IMPORTANCE-DRIVEN HIERARCHICAL STOCHASTIC RAY RADIOSITY

Jan Přikryl[†], Philippe Beakert[‡], and Werner Purgathofer[†]

[†]Institute of Computer Graphics, Vienna University of Technology
Karlsplatz 13/186/2, A-1040 Wien, Austria
e-mail: {prikryl,wp}@cg.tuwien.ac.at
[‡] Department of Computer Science, Katholieke Universiteit Leuven
Celestijnenlaan 200 A, B-3001 Leuven, Belgium
e-mail: Philippe.Bekaert@cs.kuleuven.be

ABSTRACT

In this paper we present a hierarchical Monte-Carlo radiosity algorithm driven by the view importance. The algorithm makes to possible to concentrate the computational effort on solution in the immediate environment of the observer, trading the low solution quality in invisible areas for better quality in areas that are visible for the observer. This is achieved by modifying the sampling probabilities of scene elements so that more samples are concentrated in the area of high importance and by extending the subdivision oracle function so that the subdivision is coarser in areas of low importance. This paper extends the previous work by introducing a combination of hierarchical refinement and view importance driven method for Monte-Carlo radiosity.

Keywords: Monte-Carlo, radiosity, hierarchy, hierarchical refinement, view importance, view potential

1 INTRODUCTION

Radiosity algorithms generally attempt to compute radiosity to a uniform precision throughout the whole environment. This results in globally over-solved and locally under-solved radiosity solution for most scenes [Smits92].

If we allow low accuracy of the solution in those parts of the scene, that are not directly visible and that do not influence the visible parts too much, we can spend more computation effort on parts that are directly visible. This way, we can save a lot of computational time when we are interested in illumination of only a part of a complex scene.

The principle of the method is well known: During the course of radiosity system computation, we are computing the second quantity, called *visual importance*. This quantity expresses the influence the radiosity of a particular mesh element has on the solution in the visible part of the scene.

In this paper we present a new radiosity method that combines the importance-driven Monte-Carlo radiosity and the hierarchical refinement of scene mesh elements. This way, both the memory and the time expenses of computing a single view of the scene can be reduced.

The paper is further organised as follows: in Section 2 we overview previous Monte-Carlo radiosity methods and importance-driven radiosity approaches. Hierarchical refinement for radiosity is brieifly overviewed in Section 3. In Section 4 we derive the new method. Results and comparisons of our method are presented in Section 5. Finally, in Section 6 we summarise our experiences and draw some ideas for further research.

2 MONTE-CARLO RADIOSITY

Monte-Carlo radiosity algorithms are — as well as their deterministic counterparts — based on algorithms used to compute radiative energy transport. Their advantage is that they quickly deliver solutions in that higher order interreflections are visible. Unfortunately, as with all Monte-Carlo methods, the variance of the solution drops slowly and the results suffer from noise.

Monte-Carlo radiosity solves the power form of the radiosity equation, given by

$$P_i = W_i + \rho_i \sum_{j=1}^{M} F_{ji} P_j, \qquad (1)$$

where P_i is the power of a receiving element, W_i is the self-emitted power of this element, ρ_i the reflectance of the element, F_{ji} the form-factor determining how much the power P_j radiated from element j contributes to the incoming power of the element i. The Equation (1) is solved by probabilistic simulation of photon paths traveled by photons leaving light sources in the scene.

Random-walk global illumination solution was proposed by [Patta92], followed by number of Monte-Carlo radiosity algorithms [Kelle97, Neuma97, Sbert97].

Importance-driven extensions to continuous random walks, that concentrate most of the particle paths to the region of interest, have been discussed by Pattanaik and Mudur [Patta93a, Patta93b, Patta95]. An extension of stochastic ray radiosity algorithm to take view importance into account was presented by Neumann *et al.* [Neuma96]. This method forms the basis of our algorithm.

2.1 Stochastic Ray Radiosity

Stochastic ray radiosity is a stochastic variant of the Jacobi iterative method. With stochastic ray radiosity, each iteration is performed by shooting particles from the scene elements. The expected number of particles n_i that are used to represent the power of i-th element is proportional to some probabilities $q_{i,\mathrm{pow}}$:

$$E[n_i] = N q_{i,\mathrm{pow}}, \tag{2}$$

where N is the total number of particles used in one iteration of the method.

