EUROGRAPHICS 2006 / E. Gröller and L. Szirmay-Kalos
(Guest Editors)

Volume 25 (2006), Number 3

Optimizing Control Variate Estimators for Rendering

Shaohua Fan[1], Stephen Chenney[2], Bo Hu[1], Kam-Wah Tsui[1], Yu-chi Lai[1]

[1]University of Wisconsin – Madison
[2]Emergent Game Technologies

Abstract

We present the Optimizing Control Variate (OCV) estimator, a new estimator for Monte Carlo rendering. Based upon a deterministic sampling framework, OCV allows multiple importance sampling functions to be combined in one algorithm. Its optimizing nature addresses a major problem with control variate estimators for rendering: users supply a generic correlated function which is optimized for each estimate, rather than a single highly tuned one that must work well everywhere. We demonstrate OCV with both direct lighting and irradiance-caching examples, showing improvements in image error of over 35% in some cases, for little extra computation time.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism Color, shading, shadowing, and texture G.3 [Probability and Statistics]: Probabilistic Algorithms
Keywords: *direct lighting, deterministic mixture sampling, control variates*

1. Introduction

Monte Carlo integration methods offer the most general solution to physically accurate lighting simulation: they handle near-arbitrary geometry, material properties, participatory media, etc. All Monte Carlo methods require an *estimator* that takes the information found in the samples and determines a single final value. A good estimator is unbiased and has low variance. In rendering, the unbiased property guarantees the image has on average the correct pixel values, while variance determines the noise levels in the image, or how much neighboring pixels tend to differ in value.

There are many possible estimators, each of which combines the samples in a different way to get the final answer. If we focus on unbiased estimators, then a good strategy is to choose one that minimizes variance while remaining relatively fast to compute. The most common estimator in rendering is the sample mean or an importance weighted mean. Alternatives exist, however, such as the Multiple Importance Sampling (MIS) estimator [VG95] or control variate estimators [SSSK04] (also referred to as correlated sampling).

In this paper we apply an Optimizing Control Variate (OCV) estimator to the problem of estimating irradiance integrals for direct lighting. The same basic problem is also a sub-component of many rendering algorithms, such as irradiance caching and photon-map gathering, for which we also demonstrate some results. The OCV estimator solves a small optimization problem to find a good control variate distribution given a set of samples. Unlike existing control variate methods which require a single control variate distribution for all estimates, OCV allows the distribution to vary over the scene depending on surface properties and lighting conditions. Furthermore, users are not burdened with finding an optimal correlated function; they can provide a generic parameterized function that the estimator optimizes.

OCV works with the *deterministic mixture sampling* (DMS) framework for constructing importance functions, sampling from them, and computing estimates from the samples [OZ00]. In addition to providing better estimators, DMS allows for multiple importance sampling functions to be combined in a general way. The optimizing nature of the estimator ensures that the combination of samplers performs at least as well as the best among them. In this way, OCV can be viewed as a generalization of multiple importance sampling.