Gathering multi-path: a new Monte Carlo algorithm for radiosity

Mateu Sbert, Roel Martínez and Xavier Pueyo

mateu|roel|xavier@ima.udg.es

Institut d'Informàtica i Aplicacions

Universitat de Girona

ABSTRACT

In this paper we present a new Monte Carlo technique for radiosity. This technique is shown to be the dual of the previous multi-path technique, in the same way as a gathering random walk technique is dual to a shooting one. We also show how the same global density of lines can be used to trace paths and thus to obtain estimators of radiosity simultaneously in both dual techniques. Although the gathering multi-path technique does not appear to improve the shooting technique, it can be used to compute efficiently the importance matrix.

Key Words: Rendering, Integral Geometry, Radiosity, Monte Carlo

1 INTRODUCTION

It is well known that the Radiosity system has a dual representation as a Power system, based on the reciprocity relations $A_i F_{ij} = A_j F_{ji}$. And to solve an equation system through random walk, we are provided with two methods: a direct (or gathering) method and an adjoint (or shooting) method [HH75, Rub81, DLW94]. The Radiosity system is usually solved by the direct method, and the Power system by the adjoint one. The solution for the Power system (Particle Tracing [PM92, Shi90, FP93]) is given by tracing paths from the sources (usually selected upon their power), and following those paths according to the transition probabilities given by the Form Factors. On the other hand, to implement the gathering solution of the Radiosity system means selecting each patch in turn, and tracing paths from this patch. Whenever the path hits a source, the radiosity of the patch is updated. This method is in principle highly inefficient, due both to the fact that only the origin patch is updated and to the low probability of hitting a source. But if we come back to the integral equation (Rendering Equation [Kaj86]) whose discretisation (and specialization for diffuse surfaces) originates the Radiosity system, we see that we find the gathering method under the name of distributed ray-tracing [CPC84] or also path-tracing [Kaj86]. The paths are only traced from the visible part of the scene (this would mean in our case to trace them only from visible patches) and the low probability of hitting a source is dealt with by either splitting the ray into one ray directed at the source and another stochastic one, or instead the direct illumination is spread somehow (for instance with a first shooting pass) over the scene and so we change radically the probabilities of finding a source. A (theoretical) solution has also been given to the inefficiency caused by updating only the origin patch. Rubinstein [Rub81] considers a *covering path*, a path that covers (or hits) all states (patches in our case) and is used to track paths of fixed length from all of them. Whenever the covering path hits a state, a new path is considered to begin from this state, and tracked to a predetermined fixed length. The covering path is terminated after it has visited all the states plus the fixed length (thus a single physical path conveys many different logical paths). As a mixture of non-discrete shooting and gathering we find bidirectional ray-tracing [VG95, LW93]. Shooting random walk methods have also been investigated in the limiting, non discrete case [DLW94, DW95]. Variances and complexities of shooting and gathering random walk are discussed in [Sbe97a, Sbe97d, Sbe97b].

It must be remarked that for the dual methods just discussed, the traced paths are wholly different, and so considering the density of lines that support the different segments of the paths, both densities of lines are different. In this paper we

will see how it is possible to implement both dual methods for radiosity through the same density of lines. This will be accomplished by the use of global Monte Carlo techniques [SPP95].
The organisation of this paper is as follows: In section 2 we will revise the adjoint method solution for the power system with a global density of lines presented in [SPP95, SPNP96]. In section 3 some theoretical results about the Form Factor matrix are presented, which will justify the new algorithm. In section 4 a new algorithm is presented which efficiently solves the direct solution of the Radiosity system using a global density of lines. In section 5 results are presented which show the behaviour of this new algorithm. Conclusions are presented in section 6.

2 PREVIOUS WORK

In this section we show how we can simulate random paths to solve the Power equation with a global density of lines, that is, with lines not tied to any particular patch.