The probability $q_{i,\mathrm{pow}}$ tells us how likely it is that a particle originates from the i-th element. This probability depends on the ratio of the element power P_i and the total power in the scene, P_{tot}:

$$q_{i,\mathrm{pow}} = \frac{P_i}{P_{\mathrm{tot}}}, \tag{3}$$

where for M elements in the scene we have $P_{\mathrm{tot}} = \sum_{k=1}^{M} P_k$.

In order to correctly represent P_i with n_i particles carrying some elementary power w_i, condition

$$E[n_i w_i] = P_i \tag{4}$$

has to be fulfilled — the expected value of power being shot out from the element has to be equal to its current power. This implies that

$$w_i = \frac{P_i}{N q_{i,\mathrm{pow}}} = \frac{P_{\mathrm{tot}}}{N}, \tag{5}$$

which means the elementary power is constant for all elements.

In order to avoid form-factor computation for interacting elements, particles are leaving the element surface kin directions determined by directional PDF corresponding to the $\cos\theta$ function. This results in a particle distribution which follows the form-factor value. The particle hits are then recorded on the receiver element and element power $P_{i,\mathrm{rcv}}$ due to particles received in this particular iteration is determined using Equation (1).

The method is usually used in so-called *self-correcting* version in which the power of the i-th element after the k-th iteration is given by

$$P_i^{(k)} = (1-\tau_k)\cdot P_i^{(k-1)} + \tau_k \cdot P_{i,\mathrm{rcv}}, \tag{6}$$

where τ_k is given by harmonic series, $\tau_k = 1/k$. It is usual to perform several "warming-up" iterations with $\tau_k = 1$ prior to stabilising the solution using the harmonic series.

2.2 Importance-Driven Stochastic Ray Radiosity

The importance-driven extension of stochastic ray radiosity uses the quantity dual to the power of an element, importance, to concentrate the computational effort to parts of the scene that are in the given moment visible. Importance of an element can be computed using an equation that is adjoint to Equation (1),

$$I_i = V_i + \sum_{j=1}^{M} F_{ji}\rho_j I_j. \tag{7}$$

I_i is the power-like importance of the i-th element, V_i stands for its initial, directly received importance. The value of V_i can be determined using usual approaches for computation point-to-patch form-factor.

An importance-driven iteration consists generally of two steps: First, the iteration that propagates importance is performed. The new importance values are then used to influence the propagation of power in the second step.

The power-like importance is propagated in the same way as power has been propagated in the stochastic ray radiosity method. As importance is an incident quantity, it has to be multiplied by the surface reflectance prior to shooting, which results in sampling probabilities of

$$q_{i,\mathrm{imp}} = \frac{\rho_i I_i}{\sum_{k=1}^{M}\rho_k I_k}. \tag{8}$$

The probabilities q_i,pow for power propagation depend now not only on the power of element i, but they depend on its view importance as well,

$$q_{i,\mathrm{pow}} = \frac{P_i I_i / A_i}{\sum_{k=1}^{M} P_k I_k / A_k}. \tag{9}$$

As I_i is i the power-like form of importance, it has to be converted to radiosity-like importance by dividing by the element area A_i.

It can be shown that using the total importance value for determining the element sampling probability is not an optimum choice [Phili99]. In many situations we can do better by using just the received importance,

$$q_{i,\text{pow}} = \frac{P_i(I_i - V_i)/A_i}{\sum_{k=1}^{M} P_k(I_k - V_k)/A_k}. \tag{10}$$

Until now we have been describing importance-driven radiosity scheme that uses separate steps for importance transport (stochastic ray radiosity with element sampling probabilities corresponding to Equation (8)) and for power transport (sampling probabilities after Equation (10)). However, it is possible to combine both steps into a single one in which power and importance are propagated simultaneously. Element sampling probability is then a combination of both above mentioned equations:

$$q_{i,\text{comb}} = (1-\alpha)\frac{P_i(I_i - V_i)/A_i}{\sum_{k=1}^{M} P_k(I_k - V_k)/A_k} + \alpha\frac{\rho_i I_i}{\sum_{k=1}^{M} \rho_k I_k}. \tag{11}$$

Value of $\alpha, 0 \leq \alpha \leq 1$ determines how much computational effort in the combined power-importance iteration is spent on importance propagation. The value of $\alpha = 0.1$ is often used.

