

Adaptive Sampling and Bias Estimation in Path Tracing

Rasmus Tamstorf

Henrik Wann Jensen

Graphical Communication
Technical University of Denmark
DK-2800 Lyngby, Denmark
rt@acm.org

mental images GmbH & Co. KG
Fasanenstrasse 81
D-10623 Berlin, Germany
henrik@mental.com

Abstract : One of the major problems in Monte Carlo based methods for global illumination is noise. This paper investigates adaptive sampling as a method to alleviate the problem. We introduce a new refinement criterion, which takes human perception and limitations of display devices into account by incorporating the tone-operator. Our results indicate that this can lead to a significant reduction in the overall RMS-error, and even more important that noisy spots are eliminated. This leads to a very homogeneous image quality. As most adaptive sampling techniques our method is biased. In order to investigate the importance of this problem, a nonparametric bootstrap method is presented to estimate the actual bias. This provides a technique for *bias correction* and it shows that the bias is most significant in areas with indirect illumination.

1 Introduction

Monte Carlo based global illumination methods such as path tracing [6] and bidirectional path tracing [10, 9, 17] are very powerful techniques when highly complex scenes with complex reflection models are rendered. The main disadvantage of Monte Carlo based ray tracing is the high variance of the result which is seen as noise in the rendered images. This noise can be eliminated by increasing the number of samples per pixel, but it is costly due to the slow convergence of the Monte Carlo method.

Adaptive image sampling methods tries to avoid the problem of using a fixed (high) number of samples per pixel by concentrating the samples in the difficult parts of the image. Several methods which apply adaptive sampling to regular ray tracing have been presented. The primary aim of most of these methods is to concentrate the samples along the edges in the image to reduce aliasing effects [13, 12]. In Monte Carlo based ray tracing such as path tracing these edges represents a minor problem compared to the noise seen as spots in the image. Only a few adaptive sampling techniques capable of handling this kind of noise have been presented.

Lee et al. [11] presented a method in which each pixel is sampled based on the

variance of the samples, and Purgathofer [14] presented a technique in which confidence intervals are used. The use of confidence intervals has the advantage that error bounds can be specified explicitly for each pixel. However, as demonstrated by Kirk and Arvo [8] both these methods are biased. To avoid bias, Kirk and Arvo presented a method in which an initial sample set is used only to estimate the necessary number of samples per pixel.

An alternative to using adaptive sampling is filtering in which the noisy spots are located and removed using for example a median filter [5] or a more complex energy preserving filter [15]. These methods have the advantage that it is relatively cheap to reduce the amount of noise in the picture. The accuracy of the filtered image is however difficult to predict since the blurring of noisy pixels also introduces errors in the image.

In this paper we investigate the use of adaptive sampling with path tracing. We use a confidence interval based approach similar to Purgathofer's, but with a modified formula for computing the confidence interval in which a tone-operator is included. This has the advantage that samples are concentrated in those regions where they contribute most to the final appearance of the image. Like other adaptive sampling techniques our method is biased. In order to see whether this bias is significant we estimate it using a statistical technique known as non-parametric bootstrapping.

2 Adaptive sampling

In this section we will consider adaptive sampling using confidence intervals. In addition to [14] we will present a new refinement criterion which includes the tone operator.

Adaptive sampling using confidence intervals is based on an analysis of the pixel radiance, L . This is computed by integrating the image-function, $p(x, y)$, which represents the radiance towards a viewer for any point (x, y) in the rendered image. The radiance through a pixel, \bar{L} , is thus found by :

$$\bar{L} = \int_A p(x, y) f(x, y) dx dy \quad (1)$$

where $f(x, y)$ is an appropriate anti-aliasing filter¹ and A is the support area of the filter. $p(x, y)$ is unknown and \bar{L} is estimated by a Monte Carlo method based on n random samples, $X_i, i = 1, \dots, n$:

$$\hat{\bar{L}} = \frac{1}{n} \sum_i X_i \quad (2)$$

where the samples are distributed according to the chosen filter kernel.

The idea in [14] is to continue sampling until the confidence that the true value, \bar{L} , is within a given tolerance of the estimate has reached a certain level. More formally this means that sampling continues until :

$$P\{\bar{L} \in [\hat{\bar{L}} - d; \hat{\bar{L}} + d]\} = 1 - \alpha \quad (3)$$

¹Since the final image contains one value for each pixel, the Nyquist frequency of interest for anti-aliasing is 0.5 pixel^{-1} . The attenuation of the box filter at this frequency is not very good, and it is our experience that a Gaussian filter provides visually superior images.