2.1 Some facts from Integral Geometry

Integral Geometry shows that considering a uniform global density of lines, that is, isotropic and homogeneous, and taking two surfaces embedded in that density, the conditional probability that a line having crossed the first surface will also cross the second one is equal to the Form Factor from the first surface to the second [Sbe93]. Also, if we have a surface within a convex body, the conditional probability that a line having crossed the body will also cross the surface is twice the area of the surface divided by the area of the body. We can then simulate a *global density* or distribution of lines with two aims, each based on one of the two previous facts. The first fact means that, as our random walk paths must follow the probability distribution given by the Form Factors, we can use the global distribution of lines to make a path advance from patch to patch (we can also use the global density to build the Form Factor matrix, as in [Sbe93]). Moreover, as a line may cross many surfaces at once, and can be considered to transport energy in both directions, it can advance on average n_{int} paths, where n_{int} is the average number of intersections per line. The second fact means that we can predict how many lines (on average) will cross each patch, so with each line crossing a source patch we can begin a new path transporting the correct quantity of energy (on the average). The global density of lines will be simulated through the surrounding of the scene by a sphere, and taking pairs of random points on its surface (there are also other possibilities, see [RPS51, Sbe97c]). The methods that solve radiosity through the use of a global density of lines have been called global Monte Carlo methods [SPP95], opposite to local Monte Carlo methods, that draw lines local to a patch or source. Global densities of lines have also been used by [BF89] and [Neu95]. In [Pel95] and [SF97] there is a related approach using area projection. In the next section we will review the multi-path algorithm.

2.2 Shooting multi-path method

In the above section 2.1 we have seen that it is possible to simulate a random walk using a global density of lines. Figure 1(top) represents this fact (but take into account that this is a somewhat idealized drawing: the actual lines intersect the surfaces at random points, thus the exit point is not the arriving point). This is the origin of the shooting multi-path method, exposed in [SPNP96, SPP95] (where it was just named multi-path). Next we shall look at this method.
We shall cast a predetermined number of randomly oriented global lines. Each line will produce an intersection list, and we follow this list taking into account successive pairs of patches (figure 1(bottom)). Each patch (if not emitter) has with it two quantities. One records the energy accumulated, the other one is the unshot energy. For every pair of patches along the intersection list, the first patch of the pair will transmit its unsent energy to the second patch of the pair. So the unshot energy of the first patch is reset to zero, and the two quantities at the second patch, the accumulated and the unsent energy, are incremented. In the case of a source we keep also a third quantity, the emitted energy per line exiting the source. We compute that energy in the following way: Given the number of lines we are going to cast, we compute for any light source beforehand the forecasted number of lines passing through it. The division of the total source energy by this number gives the predicted energy of one line. Then, if the first patch of a pair is a source patch, the energy transported to the second patch of the pair will also include this predicted energy portion. Also we can use as a variance reduction technique the fact that a segment of line joining a pair of patches can be used to transmit energy in both directions. For the multi-path method, as well as for all global Monte Carlo methods to work efficiently, we need first to "smooth" the scene [SPNP96], for example distributing the direct illumination sending rays from the source. After doing this the new sources are the patches that have received direct illumination, and the old sources are no more emitters (although we could leave some energy undistributed so as not to waste the lines hitting the original sources). In this way we avoid the problem of many lines not transporting energy in the initial stage of the simulation, as in

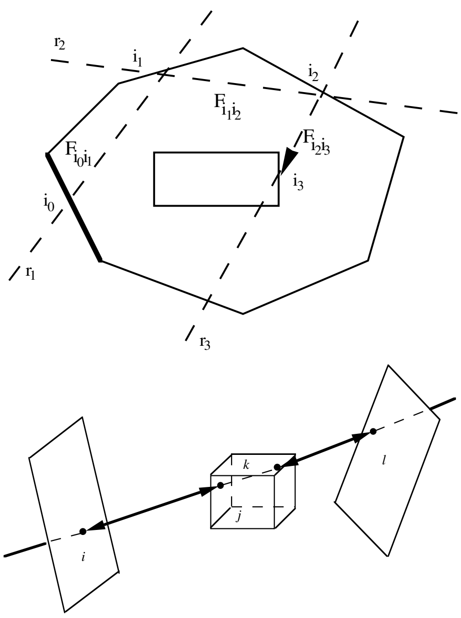

Figure 1: (top) Three global lines r_1, r_2, r_3 simulate a path from i_0 to i_3. This path is built incrementally. Between r_1 and r_2, for instance, many other lines may have been cast and made advance other paths. (bottom) The same line will transport energy from patch i to patch j, and from patch k to patch l (and viceversa). Patches i and l are on the walls of the interior scene.

