Eurographics Symposium on Rendering (2005)
Kavita Bala, Philip Dutré (Editors)

Importance Resampling for Global Illumination

Justin F. Talbot David Cline Parris Egbert

Brigham Young University

Abstract

This paper develops importance resampling into a variance reduction technique for Monte Carlo integration. Importance resampling is a sample generation technique that can be used to generate more equally weighted samples for importance sampling. This can lead to significant variance reduction over standard importance sampling for common rendering problems. We show how to select the importance resampling parameters for near optimal variance reduction. We demonstrate the robustness of this technique on common global illumination problems and achieve a 10%-70% variance reduction over standard importance sampling for direct lighting. We conclude that further variance reduction could be achieved with cheaper sampling methods.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-dimensional Graphics and Realism

1. Introduction

The goal of global illumination is the creation of physically realistic-looking images from descriptions of virtual scenes. Kajiya [Kaj86] first expressed this difficult problem as a recursive integral which he called the "Rendering Equation." More recently, Veach [Vea97] reformulated this equation as a non-recursive integral over light paths.

Since these integrals typically cannot be solved analytically, Monte Carlo integration [MU49] is commonly used to approximate them. Monte Carlo integration is a probabilistic process and is subject to variance which appears as noise in the rendered image. To combat this problem, a number of variance reduction techniques have been developed. One of the most common of these in global illumination is importance sampling.

Importance sampling refers to the general technique of carefully choosing a sampling distribution for Monte Carlo integration. It can be shown that the more closely proportional the probability density function (pdf) is to the function being integrated, the lower the variance of the Monte Carlo estimate.

To use importance sampling, it must be possible to generate samples with the distribution defined by the pdf. Two common techniques used for generating these samples are cumulative density function (CDF) inversion and rejection sampling. A third technique, Metropolis sampling, has been explored recently. In this paper we use a fourth technique, importance resampling.

Importance resampling, unlike the other techniques, generates samples that are only approximately distributed according to the desired pdf. Thus, additional care is required to preserve unbiasedness. Using importance resampling to generate the samples for importance sampling produces a variance reduction technique that we call Resampled Importance Sampling (RIS). Standard importance sampling is a special case of RIS. RIS is more robust than standard importance sampling and can reduce variance significantly.

First, we discuss previous work in Section 2. In Section 3, we describe importance resampling. We then develop the general RIS estimate and discuss its mathematical properties in Section 4. We analyze the variance of the estimate and choose robust resampling parameters to achieve near optimal variance reduction. In Section 5, we present a few case studies using RIS in global illumination. We show that importance resampling can reduce the variance and increase the robustness of Monte Carlo global illumination algorithms. We conclude and discuss possible future improvements to RIS in Section 6.

2. Background

Importance sampling is a very effective variance reduction technique for many of the problems found in global illumination. It requires generating samples from probability density functions. These samples can be generated by CDF inversion, rejection sampling, Metropolis sampling, or importance resampling.

CDF inversion is the most common sampling technique. When possible, it permits exact sampling of the pdf which provides maximum variance reduction. However, it requires being able to integrate the pdf and invert the CDF. Finding a pdf that is nearly proportional to the integral function *and* is easily CDF invertible is a very difficult problem. For some special cases in global illumination, pdfs have been found that are invertible. For example, distributions have been developed for sampling the direct lighting from spherical light sources [SWZ96] and arbitrary polygons [Arv95] and for sampling Phong [LW94] or Ward [War92] Bidirectional Reflectance Distribution Functions (BRDFs). Some recent work has focused on finding good pdfs for sampling environment maps [ARBJ03, KK03].

Rejection sampling can produce samples without the need to find or invert the CDF. However, rejection sampling is often avoided because it does not permit stratified sampling and it does not have a fixed time bound for generating a sample. Also, rejection sampling requires finding a bounding function which can be difficult to do.

Metropolis sampling [MU49] generates a Markov chain with a stationary distribution that is equal to the desired sampling distribution. The advantage of Metropolis sampling is its generality. It can be used to generate samples from even the most complex distributions. However, Metropolis sampling has a large start-up cost if unbiased samples are desired. Additionally, finding efficient transition functions, which is necessary to reduce variance, is domain specific and can be difficult. Metropolis sampling was first used in global illumination by Veach and Guibas [VG97].