The use of importance influences the size of the elementary power quanta that are being shot from a particular element. The equation (4) has to hold in this case as well. Recollecting that

$$w_{i,\text{pow}} = \frac{P_i}{N q_{i,\text{pow}}}$$

we can see that higher importance would cause more particles carrying lower energy quanta to be shot.

3 HIERARCHICAL REFINEMENT

Multi-resolution energy transport has been introduced by Hanrahan *et al.* [Hanra91]. The hierarchical radiosity algorithm makes it possible to effectively compute all the radiosity transport in the scene at a given accuracy level — some elements may interact at higher levels of the hierarchy, some of them have to be subdivided in order to fulfill the accuracy conditions. The authors limit an overall error of the computation by allowing the energy to be transported only over such links from shooter j to receiver i, that fulfills the condition

$$B_j\, F_{ji,\text{est}} < B_\varepsilon. \tag{12}$$

Here, $F_{ji,\text{est}}$ is a cheap form-factor estimate that gives upper bound of the actual form-factor, and B_ε is the user-supplied error threshold that determines the accuracy of the solution.

The idea of hierarchical refinement was extended by Smits *et al.* [Smits92] by incorporating the view importance into the refinement oracle in order to decrease the number of subdivisions in invisible areas. In this approach, a link is considered acceptable only if

$$B_j \rho_i \Psi_i\, F_{ji,\text{est}} < B_\varepsilon \psi_{\max}, \tag{13}$$

where Ψ_i is the radiosity-like form of the view importance, and $\psi_{\max}$ is a correction factor that corresponds to the maximum visible radiosity-like importance. This way the hierarchical subdivision in visible areas is comparable to that of non-importance-driven method while in invisible areas the level of subdivision is reduced.

Other possibility of using the view importance for hierarchical radiosity solutions has been shown by Bekaert and Willems [Bekae94]. They use importance to order the shooters in the course of shooting iterations so that light sources having the most influence at the region of interest are processed first. This approach was combined with that of Smiths *et al.* in [Bekae95].

3.1 Hierarchical Refinement in Monte-Carlo Radiosity

Monte-Carlo radiosity approaches have lacked a suitable hierarchical refinement approach for some period of time. Heckbert [Heckb90] and Tobler *et al.* [Toble97] proposed element subdivision schemes for continuous random-walk algorithms that use adaptive photon maps. The deficiency of these methods was the need to discard the current power stored at an element as the element was refined. The authors of the latter scheme reported 25% of recorded photon hits is typically discarded later due to element refinement.

The well-distributed ray set algorithm [Neuma97] was extended to work with the hierarchical element subdivision by Bekaert *et al.* [Bekae98]. The method is again a variant of stochastic Jacobi approach, which makes it possible to evaluate the power-equivalent of the Equation (12) oracle for every particle hitting a receiving surface with decent computational cost. In this case, an interaction has to be refined if

$$P_j \frac{A_i}{A_j}\, F_{ji,\text{est}} > P_\varepsilon. \tag{14}$$

4 IMPORTANCE-DRIVEN HIERARCHICAL STOCHASTIC RAY RADIOSITY

When looking for fast Monte-Carlo radiosity method that could be extended to hierarchical, importance-driven version, the well-distributed ray set method [Neuma97] would seem to be an excellent candidate. However, this method relies on the fact that the elementary power quantum remains approximately constant for all the patches in the scene. If the elementary quanta differ too much, the method will not converge. Our approach is therefore based on the stochastic ray radiosity [Neuma95] and its importance-driven extension [Neuma96].

Figure 1: Reference solution.

4.1 New Method

The new method is a combination of importance-driven stochastic ray radiosity as presented in Subsection 2.2 and hierarchical refinement methods presented in Subsection 3.

