

Improved Indirect Photon Mapping with Weighted Importance Sampling

László Szécsi and László Szirmay-Kalos, TU Budapest

Abstract

This paper offers a novel approach to the indirect photon mapping method. The placement of photons acting as virtual light sources is regarded as a cheap sampling scheme, allowing for the reuse of a complete shooting path at the cost of a single shadow ray. In order to counter for its shortcomings, the variance reduction technique called weighted importance sampling is applied. This allows for the extension of the indirect photon mapping method for specular settings, because the weighting function can mimic surface BRDF characteristics disregarded by the virtual light source placement. On the other hand, weighted importance sampling also helps to eliminate the “corner spikes” caused by the fact that shooting cannot mimic the geometric factor of the connection rays. Advantages and problems are examined for several weighting schemes.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Raytracing

1. Introduction

1.1. Importance sampling schemes

A global illumination algorithm should compute the average of the radiance values leaving the area visible through a pixel in the direction of the eye:

$$C_A = \frac{1}{A} \cdot \int_A L(\vec{x}, \omega_{eye}) d\vec{x},$$

where A is the area of the surface on which the average is computed. We shall call this area as the *area of interest*. According to the rendering equation, the radiance is the sum of the emission and a reflected component that can be obtained by reflecting the radiance of all points that are visible from here. Let us concentrate on the reflected component since the emission is easy to compute. The average of the reflected radiance is:

$$C_A = \frac{1}{A} \cdot \int_A \int_{\Omega} L^{in}(\vec{x}, \omega') \cdot f_r(\omega', \vec{x}, \omega_{eye}) \cdot \cos \theta_{\vec{x}}' d\omega' d\vec{x},$$

where L^{in} is the incoming radiance at point \vec{x} , f_r is the BRDF and $\theta_{\vec{x}}'$ is the angle between the surface normal and incoming direction ω' . The product of the BRDF and the cosine of the incoming angle is the *scattering density* $w(\omega', \vec{x}, \omega_{eye})$ that expresses the probability density that scattering connects directions ω_{eye} and ω' .

This integral is often evaluated by Monte-Carlo quadrature.

Classical Monte-Carlo quadrature would take M random samples with probability density $p(\vec{x}, \omega')$ and approximate the integral as follows:

$$C_A = \frac{1}{A} \cdot \int_A \int_{\Omega} L^{in}(\vec{x}, \omega') \cdot w(\omega', \vec{x}, \omega_{eye}) d\omega' d\vec{x} \approx \frac{1}{M} \cdot \sum_{n=1}^M \frac{L^{in}(\vec{x}_n, \omega_n') \cdot w(\omega_n', \vec{x}_n, \omega_{eye})}{A \cdot p(\vec{x}_n, \omega_n')}. \quad (1)$$

In order to reduce the variance, probability density p should mimic the integrand. This approach is called *importance sampling*. Sampling according to a given probability density is carried out by transforming uniformly distributed numbers provided by the pseudo or quasi random number generator. This transformation requires the inverse of the cumulative probability distribution, thus p should be analytically integrable and we should be able to compute the inverse of its integral. These requirements can be met only if p is algebraically simple, which makes it impossible to accurately mimic the integrand. On the other hand, the incoming radiance L^{in} is not available, but we have to use another Monte-Carlo estimation to approximate it.

Examining the integrand we can note that this is the product of the incoming radiance and the scattering density. Since it seems hopeless to sample according to this product, Monte-Carlo algorithms try to mimic either the incoming radiance, or the scattering probability to the eye. The first ap-

proach is followed in *gathering algorithms*, such as in path tracing, while the second is in *shooting algorithms*.

Random walks are either initiated from samples of light sources or samples of our area of interest. Walks are terminated according to the rules of Russian roulette or by some deterministic decimation scheme⁹. If a walk is to be continued, another point, visible from the previous one, is sampled according to one of the above strategies. At the visited points, this offers an estimate of either the outgoing eye potential, or the incoming radiance. However, these points usually do not coincide with light sources or our area of interest. Therefore, deterministic steps are needed to complete the light paths so that some contribution can be calculated along them. Unfortunately, the distribution of these connections will not fit into the importance sampling scheme. This makes the rendering of scenes with specular surfaces difficult.

