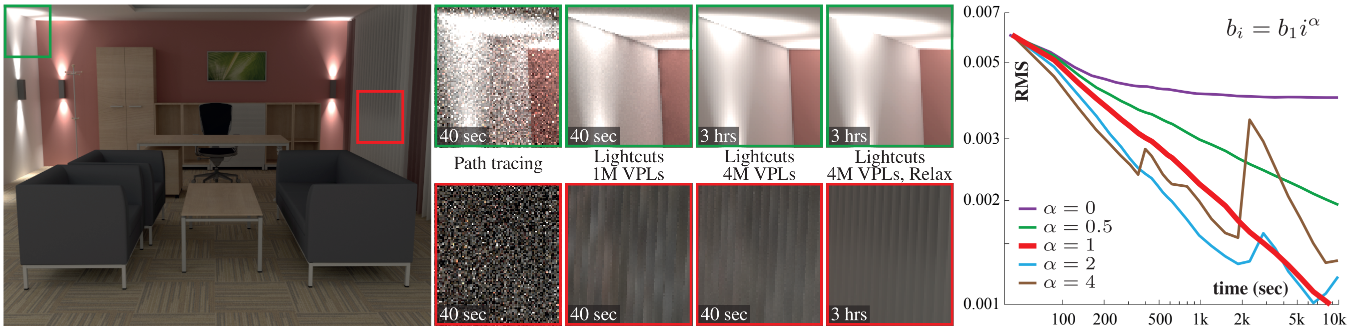

Figure 1: *We test our progressive GPU Lightcuts algorithm on a scene with strong secondary lighting and detailed geometry. In the same time, Lightcuts clearly remains superior to path tracing. However, one million virtual point lights (VPLs) cannot capture the illumination on the blinds in full detail. Increasing the number of VPLs per iteration improves quality but consumes more memory; further improvements require accumulating multiple iterations (bottom). In order to remove the VPL clamping bias, we additionally apply our progressive clamping relaxation scheme (top). The log-log plot on the right compares the convergence of different relaxation schemes, measured as RMS difference 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 using clusters of virtual point light (VPL) sources. Two of its limitations are the infinite memory footprint required for variance convergence and the bias introduced by VPL contribution clamping. We present a progressive Lightcuts algorithm, which is consistent and converges using a fixed amount of intermediate memory. 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 a different set of VPLs. This is made possible by noting that Lightcuts is essentially an unbiased adaptive VPL stratification algorithm [Walter et al. 2005].

Since variance can be infinite due to a weak singularity at corners, a common approach is to clamp the VPL 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, in turn, again increases variance, which is particularly noticeable at geometric 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 clamping necessary is closely related to the number of samples, i.e. total number VPLs, used in the estimate. We develop a progressive clamping relaxation scheme, which reduces clamping as more VPLs are accumulated to the result over the 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. We show that consistency is achieved if the variance of the individual rendering iterations increases sub-linearly. 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 singularities occur, pixel variance increases only linearly with b. Therefore, at each iteration $i > 1$, we use a clamping constant $b_i = b_1 i$, where b_1 is the initial bound. With this scheme, both the cumulative variance and bias diminish over time, ensuring the estimate consistency.

Results

We numerically compare different clamping relaxation schemes, with a clamping constant computed as $b_i = b_1 i^\alpha$ with different values for α. We measure the root mean squared (RMS) difference against a reference path tracing solution over time, plotted in Figure 1 right. Without clamping relaxation (i.e. $\alpha = 0$), after 20 minutes the image converges to a wrong solution. Both the generally conservative $\alpha = 0.5$ scheme as well as the specific $\alpha = 1$ scheme converge to the correct solution with rates $O(N^{-0.2})$ and $O(N^{-0.35})$ respectively, while the schemes with $\alpha \geq 2$ are already too aggressive.

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

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, SIGGRAPH '05, 1098–1107.

SIGGRAPH 2012, Los Angeles, California, August 5 – 9, 2012.
ISBN 978-1-4503-1435-0/12/0008