Importance resampling has been used informally in the global illumination literature. Lafortune et al. [LW95] used importance resampling to decrease the number of visibility tests necessary in bidirectional path tracing. Shirley et al. [SWZ96] suggested using resampling to improve direct lighting computations. Burke [Bur04] used importance resampling to sample the distribution of the product of a Phong BRDF model and an illuminating environment map.

Other attempts have been made to generalize importance sampling for global illumination problems. Multiple Importance Sampling (MIS), developed by Veach [Vea97], permits the use of multiple sampling distributions through careful weighting. Weighted Importance Sampling [BSW00], like RIS, uses two pdfs to reduce the variance in a Monte Carlo estimate. However, Weighted Importance Sampling is consistent, not unbiased, and doesn't permit optimizing the computation effort to further reduce variance. Combined Correlated and Importance Sampling [SSSK04] uses a pseudo-optimal combination of two variance reduction techniques. It requires computing a closed form approximation of the integral function which limits its generality.

3. Importance Resampling

Importance resampling is a common method in computational statistics for generating samples from difficult distributions. It is commonly used in sequential importance sampling and particle filtering [DdFG00]. It can also be used to generate samples from Bayesian posterior distributions [GCSR04].

Importance resampling was first described by Rubin [Rub87]. Here we briefly describe importance resampling and try to provide an intuitive explanation of why it works.

Assume we want to generate samples from a sampling distribution with pdf g, but cannot do so directly (e.g. using the CDF inversion technique) because g does not have an analytic closed form or is too complex to integrate and invert. We can, instead, generate a set of samples from a source distribution, p, weight these samples appropriately, then *resample* these samples by drawing a single sample from them with probability proportional to its weight.

Importance resampling:

1. Generate M samples ($M \geq 1$) from the source distribution p, $\mathbf{X} = \langle X_1, ..., X_M \rangle$.
2. Compute a weight for each sample, w_j.
3. Draw a single sample Y from $\mathbf{X}$ with probability proportional to $\langle w_1, ..., w_M \rangle$.

If we choose $w_j = \frac{g(X_j)}{p(X_j)}$, then the resulting sample Y will be approximately distributed according to g. The effect of the resampling step is to take samples from the source density, p, and "filter" them, so that the resulting sample, Y, has a distribution that approximates g.

We can view M, the number of samples, as a distribution interpolation variable. When $M = 1$, Y is distributed according to p. As $M \to \infty$, the distribution of Y approaches g. Typically, M must be very large to make the bias introduced by the finite M approximation negligible.

As an example, Figure 1 shows the distribution of Y for various values of M when p is uniform and $g \propto \cos(\theta) + \sin^4(6\theta)$.

4. Resampled Importance Sampling

Combining importance resampling with importance sampling produces a variance reduction technique we call Resampled Importance Sampling (RIS).

Assume we want to find the integral I of a function $f(x)$:

$$I = \int_{\Omega} f(x)\, d\mu(x)$$

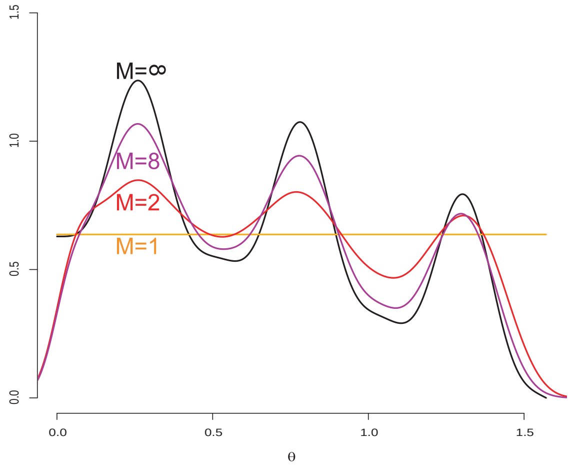

Figure 1: *Distributions resulting from importance resampling for different values of M. At $M = 1$, the distribution is p. At $M = \infty$, the distribution is g. For other values of M, the distribution interpolates between p and g, though the exact manner of interpolation is unknown. The low values on the left side for $M = 2$ and $M = 8$ are artifacts of the density estimation method.*

We also have two probability density functions. The source pdf, p, can be sampled readily, but may be a poor approximation to f. The sampling pdf, g, is a good approximation to f, but it may be unnormalized and difficult to sample. Standard importance sampling limits us to using only p. We would like to generalize importance sampling so that we can also use g to improve our estimate. RIS allows us to use g in an unbiased manner by using importance resampling to draw samples approximately from g.