As we work with power and power-like importance values, the subdivision oracle (14) has to be modified to take power-like importance into account, yielding refinement criterion

$$P_j \frac{1}{A_j} \rho_i I_i \, F_{ji,\mathrm{est}} > P_\varepsilon I_{\max}, \qquad (15)$$

where $I_{\max}$ is the correction factor that has the same function as $\psi_{\max}$ in Equation (13).

As the scene is now composed of element hierarchy, importance values can be stored either in top level patches or in leaf elements of the hierarchy. We do not allow for element subdivision due to importance transfer as importance serves only as an additional information for the radiosity algorithm and even if it is computed with higher discretisation error, it still serves its purpose well. Subdivision due to importance transport would be necessary only if we were interested in exact importance values on the mesh.

5 RESULTS

We have implemented the method outlined in the previous Sections and evaluated its convergence. As the method is importance based, we compared resulting image quality using perceptual image comparison after Mannos and Sakrison [Manno74] as described by Rushmeier *et al.* [Rushm95].

The reference solution shown in Figure 1 was obtained by running 100 iterations (640 steps) of hierarchical well-distributed ray set radiosity [Bekae98] using 433×10^6 particles and subdividing the initial 23182 patches into 308018 elements.

Figure 2: Importance driven solution after 100 iterations with 1.1×10^6 particles (48702 elements) and corresponding hierarchical solution without use of importance (324606 elements).

Figure 2 presents the result of importance-driven hierarchical method with combined power-importance propagation (upper image) and the classical hierarchical method (lower image) for 110×10^6 particles traced. The perceptual distance between solutions generated using importance and without importance is depicted in Figure 3.

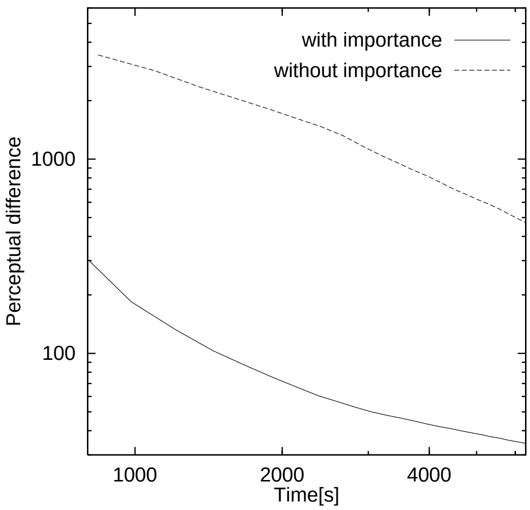

Figure 3: Image-based comparison of importance-driven and classical hierarchical stochastic ray radiosity.

Figure 4 shows the indirect importance solution and radiosity in the maze scene after 30 iterations. An interesting observation is that variance in unimportant regions is not as extremely high as with the non-hierarchical importance-driven method. The reason for this phenomenon is the importance-driven subdivision scheme: the accuracy of the solution in non-important regions is still lower, but this time the inaccuracy shows up more as the discretisation error than as variance.

Finally we were interested in influence of different minor improvements of the importance-driven method on its convergence. The questions that we wanted to answer were:

- How large is the influence of hierarchical importance storage on the solution quality?
- Does the combined power-importance iteration scheme outperform the separated power and importance iterations?
- As the element refinement during the warm-up phase of the algorithm (iterations with $\tau_k = 1$) results in quite a noisy image at the end of the warm-up, would it help if we inhibit the subdivision process until the warm-up ends?

The results of the tests are shown in Figures 5 and 6. We can see that there is just minimal difference in convergence speed for separate and combined importance propagation scheme, although the combined