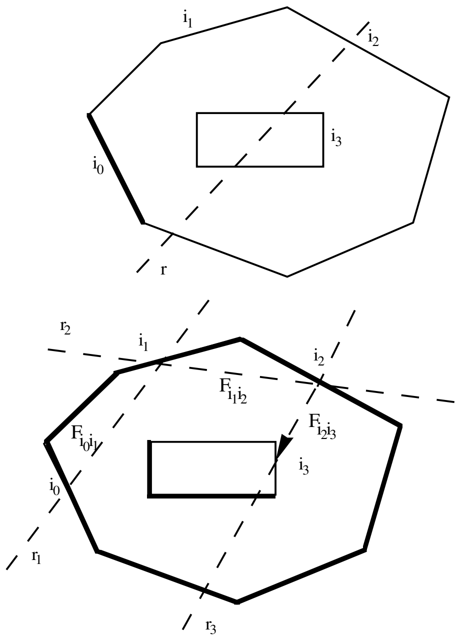

Figure 2: (top) Line r at the beginning of the simulation transports no energy. (bottom) After distributing direct illumination (note that some energy in the source i_0 is left without distributing). A single physical path $i_0 i_1 i_2 i_3$ covers logical paths $i_0 i_1 i_2 i_3$, $i_1 i_2 i_3$, $i_2 i_3$.

figure 2(top), and at the same time we use a single physical path to convey many different logical paths (as in Rubinstein's *covering paths* [Rub81]), as shown in figure 2(bottom). From this characteristics and from the fact that at any moment there are many paths alive in the simulation, we named this method multi-path.

3 SOME RESULTS ABOUT THE FORM FACTOR MATRIX

Here we present some theoretical results, obtained in [Sbe97c], which will be used in section 4 to justify the new technique presented in this paper.
Let us first remember that F is a stochastic or probability matrix, because $\forall(i,j) F_{ij} \geq 0$ and $\sum_j F_{ij} = 1$ for all i. A probability matrix is said to be aperiodic if it has no periodic state [Col74] (a state is simply a patch in the radiosity setting). A periodic state is one such that it can be visited back by a path starting from it only at multiples of the period. A probability matrix is said to be irreducible if there is always a path between any two states, that is, the probability of reaching any state starting at any other is not null. For the Form Factor matrix we can assume safely its aperiodicity, to assume its irreducibility every two patches must be reachable through a path from each other, and this is not always true. Think for instance of two separated rooms with the interconnecting door closed, or wall patches behind a drawer, or floor patches covered by furniture. But in the case of non-connecting rooms we can simply consider two different Form Factor matrices, one for each room. In the case of covered or hidden patches, we can simply suppress them and obtain a smaller and now irreducible Form Factor matrix. We have then the following propositions:

Proposition 1: *If the Form Factor matrix F is irreducible and aperiodic, then when n grows to infinity we have $(F^n)_{ij} \rightarrow \frac{A_j}{A_T}$, where A_j is the area of patch j and $A_T = \sum_{i=1}^{n_p} A_i$, for all the n_p patches in the scene.*

Proof: *Apply to F the Kolmogorov theorem (see for instance [Col74]) that states that $(F^n)_{ij} \rightarrow \pi_j$ where π is the left eigenvector of F corresponding to the eigenvalue 1 (π is called the equilibrium distribution). But from the relations $A_i F_{ij} = A_j F_{ji}$ it is trivial to verify that $\pi = \left(\frac{A_1}{A_T}, \frac{A_2}{A_T}, \cdots, \frac{A_{n_p}}{A_T}\right)$*

Proposition 2: *When the length of a random path with transition matrix F grows to infinity, the number of hits on any patch i gets proportional to $\frac{A_i}{A_T}$, and this independently of where the path started its trajectory.*

Proof: *It is simply a consequence of being* $\pi = \left(\frac{A_1}{A_T}, \frac{A_2}{A_T}, \cdots, \frac{A_{n_p}}{A_T}\right)$ *, and of the result [HH75] that the number of hits on patch* i *is proportional to* π_i*.*

4 THE NEW TECHNIQUE: GATHERING MULTI-PATH

4.1 Description and implementation

We have seen in section 2 an adjoint solution of the Power system using a global density of lines. In this section we will study a new technique, which implements the direct method solution for the Radiosity system, using the same global density of lines. This is contrary to a local Monte Carlo method solution, where the densities of lines originated in both methods, direct and adjoint, are different (taking the lines that contain the different segments that build up the path you have a density of lines). However, if a local path is allowed to follow its trajectory through the scene indefinitely, the density of local lines (from the patches) obtained approaches the one obtained with a global density of lines. This is because (proposition 1) when n grows to infinity and for any (i,j) $(F^n)_{ij} \rightarrow \frac{A_j}{A_T}$;thus the value of the limit is the fraction of total area of the scene that corresponds to patch j, and the hits on a patch are proportional to the area of a patch, as in a global density. Then a global density of lines appears as the limiting case for local densities of lines, when the path with segments laying on these lines follows an infinite trajectory (proposition 2). We call this new technique *gathering multi-path technique*, because it is the dual of shooting multi-path. That is, the paths are traced from the patches to the source (figure 3(top)).