Given the samples $\mathbf{X}$ and Y from the importance resampling process, we develop the RIS estimator as a form of weighted importance sampling:

$$\hat{I}_{ris} = \frac{1}{N}\sum_{i=1}^{N} w(\mathbf{X}_i, Y_i)\,\frac{f(Y_i)}{g(Y_i)}$$

The weighting function w must be chosen to correct for both the fact that g is unnormalized and for the fact that the density of Y only approximates g. The appropriate choice of w is surprisingly simple. It is the average of the weights computed in the resampling step:

$$w(\mathbf{X}_i, Y_i) = \frac{1}{M}\sum_{j=1}^{M} w_{ij}$$

Combining these two equations gives the RIS estimate:

$$\hat{I}_{ris} = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{f(Y_i)}{g(Y_i)}\cdot\frac{1}{M}\sum_{j=1}^{M}\frac{g(X_{ij})}{p(X_{ij})}\right) \qquad (1)$$

When $M = 1$, RIS reduces to standard importance sampling.

For the RIS estimate to be unbiased, two conditions must hold. First, g and p must be greater than zero everywhere that f is non-zero [Tal05]. Second, M and N must be greater than zero.

4.1. Variance Analysis

Since the RIS estimator is unbiased, the only error in the estimate is due to the variance:

$$V(\hat{I}_{ris}) = \frac{1}{N}\left[\frac{1}{M}(e_3 - e_2) + (e_2 - e_1)\right] \qquad (2)$$

where,

$$e_1 = E\left(\frac{f}{p}\right)^2$$

$$e_2 = E\left(\frac{f^2}{gp}\right)E\left(\frac{g}{p}\right)$$

and

$$e_3 = E\left(\frac{f^2}{p^2}\right)$$

The derivation of this equation is given by Talbot [Tal05].

Note that the domain of the variance will be the range of f. Since, in global illumination, the range of f is typically n-component spectral values, the variance will be an n-component vector.

Equation (2) is best understood by comparison with the variance of standard importance sampling:

$$V(\hat{I}_{is}) = \frac{1}{N}(e_3 - e_1)$$

The variance of standard importance sampling is the vector $e_3 - e_1$. Taking more samples inversely scales this vector (i.e. decreases the variance).

RIS splits this vector into the sum of two vectors, $e_2 - e_1$ and $e_3 - e_2$ (see Figure 2). The relative length of the vectors depends on the target distribution g. If $g \propto p$, then the length of $e_3 - e_2$ is zero. If $g \propto f$, then the length of $e_2 - e_1$ is zero.

In RIS, increasing N decreases the length of both vectors. Increasing M, however, decreases the length of just $e_3 - e_2$. The choice of which to increase clearly depends upon the relative lengths of the two vectors and the relative computational expense of increasing M versus N. The following section explores this issue.

Figure 3 shows the effect of varying M and N when used in direct lighting.

4.2. Choosing Parameters for RIS

When using RIS we can freely choose p, g, M, and N within the unbiasedness constraints given in Section 4. Clearly

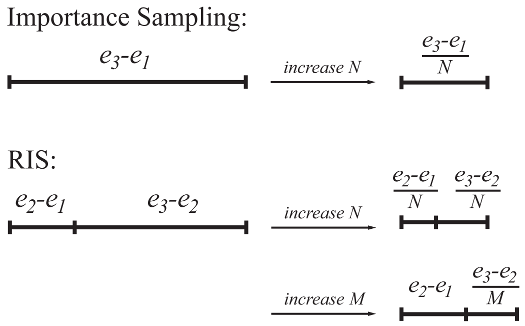

Figure 2: *The effect of increasing M or N on the variance vector of importance sampling and RIS. Increasing N shrinks the entire variance vector of both importance sampling and RIS. Increasing M in RIS shrinks just the length $e_3 - e_2$. However, this may decrease the overall variance more cheaply than increasing N.*

some choices will lead to lower variance than others. In this section we briefly discuss choosing g and p. We then formally show how to find optimal values of M and N, for fixed choices of g and p.

Equation (2) suggests three guidelines for choosing g and p. First, g should be more proportional to f than p is to f. If this is not true, standard importance sampling will have equal or lower variance than RIS. Second, g and p should be as proportional to f as possible. This directly reduces the variance. Third, g and p should be computationally cheap to sample and evaluate (in comparison to f). RIS depends upon evaluating g and p multiple times for each sample. If g and p are expensive, standard importance sampling will be more efficient.

