Noise Reduction for Progressive Photon Mapping

Zhe Fu Henrik Wann Jensen
University of California at San Diego

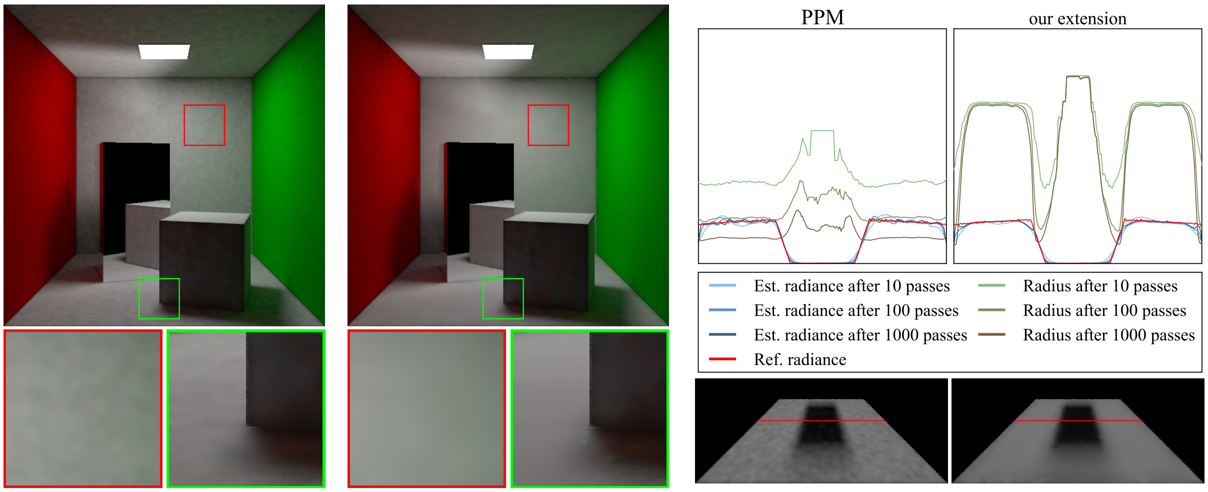

(a) Cornell box by PPM (b) Cornell box by our extension (c) 1D visualization of samples along the line on floor

Figure 1: *A scene of cornell box and a scene of a diffuse surface occluded by a small ceiling rendered by progressive photon mapping and our method. We use 10 million photons in 100 passes to render the two scenes.*

1 Introduction and Motivation

We present an adaptive noise reduction technique for progressive photon mapping, which preserves the consistency of the original algorithm. Although progressive photon mapping converges to the correct result in the limit, the image produced after a finite number of samples can suffer from both bias and noise [Hachisuka et al. 2010]. The bias can be seen as blurry lighting features while the noise can be seen as splotches.

Progressive photon mapping is primarily controlled by one parameter α that determines the rate of the radius reduction around the sample points in the scene. A large α leads to a higher acceptance of incoming photons with a slower radius reduction, while a small α accepts fewer photons and reduces the radius faster. This effectively means that α can be used to adjust the amount of bias and noise present in the scene. A larger value of α results in more bias while a smaller value causes more noise. The original progressive photon mapping algorithm uses one α value for the entire scene, which means a nearly uniform radius reduction across the scene, and consequently a similar amount of smoothing and noise.

In this paper we propose to adjust α adaptively across the scene to get a better balance between bias and noise in the rendered images. Intuitively, in some region of the scene where lighting condition changes slightly or gradually, we do not need to reduce the radius aggressively. Otherwise, these surfaces may end up with high frequency noise. On the other hand, details should be preserved by using smaller radii at locations where light intensity changes rapidly, such as shadow boundaries.

2 Adaptive α adjustment

Our goal is to pick a value for α that results in larger radii in regions with smooth lighting and a smaller radii in regions with rapid lighting changes. Hachisuka et al. [2010] showed how bias can be estimated using the laplacian of the radiance. Unfortunately, the laplacian itself is quite noisy and we found that using just the gradient of the radiance leads to a significantly more robust estimate of the rate on how lighting changes locally. Based on this observation, we have derived an empirical formula for adjusting α that uses a variable, $s = \|\nabla L(\mathbf{x}, \vec{\omega})\| / \|L(\mathbf{x}, \vec{\omega})\|$, to indicate the relative light intensity changes. Note, that the gradient is estimated as a by-product in kernel based progressive photon mapping . Given s we compute a local α value for each sample location as:

$$\alpha(\mathbf{x}) = \frac{1}{1 + e^{\beta(s-\mu)}} \tag{1}$$

where the S-shaped sigmoid function maps a small s to a large α, which in turn results in slow radius reduction and a smooth appearance. It does the opposite for a large s to preserve lighting details. It also ensures that α is between 0 and 1 in order to satisfy the convergence condition of progressive photon mapping. The two parameters μ and β in Equation 1 control the center of the sigmoid function and how steep it is, respectively. We have found that μ can be computed automatically as the 90th percentile of s of all measurement points. In other words, points with 10% largest s are considered to hold scene details, and their radii are reduced more aggressively. β is a new user-specified parameter instead of α, which serves as the effectiveness of our extension. We use $\beta = 0.7$ in our results. The extreme case with $\beta = 0$ is a standard progressive photon mapping with $\alpha = 0.5$.

Figure 1 shows our improvements over the original algorithm in two example scenes. Comparing the close-up in red boxes, our method removes most high frequency noise on the walls or floors. On the other hand, it still preserves shadow boundaries, corners, and caustics as illustrated in the green boxes. Since we greedily lower the radius reduction rate, it is roughly 30% longer to render the images compared with standard progressive photon mapping.

References

HACHISUKA, T., JAROSZ, W., AND JENSEN, H. W. 2010. A progressive error estimation framework for photon density estimation. *ACM Trans. Graph. 29*, 6 (Dec.), 144:1–144:12.

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