1.2. Indirect photon mapping

The indirect photon mapping is also called *instant radiosity*⁴ or the virtual light sources method⁹. It is a shooting type algorithm, where all the points visited by shooting walks and their respective power values are stored. Thereafter, in order to calculate pixel values one by one, the points on the surfaces visible through the pixels are connected to all of these photon hits in a deterministic manner, as depicted in Figure 1. These hits act like abstract, point-like light sources. Therefore, they are often referred to as *virtual light sources*. Even in the diffuse case, the deterministic step produces obvious artifacts. Hits near a surface element have more influence on its incoming radiance. Importance sampling would choose paths through near surface points with a high probability, and scale down their contribution appropriately. In the deterministic sampling scheme, however, it is relatively rare that a virtual light source is near the surface point. In turn, when it happens, its contribution is not scaled down. In the image this will appear as bright spots near the virtual light sources, and lack of proper illumination in between (Figure 2). If all




Figure 1: Indirect photon mapping.

surfaces are considered to be diffuse, the estimator for the radiance in the indirect photon mapping method will be the




Figure 2: Diffuse scene rendered by indirect photon mapping.

following:

$$L^{out}(\vec{x}, \omega_{eye}) = \sum_{\vec{y}} \frac{a_{\vec{x}} \cos \theta_{\vec{x}} \cos \theta'_{\vec{y}}}{\pi |\vec{x} - \vec{y}|^2} \frac{a_{\vec{y}}}{\pi} \Phi_{\vec{y}}^{in}, \quad (2)$$

where a is the local albedo, summation is performed for every virtual light source \vec{y} , and $\Phi_{\vec{y}}^{in}$ is the incoming power provided by the shooting walk algorithm. The term $\frac{a_{\vec{y}}}{\pi} \Phi_{\vec{y}}^{in}$ is the radiant exitance characteristic for a virtual light source.

It is straightforward to extend the technique to handle specular surfaces⁹. If the incoming direction of the radiance at the shooting walk nodes is also stored, the local BRDF may be evaluated. However, in case of a highly specular, visible surface, the deterministic connections may or may not coincide with the directions where the scattering density is high. Therefore, variance of the estimator will be high. This means that large specular highlights will appear where the virtual light sources suggest, but nothing in between (Figure 3). For the general case, the radiance is estimated by the sum:

$$L^{out}(\vec{x}, \omega_{eye}) = \sum_{\vec{y}} \frac{w(\omega_{\vec{y} \rightarrow \vec{x}}, \vec{x}, \omega_{eye}) w(\omega_{in}, \vec{y}, \omega_{\vec{y} \rightarrow \vec{x}})}{|\vec{x} - \vec{y}|^2} \Phi_{\vec{y}}^{in}. \quad (3)$$

The direction of the incoming radiance ω_{in} is also to be stored for every photon hit.

If the indirect photon mapping method is viewed as a gathering type algorithm, the weak point is revealed. As the next point of the walk is chosen, continuing from the surface point visible from the pixel, we are limited to the enumeration of the stored virtual light sources. However, as discussed above, it would be beneficial to choose a direction according




Figure 3: Specular scene rendered by indirect photon mapping.

to a probability distribution that mimics the scattering probability, and thus, a surface point whose distribution mimics

$$f_r(\omega_{\vec{y} \rightarrow \vec{x}}, \vec{x}, \omega_{eye}) \cdot \frac{\cos \theta_{\vec{x}} \cos \theta'_{\vec{y}}}{|\vec{x} - \vec{y}|^2}.$$

This means, we have a sampling scheme that has some shortcomings, and a probability distribution that we know would be better. This problem has been observed in other bi-directional path tracing ^{5,8} and in classical photon mapping ³. In bi-directional path tracing the problem has been solved by multiple importance sampling ⁷, while in photon mapping by multi-step final gathering ².

In indirect photon mapping, however, we usually do not want to create random gathering paths including more than one rays, since that would destroy the efficiency of the algorithm (an exception is the hit of ideal mirrors or refracting surfaces, when we cannot do anything else). Thus indirect photon mapping requires another solution that do not applies additional paths for final gathering but uses only those paths that have been created by photon shooting. This is exactly the case where weighted importance sampling can be effectively applied.

Weighted importance sampling is a rather old method ⁶ which has received just little attention in rendering so far. The exception is the pioneering paper ¹, which applied this technique in stochastic iteration of the radiosity equation. However, we believe that this technique has more potential in other algorithms, and particularly in the non-diffuse setting. In the next section we review the theory of weighted importance sampling.

