TECHNISCHE UNIVERSITÄT
KAISERSLAUTERN

Postfach 3049 · D-67653 Kaiserslautern

Interner Bericht

On the Efficiency of Multiple Importance Sampling: Robust Instant Global Illumination

T. Kollig, A. Keller

Technical Report 328/04
Kaiserslautern University of Technology

FACHBEREICH
INFORMATIK

On the Efficiency of Multiple Importance Sampling: Robust Instant Global Illumination

T. Kollig, A. Keller

328/04

Technische Universität Kaiserslautern
AG Numerische Algorithmen
Postfach 30 49
67653 Kaiserslautern
Germany

Februar 2004

Herausgeber: AG Numerische Algorithmen
Leiter: Professor Dr. S. Heinrich

On the Efficiency of Multiple Importance Sampling: Robust Instant Global Illumination

T. Kollig[1] and A. Keller[2]

[1] University of Kaiserslautern, Germany
[2] University of Ulm, Germany

Abstract

Approximating illumination by point light sources, as done in many professional applications, suffers from the problem of the weak singularity: Numerical exceptions caused by the division by the squared distance between the point light source and the point to be illuminated must be avoided. Multiple importance sampling overcomes these problems by combining multiple sampling techniques by weights. Such a set of weights is called a heuristic. So far the estimators resulting from a heuristic only have been analyzed for variance. Since the cost of sampling is not at all constant for different sampling techniques, it is possible to find more efficient heuristics, even though they may have higher variance. Based on our new stratification heuristic, we present a robust and unbiased global illumination algorithm. By numerical examples, we show that it is more efficient than previous heuristics. The algorithm is as simple as a path tracer, but elegantly avoids the problem of the weak singularity.

1. Introduction

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

Based on our new *stratification heuristic*, we present a robust, faster, and very simple global illumination algorithm that is suited for production as well as interactive rendering. The new algorithm is easily implemented in any ray tracing system and benefits from orthogonal techniques as e.g. the photon map.

2. The new Global Illumination Algorithm

We first introduce the new global illumination algorithm (see figures 1 and 5) in order to illustrate its simplicity:

1. **Generation of Point Light Sources:** Identical to the preprocessing of the instant radiosity algorithm [Kel97] or a very sparse global photon map [Jen01], a set $(y_j, L_j)_{j=0}^{M-1}$ of M point light sources is created. In terms of bidirectional path tracing this corresponds to tracing light paths starting from the lights and storing all the points $y_j \in \mathbb{R}^3$ of incidence with the radiance 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 the contributions for the current point x_i until the eye path is terminated:

   a. **Shadows:** The contribution of the j-th point light source is

   $$w(x_i, y_j) f_r(\omega_{x_i,y_j}, x_i, \omega_i) G(x_i, y_j) V(x_i, y_j) L_j \quad , \qquad (1)$$

   where

   $$w(x_i, y_j) := \begin{cases} 1 & G(x_i, y_j) < t \\ 0 & \text{else} \end{cases} \quad . \qquad (2)$$

   The direction ω_{x_i,y_j} points from x_i to y_j, ω_i is the direction from where x_i has been hit, and the visibility $V(x_i, y_j)$ is either 1, if the two points are mutually visible, or zero otherwise. Obviously a shadow ray for determining the visibility only has to be traced if the geometry term

   $$G(x_i, y_j) := \frac{\cos^+ \theta_{x_i} \cdot \cos^+ \theta_{y_j}}{\|x_i - y_j\|_2^2} < t \qquad (3)$$

   for some threshold parameter $t > 0$. By appropriately choosing the threshold t numerical exceptions caused

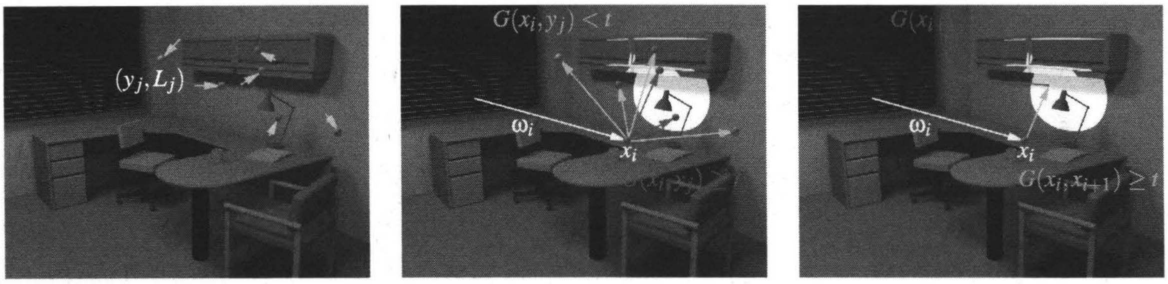

1. Generation of point light sources | 2.a. Robust shadowing | 2.b. Shading by scattering

Figure 1: *Principle steps of the unbiased robust global illumination algorithm as a simple comic: 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 threshold $t = 0.5$. In step 2.a. shadow rays towards the point lights outside this domain are traced. This is robust, because numerical exceptions by the inverse squared distance in G cannot occur. In order to be unbiased, step 2.b. 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.*