As with the shooting multi-path technique, we must first smooth the scene, for instance distributing direct illumination, so that the technique becomes efficient, that is, the probability that a path hits a source is high (figure 3(bottom)). On the other hand, we also want to use one physical path to track many logical paths, as in the multi-path technique. To do this we must track every patch a path hits. Whenever a path hits a patch, this patch is added to the list of visited patches for this path. Whenever a path hits a source, the emissivity of the source travels back through the list of patches in the path updating each radiosity in turn. With every patch updated the emissivity transferred to the next patch in the list decreases by the amount of the reflectivity of the previously updated patch. So, for each path only the last ten or fifteen visited patches need to be stored (a threshold can be pre-established upon which to prune the list of visited patches). As in the multi-path method, we have many paths "frozen" at the same moment. When a global line is cast, some of

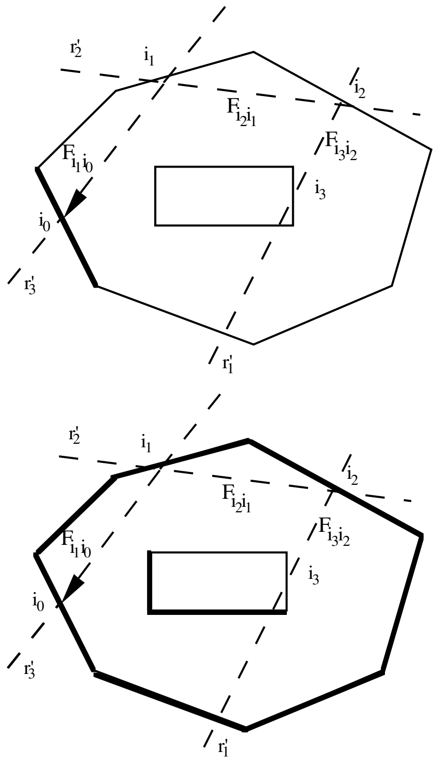

Figure 3: (top) Three global lines r'_1, r'_2, r'_3 simulate a path from i_3 to i_0. This path is built incrementally. Between r'_1 and r'_2, for instance, many other lines may have been cast and made advance other paths. (bottom) After distributing direct illumination (note that some energy in the source i_0 is left without distributing). A single physical path $i_3 i_2 i_1 i_0$ covers logical paths $i_3 i_2 i_1 i_0$, $i_2 i_1 i_0$, $i_1 i_0$.

these paths will advance. We will store each path with the last visited patch. So the total number of paths in any moment of the simulation is equal to the number of patches (although of course at the beginning most of them will be empty). When a line crosses two patches i and j, this line is considered to advance the path stored with i from i to j and the path stored with j from j to i, so the patches are added to the paths, and after this the paths are interchanged (so that they get stored in the respective last visited patches). Now we process the paths to transmit emissivities in case patches i or j are emitters. This process is represented in figures 4,5,6 and 7. In figure 4 we see the two paths stored with patches i and j (we implement it with an array of pointers with a range covering all patches). In figure 5 a global line has

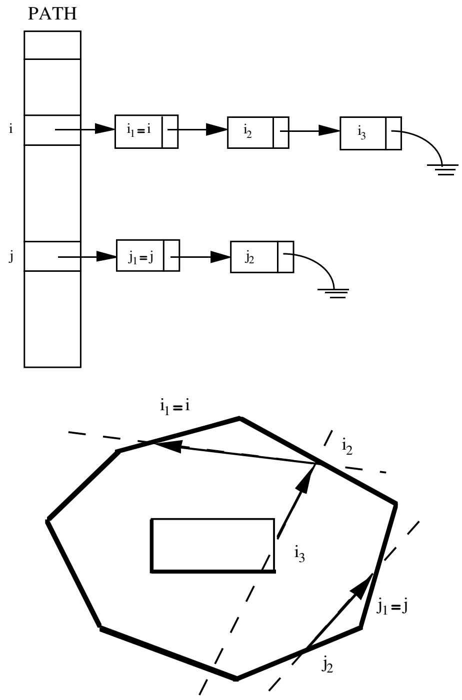

Figure 4: (top) The paths stored with patches i and j. (bottom) Another view of the same paths.

