
Interner Bericht

Efficient Importance Sampling Techniques for the
Photon Map

Alexander Keller, Ingo Wald

302/00

Fachbereich Informatik

Universität Kaiserslautern · Postfach 3049 · D-67653 Kaiserslautern

Efficient Importance Sampling Techniques for the Photon Map

Alexander Keller, Ingo Wald

302/00

Universität Kaiserslautern
AG Numerische Algorithmen
Postfach 30 49
67653 Kaiserslautern
Germany

Januar 2000

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

Efficient Importance Sampling Techniques for the Photon Map

Alexander Keller and Ingo Wald
Kaiserslautern University

Abstract. The photon map provides a powerful tool for approximating the irradiance in global illumination computations independent from geometry. By presenting new importance sampling techniques, we dramatically improve the memory footprint of the photon map, simplify the caustic generation, and allow for a much faster sampling of direct illumination in complicated models as they arise in a production environment.

1 Introduction

As introduced in [Jen96a, JC95b] the photon map algorithm generates an approximation of irradiance by storing the collisions of a random walk simulation of radiance. Its basic implementation is strikingly simple and the independence of scene geometry is a big advantage as compared to finite element approaches. In the sequel we present three new algorithms, improving the current photon map techniques:

1. **Importance driven photon deposition:** Originally the photon map is generated using the von Neumann-Ulam scheme, i.e. a pure forward simulation of the particle nature of light. In [PP98] an importance driven method for generating photon maps has been introduced, which directed the photon paths into the areas of high visual importance. The improvements were hardly visible from the images, and the core problem remained: Photons are stored all over the scene, even in regions, in which the importance is almost zero. In addition the energy of the photons varies due to the weighting by the emission and scattering probabilities, increasing the variance of the estimate.
Instead of controlling emission and scattering as in [PP98], we control the deposition of the photons, resulting in a more precise photon map, reduced memory requirements and faster photon map queries.
2. **Automated generation of caustics:** Caustics, which are generally speaking the lighting effects that arise from light hitting a diffuse surface coming from a specular surface, require a photon map with a much higher resolution than the global photon map used for ambient lighting. In [Jen96b, JC95b] projection maps were used to shoot an increased number of photons into the direction of caustics generators. This however only captures caustics caused by direct lighting, whereas caustics from indirect illumination have to be added by massive user intervention. By extending our approach of importance driven photon deposition, we enable the automatic generation of all visually important caustic effects.
3. **Faster direct illumination computations:** Since the direct visualization of the photon map often is too blurry, the calculation of direct illumination is separated from the calculation of indirect illumination [Jen96a]. This has the advantage that the main contribution to an image is sampled very accurately, while only

the minor effect of indirect ambience is estimated from the photon map. The disadvantage, however, is that the sampling effort depends on the number of light sources, which can be large.

We present an importance sampling scheme, which enables a more accurate and faster computation of direct illumination in settings with a large number of occluded light sources.

2 Importance driven Photon Map Generation

The basic idea is to control the deposition of photons by visual importance. Using a representation of the visual importance by an importon map as introduced in [PP98], by a probabilistic decision we store photons only in areas of high visual impact. It such is possible to obtain the same quality as in previous approaches, but at reduced memory cost and increased query speed.

2.1 Probabilistic Photon Deposition

The photon map can be seen as a snapshot of a random walk simulation of radiance transport from the light sources. A convenient algorithm for the collision estimate with Russian Roulette path termination including deterministic low discrepancy sampling has been derived in [KMS94]. The simulation yields a cloud of particles \mathcal{P} , where each particle $k \in \mathcal{P}$ consists of its point of incidence x_k , its direction of incidence $\vec{\omega}_k$ and its irradiance L_k . Using the Dirac- δ -distribution the particle k is localized by

$$\delta_k(x, \vec{\omega}) := \delta(x - x_k) \cdot \delta(\vec{\omega} - \vec{\omega}_k).$$