by the denominator of G efficiently are avoided. The positive cosine $\cos^+ \theta_{x_i}$ is the scalar product between the unit direction incident in x_i and the surface normal in x_i, which is set to zero, if the cosine is less than zero (analogous for $\cos^+ \theta_{y_j}$). Light sources that are hit directly contribute $L_e(x_i, -\omega_i)$ in addition.

b. **Scatter:** In order to account for the contribution of the cases $G(x_i, y_j) \geq t$, i.e. for all techniques, which had guaranteed zero weight in the previous step, we trace only one ray into a random direction yielding the next vertex x_{i+1} along the eye path on the scene surface S. If $G(x_i, x_{i+1}) < t$ then the eye path is terminated, otherwise i is incremented and we continue with step (a), whose result is attenuated by the color of the bidirectional reflectance distribution function f_r.

2.1. Contribution

As we will show in the next section, our new algorithm can be considered as a variant of bidirectional path tracing, especially a variant of stratified sampling. It has important advantages:

1. Instead of variance, we will analyze the more important efficiency: Despite potentially higher variance, the new scheme is more efficient than any known bidirectional path tracing algorithm, but as robust and unbiased. This results in shorter rendering times for the same image quality.
2. The actual implementation exposes the simplicity of a path or distribution ray tracer, thus it easily is incorporated into any production ray tracing software. Moreover even realtime ray tracers [WKB*02, BWS03] can be improved with the new technique.
3. Opposite to similar approaches like e.g. instant radiosity [Kel97], the numerical problem of weak singularities caused by the inverse squared distance in the geometry term G is avoided in a very elegant way, which keeps the algorithm unbiased and makes it numerically robust.

Besides the clarity of the unbiased robust algorithm, it is obvious how to extend the method by orthogonal techniques (see section 4.2) like e.g. the photon map for caustics [Jen01], importance sampling for point light sources [War91, Kel98, WBS03], or the discontinuity buffer [Kel98] to only name a few.

3. Analysis of the Algorithm

As we will show in the subsequent analysis, the simple algorithm belongs to the class of bidirectional path tracers and such is based on profound theory. The key to the efficiency of our new heuristic is to consider the support of the single sampling techniques, which never has been done before: Although our new heuristic may not have optimal variance, it is more efficient than other heuristics (see section 4.1). In fact the new algorithm is a stratification technique.

We briefly recall the basics of bidirectional path tracing (for the full details see [VG95, Vea97]) and then introduce the new algorithm as the *stratification heuristic*.

3.1. Bidirectional Path Tracing

For bidirectional path tracing, the global illumination problem has to be formulated as a path integral. We use the product area measure $d\mu_\ell(\bar{x}) = \prod_{i=0}^{\ell} dx_i$ on the surface S of the scene for all paths $\bar{x}$ of length ℓ in the path space

$$\mathcal{P}_\ell = \{\bar{x} = x_0 x_1 \ldots x_\ell \mid x_i \in S\} .$$

The radiance contribution

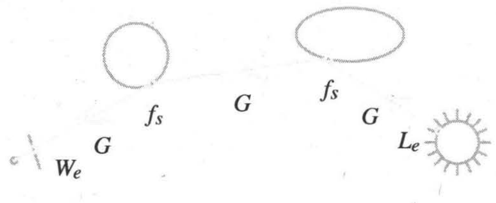

$$f(\bar{x}) = W_e \cdot G \cdot f_r \cdot G \cdot f_r \cdot G \cdot L_e$$

of a path $\bar{x} \in P_\ell$ is the product of the radiance L_e of the light source, the sensor responsitivity W_e, the bidirectional reflectance distribution functions f_r at the vertices, and the geometry terms G between all pairs of subsequent vertices of the path. The path integral formulation of the color I of a pixel then is

$$I = \sum_{\ell=1}^{\infty} \int_{\mathcal{P}_\ell} f(\bar{x}) d\mu_\ell(\bar{x}) = \int_{\mathcal{P}} f(\bar{x}) d\mu(\bar{x}) \ .$$

The bidirectional path tracing algorithm generates samples of the path space by first generating a light and an eye path. Then all vertices of both paths are deterministically connected by shadow rays (dashed lines) in order to determine the contribution.

