
Illumination in the Presence of Weak Singularities

Thomas Kollig¹ and Alexander Keller²

¹ Dept. of Computer Science, Kaiserslautern University of Technology, 67653
Kaiserslautern, Germany
`kollig@informatik.uni-kl.de`

² Dept. of Computer Science, University of Ulm, 89069 Ulm, Germany
`alexander.keller@uni-ulm.de`

Summary. Approximating illumination by point light sources, as done in many professional applications, allows for efficient algorithms, but suffers from the problem of the weak singularity: Besides avoiding numerical exceptions caused by the division by the squared distance between the point light source and the point to be illuminated, the estimator should be unbiased and of finite variance. We first illustrate that the common practice of clipping weak singularities to a reasonable value yields clearly visible bias. Then we present a new global illumination algorithm that is unbiased and as simple as a path tracer, but elegantly avoids the problem of the weak singularity. In order to demonstrate its performance, the algorithm has been integrated in an interactive global illumination system.

1 Introduction

Simulating light transport in a physically correct way has become a mainstream feature in movie productions and interactive rendering systems. On the one hand many approximations are used to make the algorithms simpler, faster, and more numerically robust. On the other hand unbiased approaches like e.g. bidirectional path tracing or the Metropolis light transport algorithm are too complicated for use in professional productions and not sufficiently efficient.

Based on the popular approximation instant radiosity [Kel97], we present an unbiased, robust, and very simple global illumination algorithm that is used for production as well as interactive rendering. The new algorithm is easily implemented in any ray tracing system and exposes many advantages over previous techniques, as we will illustrate in the sequel.

2 Avoiding Bias Caused by Bounding

Before we introduce our new global illumination algorithm, we need to explain an important observation for the example problem of direct illumination by an area light source.

The direct illumination is the radiance

$$L_r(x, \omega_r) = \int_A f_r(\omega, x, \omega_r) G(x, y) V(x, y) L_e(y, -\omega) dy$$

reflected in x into direction ω_r , which is the integral over the surface A of the light source. Its radiance L_e from y towards x , i.e. into direction $-\omega$, is attenuated by the geometry term G and the bidirectional reflectance distribution function f_r , which represents the physical surface properties. The visibility is accounted for by $V(x, y)$, which is 1 if x and y are mutually visible and zero otherwise.

It is common practice to use a bounded geometry term

$$G'(x, y) := \min\{G(x, y), b\}$$

for some bound $b > 0$ instead of the correct expression

$$G(x, y) := \frac{\cos^+ \theta_x \cdot \cos^+ \theta_y}{\|x - y\|_2^2} ,$$

where the positive cosine $\cos^+ \theta_x$ is the scalar product between the unit direction of $y - x$ and the surface normal in x , which is set to zero, if the cosine is less than zero (analogous for $\cos^+ \theta_y$).

The obvious reason for uniformly bounding G by some $b > 0$ is to avoid infinite variance from the weak singularity, which can be caused by an arbitrarily small Euclidean distance $\|x - y\|_2$, i.e. when the point x to be lit and the sampled point y on the light source are very close.

2.1 Consequences of Bounding the Integrand

Almost any rendering software approximates the direct illumination $L_r(x, \omega_r)$ by

$$\begin{aligned} L'_r(x, \omega_r) &= \int_A f_r(\omega, x, \omega_r) G'(x, y) V(x, y) L_e(y, -\omega) dy \\ &\approx \frac{|A|}{N} \sum_{i=0}^{N-1} f_r(\omega_i, x, \omega_r) G'(x, y_i) V(x, y_i) L_e(y_i, -\omega_i) \end{aligned} \quad (1)$$

using Monte Carlo integration or a variant of it. Here y_i are uniformly distributed samples on the area A of the light source and ω_i is the unit vector pointing from x to y_i .