We now modify the discrete density approximation of the irradiance L_i by introducing an arbitrary probability $p(x, \vec{\omega})$:

$$\begin{aligned} L_i(x, \vec{\omega}) &\approx \sum_{k \in \mathcal{P}} \delta_k(x, \vec{\omega}) L_k = \sum_{k \in \mathcal{P}} \delta_k(x, \vec{\omega}) L_k \frac{p(x, \vec{\omega})}{p(x, \vec{\omega})} \\ &= \sum_{k \in \mathcal{P}} \delta_k(x, \vec{\omega}) \frac{L_k}{p(x, \vec{\omega})} \int_0^1 \chi_{(0, p(x, \vec{\omega})]}(t) dt, \end{aligned}$$

where χ is the characteristic function of the interval $(0, p(x, \vec{\omega})]$. This transformation very much resembles the Russian Roulette absorption mechanism [KMS94, AK90], and similarly the integral now can be evaluated by a one-sample Monte Carlo integration, yielding

$$p(x, \vec{\omega}) = \int_0^1 \chi_{(0, p(x, \vec{\omega})]}(t) dt \approx \chi_{(0, p(x, \vec{\omega})]}(\xi) = \begin{cases} 1 & \xi \leq p(x, \vec{\omega}) \\ 0 & \text{else} \end{cases}$$

For the actual implementation this means to simply insert the random decision, whether to store or discard the photon. It such is possible to concentrate the photons in regions of high probability $p(x, \vec{\omega})$, since only the fraction of photons is stored for which $\xi \leq p(x, \vec{\omega})$ for some uniform random variable $\xi \in (0, 1]$. This is illustrated in figure 1, where we first show the standard photon map, then the importons and finally the same amount of photons deposited using the importance information. Figure 2 shows the improved photon distribution that is obtained by tracing more trajectories in the preprocessing step but storing the same number of photons as without the importance driven deposition.




Fig. 1. Photon map in a scene of 10x10 rooms (each with one light source) seen from the top. a) Standard photon map without importance, b) distribution of visual importance, i.e. the importons, and c) photon map generated using visual importance controlled deposition.




Fig. 2. 10x10 rooms lit by 100 light sources. In the top row $6 \cdot 10^5$ photons have been stored, whereas the bottom row $6 \cdot 10^6$ photons have been deposited. The left column directly visualizes the standard photon map, while the right column illustrates the improvement by importance driven photon deposition.




Fig. 3. Caustics and close-up at a) ratio $q = 1$, b) $q = \frac{1}{5}$, and c) $q = \frac{1}{20}$. Note that the amount of photons in the global photon map remains about unchanged.

2.2 Automatic Caustics Generation

In [JC95b] projection maps are used to shoot an increased number of photons in the directions of the solid angle covered by caustic generating objects. This approach covers direct caustics, i.e. $L[S]^+D$ paths. However secondary light sources (like e.g. bright diffuse reflections as often encountered in architectural visualizations) might cause indirect caustics which are not simulated by the previous approach. In order to automatically generate the very detailed caustic photon map and the less detailed global photon map [JC98], we partition the set of all photons \mathcal{P} into two sets

$$\sum_{k \in \mathcal{P}} \delta_k(x, \vec{\omega}) L_k = \sum_{k \in \mathcal{P}_{cau}} \delta_k(x, \vec{\omega}) L_k + \sum_{k \in \mathcal{P} \setminus \mathcal{P}_{cau}} \delta_k(x, \vec{\omega}) L_k,$$

where \mathcal{P}_{cau} contains all photons scattered by a non-diffuse surface¹, i.e. a potential caustic generator. Using the integral transformation from the previous section with the constant probability q , we approximate

$$\begin{aligned} L_i(x, \vec{\omega}) &\approx \sum_{k \in \mathcal{P}_{cau}} \delta_k(x, \vec{\omega}) L_k \frac{p(x, \vec{\omega})}{p(x, \vec{\omega})} + \sum_{k \in \mathcal{P} \setminus \mathcal{P}_{cau}} \delta_k(x, \vec{\omega}) L_k \frac{q \cdot p(x, \vec{\omega})}{q \cdot p(x, \vec{\omega})} \\ &\approx \sum_{k \in \mathcal{P}_{cau}} \delta_k(x, \vec{\omega}) \frac{L_k}{p(x, \vec{\omega})} \cdot \chi_{(0, p(x, \vec{\omega}))}(\xi_k) \\ &\quad + \sum_{k \in \mathcal{P} \setminus \mathcal{P}_{cau}} \delta_k(x, \vec{\omega}) \frac{L_k}{q \cdot p(x, \vec{\omega})} \cdot \chi_{(0, q \cdot p(x, \vec{\omega}))}(\xi_k), \end{aligned}$$