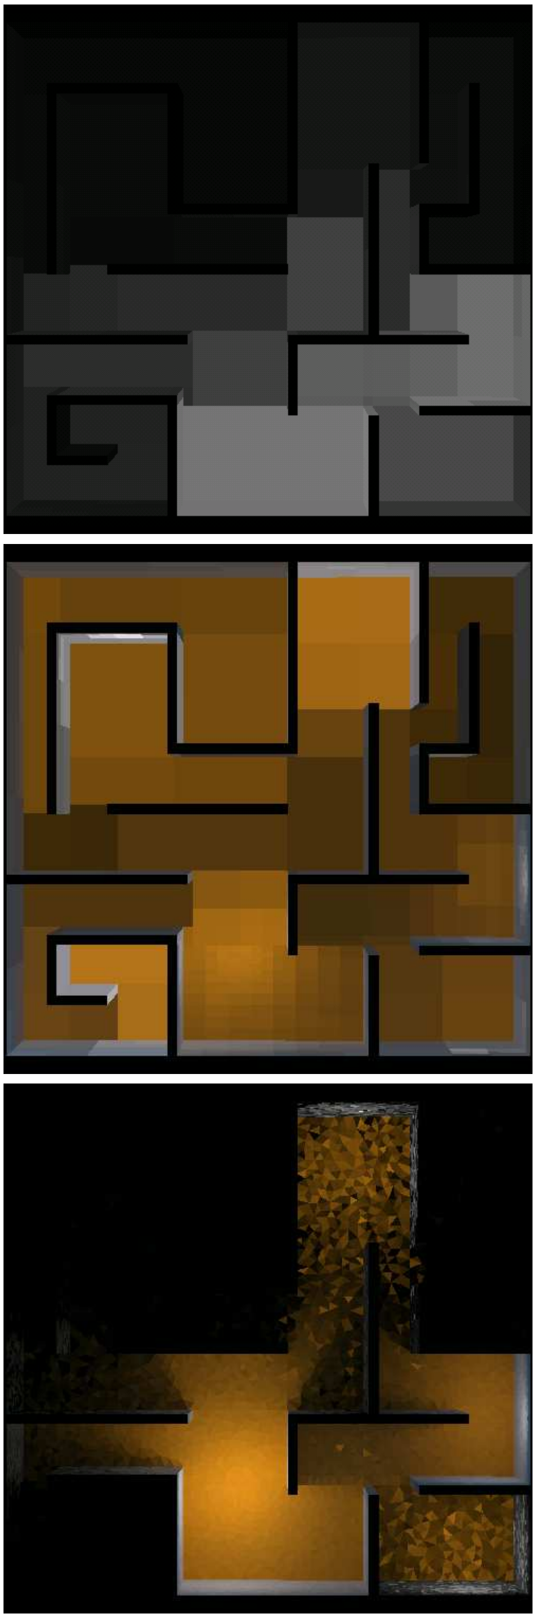

Figure 4: Importance solution, corresponding importance-driven hierarchical radiosity solution and importance driven solution on fixed size elements for the maze scene without furniture.

scheme scores a bit better. An interesting observation is that the hierarchical importance storage combined with separate propagation scheme is even better than the combined propagation scheme. All other improvements exhibit roughly the same convergence rate.

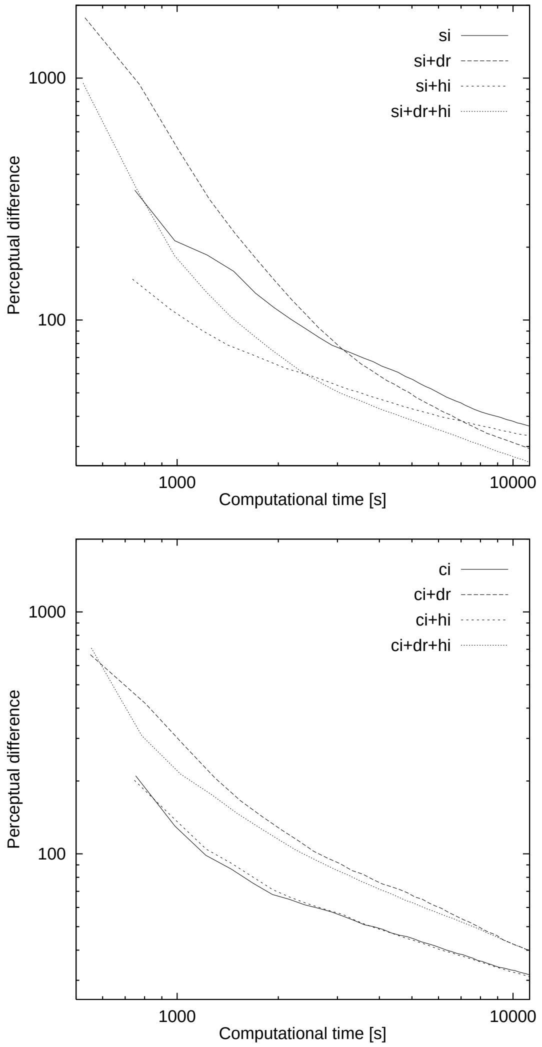

Figure 5: Influence of different improvements in the algorithm on its overall convergence. Method description: idr – separate importance and power propagation, cidr – combined power-importance iterations, dr – no refinement during the warm-up phase, hi – hierarchical importance.