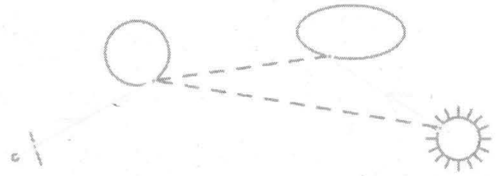

Consequently a path of length ℓ can be generated by ℓ techniques with the associated probability density functions

$$p_{\ell,0}, p_{\ell,1}, \ldots, p_{\ell,\ell-1} \ : \ \mathcal{P}_\ell \ \rightarrow \ \mathbb{R}_0^+$$

that are determined by which vertices deterministically are connected.

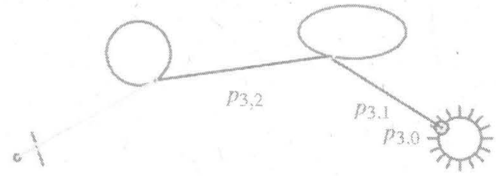

Combining all the path generation techniques by a set of weights

$$w_{\ell,0}, w_{\ell,1}, \ldots, w_{\ell,\ell-1} \ : \ \mathcal{P}_\ell \ \rightarrow \ \mathbb{R}_0^+ \ ,$$

i.e. a heuristic, yields the multiple importance sampling estimator

$$\int_{\mathcal{P}_\ell} f(\bar{x}) d\mu_\ell(\bar{x}) \approx \frac{1}{n} \sum_{j=1}^{n} \sum_{i=0}^{\ell-1} w_{\ell,i}(\bar{x}_{i,j}) \frac{f(\bar{x}_{i,j})}{p_{\ell,i}(\bar{x}_{i,j})} \ , \qquad (4)$$

where $\bar{x}_{i,j}$ is generated according to the density $p_{\ell,i}$. It can compensate for the weaknesses of a single technique, like e.g. the overmodulation of the geometry term G. The estimator (4) is unbiased, if for all possible paths $\bar{x} \in \mathcal{P}_\ell$ with contribution $f(\bar{x}) \neq 0$ the weights sum up to one, i.e.

$$\sum_{i=0}^{\ell-1} w_{\ell,i}(\bar{x}) = 1 \ ,$$

and if a path cannot be generated by a technique then the weight must be zero, i.e. $w_{\ell,i}(\bar{x}) = 0$ for $p_{\ell,i}(\bar{x}) = 0$.

Two such heuristics are the power heuristic [VG95, Vea97] with the weights

$$w_{\ell,i}^{\text{pow}}(\bar{x}) := \frac{p_{\ell,i}^{\beta}(\bar{x})}{\sum_{i=0}^{\ell-1} p_{\ell,i}^{\beta}(\bar{x})}$$

for $1 \leq \beta \leq \infty$ and the balance heuristic resulting from the special case $\beta = 1$. Compared to the optimal heuristic estimator η^{opt} [Vea97], the variance of a power heuristic estimator η^{pow} is bounded by

$$\sigma^2\left(\eta^{\text{pow}}\right) \leq c \cdot \sigma^2\left(\eta^{\text{opt}}\right) + \frac{1}{n} \cdot \mathbf{E}\left(\eta^{\text{opt}}\right) \ , \qquad (5)$$

where $c = 1$ for $\beta = 1$ and $c = \frac{1}{2}(1 + \sqrt{\ell})$ for $\beta = 2$. This means that the balance heuristic is asymptotically optimal.

The big disadvantage of these heuristics is that all densities have to be evaluated in order to determine the weights and all samples have to be generated even if their weighted contribution is rather small.

3.2. The Stratification Heuristic

Shooting shadow rays involves the evaluation of the geometry term $G(x, y)$ as defined in equation (3). If the two points x and y lie closely together, i.e. $\|y - x\|_2^2$ becomes very small, then numerical problems arise from the geometry term by either overmodulation or division by zero. Therefore various hacks including clipping the value to a plausible range or adding small values to the denominator are used to avoid these numerical problems.

Multiple importance sampling (4), as used in bidirectional path tracing, robustly avoids the problem of the weak singularity by suited heuristics. The resulting algorithms are unbiased and do not rely on the just mentioned hacks. However, all heuristics known so far require to evaluate all techniques, which make bidirectional path tracing rather inefficient and even complicated to implement.