2. Weighted Importance Sampling

Suppose that integral $F = \int_V f(\mathbf{z}) d\mathbf{z}$ needs to be evaluated by Monte-Carlo quadrature. The classical Monte-Carlo approach would compute the following sum:

$$F = \int_V f(\mathbf{z}) d\mathbf{z} \approx \frac{1}{M} \cdot \sum_{n=1}^M \frac{f(\mathbf{z}_n)}{p(\mathbf{z}_n)},$$

where p is the sampling density, which should mimic integrand f . In practical cases p cannot mimic f accurately and be appropriate for sample generation at the same time. Weighted importance sampling ^{6,1} attacks this problem by working with two probability densities simultaneously. Suppose we have probability density $g(\mathbf{z})$ that is quite good in mimicking integrand f but we are unable to sample according to this density due to its algebraic complexity. On the other hand, we also have another probability density $p(\mathbf{z})$ which is possibly poorer in mimicking f but is appropriate to construct a sampling scheme. Weighted importance sampling proposes the following quadrature formula to estimate the integral:

$$\int_V f(\mathbf{z}) d\mathbf{z} \approx \frac{\sum_{n=1}^M f(\mathbf{z}_n)/p(\mathbf{z}_n)}{\sum_{n=1}^M g(\mathbf{z}_n)/p(\mathbf{z}_n)},$$

where samples \mathbf{z}_n are obtained with probability density p . In order to demonstrate that this quadrature is asymptotically equivalent to the original Monte-Carlo quadrature, let us divide both the numerator and the denominator by the number of samples M :

$$\frac{\frac{1}{M} \cdot \sum_{n=1}^M f(\mathbf{z}_n)/p(\mathbf{z}_n)}{\frac{1}{M} \cdot \sum_{n=1}^M g(\mathbf{z}_n)/p(\mathbf{z}_n)}.$$

The new numerator is the same as the original integral quadrature, thus it converges to the integral. The denominator, on the other hand, is the Monte-Carlo estimate of integral $\int_V g(\mathbf{z}) d\mathbf{z}$. Since g is a probability density function, its integral is 1, thus the new quadrature converges to the same value as the old quadrature.

The question is whether or not this new estimate is better than the old one. This depends on whether or not density g is better in mimicking f than p . The formal analysis ⁶ results in the following formulae. The mean square error after M samples obtained with weighted importance sampling, including both the bias and the variance, is

$$\varepsilon_{WIS} \approx \frac{1}{M} \cdot \int_V \left(\frac{f(\mathbf{z})}{p(\mathbf{z})} - F \cdot \frac{g(\mathbf{z})}{p(\mathbf{z})} \right)^2 \cdot p(\mathbf{z}) d\mathbf{z}. \quad (4)$$

For the sake of comparison, we also present the mean square error of the classical Monte-Carlo estimator applying standard importance sampling:

$$\varepsilon_{CMC} \approx \frac{1}{M} \cdot \int_V \left(\frac{f(\mathbf{z})}{p(\mathbf{z})} - F \right)^2 \cdot p(\mathbf{z}) d\mathbf{z}. \quad (5)$$

Instead of repeating the proof, we provide an intuitive explanation. Suppose that p is poor in sampling a particular sub-domain \mathcal{S} . It does not generate samples here as frequently as would be required by the large integrand values. If classical Monte-Carlo method is used, when we are lucky enough to generate a sample in sub-domain \mathcal{S} , we get a large integrand value that is divided by a small probability, which results in a huge term in the average approximating the integral. These infrequent but huge values are responsible for large fluctuations. However, if this method is used with a probability density g according to weighted importance sampling, then the approximating sum is also divided by the sum of $g(\mathbf{z}_n)/p(\mathbf{z}_n)$ terms. When we are not lucky to sample the important regions, g will also be small, thus the denominator will be smaller than one. Dividing by the denominator, the approximation will be scaled up. However, when the sample is in important sub-domain \mathcal{S} , the Monte-Carlo estimate and the denominator will be increased simultaneously by a larger value. Since the fluctuations of the numerator and the denominator are thus synchronised, the fluctuation of their ratio is decreased.

However, weighted importance sampling may have not only advantages but disadvantages as well. It has a small bias, which disappears quickly. On the other hand, if p is already good enough to mimic integrand f , then the numerator will be stable. The fluctuation of the denominator around 1 appears just as an additional noise.