The consequence of uniformly bounding the integrand is an exponential decay of the error probability. In the case of pure Monte Carlo integration Hoeffding's inequality yields the probability

$$\text{Prob} \left(\left| \int_{[0,1)^s} f(x) dx - \frac{1}{N} \sum_{i=0}^{N-1} f(x_i) \right| \geq \epsilon \right) \leq 2e^{-\frac{1}{4\tau^2} N \epsilon^2}$$

of an integration error more than an arbitrary threshold $\epsilon > 0$ if it is possible to uniformly bound $|f(x) - I| < \tau$ for almost all $x \in [0, 1)^s$, where I is the integral of f . Since most Monte Carlo rendering algorithms bound the samples before averaging, the previous formula explains the observed fast convergence.

By the fast convergence visible artifacts rapidly disappear and the images look nice. However, the estimator is biased, i.e. does not converge to the desired value $L_r(x, \omega_r)$, and important visible contributions of the illumination are missing, as can be seen from the differences between Figs. 1a) and e). Obviously the bias is especially high in the vicinity of concave geometry such as the curtains and the fine geometry of the chairs. In fact the bias introduced by bounding the geometry term cannot be ignored.

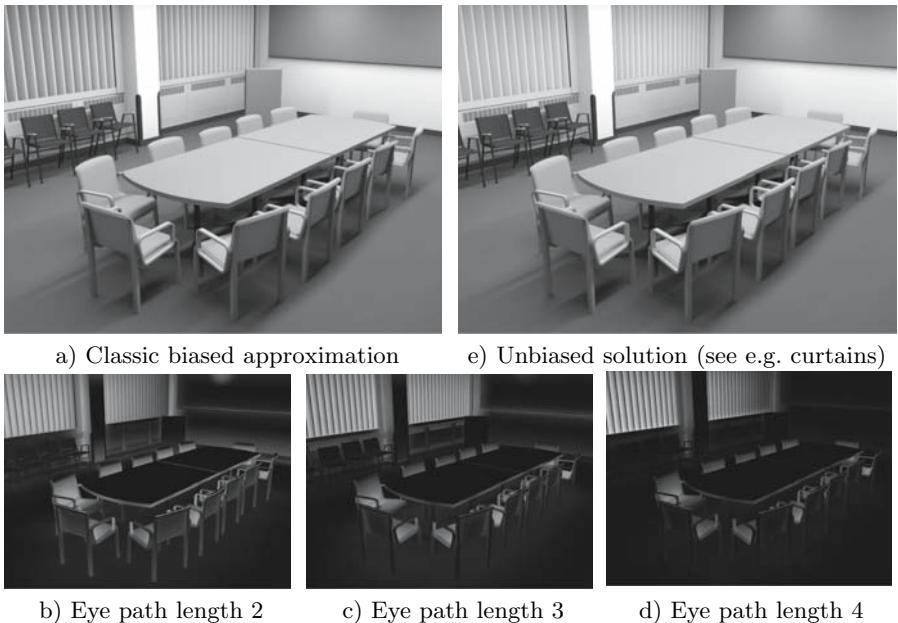


Fig. 1. Looking at the curtains, the bias between the true global illumination in (e) and the classic approximation in (a) is clearly visible as a difference in brightness. The images (b) to (d) show the missing contributions according to the eye path length. For display the images (b) to (d) have been amplified by a factor of 3, 9, and 27, respectively. The bias of the classic approximation is clearly located in regions of concave geometry.

2.2 Unbiased Robust Estimator

It is favorable to preserve the fast convergence of the estimator (1), since it actually contributes most of the illumination and exposes low variance. The so-called bias, i.e. the difference between the desired and the computed integral is

$$\begin{aligned}
& L_r(x, \omega_r) - L'_r(x, \omega_r) \\
&= \int_A L_e(y, -\omega) f_r(\omega, x, \omega_r) V(x, y) \max\{G(x, y) - b, 0\} dy \\
&= \int_A L_e(y, -\omega) f_r(\omega, x, \omega_r) V(x, y) \frac{\max\{G(x, y) - b, 0\}}{G(x, y)} G(x, y) dy \\
&= \int_{S^2} L_e(h(x, \omega), -\omega) \frac{\max\{G(x, h(x, \omega)) - b, 0\}}{G(x, h(x, \omega))} f_r(\omega, x, \omega_r) \cos^+ \theta_x d\omega. \quad (2)
\end{aligned}$$

Changing the domain of integration to the unit sphere S^2 requires the ray tracing function $h(x, \omega)$, which returns the first surface point hit when shooting a ray from x into direction ω .