Taking a closer look at what is really necessary to avoid the problems caused by the weak singularity reveals that

- in fact technique $p_{\ell,\ell-1}$ is sufficient, if $G(x_1, x_2) < t$,
- otherwise we have to proceed with technique $p_{\ell,\ell-2}$.

Continuing this scheme along the eye paths yields the

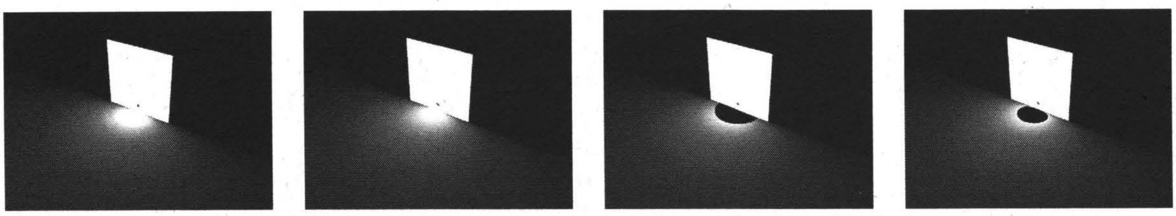

a. clipped illumination b. balance heuristic c. stratification by distance d. stratification heuristic

Figure 2: *Effect of the stratification for a single point light source (red dot): The a.) clipped overmodulation of the geometry term as present in the instant radiosity algorithm, the b.) attenuation by the weight of the balance heuristic, removing the weak singularity by c.) thresholding the distance to the point light source. Thresholding the geometry term d.) mimics best the integrand. A threshold $t = 0.5$ was used.*

weights

$$\begin{aligned} w_{\ell,\ell-1}(\bar{x}) &:= w(x_1,x_2) := \begin{cases} 1 & G(x_1,x_2) < t \\ 0 & \text{else} \end{cases} \\ w_{\ell,\ell-2}(\bar{x}) &:= (1 - w(x_1,x_2)) \cdot w(x_2,x_3) \\ &\vdots \end{aligned}$$

of our new *stratification heuristic*. Obviously the weights sum up to one as required for multiple importance sampling. Admittedly, it looks complicated now. However, it proves that the simple algorithm from section 2 is a bidirectional path tracing algorithm and consequently is unbiased.

Contrary to previous heuristics, especially the similar looking maximum heuristic, our new heuristic does not require to sample from all techniques before being able to determine the weights and consequently is more efficient.

In fact the new heuristic is a stratification method, where the technique to use is found by thresholding the geometry term. The cost of determining the stratum we are in is negligible, because parts of the integrand are used that have to be evaluated anyhow.

3.2.1. Choice of the Stratification

As illustrated in figure 1, the stratification heuristic divides the integration domain around any point x into a kind of near and far field. This is best seen from the radiance integral equation

$$\begin{aligned} &L(x,\omega') \\ &= L_e(x,\omega') + \int_\Omega f_r(\omega',x,\omega)L(y,-\omega)\cos^+\theta_x d\omega \\ &= L_e(x,\omega') \\ &\quad + \int_\Omega w(x,y) f_r(\omega',x,\omega)L(y,-\omega)\cos^+\theta_x d\omega \qquad (6) \\ &\quad + \int_\Omega (1-w(x,y)) f_r(\omega',x,\omega)L(y,-\omega)\cos^+\theta_x d\omega, (7) \end{aligned}$$

where $y = h(x,\omega)$ is the closest point hit from x, when shooting a ray into direction ω: The resulting first integral (6) is safely evaluated using point light sources as in step 2.a. of the algorithm, while the second integral (7) has to be evaluated by sampling the hemisphere as in step 2.b. in order to avoid numerical exceptions.

The shape of the strata is determined by the weight function (2) as illustrated in figure 2. Instead of just thresholding the inverse squared distance, which causes the integrand to be weakly singular, we threshold the whole geometry term G in order to closely mimic the integrand. The threshold parameter t then determines the size of the strata (see figure 3). Choosing $0 \le t < \infty$ avoids infinite variance to be caused by the geometry term G.

In the context of parametric integration, Heinrich [Hei00] proposed an optimal algorithm for the Monte Carlo approximation of weakly singular operators: For smooth function classes his algorithm used a similar stratification idea to separate the weak singularity.

4. Discussion and Numerical Evidence

The new algorithm units many seemingly isolated techniques in a simple way: Bidirectional path tracing, stratification, ambient occlusion and local illumination environments, and final gathering and secondary final gathering are all intrinsic.

On the one hand the scattering step 2.b of the algorithm could be considered as a secondary final gather [Chr99]. On the other hand it can be interpreted in the context of ambient occlusion techniques [IKSZ03, Neu03]. However, 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].