joined patches i and j, so we must advance and interchange the respective paths, which are again shown in figure 6. Figure 7 shows the process of the path pending from j, supposing j is an emitter patch. The variables used are defined in figure 8. The algorithm for the gathering multi-path technique is given in figure 9 and 10.

4.2 Comparison with shooting solutions

The gathering multi-path estimator can be shown to have the same variance as the shooting multi-path [Sbe97c]. As both techniques can be simulated with the same density of lines, the optimal combination of them is with weight 1/2 each.

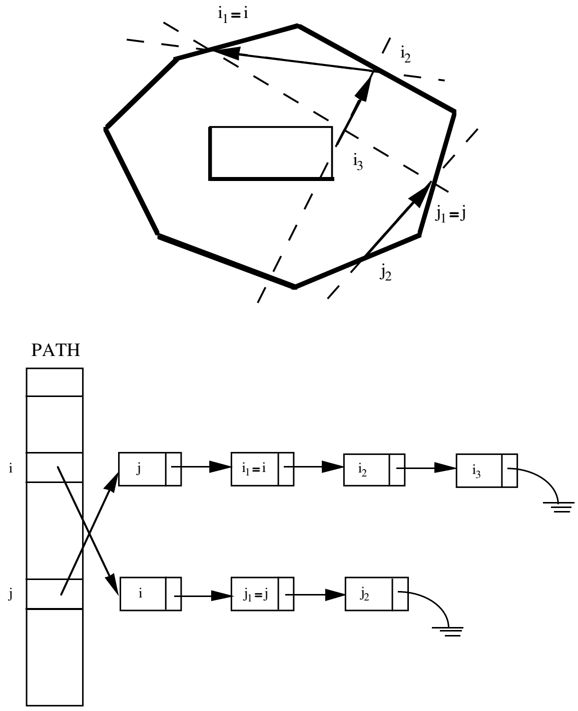

Figure 5: (top) A global line connects patches i and j.(bottom) Patch j is added to the path pending from patch i, patch i is added to the path pending from patch j and the paths are interchanged.

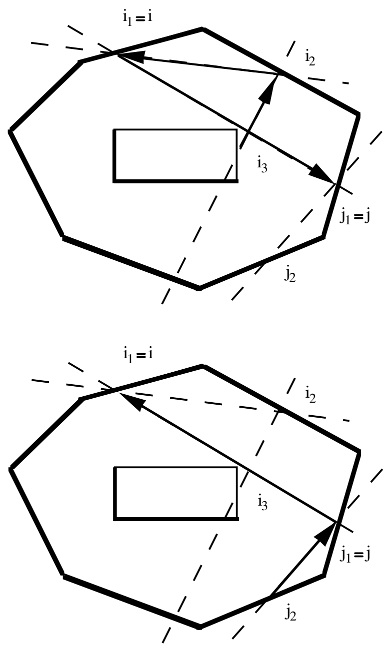

Figure 6: (top) The enlarged path pending from patch j. (bottom) The enlarged path pending from patch i.

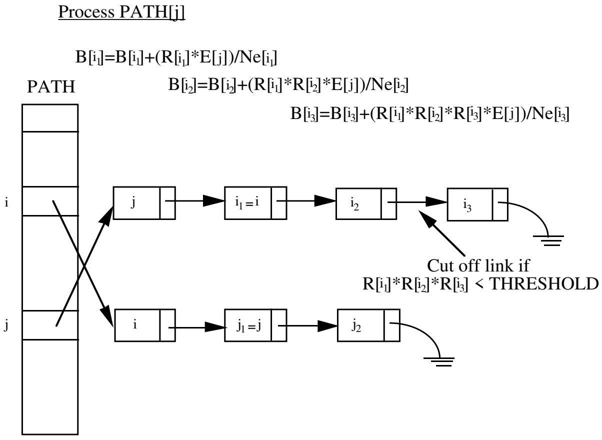

Figure 7: Processing the path from j. The same would be done with the path from i.