Thus we can conclude that weighted importance sampling should be used carefully, since it can reduce and increase the variance depending on the mimicking capabilities of the two probability densities. Let us examine the difference of the mean square errors of the estimator of classical Monte-Carlo and the estimator of weighted Monte-Carlo:

$$\begin{aligned} \epsilon_{CMC} - \epsilon_{WIS} = \\ \frac{1}{M} \cdot \int_V F \cdot \left(\frac{g(\mathbf{z})}{p(\mathbf{z})} - 1 \right) \cdot \left(\frac{2f(\mathbf{z})}{p(\mathbf{z})} - F \left(\frac{g(\mathbf{z})}{p(\mathbf{z})} + 1 \right) \right) \cdot p(\mathbf{z}) \, d\mathbf{z} \end{aligned}$$

In order to get improvement, this difference should be positive. Note that the integrand of the error difference is a product where the first factor cannot be negative, but the second and the third factors can. Thus improvement is guaranteed if the second and the third terms change their signs simultaneously. There are two cases. The quadrature overestimates in \mathbf{z} if $f(\mathbf{z})/p(\mathbf{z}) > F$. On the other hand, the quadrature underestimates in \mathbf{z} if $f(\mathbf{z})/p(\mathbf{z}) < F$.

In case of overestimation the new probability density g should meet the following requirement in order to make the integrand of the error difference positive:

$$\frac{f(\mathbf{z})}{F} \geq g(\mathbf{z}) > p(\mathbf{z}).$$

It means that g should also result in overestimation but its level should be decreased. This statement can be proven by checking that in this case both the second and the third factors are positive.

In case of underestimation, the requirement of the integrand of the error difference being positive is

$$\frac{f(\mathbf{z})}{F} \leq g(\mathbf{z}) < p(\mathbf{z}),$$

thus the level of underestimation should be decreased. In this case the second and third factors are both negative.

3. Application of weighted importance sampling for indirect photon mapping

In the classic indirect photon mapping algorithm, we solve the radiance equation for a pixel using multi-dimensional samples, or light paths, generated using the shooting walk method. However, a large number of effects, most of them directly affecting image quality would better be handled using samples generated with the gathering walk sampling scheme. This means, that when we evaluate the radiance estimator for a pixel, we may weight samples generated by one method using the probabilities of the sample being generated by the other one.

However, it is not obvious how the deterministic formula given in Equation 3 can be interpreted as a Monte-Carlo integration of the form specified by Equation 1. The division with the generation probability of the sample is hidden in the term Φ_y^{in} , the incoming radiance estimated by the shooting walk algorithm.

3.1. Single step weighting

Shooting creates a radiance estimate for the entire scene, represented by the virtual light sources. When we calculate the incoming radiance using the indirect photon mapping method, we actually sample radiance in a limited way. The scattering probability to the eye would mimic transfer better, and could therefore be used as a weighting probability $g_1(\vec{y})$. However, we need to know the probability $p_1(\vec{y})$ of a surface point being visited by the shooting. Because of the importance sampling approach of shooting, this probability is proportional to the radiance, estimated by the virtual light sources themselves. Therefore, it is possible to derive an approximation for $p_1(\vec{y})$ from their density distribution. If a proximity search for the nearest n nodes of the total N is carried out, and the area of the surface element within which they are located is A , then:

$$p_1(\vec{y}) = \frac{n}{AN}$$

The weighting probability is that of a gathering step:

$$g_1(\vec{y}) = \frac{w(\omega_{eye}, \vec{x}, \omega_{\vec{y} \rightarrow \vec{x}}) \cos \theta'_y}{a(\vec{x}, \omega_{eye}) |\vec{x} - \vec{y}|^2}$$

That makes the final estimator:

$$L_1^{out}(\vec{x}, \omega_{eye}) = \frac{\sum_{\vec{y}} \frac{w(\omega_{eye}, \vec{x}, \omega_{\vec{y} \rightarrow \vec{x}}) w(\omega_{in}, \vec{y}, \omega_{\vec{y} \rightarrow \vec{x}}) \Phi_y^{in}}{|\vec{x} - \vec{y}|^2}}{\sum_{\vec{y}} \frac{w(\omega_{eye}, \vec{x}, \omega_{\vec{y} \rightarrow \vec{x}}) \cos \theta'_y}{a(\vec{x}, \omega_{eye}) |\vec{x} - \vec{y}|^2} \frac{1}{p_1(\vec{x})}}$$

The similarity of the numerator and the denominator is obvious. The disturbing high-variance terms disregarded by the original sampling scheme, namely the scattering probability to the eye and the distance to the virtual light source appear in both. Singularities cancel out, removing extreme values and associated image artifacts. However, $p_1(\vec{y})$ is only an approximation. In the image, the former bright singularities will be traded for far less intense light and dark spots (Figures 4 and 5).