To summarize, the new algorithm is unbiased, because it can be regarded as bidirectional path tracing algorithm, especially a stratification method. The algorithm is also robust, because no numerical exceptions can be caused by the weak

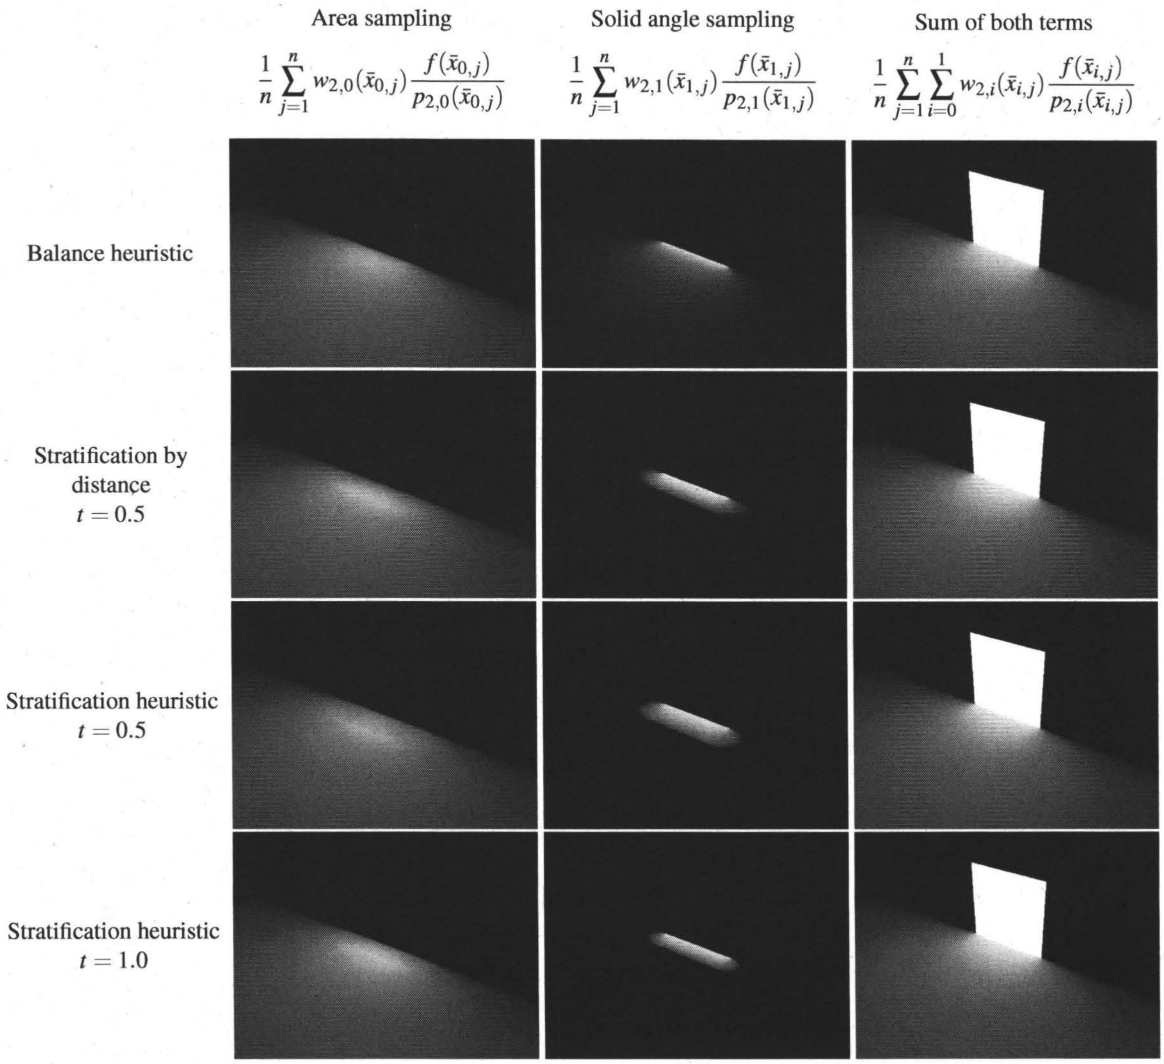

Figure 3: *For path length $\ell = 2$ we compare bidirectional path tracing heuristics in a simple scene with a light source on a plane. Each row represents a heuristic. The columns show the weighted contributions of sampling the light source by area and solid angle, while the third column shows the sum of both terms and the light source.*

singularity. Like instant radiosity [Kel97] or similar interactive global illumination systems [WKB*02, BWS03], it uses efficient area sampling by point light sources, but singularities are not clipped and averaged away in a biased way. This results in an increased image quality.