We now derive a heuristic for choosing near optimal values for M and N. The values are chosen to minimize the overall variance of the RIS estimate under a fixed computation time constraint and given fixed choices of p and g.

If the total execution time available for computing $\hat{I}_{ris}$ is T, then we have the following constraint:

$$T = MNT_X + N(r(M) + T_Y)$$

where T_X is the time necessary to draw a sample from p and compute its weight. The function r is the time necessary to perform the resampling step. T_Y is the time necessary to compute Equation (1), given that the resampling has already been done.

There are a number of extremely efficient techniques for drawing samples from discrete distributions, so, in practice, $r(M)$ will be negligible. Ignoring $r(M)$ gives the constraint:

$$T = MNT_X + NT_Y$$

As we noted earlier, the variance resulting from Equation (2) will be an n-component vector. To minimize the variance we must choose a real-valued length function over the spectral vectors. We use the l^2-norm. Perceptually-based measures could also be used.

Using this constraint, we minimize Equation (2) by substitution to find a near optimal value of M:

$$M = \sqrt{\frac{|e_3 - e_2|}{|e_2 - e_1|}\frac{T_Y}{T_X}} \quad (3)$$

where the vertical bars represent the chosen length function.

To ensure that $M \geq 1$ and $N \geq 1$, which is necessary for the resampling process, we first clamp the resulting $M \geq 1$. We then solve for N

$$N = \frac{T}{MT_X + T_Y}$$

and clamp the resulting $N \geq 1$. We solve again for M using the clamped N:

$$M = \frac{\frac{T}{N} - T_Y}{T_X}$$

Finally, when sampling, we probabilistically take the floor or the ceiling of M and N such that the expected value remains the same.

As should be expected, the optimal values of M and N are functions of the variance and the execution time of the two portions of the estimate. Figure 3 shows the optimal values of M and N computed using Equation (3) in a direct lighting application.

4.3. Robust Approximations of M and N

In practice, the true optimal values of M and N cannot be computed since Equation (3) relies on e_1, e_2, e_3, T_X, and T_Y which are unknown. If we were to estimate the unknown parameters, we could compute approximate values for M and N. Unfortunately, computing all the unknown parameters can be difficult or very time consuming. In this section we introduce a robust approximation of Equation (3) that only requires estimates for T_X and T_Y. These values are very simple to compute in global illumination applications.

If we only estimate T_X and T_Y, a provably robust heuristic for choosing M and N is to assign equal time to both stages of the resampling process:

$$MNT_X = NT_Y$$

This is equivalent to using:

$$M = \frac{T_Y}{T_X} \quad (4)$$

in place of Equation (3) in the previous section. We use M_r and N_r to denote the robust values that result from Equation (4).

Within the overall time constraint, no other MN can be

Figure 3: *Dragon sampled with a single primary ray and using RIS with different values of M and N to compute direct lighting (in equal time). On the left, $N = 20$, $M = 1$ and on the right $N = 1$, $M = 60$. The right image has less variance except where visibility is a major component of the variance. The lower image shows the optimal ratio of M and N. Green (lighter) pixels correspond to larger M, red (darker) pixels to larger N.*

more than twice $M_r N_r$. Also, no other N can be more than twice N_r. With Equation (2), this implies that RIS using M_r and N_r will have no more than twice the variance of RIS using the true optimal values of M and N.

Although the resulting bound on the variance is poor, we have found that this approximation works very well in practice. It is very cheap to compute and it avoids the sometimes severe noise and artifacts associated with Monte Carlo estimates of e_1, e_2, and e_3.

Finally, we note that if we cannot estimate T_X and T_Y, a robust default choice is $T_X = T_Y$. Plugging this into Equation (4) results in $M = 1$, which is standard importance sampling.

5. Applications

In this section we demonstrate some of the properties of RIS when applied to different problems in global illumination.

5.1. Sampling the Direct Lighting

Using RIS for direct lighting calculations has already been addressed, informally, by Shirley et al. [SWZ96] and Burke [Bur04]. Here we generalize their approaches and show how to approximate the optimal values of M and N for direct lighting.

The direct lighting equation is

$$f = F_s G V L_e$$

where F_s is a BRDF term, G is a geometry term, V is a binary visibility term, and L_e is the emitted light.