But unfortunately, we have found in our tests that there is no gain in efficiency (see section 5). This leads to the strong suspicion that both estimators are *assymptotically* a single one seen from two different points of view. As the gathering multi-path needs more storage (to keep the paths, $l_{ave}\, 2\, sizeof(long\_int)$ for each patch) than the shooting one, the question arises as what practical utility brings the new technique.

```
Ne[1..number_of_patches] = Vector holding
the number of expected lines through the patches
/*precomputed at the beginning*/
B[1..number_of_patches] = Vector holding
the radiosities of the patches
PATH[1..number_of_patches] = Vector holding
pointers to paths
```

Figure 8: Global arrays used in the implementation.

```
Initialize radiosities to the emissivities.
Initialize PATH[1..number_of_patches]
While all global lines not cast
   Cast a line
   Build sorted intersection list
   For each consecutive pair (i, j) of patches in the list
      Add patch i in front of PATH[j]
      Add patch j in front of PATH[i]
      Swap (PATH[i], PATH[j])
      Process PATH[i]
      Process PATH[j]
   End for
End while
```

Figure 9: Gathering multi-path algorithm. The procedure to process a path is in figure 10

One answer is as follows: If we need to compute the Importance matrix (for instance to simulate a biased by importance random walk [PM93], [DLW94]), only the gathering multi-path is useful, as it keeps track of what source comes the energy from. And, if we want to obtain the same importance matrix with a local shooting Monte Carlo technique, such as the best finite path length shooting one, we obtain in [Sbe97c] a relative efficiency of n_{int} (this is because each patch must be a source). This comparison is based on the Mean Square Errors, and thus it is made on an average on all patches basis.

```
If first patch in CURRENT_PATH is a source
  AR ← 1.0 /* initialize accumulated reflectivity*/
  E ← Emissivity of the first patch
  current_patch ←next patch in CURRENT_PATH
  While not at the end of the CURRENT_PATH
     AR ← AR * reflectivity_of_current_patch
     B[current_patch]B[current_patch]+
                     (AR * E/Ne[current_patch])
     If AR < THRESHOLD
     prune CURRENT_PATH
     /*set to nil next patch of CURRENT_PATH*/
     End if
     current_patch ← next patch
     in CURRENT_PATH
  End while
End if
```

Figure 10: Procedure to process the *CURRENT_PATH*.

4.3 Comparative cost with a biased infinite gathering random walk solution

It is straightforward to compare the gathering multi-path with the infinite path length gathering random walk estimator (see [Sbe97d]), when each patch is selected with probability $p_i = \frac{A_i}{A_T}$ (suppose for instance we want as before compute the Importance matrix). This is because both techniques use the same estimator, only the way to obtain the paths is different. In fact the variances are the same (save a negligible term, see [Sbe97c]). In our new technique, the simulation is achieved trough the casting of N global lines. As each global line hitting a patch starts a new path from it, and a global line intersects an average of n_{int} patches, N global lines will originate a total number of Nn_{int} paths. So, with a cost of tracing N global lines we have obtained Nn_{int} paths with average lenth l_{ave}. The paths that start from a given patch j, N_j, are proportional to the area of the patch.
Now, in the local Monte Carlo gathering method strategy (with $p_i = \frac{A_i}{A_T}$) we have to cast rays also according to the patch area. To obtain the same variance we trace from each patch j the same number of paths N_j that with the global strategy. If we cut the infinite path with the same threshold with which we cut the global paths, we have to trace a total of Nn_{int} paths with average lenth l_{ave}. This makes a cost of $Nn_{int}l_{ave}$ local lines. Considering a similar cost for both a local and global line (supposing negligible the cost of sort-

ing the intersections list), as the variances are the same in both cases, we have the result that the relative efficiency of both techniques is $n_{int}l_{ave}$ (the actual efficiency could be lesser for the smallest patches, as the actual number of paths exiting them in the global method could deviate significantly from the expected value). As n_{int} depends on the complexity of the scene, and l_{ave} depends on the average reflectivity, the more complex the scene and the highest the reflectivity, the highest relative efficiency we will obtain. The comparison here is based directly upon variances and not upon an average as in previous section.