4.1. Efficiency and Time Complexity

Analyzing only the variance $\sigma^2(\eta)$ of a Monte Carlo estimator η ignores the cost of realizing the estimator. Therefore a complete analysis must consider the time complexity

$$\sigma^2(\eta) \cdot T(\eta),$$

whose inverse is the efficiency. $T(\eta)$ is the time of one realization of η. The variance of our new estimator may not optimal in the sense of (5), however, for several reasons its realization time is much smaller:

Cheaper rays: Equating the geometry term G and the threshold t allows one to bound the maximum length

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

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

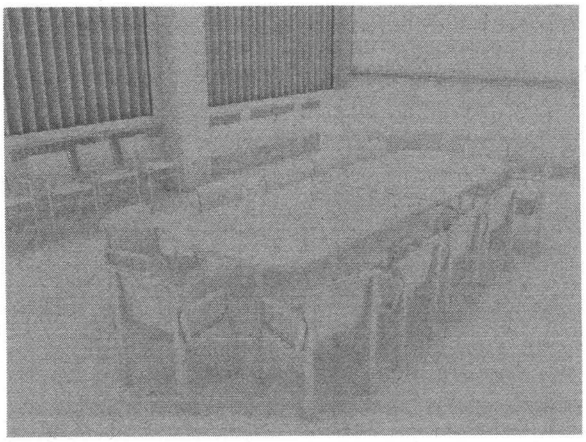

Figure 4: *From green to red the color indicates the average eye path length per pixel. While green corresponds to the minimum length of one, pure red indicates a maximal length of* 4.657*. The maximum path length considered was* 5*.*

have to be traversed. Without our heuristic it would be only possible to bound shadow rays.

Shorter eye paths: Compared to previous bidirectional path tracing heuristics, the eye path length of the stratification heuristic is shorter on the average. Thus less rays have to be traced and shaded as can be seen in figures 4 and 5.

Less shadow rays: The shorter average eye path length directly results in a reduced number of shadow rays to be shot and a consequently higher data locality. In addition many shadow rays do not need to be traced, because of the threshold (3) on the geometry term.

Intrinsic cache coherence: Only in the vicinity of concave corners the eye path length slightly increases as illustrated in figure 4. Then the ray length is short, which implies that most 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.

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 than previous bidirectional path tracing algorithms and in addition 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.

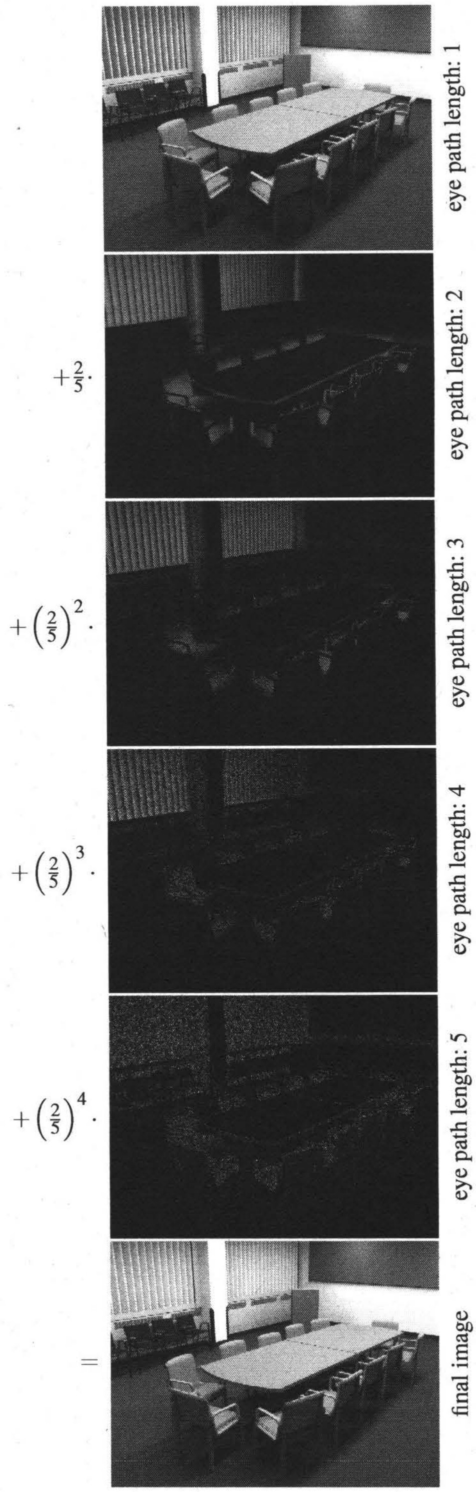

Figure 5: *Visualization of the contributions made by the different strata for* $t = 0.5$*, which are identified by the eye path length. Note that the images have been amplified for display, i.e. the contribution vanishes in an exponential way.*