We need to choose a density g that approximates f and is cheap to compute. Shirley et al. and Burke chose

$$g = F_s G L_e$$

dropping the visibility term which is usually the most expensive part of lighting computation.

In the RIS framework we can see that it would be equally valid to approximate f in any other way. For example, as the computational expense of evaluating F_s or L_e increases, due to more physically realistic surface models or to the calculation of these terms in complex shader programs, it may be more efficient to use $g = GV$. Also, if any of the dropped terms can be approximated efficiently, the approximation should be used to improve g. In our implementation, computing the visibility is still the most expensive operation so we follow Shirley et al. and Burke and set $g = F_s G L_e$.

Since Shirley et al. and Burke have already shown that resampling can be effective in direct lighting, the rest of this section shows how to approximate the optimal values of M and N. We show that using these approximate values can reduce the variance past that which is possible with standard importance sampling alone.

Estimating e_1, e_2, and e_3, in order to use Equation (3) is too computationally expensive. Instead we will use the approximation given by Equation (4) to compute robust values of M and N. As described in Section 4.3, we must first approximate T_X and T_Y. To do this, we cast a few thousand primary rays. We then track the time necessary to compute the direct lighting at the first hit point. T_X is the average time necessary to sample the light source and compute g. T_Y is the average time to check the visibility. The time necessary to estimate these values is negligible.

Across scenes of similar complexity, the values of T_X and T_Y will probably be quite stable. Thus, these values could be precomputed for a particular renderer implementation. If precomputed, robust RIS would require absolutely no extra computation time over standard importance sampling.

The images in Figure 4 show a dragon lit by two polygonal light sources and an environment map. The left image uses standard importance sampling. The right uses the robust values of M and N computed from estimated T_X and T_Y values. In this scene we achieve a 70% reduction in overall variance compared to just using standard importance sampling.

This variance reduction is scene dependent. If the environment map is removed from the scene, RIS only gives a 10% variance reduction. This difference is largely due to the fact

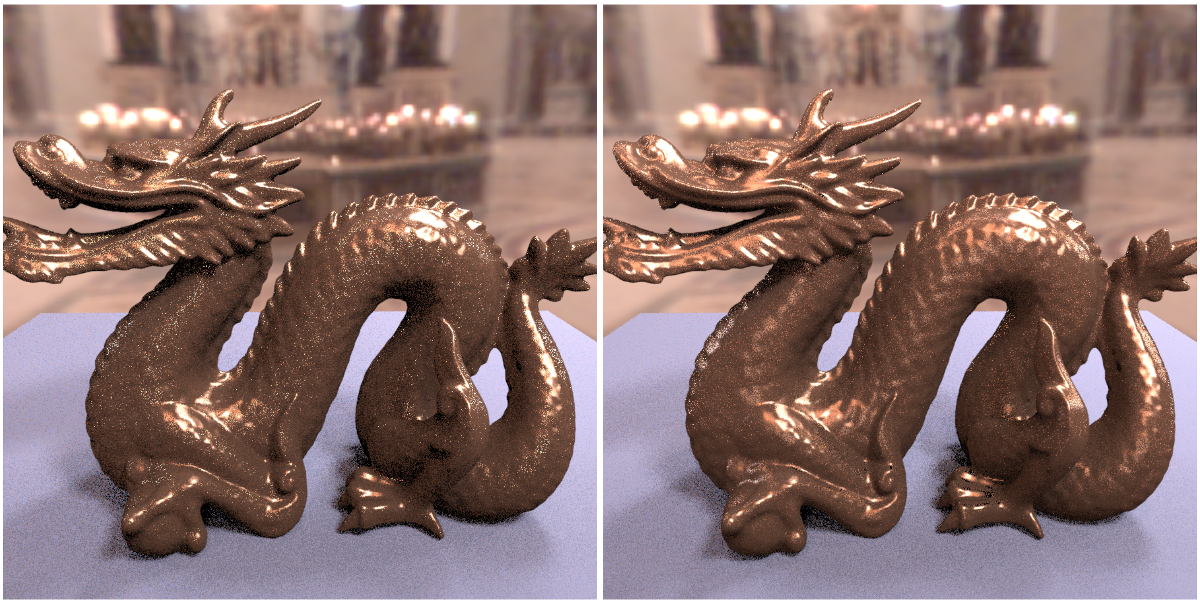

Figure 4: *Robust RIS for direct lighting. Both images are rendered using 20 primary rays and approximately equal computation times. The left image uses standard importance sampling,* $M = 1$, $N = 20$, *for the direct lighting. The right image uses the computed robust values from Equation (4),* $M = 15.75$, $N = 10.64$. *There is a 70% reduction in variance.*

that the environment map introduces a lot of variance into the L_e term. This term is included in g and thus RIS can reduce its variance. Without the environment map, much of the variance is due to visibility. Since visibility is not included in g, RIS cannot reduce its variance.