¹Caustics could be characterized more precisely by investigating focal properties of wavefronts and surface derivatives as used in e.g. [Ige99].

where the ξ_k are independent, uniform random variables on $[0, 1)$. The improved quality of the automatically generated caustic photon map for different choices of the probability q can be seen in figure 3.

2.3 The Algorithm

The implementation of the ideas of the previous sections is obtained by replacing the standard Store-call for photon k by the following code fragment:

```

if ( $k \in \mathcal{P}_{cau}$ )
{
    if ( $\xi_k < p(x_k)$ )
        CausticMap->Store( $x_k, \vec{\omega}_k, \frac{L_k}{p(x_k)}$ );
}
else if ( $\xi_k < q \cdot p(x_k)$ )
    GlobalMap->Store( $x_k, \vec{\omega}_k, \frac{L_k}{q \cdot p(x_k)}$ );

```

If the photon under consideration could generate a caustic, it is stored in the detailed caustic photon map if the probability induced by visual importance is high enough. Otherwise only the fraction q of the photons is stored in the global photon map, which in consequence is less detailed. This of course introduces the overhead of tracing $\frac{1}{q}$ times more trajectories than required to fill the global photon map. Since the computation of $p(x_k)$ is very expensive, it is more efficient to use two random variables $\xi_{k,1}, \xi_{k,2}$, first testing whether $\xi_{k,1} < p$ and then if necessary whether $\xi_{k,2} < p(x_k)$. This indicates a third alternative: Only the fraction of q trajectories is used for the global photon map. Using a total of N trajectories, we use the first $\lfloor q \cdot N \rfloor$ trajectories to fill the caustic and global photon map by only testing for $\xi_k < p(x_k)$ and then use the remaining trajectories $\lfloor q \cdot N \rfloor + 1 \dots N$ to complete the caustic photon map, yielding the most efficient implementation.

The Deposition Probability $p(x)$. While the ratio q of collisions to be stored in the global and the caustic photon map is constant, the probability p is determined from the visual importance. Although we could use $p(x, \vec{\omega})$, we only use a probability $p(x)$ which solely depends on the location x and is faster to determine. Identical to the approach in [PP98], we shoot importons from a general camera sensor into the scene. The probability $p(x)$ then is determined by performing a query for the importons. Since the photon map performs best, if the photons are of about equal irradiance, the collected visual importance is scaled by a factor f and clipped if above 1, where f is chosen so that $p(x) = 1$ in the areas of high visual importance, leaving the irradiance L_k unchanged. Only for the regions where $p(x_k) < 1$, i.e. for regions of almost no visual impact, the irradiance grows to $\frac{L_k}{p(x_k)}$ increasing the variance of the irradiance approximation. The scaling factor f can be estimated by tracing some test rays from the camera into the scene sensing for the visual importance. Choosing the reciprocal value of the minimum of these queries to be f fulfils the requirements.

Numerical problems can arise in regions where $\frac{L_k}{p(x_k)}$ becomes very large due to a small deposition probability $p(x_k)$. This can be remedied in two ways by either bounding $p(x_k)$ from below by some $\epsilon > 0$, or by discarding all photons with $p(x_k) <$

ϵ , which although hardly perceivable of course yields a biased algorithm.

Discarding photons requires to trace more trajectories, i.e. having a more expensive preprocessing step, in order to obtain the same number of photons as in the pure forward simulation, but is paid off by the much better photon distribution in the regions of high visual importance at the same memory cost, which is illustrated in figure 2. On the other hand storing photons only in areas of high visual impact can dramatically reduce the memory footprint of the photon maps by preserving their quality. Assuming the scale f of visual importance to be determined by the above procedure, the compact set of parameters (N, q) of the number N of photons to trace, and the ratio q of the global and caustic photons of the new importance driven photon map generation algorithm is intuitive and increases usability since no manual intervention is required for complete caustics generation.