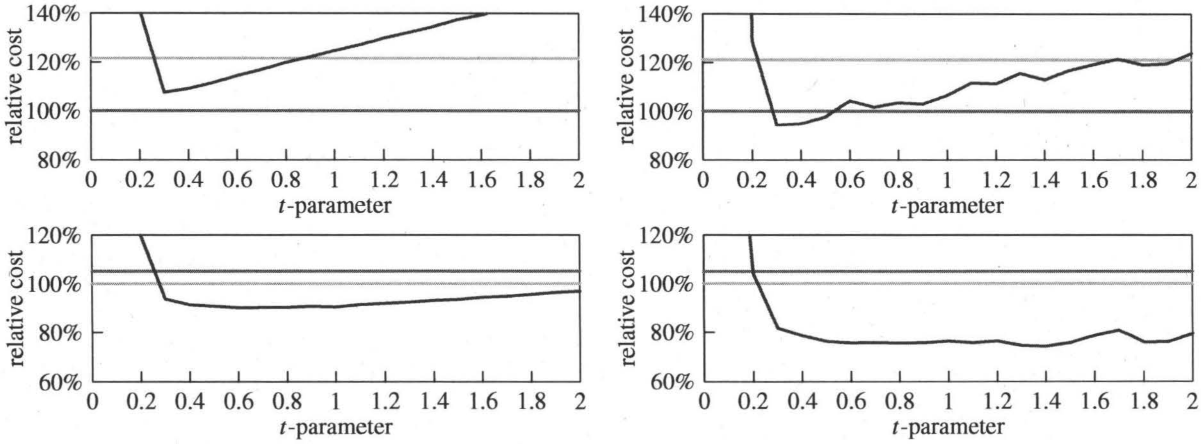

Figure 6: *Comparison of the relative rendering time of the power heuristic with $\beta = 2$ in red, balance heuristic in green, and stratification heuristic in blue at identical image quality. The top row shows the measurements for the office scene and the bottom row for the conference room. Timings on the left are pure Monte Carlo, while on the right we used the more efficient interleaved sampling, i.e. the method of dependent tests. While it is not clear, whether the balance or power heuristic is better, the stratification heuristic reliably outperforms the classical heuristics.*

4.1.1. Numerical Evidence

In figure 6 the relative running time is compared at identical image quality for the stratification heuristic, the power heuristic, and the balance heuristic. Obviously the performance of the power or balance heuristic depends on the scene and none of them is a clear favorite.

For the experiments in the left column of figure 6 we used a bidirectional path tracer, where for each sample an independent light and eye path have been generated. Due to the shorter eye paths, the stratification heuristic is competitive, but does not always beat the classical heuristics. For the measurements in the right column, the bidirectional path tracer has been improved by dramatically reducing the number of lights paths by interleaved sampling [KH01]. While the realization cost now has been decreased by the same constant time for all heuristics, clearly the stratification heuristic benefits most from reusing light paths and reliably is faster than the classical heuristics.

4.1.2. Choice of the Threshold t

As can be seen, the performance of the stratification heuristic depends on the choice of the threshold parameter t. Unfortunately, this parameter depends on the given scene, which can be easily seen by a short thought experiment: By shrinking or enlarging a scene also the stratification of the path space and thus the threshold t has to be resized. Note that this is intrinsic to heuristics, as e.g. the power heuristic, whose weights are based on the ratio of the probability densities. Nevertheless our experiments indicate that usually $t \approx 0.3$ results in the fastest rendering times.

4.2. Extensions

For clarity and brevity we only introduced the basic algorithm in section 2. For completeness we now indicate the obvious extensions by orthogonal techniques:

Efficient multidimensional sampling: In order to verify the convergence rates the images have been computed using pure random sampling and interleaved sampling [KH01]. Efficiency can be further increased by applying quasi-Monte Carlo and randomized quasi-Monte Carlo sampling methods [Kel02, KK02a, KK02b].

Caustics: Some caustic paths cannot be captured efficiently by any path tracing algorithm [KK02a], however, these easily are complemented by a caustic photon map.

Shadow computation: The techniques of Ward [War91] 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.

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.