6 CONCLUSIONS AND FUTURE WORK

We presented the hierarchical importance-driven stochastic ray radiosity algorithm which consists of straightforward combination of the importance-driven stochastic ray radiosity and of the hierarchical refinement method for Monte-Carlo radiosity. We did not use the faster converging well-distributed ray set algorithm for this purpose as it turns out that the

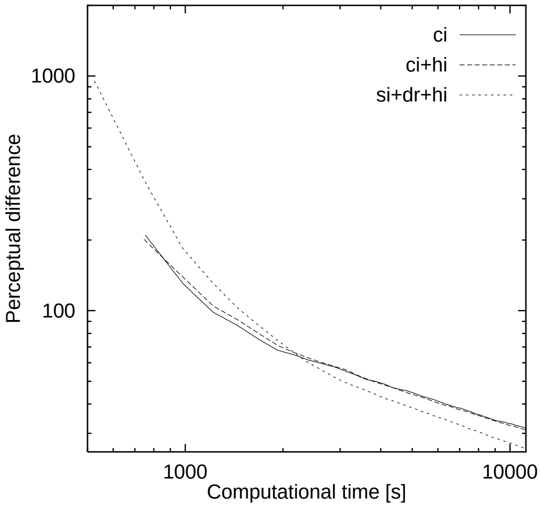

Figure 6: Best methods selected from Figure 5. Methods descriptions are the same as above.

"well-distributed ray set" would have to be built from scratch if the importance of elements changes. We are currently investigating possible workarounds.

As one can suppose, the method we presented indeed converges faster than its classical counterpart. In our test scene, the importance driven solution needed only 15% of the memory that was occupied by elements produced by the non-importance-driven method. From the further experiments we could see that the fastest convergence can be achieved by using either the separate importance propagation scheme combined with deferred element refinement and hierarchical importance storage, or using the combined power-importance propagation scheme. As we could see in our measurements, it probably does not pay off to store the importance in the hierarchy: the convergence rate of the method using hierarchical importance was in the best case approximately equal to the convergence rate of the method that stores importance values at top-level elements. Our tests did not show that switching the hierarchical refinement off during the warming-up phase would have any significant influence at the convergence of the method, although some minor improvements have been observed.

The method is still far from optimal. Possible improvements, besides above mentioned extending to well-distributed ray set, include use of linear basis functions and a robust perceptually driven subdivision oracle.

Acknowledgements

The above presented ideas have been implemented within the RenderPark software from K.U.Leuven [Ren]. This work has been supported by Austrian “Fonds zur Förderung der wissenschaftlichen Forschung” under project number P11545-MAT and by a joint Czech-Austrian scientific collaboration funding under project number 1999/17. The second author acknowledges support of Flemish fund for scientific research under grant number FWO-Vl #G.0105.96.

References

[Bekae94] Philippe Bekaert and Yves D. Willems. A Progressive Importance-Driven Rendering Algorithm. In *Proceedings of the Tenth Spring School on Computer Graphics '94*, pages 58–67, Comenius University, Bratislava, Slovakia, June 1994.

[Bekae95] Philippe Bekaert and Yves D. Willems. Importance-Driven Progressive Refinement Radiosity. In Hanrahan and Purgathofer [Hanra95], pages 316–325.

[Bekae98] Philippe Bekaert, László Neumann, Attila Neumann, Mateu Sbert, and Yves D. Willems. Hierarchical Monte Carlo radiosity. In *Rendering Techniques '98 (Proceedings of the Ninth Eurographics Workshop on Rendering)*, pages 259–268. Eurographics, Springer-Verlag/Wien, 1998.