In order to obtain an unbiased estimate of the direct illumination $L_r(x, \omega_r)$, we use the estimator (1) and add an estimate of the above equation (2). Applying importance sampling according to the density $f_r(\omega, x, \omega_r) \cdot \cos^+ \theta_x$, the integrand becomes bounded [Shr66], too.

Although the method seems simple, it never before has been used to single out the weak singularities contained in the geometry term G . There are several advantages to this approach: Bounding the integrand in (1) does not add new discontinuities and consequently variance is not increased. Since both integrands are bounded, the variance remains finite and the estimate is numerically robust.

In the context of parametric integration, Heinrich [Hei00] proposed an optimal method for the Monte Carlo approximation of weakly singular operators: For smooth function classes his algorithm used a stratification idea to separate the weak singularity. This is related to our approach, however, introduces more discontinuities to the integrand as compared to bounding.

2.3 Choice of the Bound

Obviously, the radiance L_r is estimated in an unbiased way for any choice of $0 \leq b < \infty$. However, most renderers implicitly are using a fixed bound without compensating the bias (2). Choosing

$$b \equiv b(\omega, x, \omega_r) = \frac{c}{f_r(\omega, x, \omega_r)} \quad (3)$$

allows one to use the very efficient estimator (1) as long as $G(x, y) f_r(\omega, x, \omega_r) \leq c$. Using importance sampling as indicated in the previous section, the transport kernel in (2) then is bounded by 1.

By c we can adjust the efficiency, i.e. how much of the estimate is obtained by sampling the area of the light source and how much is contributed by importance sampling of the solid angle. Choosing c independent of the scene geometry, the contributions of the estimators for (1) and (2) depend on the scale of the geometry, which is hidden in the geometry term G .

For $c = \frac{1}{|A|}$ no sample from either (1) or (2) can be larger than the source radiance L_e , i.e. the radiance is never amplified but only attenuated. Consequently the contributions of the integration over the light source as well as the solid angle can contribute about the same noise level at maximum. In addition the bound becomes independent of the scale, since both G and A contain the scale of the scene.

Since usually the radiance is vector valued, i.e. it comprises components for red, green, and blue, in fact the maximum norm $\|f_r(\omega, x, \omega_r)\|_\infty$ should be used in the denominator of b .

3 The New Global Illumination Algorithm

Using the observation from the previous section, it is simple to construct a global illumination algorithm that is unbiased and numerically robust. We just need to compensate for the bias of popular approaches like instant radiosity [Kel97] or successor approaches to interactive global illumination [WKB⁺02, BWS03]. The procedure to compute a local solution from the radiance Fredholm integral equation of the second kind

$$L(x, \omega_r) = L_e(x, \omega_r) + \int_{S^2} L(h(x, \omega), -\omega) f_r(\omega, x, \omega_r) \cos^+ \theta_x d\omega$$

is illustrated in Fig. 2:

1. **Generation of point light sources:** Identical to the preprocessing of the instant radiosity algorithm [Kel97] or a very sparse global photon map [2], a set $(y_j, L_j)_{j=0}^{M-1}$ of M point light sources is created. This corresponds to tracing paths starting at the lights and storing all the points $y_j \in \mathbb{R}^3$ of incidence with their power L_j .
2. **Shading:** Similar to a path tracer, an eye path is started from the lens incident in point x_0 from direction ω_0 . Starting with $i = 0$, we sum up three contributions for the current point x_i until the eye path is terminated:
 - a) **Light sources:** that are hit contribute their emission $L_e(x_i, -\omega_i)$.
 - b) **Illumination:** The contribution of the j -th point light source is

$$f_r(\omega_{x_i, y_j}, x_i, \omega_i) G'(x_i, y_j) V(x_i, y_j) L_j \quad , \quad (4)$$

where the direction ω_{x_i, y_j} points from x_i to y_j and ω_i is the direction from where x_i has been hit. Note that G' is the bounded version of the geometry term G .