5. Results and Conclusion

By considering efficiency, i.e. both variance and realization time, we introduced a faster global illumination algorithm. It has been shown that the new Monte Carlo algorithm in fact is a stratification method that as well can be considered as a multiple importance sampling heuristic. Consequently it is an unbiased and robust way to get rid of the problem of the weak singularity. Its 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.

Finally the combination of occlusion maps and shadow buffering by the stratification heuristic 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 the paper.

References

[BWS03] BENTHIN C., WALD I., SLUSALLEK P.: A Scalable Approach to Interactive Global Illumination. *Computer Graphics Forum 22*, 3 (2003), 621–629. 2, 5, 7

[Chr99] CHRISTENSEN P.: Faster Photon Map Global Illumination. *Journal of Graphics Tools 4*, 3 (1999), 1–10. 4

[FBG02] FERNANDEZ S., BALA K., GREENBERG D.: Local Illumination Environments for Direct Lighting Acceleration. In *Rendering Techniques 2002 (Proc. 13th Eurographics Workshop on Rendering)* (2002), Debevec P., Gibson S., (Eds.), pp. 7–13. 6

[Hei00] HEINRICH S.: Monte Carlo Approximation of Weakly Singular Operators. Talk at the Dagstuhl Seminar 00391 on Algorithms and Complexity for Continuous Problems, Sept. 2000. 4

[IKSZ03] IONES A., KRUPKIN A., SBERT M., ZHUKOV S.: Fast, Realistic Lighting for Video Games. *IEEE Computer Graphics and Applications 23*, 3 (2003), 54–64. 4

[Jen01] JENSEN H.: *Realistic Image Synthesis Using Photon Mapping*. AK Peters, 2001. 1, 2

[Kel97] KELLER A.: Instant Radiosity. In *SIGGRAPH 97 Conference Proceedings* (1997), Annual Conference Series, pp. 49–56. 1, 2, 5

[Kel98] KELLER A.: *Quasi-Monte Carlo Methods for Photorealistic Image Synthesis*. Ph.D. thesis, Shaker Verlag Aachen, 1998. 2, 7

[Kel02] KELLER A.: *Beyond Monte Carlo – Course Material*. Interner Bericht 320/02, University of Kaiserslautern, 2002. Lecture at the Caltech, July 30th – August 3rd, 2001. 7

[KH01] KELLER A., HEIDRICH W.: Interleaved Sampling. In *Rendering Techniques 2001 (Proc. 12th Eurographics Workshop on Rendering)* (2001), Myszkowski K., Gortler S., (Eds.), Springer, pp. 269–276. 7

[KK02a] KOLLIG T., KELLER A.: Efficient bidirectional path tracing by randomized quasi-Monte Carlo integration. In *Monte Carlo and Quasi-Monte Carlo Methods 2000*, Niederreiter H., Fang K.,, Hickernell F., (Eds.). Springer, 2002, pp. 290–305. 7

[KK02b] KOLLIG T., KELLER A.: Efficient Multidimensional Sampling. *Computer Graphics Forum 21*, 3 (Sept. 2002), 557–563. 7

[Neu03] NEULANDER I.: Image-Based Diffuse Lighting using Visibility Maps. In *Proceedings of the SIGGRAPH 2003 Conference on Sketches & Applications* (2003), pp. 1–1. 4

[Vea97] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1997. 2, 3

[VG95] VEACH E., GUIBAS L.: Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *SIGGRAPH 95 Conference Proceedings* (1995), Annual Conference Series, pp. 419–428. 2, 3

[War91] WARD G.: Adaptive Shadow Testing for Ray Tracing. In *2nd Eurographics Workshop on Rendering* (Barcelona, Spain, 1991). 2, 7

[WBS03] WALD I., BENTHIN C., SLUSSALEK P.: Interactive Global Illumination in Complex and Highly Occluded Environments. In *Rendering Techniques 2003 (Proc. 14th Eurographics Workshop on Rendering)* (2003), Christensen P., Cohen-Or D., (Eds.), pp. 74–81. 2, 7

[WBWS01] WALD I., BENTHIN C., WAGNER M., SLUSSALEK P.: Interactive Rendering with Coherent Ray Tracing. *Computer Graphics Forum 20*, 3 (2001), 153–164. 6

[WKB*02] WALD I., KOLLIG T., BENTHIN C., KELLER A., SLUSALLEK P.: Interactive Global Illumination using Fast Ray Tracing. In *Rendering Techniques 2002 (Proc. 13th Eurographics Workshop on Rendering)* (2002), Debevec P., Gibson S., (Eds.), pp. 15–24. 2, 5, 7