[Hanra91] Pat Hanrahan, David Salzman, and Larry Aupperle. A Rapid Hierarchical Radiosity Algorithm. In *Computer Graphics Proceedings*, Annual Conference Series, pages 197–206. ACM SIGGRAPH, 1991.

[Hanra95] Patrick M. Hanrahan and Werner Purgathofer, editors. *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*. Eurographics, Springer-Verlag/Wien, 1995.

[Heckb90] Paul Heckbert. Adaptive Radiosity Textures for Bidirectional Ray Tracing. In *Computer Graphics Proceedings*, Annual Conference Series, pages 145–154. ACM SIGGRAPH, 1990.

[Kelle97] Alexander Keller. Instant Radiosity. In *Computer Graphics Proceedings*, Annual Conference Series, pages 49–56. ACM SIGGRAPH, 1997.

[Manno74] James L. Mannos and David J. Sakrison. The Effect of a Visual Fidelity Criterion on the Encoding of Images. *IEEE Transactions of Information Theory*, 20(4):525–536, July 1974.

[Neuma95] László Neumann, Werner Purgathofer, Robert F. Tobler, Attila Neumann, Pavol Eliáš, Martin Feda, and Xavier Pueyo. The Stochastic Ray Method for Radiosity. In Hanrahan and Purgathofer [Hanra95], pages 206–218.

[Neuma96] Attila Neumann, László Neumann, Philippe Bekaert, Yves Willems, and Werner Purgathofer. Importance-Driven Stochastic Ray Radiosity. In *Rendering Techniques '96 (Proceedings of the Seventh Eurographics Workshop on Rendering)*, pages 111–122. Eurographics, Springer-Verlag/Wien, 1996.

[Neuma97] László Neumann, Attila Neumann, and Philippe Bekaert. Radiosity with Well Distributed Ray Sets. In *Eurographics '97 Conference Proceedings*, volume 16 of *Computer Graphics Forum*, pages C–261–C–270. Eurographics, September 1997.

[Patta92] Sumanta N. Pattanaik and S. P. Mudur. Computation of the Global Illumination by Monte Carlo Simulation of the Particle Model of Light. In *Third Eurographics Workshop on Rendering*, pages 71–83, Bristol, May 1992. Eurographics.

[Patta93a] Sumanta N. Pattanaik and S. P. Mudur. Efficient Potential Equation Solutions for Global Illumination Computation. *Computers and Graphics*, 17(4):387–396, July–August 1993.

[Patta93b] Sumanta N. Pattanaik and S. P. Mudur. The Potential Equation and Importance in Illumination Computations. *Computer Graphics Forum*, 12(2):131–136, June 1993.

[Patta95] Sumanta N. Pattanaik and S. P. Mudur. Adjoint Equations and Random Walks for Illumination Computation. *ACM Transactions on Graphics*, 14(1):77–102, January 1995.

[Phili99] Philippe Bekaert. *Hierarchical and Stochastic Algorithms for Radiosity*. PhD thesis, Katholieke Universiteit Leuven, @dec 1999.

[Ren] RenderPark system. Available from `http://www.cs.kuleuven.ac.be/cwis/research/graphics/RENDERPARK/`.

[Rushm95] Holly Rushmeier, Greg Ward, Christine D. Piatko, Phil Sanders, and Bert Rust. Comparing Real and Synthetic Images: Some Ideas about metrics. In Hanrahan and Purgathofer [Hanra95], pages 82–91.

[Sbert97] Mateou Sbert. *The Use of Global Random Directions to Compute Radiosity – Global Monte-Carlo Techniques*. PhD thesis, Universitat de Girona, 1997.

[Smits92] Brian E. Smits, James R. Arvo, and David H. Salesin. An Importance-Driven Radiosity Algorithm. In *Computer Graphics Proceedings*, Annual Conference Series, pages 273–282. ACM SIGGRAPH, 1992.

[Toble97] Robert F. Tobler, Alexander Wilkie, Martin Feda, and Werner Purgathofer. A Hierarchical Subdivision Algorithm for Stochastic Radiosity Methods. In *Rendering Techniques '97 (Proceedings of the Eigtht Eurographics Workshop on Rendering)*, pages 193–204. Eurographics, Springer-Verlag/Wien, 1997.