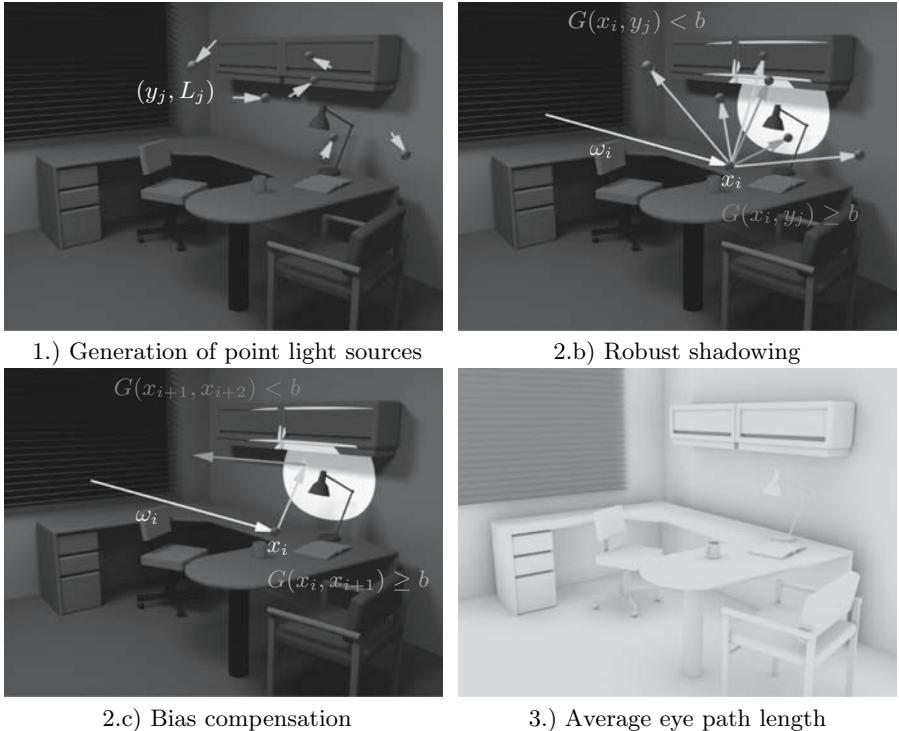


Fig. 2. Principal steps of the unbiased robust global illumination algorithm: In the first step a set $(y_j, L_j)_{j=0}^{M-1}$ of point light sources is generated. Hitting x_i from direction ω_i in the second step, the highlighted areas show the domain, where the geometry term $G(x_i, \cdot)$ is below the bound $b = 0.5$. In step 2.b shadow rays towards the point lights are traced. This is robust, because numerical exceptions by the inverse squared distance in G cannot occur due to bounding. In order to be unbiased, step 2.c continues the eye path from x_i by scattering a ray. While the ray hits the domain, we continue with step 2.a, otherwise the eye path is terminated. Image 3.) shows the average eye path length as a gray image, where the maximum path length considered was 5. Darker areas in the image indicate longer eye paths. The resulting image clearly resembles images obtained by ambient occlusion, i.e. concave corners are darker.

c) **Bias compensation:** Because the weak singularity was avoided by bounding the integrand, we have to account for the bias. Therefore we trace a ray into a random direction yielding the next vertex x_{i+1} along the eye path on the scene surface S . If $G(x_i, x_{i+1}) < b$ then this contribution has already been accounted for in the previous step and the eye path is terminated. Otherwise i is incremented and we continue with step 2.a, whose result has to be attenuated by the product of

the bidirectional reflectance distribution function f_r and $\frac{G(x_i, x_{i+1}) - b}{G(x_i, x_{i+1})}$ because of the derivation in equation (2).

The evolution of an image with the eye path length can be seen in Fig. 1, where Figs. 1a) – d) show the contribution of the eye path length $i = 1, \dots, 4$ and Fig. 1e) the sum of the contributions computed by our algorithm. Note that the contributions have been amplified by $3^{(i-1)}$ for display, i.e. the bias decays exponentially over the eye path length in our new algorithm.