Figure 4: Diffuse scene rendered without weighting and using single step weighting in 14 seconds on an Athlon 950 Mhz computer




Figure 5: Specular scene rendered without weighting and using single step weighting in 20 seconds.

3.2. Full path weighting

A sample produced by shooting is a light path, for which the generation probability is known exactly. Although surface points used in the previous approach are also produced, we can only approximate their generation probabilities. Therefore, we would need a weighting probability distribution over the domain of light paths, not surface points. The recursive shooting walk scheme provides such a distribution. Whenever we calculate the contribution of a light path represented by a virtual light source, we can determine the probability with which gathering would have generated the path.

The probability of generating the path that ends in the

virtual light source at \vec{z}_{l-1} by shooting will be denoted by $p_2(\omega_1, \dots, \omega_{l-2}, \vec{z}_{l-1})$.

$$p_2(\vec{z}_0, \omega_1, \dots, \omega_{l-2}) = p_{light}(\vec{z}_0, \omega_1) \cdot \prod_{i=1}^{l-2} \frac{w(-\omega_i, \vec{z}_i, \omega_{i+1}) \cos \theta_{\vec{z}_{i+1}}}{|\vec{z}_i - \vec{z}_{i+1}|^2}$$

where l is the number of surface nodes, z_i denotes consecutive nodes of the path and ω_i denotes consecutive directions. z_0 is the starting point and $w(-\omega_0, \vec{z}_0, \omega_1)$ can be interpreted as the emittance distribution of the light source (see Figure 6 for nomenclature). This probability can be calculated for every virtual light source during shooting. The probability of




Figure 6: Full path weighting.

generating the same path using gathering would be:

$$g_2(\vec{z}_l, -\omega_{l-1}, \dots, -\omega_2) = \prod_{i=1}^{l-2} \frac{w(\omega_{i+2}, \vec{z}_{i+1}, -\omega_{i+1}) \cos \theta'_{\vec{z}_i}}{|\vec{z}_i - \vec{z}_{i+1}|^2}$$

In order to be able to calculate $g_2(\vec{x})$, we would need to store information about the entire shooting path. However, only the ratio $\frac{g_2(\vec{z}_l, -\omega_{l-1}, \dots, -\omega_2)}{p_2(\vec{z}_0, \omega_1, \dots, \omega_{l-2})}$ is needed. The ratio of the two factors within the numerator and the denominator that correspond to the same path section, without the distance terms that obviously cancel out, is:

$$\frac{f_r(\omega_{i+2}, \vec{z}_{i+1}, -\omega_{i+1}) \cos \theta'_{\vec{z}_i} \cos \theta'_{\vec{z}_{i+1}}}{f_r(-\omega_i, \vec{z}_i, \omega_{i+1}) \cos \theta'_{\vec{z}_i} \cos \theta_{\vec{z}_{i+1}}} = \frac{f_r(\omega_{i+2}, \vec{z}_{i+1}, -\omega_{i+1})}{f_r(-\omega_i, \vec{z}_i, \omega_{i+1})}$$

The product of these ratios, and the terms for the first and the last segment that do not cancel out give:

$$\begin{aligned} & \frac{g_2(-\omega_{l-1}, \dots, -\omega_2, \vec{z}_1)}{p_2(\omega_1, \dots, \omega_{l-2}, \vec{y})} = \\ & = \frac{w(\omega_{eye}, \vec{z}_l, -\omega_l) \cos \theta'_{\vec{z}_{l-1}} f_r(-\omega_{l-1}, \vec{z}_{l-1}, \omega_l)}{|\vec{z}_l - \vec{z}_{l-1}|^2} \cdot \\ & \quad \frac{|\vec{z}_0 - \vec{z}_1|^2}{p_{light}(\vec{z}_0, \omega_1) \cos \theta_{\vec{z}_1} f_r(\omega_2, \vec{z}_1, -\omega_1)} \end{aligned}$$

The first part of the product can be computed knowing the incoming direction of the photon hits, and the second part can be calculated for every virtual light source during shooting. A notable special case is, when $l = 2$, and ω_2 is the direction pointing from the photon hit towards the surface point visible from the pixel. For these paths, $w(\omega_2, \vec{z}_1, -\omega_1)$ can not be calculated for primary photon hits in advance.