5 RESULTS

Here we present some results to compare shooting and gathering multi-path. For the first test scene, figure 11(bottom) with a unit source in the first face (patches 1-9), we made 100 runs of 10000 global lines for the combination of both estimators (cutting the paths of gathering with a threshold of 0.001). On the y-axis of figure 11(top) is represented the square error for a single line, obtained averaging for all the runs and afterthat multiplying by 10000. We see that there is no gain combining both estimators. Other results are shown in figure 12. We used the same amount of local (for the first shot) and global lines in both images (about 2 milions for both local and global and 10 minutes of processing time on a HP 735 workstation), but the first was obtained only with gathering multi-path, and the second with the combined estimator (cutting the paths of gathering with a threshold of 0.001). The MSE for both images was the same, and they look the same quality. Thus again there is no improvement combining the two estimators.

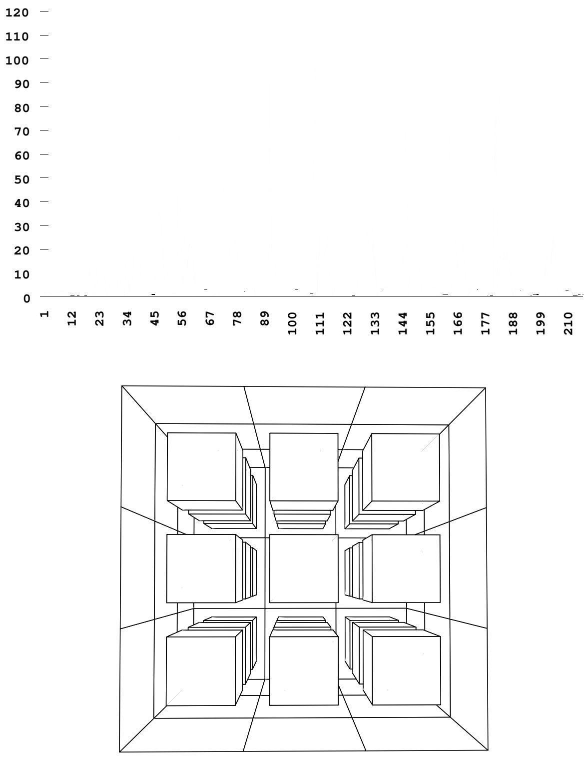

Figure 11: (top) Comparison of the experimentally obtained square errors (on y axis) for the shooting multi-path (plotted as a continous line) and for the combined estimator (plotted as a discontinous line), for the 216 patches (on x axis) scene of (bottom). A source with emittance 1 covers the first face.

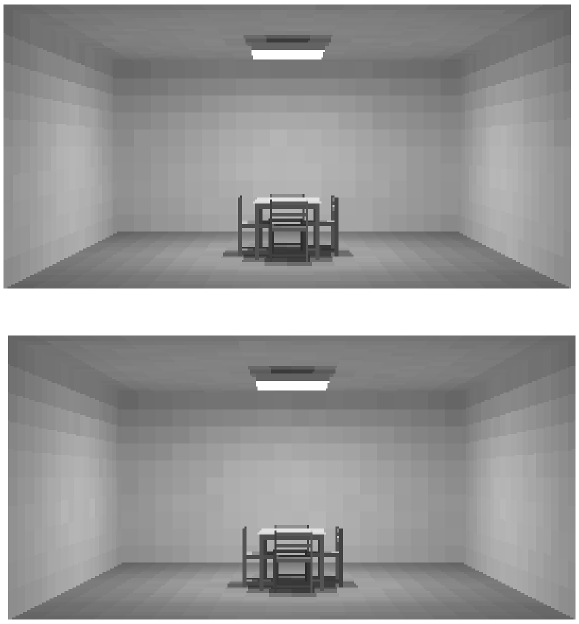

Figure 12: (top) Image obtained with the gathering multi-path method. (bottom) Image obtained combining the shooting and gathering multi-path, with the same number of lines as in the top figure. No Gouraud shading is done.

6 CONCLUSIONS AND FUTURE RESEARCH

A new global Monte Carlo technique has been presented that consists in the dual of the multi-path technique. Both techniques appear to be *assymptotically* a single estimator. Although the memory requirement of the new technique is higher, it has a practical utility in computing the Importance matrix. In this case the relative efficiency against the best shooting local random walk technique is of n_{int}, that is, the average number of intersections a random line has with the scene. And of $n_{int}l_{ave}$ when comparing to a biased infinite gathering local random walk estimator, where l_{ave} is the average length of the cut off path. Future research will be directed towards studying the applicability of quasi Monte Carlo improvements as in [SFNC97], and to integrate this technique as a first pass into importance driven techniques.