3.1 Numerical Comparison to Bidirectional Path Tracing

Our new technique can be formulated as a heuristic for multiple importance sampling [VG95, Vea97] and consequently belongs to the class of bidirectional path tracing algorithms. Although this larger mathematical framework is not required for the derivation, it is interesting to compare the efficiency of our method to the classical techniques.

The New Algorithm as a Heuristic for Bidirectional Path Tracing

Using the notions as defined in [KK02a], our algorithm computes a path integral

$$\sum_{\ell=1}^{\infty} \int_{\mathcal{P}_{\ell}} f(\bar{x}) d\mu(\bar{x}) \approx \sum_{\ell=1}^{\infty} \frac{1}{N} \sum_{j=0}^{N-1} \sum_{i=0}^{\ell-1} w_{\ell,i}(\bar{x}_{\ell,i,j}) \frac{f(\bar{x}_{\ell,i,j})}{p_{\ell,i}(\bar{x}_{\ell,i,j})}$$

by multiple importance sampling. \mathcal{P}_{ℓ} is the path space containing all transport paths $\bar{x} = x_0 x_1 \cdots x_{\ell-1}$ of length ℓ and f is the measurement contribution function, which contributes the radiance of the path \bar{x} . For a fixed path length ℓ there are ℓ techniques to generate it by a corresponding probability density function $p_{\ell,i}$. The path $\bar{x}_{\ell,i,j}$ has been generated by sampling the probability density $p_{\ell,i}$. With these definitions our new algorithm results in the weights

$$\begin{aligned} w_{\ell,\ell-1}(\bar{x}) &= \min \left\{ 1, \frac{b}{G(x_{\ell-2}, x_{\ell-1})} \right\} \\ w_{\ell,\ell-2}(\bar{x}) &= \min \left\{ 1, \frac{b}{G(x_{\ell-3}, x_{\ell-2})} \right\} \cdot (1 - w_{\ell,\ell-1}(\bar{x})) \\ &\vdots \\ w_{\ell,1}(\bar{x}) &= \min \left\{ 1, \frac{b}{G(x_0, x_1)} \right\} \cdot (1 - w_{\ell,2}(\bar{x})) \cdot \dots \cdot (1 - w_{\ell,\ell-1}(\bar{x})) \\ w_{\ell,0}(\bar{x}) &= (1 - w_{\ell,1}(\bar{x})) \cdot \dots \cdot (1 - w_{\ell,\ell-1}(\bar{x})) \quad , \end{aligned} \quad (5)$$

which obviously fulfill $\sum_{i=0}^{\ell-1} w_{\ell,i}(\bar{x}) = 1$ for any path \bar{x} of length ℓ as required for an unbiased estimator [KK02a].

Numerical Evidence for the Increased Efficiency

Since our method is unbiased, for a sufficient number of samples we obtain images without artifacts. At too low sampling rates noise, blossoming, and sharp shadow boundaries can be visible. One might think that this noise is caused by step 2.a of the algorithm, which however rarely happens. It is more likely that the point lighting after scattering in step 2.c contributes noise. Sharp shadow boundaries become visible, if one set of point light sources is used for the whole image in step 2.b. If furthermore point light sources are located in concave geometry, it can happen that the close-by geometry is brightly lit, which we call blossoming. Using a different set of point light sources for adjacent pixels (uncorrelated sampling, e.g. interleaved sampling [Kel03, KH01]) the latter two artifacts are turned into noise. All these artifacts, however, are bounded as proved in Sect. 2.3 and thus rapidly average out during Monte Carlo integration.

The choice of c balances the artifacts at low sampling rates and thus controls the efficiency of the algorithm: The larger c , the more artifacts are caused by the point light sources, the smaller c the more noise from scattering becomes noticeable.

Known heuristics from multiple importance sampling, as e.g. the power heuristic [VG95, Vea97], are ratios of probability densities and therefore independent of the scale of a given scene. For area light sources we obtain this property for $c = \frac{1}{|A|}$ as previously mentioned. For global illumination, which includes indirect illumination, however, this choice no longer is obvious.