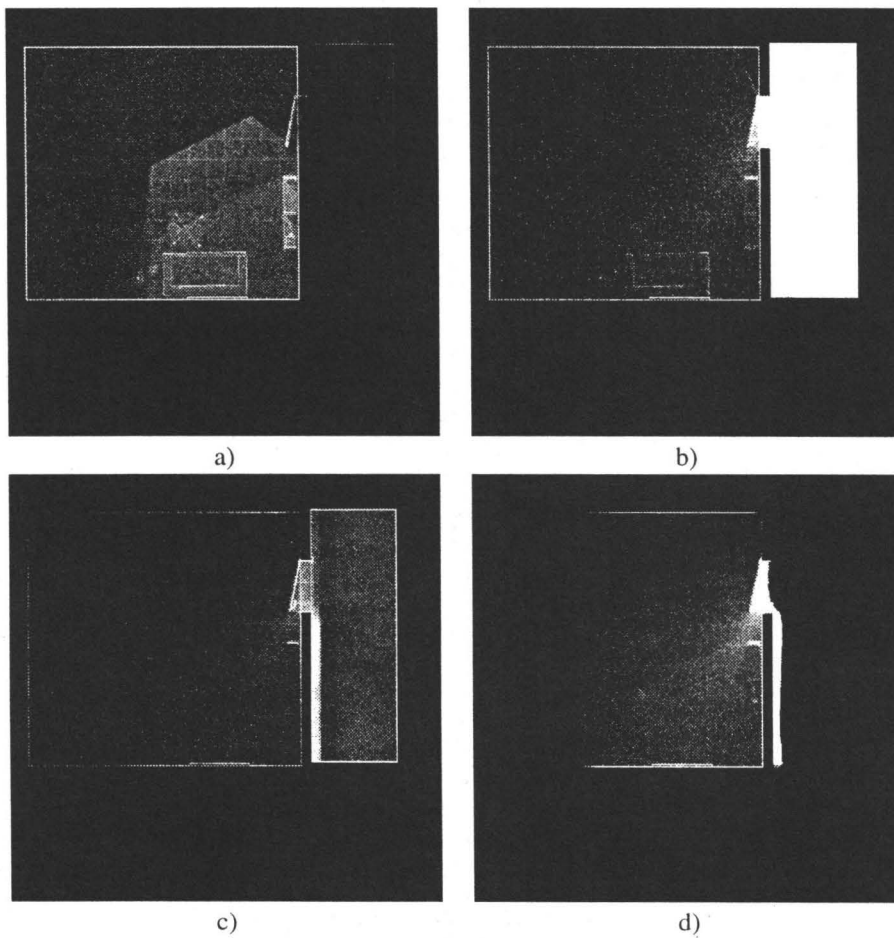


Fig. 4. The very problematic setting of a room being lit by a door slit seen from atop. The room on the left has no light sources but the camera within; the room on the right contains one light source. a) The importance distribution, b) the standard photon map, c) the unbiased importance driven photon map, and the biased version using a minimum cutoff probability ϵ .

Remaining Problems. Our approach is complementary to [PP98], where the emission rates and shooting directions are sampled from the visual importance. In a combination the advantages of our algorithm, especially the improved memory requirements, would transfer. In settings of multiple light sources, however, the algorithm of [PP98] causes problems: The importance of a light source is determined by a small random walk, and the fraction of trajectories started from one light source is given by the ratio of its importance to the total sum of importances. In consequence the photons are of very different irradiance. If now a photon of an unimportant light source scatters into a region of high visual importance, it increases the variance of the photon map estimates. This effect is further amplified by the importance sampling of scattering directions. The problem is not encountered by exclusively using our approach, since the photons in the important regions remain of about equal irradiance as explained in the previous section.

Solely selecting the scattering directions by visual importance as in [PP98] can be problematic, too. This is illustrated in figure 4, where several photon distributions of a room lit by a doorslit are compared. The standard approach would deposit the majority of the photons in the visually unimportant part of the scene, since the small doorslit is hard to hit, even when using importance sampling. The probability controlled deposition requires much more trajectories to be traced, but stores the photons only where they are important, yielding a better basis for the successive rendering step.