5.2. Robust Sampling of BRDFs

In the previous section we showed that RIS can produce better results than standard importance sampling. In this section we demonstrate that RIS can be more robust than importance sampling. Specifically, we show that RIS can be applied to a wide class of integration problems *without* change. Importance sampling, on the other hand, requires a special case for each specific problem. We use the example of Bidirectional Reflectance Distribution Function (BRDF) sampling to demonstrate.

BRDFs represent the relationship between incident and exitant light at a surface. Traditionally, BRDFs have been sampled with special case distributions developed for each specific BRDF model. Implementing all of these in a global illumination renderer can be very time consuming. This approach can also become unwieldy when the parameters of the BRDF are allowed to vary spatially (as in Bidirectional Texture Distribution Functions). Furthermore, if the BRDF is specified using a shading language, automatically creating a distribution would be difficult.

We would like to find a more robust solution. Ideally, it will improve the sampling of any BRDF model whether or not a good distribution is available for importance sampling. Resampled Importance Sampling can do this.

When sampling BRDFs, we want to sample

$$f = F_s(x, x', x'')\cos(\theta) L_i(x', x'')$$

where F_s is the value of the BRDF at point x', θ is the angle between the normal at x' and the vector $\overrightarrow{x'x''}$ and L_i is the light incident on point x' from x''. In the rest of the discussion we will drop the parameters to these functions.

To use RIS we need to choose a function g that is both closely proportional to f and is inexpensive to compute. Here we will take a general approach and choose:

$$g = F_s \cos(\theta)$$

This choice is very general since we have only assumed that the BRDF can be evaluated, which is required for Monte Carlo integration anyway. Since this choice of g works for any BRDF, we can implement RIS once for all BRDF sampling.

In the following examples, we recognize that standard importance sampling could perform much better than the results we show. However, that would require a specialized sampling distribution for each case. Our goal is to show that RIS is more robust since a single implementation can dramatically improve the sampling of very different BRDF models.

Figure 5 shows two pairs of spheres each sampled with p equal to a uniform distribution over the hemisphere. The first two spheres have a diffuse BRDF. Our chosen p matches the BRDF exactly, but does not take into account the cosine of the angle with the normal. RIS reduces the variance significantly. The second two spheres have a Cook-Torrance BRDF. In this case, p is a very poor sampling density. Nevertheless, RIS still manages to dramatically improve the sampling quality. We emphasize that no changes are made to the RIS implementation for either BRDF.

In these examples we used a uniform distribution for p for simplicity. In practice, a better default density would be a cosine-weighted hemisphere with lobes in the reflective and retroreflective directions. In cases where a specialized distribution is available for a BRDF model, that can be used instead of the default density.

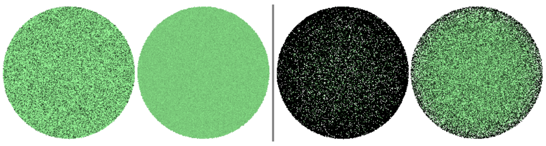

Figure 5: *Uniformly-lit spheres sampled with a uniform hemispherical distribution. The first two are perfectly diffuse and the second two use a Cook-Torrance BRDF. The first sphere in each pair is rendered without resampling ($N = 1, M = 1$). The second sphere in each pair is rendered with RIS ($N = 1, M = 20$). RIS greatly reduces the variance independent of the BRDF model used.*

6. Conclusions and Future Work

We have presented a simple explanation of importance resampling. We have shown how to use importance resampling as the sample generation technique for importance sampling. The resulting variance reduction technique, Resampled Importance Sampling, is a generalization of standard importance sampling.

We have shown how to compute the optimal resampling parameters, M and N, and how to choose robust approximations of M and N that require significantly less computation time. We have shown that RIS increases the robustness of importance sampling when used in some global illumination problems and we have achieved a 10%-70% variance reduction for the direct lighting in a complex scene.