Therefore we numerically analyzed the efficiency, i.e. the reciprocal of the variance multiplied by the rendering time. In Fig. 3 the relative running time of the balance, power, and maximum heuristic is compared to our new algorithm at identical image quality. We used a bidirectional path tracer that

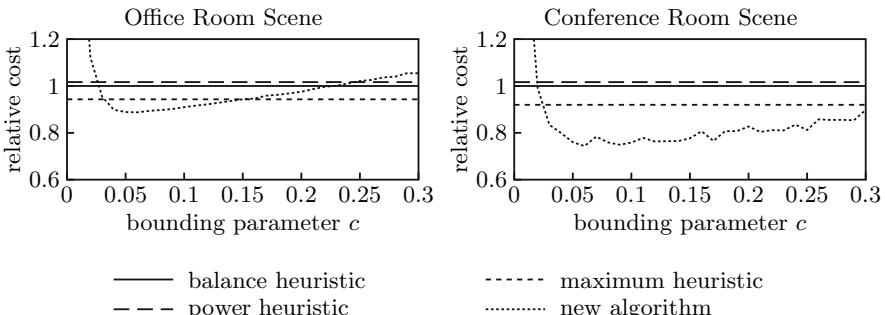


Fig. 3. Comparison of the relative rendering time of the power heuristic with $\beta = 2$, the balance heuristic, the maximum heuristic, and our new algorithm at identical image quality. We used the more efficient interleaved sampling, i.e. the method of dependent tests. For a wide choice of the bounding parameter c our new algorithm reliably outperforms the classical techniques saving up to 20% of the rendering time.

has been improved by dramatically reducing the number of light paths by interleaved sampling [KH01], which decreases the realization cost by the same constant amount of time for all heuristics. Although our algorithm may not have optimal variance, it is more efficient than other heuristics for a broad range of the parameter c .

Efficient Implementation

Contrary to classical bidirectional path tracing techniques [VG95, Vea97], the algorithm from section 3 is easily implemented in the standard shader concept of industrial ray tracing software. A complicated implementation of the balance, power, or maximum heuristic is not required, because our algorithm just is a double `for-loop` over the eye path length and the number of point lights. This has been a big advantage for the acceptance in a production environment. The observed increased efficiency has several reasons:

Cheap rays: Equating the geometry term G and the bound b allows one to bound the maximum length

$$r(x) \leq \sqrt{\frac{\cos^+ \theta_x}{b}}$$

of the eye rays, where we used $\cos^+ \theta_{y_i} \leq 1$. This distance often is much shorter than the obvious bound determined by the bounding box of the scene. Consequently the amount of geometry loaded into the caches remains much smaller and less voxels of the acceleration data structures have to be traversed.

Short eye paths: Compared to previous bidirectional path tracing heuristics, the eye path length of our new method is shorter on the average. Thus less rays have to be traced and shaded as can be seen in Figs. 1 and 2.

Less shadow rays: The short average eye path length directly results in a moderate number of shadow rays to be shot and consequently a higher data locality.

One might argue that the maximum heuristic in bidirectional path tracing also allows one to avoid the shooting of shadow rays. This is true, however, most of the possible path weights have to be computed in order to determine, whether or not to omit a shadow ray. The second disadvantage of the maximum heuristic is that it introduces discontinuities in the integrands around the weak singularity, which in our approach we explicitly avoided in order to obtain a lower noise level.

Intrinsic cache coherence: Only in the vicinity of concave corners the eye path length slightly increases as illustrated in Fig. 2. Then the ray length is short, which implies that most of the geometry already is in the processor cache. This corresponds to the idea of local illumination environments [FBG02], however, our method is unbiased and implicit, i.e. does not require an extra implementation for cache locality. Working with point

light sources, the shadow rays can be traced as bundles originating from one point [WBWS01]. Because shadow rays only access the scene geometry and do not require shader calls, less cache memory is required for shader data.

Compared to bidirectional path tracing, the cost of the light path generation remains the same for our new algorithm. However, since eye rays are cheaper, eye paths are shorter, and less shadow rays have to be shot, the new algorithm is more efficient and in addition it benefits much more from speedups in tracing rays. Because cache requirements are minimal, the efficient use of the processor cache is intrinsic to our algorithm and does not require extra care while coding.