In figures 4c and 4d we see the photon distribution of the importance driven photon map generation. Due to energy bleeding [KW00] the importance shines through the walls and photons are deposited in area of actually no importance, i.e. behind the wall. This could be avoided by using the Metropolis Light Transport algorithm [VG97, PKK00] to deposit photons by the techniques of the previous sections. This however indicates that using MLT alone would yield a more sophisticated algorithm.

3 Efficient Direct Illumination Computation

The efficient computation of direct illumination still is a core problem of photorealistic image synthesis. Exact shadow computations require a lot of shadow rays to be shot, especially in environments with a large number of light sources. The performance of shadow caches is not worth the effort for current scenes, since the coherence is destroyed by the fine tessellation or the incoherent secondary rays of global illumination algorithms [Smi98]. Approaches as in [HDG99] indicate that nevertheless coherence can be exploited, but fail for secondary rays, general scene geometry and due to a substantial memory footprint. A very powerful approach is due to Ward [War91], where the most important light sources are sampled first and shadow rays are saved by estimating the contribution of the less important light sources. The approach has been considered further by Shirley in [SWZ96], motivated by an article of Kirk on unbiased sampling [KA91]. Similar to the observations in [HDG99], in [PMS⁺99] it has been shown that reasonable approximations can be obtained by shooting only one shadow ray in the region of penumbra. Concerning the photon map, Jensen developed the technique of shadow photons [JC95a], where however the question where to deposit shadow photons remains an open problem or is impractical due to excessive memory consumption. Even assuming this problem to be solved satisfactory, the shadow estimates are too imprecise and fail for finely structured shadow details.

We now combine the idea of Ward of partitioning the set of lightsources \mathcal{L} into two sets, defining probabilities for the lightsources in both sets like Shirley and sampling from these sets in an unbiased way similar to Arvo to a unified concept based on the ideas already presented in [Kel98]: We extend the photon data structure by storing the




Fig. 5. Direct lighting calculations a) without, b) unbiased with importance, and c) biased with importance information generated from the photon map. Images a) and b) took the same time to compute. The noise in the rear of b) demonstrates the overmodulation effect due to a too small importon query radius, which is eliminated by omitting to sample from $\mathcal{L} \setminus \mathcal{L}_{con}$ in c).

light source identification with each primary photon, i.e. each photon that directly came from a light source. For the point to be illuminated directly, we perform a photon map query. The collected set of primary photons now very roughly indicates which light source may be visible from the query point. The light sources, whose identification is not among the collected photons are most probably occluded. So we have the two sets \mathcal{L}_{con} of potentially contributing light sources and $\mathcal{L} \setminus \mathcal{L}_{con}$ of most probably occluded light sources.

We need the stochastic evaluation of sums, which is derived by the integral transformation

$$\begin{aligned} \sum_{j=0}^{M-1} s_j &= \int_0^M \left(\sum_{j=0}^{M-1} s_i \chi_{[j,j+1)}(x) \right) dx \\ &\approx \frac{M}{S} \sum_{i=0}^{S-1} \left(\sum_{j=0}^{M-1} s_i \chi_{[j,j+1)}(Mx_i) \right), \end{aligned} \quad (1)$$

where S random samples $x_0, \dots, x_{S-1} \in [0, 1)$ are used to evaluate the sum of M summands s_i . We now fix the amount of S_{con} and $S_{non-con}$ shadow rays, which is independent of the actual number of light sources, in order to evaluate the lighting contributions of \mathcal{L}_{con} and $\mathcal{L} \setminus \mathcal{L}_{con}$, respectively.