More work needs to be done to find good choices for g and p. Unexpectedly, drawing samples from p accounts for the majority of T_X. A lot of work has gone into speeding up tracing rays [Wal04], but not much work has gone into finding faster sample generation techniques. Since RIS uses M samples at a time, techniques that generate samples in parallel (perhaps using SSE) could be very useful.

In Equation (3), we used the Euclidean lengths of $e_3 - e_2$ and $e_2 - e_1$. This could be improved upon by using a perceptual distance measure.

Burke [Bur04] has done some initial work on combining RIS and MIS. This work needs to be extended to a more general approach.

Acknowledgements

The authors would like to thank the Stanford University Computer Graphics Laboratory for the dragon model, Paul Debevec for the environment map, and Amy Talbot for her careful editorial work.

References

[ARBJ03] AGARWAL S., RAMAMOORTHI R., BELONGIE S., JENSEN H. W.: Structured importance sampling of environment maps. *ACM Trans. Graph. 22*, 3 (2003), 605–612.

[Arv95] ARVO J.: Stratified sampling of spherical triangles. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques* (1995), ACM Press, pp. 437–438.

[BSW00] BEKAERT P., SBERT M., WILLEMS Y. D.: Weighted importance sampling techniques for Monte Carlo radiosity. In *Proceedings of the 11th Eurographics Workshop on Rendering* (2000), Eurographics, pp. 35–46.

[Bur04] BURKE D.: *Bidirectional Importance Sampling for Illumination from Environment Maps*. Master's thesis, University of British Columbia, 2004.

[DdFG00] DOUCET A., DE FREITAS N., GORDON N. (Eds.): *Sequential Monte Carlo Methods in Practice*. Springer Verlag, 2000.

[GCSR04] GELMAN A., CARLIN J. B., STERN H. S., RUBIN D. B.: *Bayesian Data Analysis*, 2 ed. Chapman and Hall/CRC, London, 2004.

[Kaj86] KAJIYA J.: The rendering equation. In *Proceedings of SIGGRAPH 1986* (1986), Evans D. C., Athay R. J., (Eds.), Computer Graphics Proceedings, Annual Conference Series, ACM, ACM Press / ACM SIGGRAPH, pp. 143–150.

[KK03] KOLLIG T., KELLER A.: Efficient illumination by high dynamic range images. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2003), Eurographics Association, pp. 45–50.

[LW94] LAFORTUNE E. P., WILLEMS Y. D.: *Using the Modified Phong BRDF for Physically Based Rendering*. Tech. Rep. CW197, Leuven, Belgium, 1994.

[LW95] LAFORTUNE E. P., WILLEMS Y. D.: Reducing the Number of Shadow Rays in Bidirectional Path Tracing. In *Proceedings of the Winter School of Computer*

Graphics and CAD Systems '95 (Plzen, Czech Republic, 1995), Skala V., (Ed.), University of West Bohemia, pp. 384–392.

[MU49] METROPOLIS N., ULAM S.: The Monte Carlo method. *Journal of the American Statistical Association 44*, 247 (1949), 335–341.

[Rub87] RUBIN D. B.: A noniterative sampling/importance resampling alternative to the data augmentation algorithm for creating a few imputations when fractions of missing information are modest: the SIR algorithm. Discussion of Tanner and Wong. *Journal of the American Statistical Association 82* (1987), 543–546.

[SSSK04] SZÉCSI L., SBERT M., SZIRMAY-KALOS L.: Combined correlated and importance sampling in direct light source computation and environment mapping. *Comput. Graph. Forum 23*, 3 (2004), 585–594.

[SWZ96] SHIRLEY P., WANG C., ZIMMERMAN K.: Monte Carlo techniques for direct lighting calculations. *ACM Trans. Graph. 15*, 1 (1996), 1–36.

[Tal05] TALBOT J.: *Importance Resampling for Global Illumination*. Master's thesis, Brigham Young University, 2005.

[Vea97] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1997.

[VG97] VEACH E., GUIBAS L. J.: Metropolis light transport. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1997), ACM Press/Addison-Wesley Publishing Co., pp. 65–76.

[Wal04] WALD I.: *Realtime Ray Tracing and Interactive Global Illumination*. PhD thesis, Computer Graphics Group, Saarland University, 2004. Available at http://www.mpi-sb.mpg.de/~wald/PhD/.

[War92] WARD G. J.: Measuring and modeling anisotropic reflection. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques* (1992), ACM Press, pp. 265–272.