3.2 Extensions

Our approach is a very general mathematical concept and unifies many seemingly isolated techniques in a simple way: Russian roulette, ambient occlusion and local illumination environments, and final gathering and secondary final gathering are all intrinsic. In order to further increase efficiency, the algorithm can easily be complemented by the following, orthogonal techniques:

Efficient multidimensional sampling: For the sake of clarity, we based all explanations on arguments using pure random sampling. It is straightforward to improve the efficiency by quasi-Monte Carlo and randomized quasi-Monte Carlo sampling methods [Kel02, KK02a, KK02b]. The big advantage of our approach is that additional discontinuities, which could have harmed the performance of stratified sampling patterns, are explicitly avoided.

Shadow computation: The techniques of Ward [War91], Keller [Kel98], and Wald [WBS03] can be used for reducing the number of shadow rays. The shadows also could be computed using various algorithms on graphics hardware. Due to the rapidly decaying contribution of longer eye paths, it is also possible to reduce the number of point light sources used.

Discontinuity buffer: It is straightforward to apply the discontinuity buffer [Kel98] for faster but biased anti-aliasing.

Non-blocking parallelization: Our method is a Monte Carlo algorithm and as such trivial to parallelize. By the high coherency our algorithm in addition benefits from realtime ray tracing architectures as introduced in [WKB⁺02, BWS03] and improves their image quality.

Finally, it is known that some caustic paths cannot be captured efficiently by any path tracing algorithm [KK02a], however, these are easily complemented by a caustic photon map [Jen01].

As shown in several papers [WKB⁺02, BWS03], approximations to global illumination can be computed at interactive frame rates. In our implementation the unbiased solution requires roughly up to 30% more computing power as compared to the biased version.

3.3 Interpretation of the Bias Compensation Step

The bias compensation step 2.c of the algorithm could be considered as a secondary final gathering [Chr99] as well as Russian roulette for an unbiased path termination. A third interpretation is available in the context of ambient occlusion techniques [IKSZ03, Neu03]: In Fig. 2 the average path length of our eye paths is displayed as a gray image, where pixels are darker as the eye path becomes longer. Because paths are terminated whenever the geometry term is below the bound b , this image in fact looks like an image computed by the ambient occlusion technique. In Sect. 2 we thus provided the missing mathematical facts for why ambient occlusion works so fine and is justified.

Moreover, by our technique, we do not only scan the hemisphere around one point, but the whole vicinity that can be reached by short paths. This completely removes the problem of blurry patterns in concave corners as it may occur with final gathering [Chr99].

4 Conclusion

We presented a new mathematical framework for robustly computing integrals of integrands containing weak singularities without any numerical exceptions. Based on this concept we derived a robust algorithm for computing local solutions of a Fredholm integral equation of the second kind. In the context of computer graphics our approach is more general than ambient occlusion and secondary final gathering. The implementation exposes the simplicity of a path tracer and the resulting images do not show the artifacts of current state-of-the-art rendering techniques, since the algorithm is unbiased. Compared to other unbiased techniques like classical bidirectional path tracing, our method is more efficient and easily implemented in professional rendering software systems.

The method of how the integrands are bounded allows one to efficiently apply hierarchical Monte Carlo methods [Hei00,5]. In future research, we also will investigate how to determine the constant c other than by numerical experiments. Finally the combination of occlusion maps and shadow buffering by our new method can yield more realistic hardware rendering algorithms.

Acknowledgements

The first author has been funded by the Stiftung Rheinland- Pfalz für Innovation. The paper has been dedicated to Anneliese Oeder. The authors thank Regina Hong for proofreading.

