Radiance Filtering for Interactive Path Tracing

Karsten Schwenk and Timm Drevensek (Fraunhofer IGD)

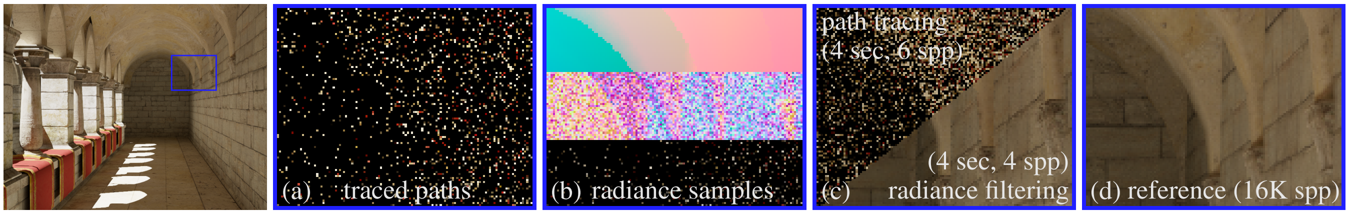

Figure 1: *For each sampled path* (a) *we extract a radiance sample representing the incident indirect radiance. These samples consist of position, direction, and radiance* (b)*. When shading a pixel, we apply its BRDF also to radiance samples of neighboring pixels (weighted by a Gaussian kernel in world-space). Reusing radiance samples reduces variance significantly without losing detail in geometry or texture. Our algorithm can provide reliable previews of global illumination with only a few samples per pixel* $(c + d)$.

1 Introduction

We present a method for noise reduction that is especially tailored to interactive progressive path tracing (PT). The idea is to exploit spatial coherence in the image and reuse information from neighboring pixels. However, in contrast to image filtering techniques (e.g. [Schwenk et al. 2012]), we do not simply filter pixel values or samples of outgoing radiance. Instead, we only reuse the incident indirect radiance of neighboring pixels in a radiance estimation step with a shrinking kernel similar to stochastic progressive photon mapping (SPPM) [Hachisuka and Jensen 2009]. This novel approach significantly reduces the variance in indirect lighting without blurring details in geometry or texture. In equal time comparisons we often achieve higher quality than previous approaches. The primary use case of our algorithm is to provide fast, reliable previews of global illumination. It is also consistent, retains the conceptual simplicity of PT, is orthogonal to importance and stratified sampling, and is easy to integrate into existing renderers.

2 Our Method

The algorithm is sketched in Fig.1 and described in greater detail in the accompanying poster, which also shows some preliminary results. In this document, we will focus on the unique properties of our method and how it compares to related work.

The basic concept of our method is to treat radiance samples of neighboring pixels as independent realizations of the same random variable. The variance of this variable (and thus the noise in the image) can be reduced by taking a weighted average of independent realizations (i.e. by filtering). In general neighboring samples represent different variables, so this assumption is only approximately true and filtering means trading noise for bias. Many approaches filter pixel values and try to limit bias by using some variant of (cross) bilateral filtering, which reduces the influence of strongly biasing pixels. The biggest problem of these approaches is that in most practical situations the bilateral filter will either be too broad and blur fine details in geometry or texture, or it will be too sharp and will not reduce variance enough. Another problem is that the bilateral filter does not handle sharp antialiased edges correctly, because the pixel value is a linear combination of the regions adjacent to the edge and not present in the undiscretized signal itself. Addressing these two issues was the primary motivation for our work.

The key observation is that the unwanted portion of the variance in the outgoing radiance is due to the incident radiance, not the BRDF. So instead of filtering pixel values, we try to reduce variance by only averaging samples of the incident radiance L_i. In the usual notation our radiance estimate is not

$$\widehat{L_o}(x,\omega_o) = \sum_j w_j f_r(x_j,\omega_{i,j},\omega_{o,j}) L_i(x_j,\omega_{i,j})(n_j \cdot \omega_{i,j}),$$

but

$$\widehat{L_o}(x,\omega_o) = \sum_j w_j f_r(x,\omega_{i,j},\omega_o) L_i(x_j,\omega_{i,j})(n \cdot \omega_{i,j}),$$

where the w_j are normalized weights for the samples inside the kernel. Note that with our estimate x, n, and ω_o are taken from the current shading point, not the neighbors, which reduces blur in the factors f_r and $(n \cdot \omega_{i,j})$. However, shadows and sharp glossy reflections are still blurred, as they are included in L_i. Usually we only filter indirect illumination, which is the primary source of noise and can be expected to be reasonably smooth. Our method averages before the pixel reconstruction filter is applied, so we handle antialiased pixels correctly, which is another advantage over image filtering. A disadvantage is the slightly higher overhead due to repeated BRDF evaluations.

The main difference to SPPM is that we distribute radiance samples in image-space during PT, not by a separate photon tracing path. The advantage is that we can expect a sample density of 1 sample per pixel and can adapt the kernel to aim to collect as many samples as needed to reach a user-defined threshold on variance. The assumption that neighboring pixels contain relevant radiance samples breaks down in the presence of geometric edges and complex perfect specular objects, but in most practical cases there will be at least some spatial coherence in the image. Early in the rendering process kernel sizes will be relatively large and will produce the typical artifacts known from PM (e.g. light leaks). However, as the path traced image converges, the variance will decrease and eventually the image-space size of the kernel will drop below the pixel size, at which point our algorithm reduces to standard PT (the sum will only include the sample for the current pixel). This makes our algorithm consistent. A disadvantage in comparison with SPPM is that we inherit the weaknesses of PT with respect to SDS paths, where PM is clearly the superior algorithm.

What was presented here is still work in progress. With future work we plan to improve the performance with complex perfect specular objects. We also want to evaluate our algorithm with bidirectional path tracing, metropolis light transport, depth of field, and motion blur. The last two effects are a further weakness of image filtering.

References

Hachisuka, T., and Jensen, H. W. 2009. Stochastic progressive photon mapping. In *ACM SIGGRAPH Asia 2009 papers*, SIGGRAPH Asia '09, 141:1–141:8.

Schwenk, K., Kuijper, A., Behr, J., and Fellner, D. 2012. Practical noise reduction for progressive stochastic ray tracing with perceptual control. *IEEE CG&A*.

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