Although the previous formulae appear to be convincing, they do not consider occlusions. Whenever the last, deterministic connection of a gathering walk failed, because the light source was occluded, the resulting path would not possibly be generated by shooting. Therefore, a portion of $g_2(\vec{z}_l, -\omega_{l-1}, \dots, -\omega_2)$ is not sampled by $p_2(\vec{z}_0, \omega_1, \dots, \omega_{l-2})$. This means that the expected value of the denominator of the weighted importance sampling formula will not be one. However, the method works for special scenes where the light source is visible from every surface point, like the ones in Figures 7 and 8.




Figure 7: Scene with unoccluded light source rendered without weighting and with full path weighting in 14 seconds

In order to attack the problem of occluded scenes, several unbiased and biased alternatives are available. It is possible to provide a modified gathering-type probability distribution that does not include paths never generated by shooting. Firstly, a gathering walk could always be continued when the light source sample to be reached is not visible from its end point. Secondly, the gathering walk could fully exclude Russian roulette and the deterministic step, and go on until it reaches a light source. The formulae are easily modified to include new continuation probabilities. Although both methods assure that the expected value of the estimator is correct




Figure 8: Scene with specularly and caustic effects rendered without weighting and with full path weighting in 14 seconds.

(these methods are unbiased), they assign high generation probabilities to longer paths. This distortion increases the variance of the estimator that appears on the image as misplaced shadows (Figure 9).




Figure 9: Scene rendered using modified gathering probability for weighting.

Biased methods, on the other hand, aim at the approximate normalization of the density g in the subdomain which can be generated by p . In particular it means that we have to approximate the probability that a gathering path ends in a light source, or, alternatively, arrives in a point from which the light source is visible. Note that this probability has usually low variation over the pixels, thus very few gathering paths can result in accurate estimates. We propose to compute this probability separately for those set of pixels in which the same object is seen.

4. Conclusion

Weighted importance sampling is a promising tool to improve the quality of cheap sampling schemes that are not good enough in mimicking the integrand. In fact, we sample according to the cheap scheme, but pretend that we got the samples with a more accurate and therefore more expensive scheme.

In this paper we presented a work in progress, in which the problems of deterministic connections in random paths have been eliminated by weighting the samples by the probability density of gathering walks. Although the basic idea is simple, the combination of the two probabilities is not trivial. The problem is that both shooting and gathering tend to sample zero contribution paths, and the sets of these irrelevant paths are different in shooting and gathering. One way of solving this is to consider only the last, and the most problematic reflection of the paths.

The second approach is to assume a gathering density which cannot generate paths which cannot be obtained with shooting. However, such densities are not as good quality as the original path tracing applying Russian roulette, which degrades the performance. Speaking of the images generated, these methods generally lack at creating proper shadows. The viewer could feel that misplaced lights are traded for misplaced shadows. This is because weighting ignores virtual light sources not visible from a surface point. It would be desirable to find a mathematically reasonable way to make use of the occlusion information carried by these virtual light sources.

Finally, we have the possibility to normalize the probability densities in those subregions which can be obtained by the cheap sampling schemes. This is left for the future work.

References

1. Philippe Bekaert, Mateu Sbert, and Yves D. Willems. Weighted importance sampling techniques for monte carlo radiosity. In B. Peroche and H. Rushmeier, editors, *Rendering Techniques 2000 (Proceedings of the Eleventh Eurographics Workshop on Rendering)*, pages 35–46, New York, NY, 2000. Springer Wien. 3
2. P. Christensen. Faster photon map global illumination. *Journal of Graphics Tools*, 4(3):1–10, 2000. 3
3. H. W. Jensen. Global illumination using photon maps. In *Rendering Techniques '96*, pages 21–30, 1996. 3
4. A. Keller. Instant radiosity. *Computer Graphics (SIGGRAPH '97 Proceedings)*, pages 49–55, 1997. 2
5. E. Lafortune and Y. D. Willems. Bi-directional path-tracing. In *Compugraphics '93*, pages 145–153, Alvor, 1993. 3
6. M. Powell and J. Swann. Weighted importance sampling — a monte-carlo techniques for reducing variance. *Inst. Maths. Applics.*, 2:228–236, 1966. 3
7. E. Veach. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, http://graphics.stanford.edu/papers/veach_thesis, 1997. 3
8. E. Veach and L. Guibas. Bidirectional estimators for light transport. In *Computer Graphics (SIGGRAPH '95 Proceedings)*, pages 419–428, 1995. 3
9. I. Wald, C. Benthin, and P. Slussalek. Interactive global illumination using fast ray tracing. In *13th Eurographics Workshop on Rendering*, 2002. 2