References

- [BWS03] C. Benthin, I. Wald, and P. Slusallek, *A Scalable Approach to Interactive Global Illumination*, Computer Graphics Forum **22** (2003), no. 3, 621–629.
- [Chr99] P. Christensen, *Faster Photon Map Global Illumination*, Journal of Graphics Tools **4** (1999), no. 3, 1–10.
- [FBG02] S. Fernandez, K. Bala, and D. Greenberg, *Local Illumination Environments for Direct Lighting Acceleration*, Rendering Techniques 2002 (Proc. 13th Eurographics Workshop on Rendering) (P.Debevec and S. Gibson, eds.), 2002, pp. 7–13.
- [Hei00] S. Heinrich, *Monte Carlo Approximation of Weakly Singular Operators*, Talk at the Dagstuhl Seminar 00391 on Algorithms and Complexity for Continuous Problems, Sept. 2000.
- [IKSZ03] A. Iones, A. Krupkin, M. Sbert, and S. Zhukov, *Fast, Realistic Lighting for Video Games*, IEEE Computer Graphics and Applications **23** (2003), no. 3, 54–64.
- [Jen01] H. Jensen, *Realistic Image Synthesis Using Photon Mapping*, AK Peters, 2001.
- [Kel97] A. Keller, *Instant Radiosity*, SIGGRAPH 97 Conference Proceedings, Annual Conference Series, 1997, pp. 49–56.
- [Kel98] ———, *Quasi-Monte Carlo Methods for Photorealistic Image Synthesis*, Ph.D. thesis, Shaker Verlag Aachen, 1998.
- [Kel01] ———, *Hierarchical Monte Carlo Image Synthesis*, Mathematics and Computers in Simulation **55** (2001), no. 1-3, 79–92.
- [Kel02] A. Keller, *Beyond Monte Carlo – Course Material*, Interner Bericht 320/02, University of Kaiserslautern, 2002, Lecture at the Caltech, July 30th – August 3rd, 2001.
- [Kel03] A. Keller, *Strictly Deterministic Sampling Methods in Computer Graphics*, SIGGRAPH 2003 Course Notes, Course #44: Monte Carlo Ray Tracing (2003).
- [KH01] A. Keller and W. Heidrich, *Interleaved Sampling*, Rendering Techniques 2001 (Proc. 12th Eurographics Workshop on Rendering) (K. Myszkowski and S. Gortler, eds.), Springer, 2001, pp. 269–276.
- [KK02a] T. Kollig and A. Keller, *Efficient Bidirectional Path Tracing by Randomized Quasi-Monte Carlo Integration*, Monte Carlo and Quasi-Monte Carlo Methods 2000 (H. Niederreiter, K. Fang, and F. Hickernell, eds.), Springer, 2002, pp. 290–305.
- [KK02b] ———, *Efficient Multidimensional Sampling*, Computer Graphics Forum **21** (2002), no. 3, 557–563.
- [Neu03] I. Neulander, *Image-Based Diffuse Lighting using Visibility Maps*, Proceedings of the SIGGRAPH 2003 Conference on Sketches & Applications, 2003, pp. 1–1.
- [Shr66] Y. Shreider, *The Monte Carlo Method*, Pergamon Press, 1966.
- [Vea97] E. Veach, *Robust Monte Carlo Methods for Light Transport Simulation*, Ph.D. thesis, Stanford University, 1997.
- [VG95] E. Veach and L. Guibas, *Optimally Combining Sampling Techniques for Monte Carlo Rendering*, SIGGRAPH 95 Conference Proceedings, Annual Conference Series, 1995, pp. 419–428.

- [War91] G. Ward, *Adaptive Shadow Testing for Ray Tracing*, 2nd Eurographics Workshop on Rendering (Barcelona, Spain), 1991.
- [WBS03] I. Wald, C. Benthin, and P. Slusallek, *Interactive Global Illumination in Complex and Highly Occluded Environments*, Rendering Techniques 2003 (Proc. 14th Eurographics Workshop on Rendering) (P. Christensen and D. Cohen-Or, eds.), 2003, pp. 74–81.
- [WBWS01] I. Wald, C. Benthin, M. Wagner, and P. Slusallek, *Interactive Rendering with Coherent Ray Tracing*, Computer Graphics Forum **20** (2001), no. 3, 153–164.
- [WKB⁺02] I. Wald, T. Kollig, C. Benthin, A. Keller, and P. Slusallek, *Interactive Global Illumination using Fast Ray Tracing*, Rendering Techniques 2002 (Proc. 13th Eurographics Workshop on Rendering) (P. Debevec and S. Gibson, eds.), 2002, pp. 15–24.