If $S_{con} < |\mathcal{L}_{con}|$, we use (1) to randomly select S_{con} locations on the $|\mathcal{L}_{con}|$ light sources for the direct illumination computations. Otherwise the variance can be further reduced by stratifying the total of S_{con} directly to the light sources in \mathcal{L}_{con} by their emission area². Concerning the light sources which probably do not contribute, the minimal choice $S_{non-con} = 1$ guarantees an unbiased algorithm, i.e. we shoot one shadow ray to a randomly selected location on a randomly selected light source from $\mathcal{L} \setminus \mathcal{L}_{con}$. From (1) we infer that this sample is scaled by $\frac{|\mathcal{L} \setminus \mathcal{L}_{con}|}{1}$ and in consequence any unoccluded light source belonging to $\mathcal{L} \setminus \mathcal{L}_{con}$ can cause a very high variance. It such is very important to select \mathcal{L}_{con} securely, i.e. to use a large radius in the photon map query, in order to guarantee an almost zero probability of missing any contributing

²Note that the composition method can raise variance. It such is more efficient to use biased poststratification estimators [KW86].

light source. In that case, however, omitting that single shadow ray is biased, but hardly noticeable as shown in figure 5c.

This basic technique is illustrated in figure 5, where we compare sampling the light sources with and without importance. Image 5b clearly shows much less noise than the direct computation without importance function in 5a at the same computational cost. We intentionally chose a little too small radius for the importon map queries in order to demonstrate the punctual overmodulations, which have been explained in the previous section.

This importance sampling scheme is easily added to an existing photon map implementation. It is much simpler than the spatial partitioning scheme from [SWZ96] and more exact than the global visibility estimates in [War94]. However it can be completed by using Ward's sorting algorithm [War94] for the case $S_{con} < |\mathcal{L}_{con}|$, where then the S_{con} samples can be distributed to the brighter light sources. In the other case the samples are distributed by the importance of the light source, which is computed identical to [PP98]. The user so only needs to specify the overall shadow ray number and does not need to take care of the sample rate for the single light sources.

The technique also easily transfers to the bidirectional path tracing algorithm [LW93, VG94, VG95], where it is used to save shadow or connection rays. Then the photon map would serve as a point approximation of the radiance like in [Kel97]. This would also improve the bidirectional mutations of the MLT [VG97] algorithm.

4 Conclusion

We presented new importance sampling techniques for the photon map, which result in a reduced memory footprint and increased rendering efficiency. Our variant of the importance driven generation of the photon map controls the deposition, which is more efficient than controlling the scattering directions. It also allows to automatically generate the complete caustic photon map at the required high resolution. Finally we presented an efficient sampling scheme for environments with a large number of light sources of which a reasonable fraction is occluded. These importance sampling techniques reduce the required user interaction for e.g. fine tuning of the emission and sampling rates for the light sources. This automatism increases usability of the photon map algorithms by replacing the set of parameters by an intuitive one requiring less user intervention in the actual computations. Although the techniques presented here are major improvements over current algorithms, open problems remain (see also [KW00]):

1. The sampling of the direct illumination efficiently detects occlusions. However it does not save shadow rays, if a large number of the light sources is visible. Spatial coherence such needs to be exploited.
2. Photon map implementations typically sample the light sources directly. The current trend in design however is to use more and more indirect illumination. Due to the lack of direct illumination the photon map is visualized directly exposing its blurry artifacts.

Ward partially dealt with these issues in his Radiance rendering system [War94, LS98] by reclassifying indirect lights as light sources and performing irradiance interpolation [LS98]. These techniques however are computationally expensive. Simpler and more robust techniques for these problems are subject of our future research and will be based on the photon map mechanism for representing discrete densities.

Acknowledgement

The authors like to thank Gerald Maitshke from Compaq Computers for supporting this work by the donation of an DEC Alpha workstation.

References

- AK90. J. Arvo and D. Kirk. Particle Transport and Image Synthesis. In *Computer Graphics (SIGGRAPH 90 Conference Proceedings)*, pages 63 – 66, 1990.
- HdG99. D. Hart, P. Dutré, and D. Greenberg. Direct Illumination with Lazy Visibility Evaluation. In *SIGGRAPH 99 Conference Proceedings*, Annual Conference Series, pages 147–154, 1999.
- Ige99. H. Igehy. Tracing Ray Differentials. In *SIGGRAPH 99 Conference Proceedings*, Annual Conference Series, pages 179–186, 1999.
- JC95a. H. Jensen and N. Christensen. Efficiently Rendering Shadows Using the Photon Map. In H. Santo, editor, *Edugraphics + Compugraphics Proceedings*, pages 285–291. GRASP- Graphic Science Promotions & Publications, 1995.
- JC95b. H. Jensen and N. Christensen. Photon Maps in Bidirectional Monte Carlo Ray Tracing of complex Objects. *Computer and Graphics*, 19(2):215–224, 1995.
- JC98. H. Jensen and P. Christensen. Efficient Simulation of Light Transport in Scenes with Participating Media using Photon Maps. In Michael Cohen, editor, *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pages 311–320. ACM SIGGRAPH, Addison Wesley, July 1998.
- Jen96a. H. Jensen. Global Illumination Using Photon Maps. In *Rendering Techniques '96 (Proc. 7th Eurographics Workshop on Rendering)*, pages 21–30. Springer, 1996.
- Jen96b. H. Jensen. Rendering Caustics on Non-Lambertian Surfaces. In *Proc. Graphics Interface*, pages 116–121. Morgan Kaufmann, 1996.
- KA91. D. Kirk and J. Arvo. Unbiased Sampling Techniques for Image Synthesis. *Computer Graphics (ACM SIGGRAPH '91 Proceedings)*, 25(4):153–156, July 1991.
- Kel97. A. Keller. Instant Radiosity. In *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 49–56, 1997.
- Kel98. A. Keller. *Quasi-Monte Carlo Methods for Photorealistic Image Synthesis*. Ph.D. thesis, Shaker Verlag Aachen, 1998.
- KMS94. A. Kersch, W. Morokoff, and A. Schuster. Radiative Heat Transfer with Quasi-Monte Carlo Methods. *Transport Theory and Statistical Physics*, 7(23):1001–1021, 1994.
- KW86. M. Kalos and P. Whitlock. *Monte Carlo Methods, Volume I: Basics*. J. Wiley & Sons, 1986.
- KW00. A. Keller and I. Wald. An Analysis of the Photon Map. Interner Bericht, University of Kaiserslautern, 2000.
- LS98. G. Larson and R. Shakespeare. *Rendering with Radiance*. Morgan Kaufmann Publishers, 1998.
- LW93. E. Lafortune and Y. Willems. Bidirectional Path Tracing. In *Proc. 3rd International Conference on Computational Graphics and Visualization Techniques (Compugraphics)*, pages 145–153, 1993.
- PKK00. M. Pauly, T. Kollig, and A. Keller. Metropolis Light Transport for Participating Media. Interner Bericht, University of Kaiserslautern, 2000.
- PMS⁺99. S. Parker, W. Martin, P. Sloan, P. Shirley, B. Smits, and C. Hansen. Interactive Ray Tracing. In *Symposium on Interactive 3D Graphics*, pages 119–126. ACM SIGGRAPH, 1999.
- PP98. I. Peter and G. Pietrek. Importance driven Construction of Photon Maps. In *Rendering Techniques '98*, pages 269–280, 1998.
- Smi98. B. Smits. Efficiency Issues for Ray Tracing. *Journal of Graphics Tools*, 3(2):1–14, 1998.

- SWZ96. P. Shirley, C. Wang, and K. Zimmerman. Monte Carlo Techniques for Direct Lighting Calculations. *ACM Trans. Graphics*, 15(1):1–36, 1996.
- VG94. E. Veach and L. Guibas. Bidirectional Estimators for Light Transport. In *Proc. 5th Eurographics Workshop on Rendering*, pages 147 – 161, Darmstadt, Germany, June 1994.
- VG95. E. Veach and L. Guibas. Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 419–428, 1995.
- VG97. E. Veach and L. Guibas. Metropolis Light Transport. In Turner Whitted, editor, *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series, pages 65–76. ACM SIGGRAPH, Addison Wesley, August 1997.
- War91. G. Ward. Adaptive Shadow Testing for Ray Tracing. In *2nd Eurographics Workshop on Rendering*, Barcelona, Spain, 1991.
- War94. G. Ward. The RADIANCE Lighting simulation and Rendering System. In *Computer Graphics*, pages 459 – 472, July 1994.