ACKNOWLEDGEMENTS

Many thanks to Werner Purgathofer and Gonzalo Besuievski for their useful comments. This project has been funded in part with grant number TIC 95-0614-C03-03 of CICYT of the Spanish Government, and with Austrian-Spanish Joint-Action HU96-037.

REFERENCES

[BF89] Chris Buckalew and Donald Fussell. Illumination networks: Fast realistic rendering with general reflectance functions. In *Computer Graphics (ACM SIGGRAPH '89 Proceedings)*, pages 89–98, July 1989.

[Col74] Robert Coleman. *Stochastic processes.* George Allen&Unwin Ltd, London, 1974.

[CPC84] Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. In *Computer Graphics (ACM SIGGRAPH '84 Proceedings)*, pages 137–145, July 1984.

[DLW94] P. Dutre, E. Lafortune, and Y. D. Willems. A mathematical framework for global illumination algorithms. In *Proceedings of the Winter School of Computer Graphics and CAD Systems '94*, pages 75–84, Plzen, Czech Republic, January 1994.

[DW95] Philip Dutre and Yves D. Willems. Potential-driven monte carlo particle tracing for diffuse environments with adaptive probability density functions. In *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 306–315, New York, 1995. Springer-Verlag.

[FP93] Martin Feda and Werner Purgathofer. Progressive ray refinement for monte carlo radiosity. In *Fourth Eurographics Workshop on Rendering*, pages 15–26, Paris, France, June 1993.

[HH75] J.M. Hammersley and D.C. Handscomb. *Monte Carlo Methods.* Methuen & Co Ltd., London, second edition, 1975.

[Kaj86] James T. Kajiya. The rendering equation. In *Computer Graphics (ACM SIGGRAPH '86 Proceedings)*, pages 143–150, August 1986.

[LW93] Eric P. Lafortune and Yves D. Willems. Bi-directional path tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)*, pages 145–153, Alvor, Portugal, December 1993.

[Neu95] Laszlo Neumann. Monte carlo radiosity. *Computing*, 55(1):23–42, 1995.

[Pel95] M. Pellegrini. Monte carlo approximation of form factors with error bounded a priori. In *Eleventh ACM Symposium on Computational Geometry*, Vancouver, Canada, June 1995.

[PM92] S. N. Pattanaik and S. P. Mudur. Computation of global illumination by monte carlo simulation of the particle model of light. In *Third Eurographics Workshop on Rendering*, pages 71–83, Bristol, UK, May 1992.

[PM93] S. Pattanaik and S. Mudur. Efficient potential equation solutions for global illumination computation. *Computers & Graphics*, 17(4):387–396, 1993.

[RPS51] J. Rey-Pastor and Luis A. Santaló Sors. *Geometría Integral.* Espasa Calpe, Madrid, 1951.

[Rub81] Reuven Y. Rubinstein. *Simulation and the Monte Carlo method.* John Wiley & Sons, New York, 1981.

[Sbe93] Mateu Sbert. An integral geometry based method for fast form factor computation. *Computer Graphics Forum (Eurographics '93)*, 12(3):C409–C420, September 1993.

[Sbe97a] Mateu Sbert. Error and complexity of random walk monte carlo radiosity. *IEEE Transactions on Visualization and Computer Graphics*, 3(1), March 1997.

[Sbe97b] Mateu Sbert. Optimal source selection in shooting random walk monte carlo radiosity. In *Proceedings of Eurographics'97*, Budapest, September 1997.

[Sbe97c] Mateu Sbert. *The use of global random directions to compute radiosity. Global Monte Carlo methods.* Ph.D. thesis, Universitat Politècnica de Catalunya, Barcelona, March 1997. Available in *http://ima.udg.es/˜ mateu*

[Sbe97d] Mateu Sbert. Variance of two infinite path length random walk radiosity estimators. In *Proceedings of WSCC'97*, Plzen, February 1997. To appear in *Computer & Graphics.*

[Shi90] Peter Shirley. A ray tracing method for illumination calculation in diffuse-specular scenes. In *Proceedings of Graphics Interface '90*, pages 205–212, Toronto, Ontario, May 1990. Canadian Information Processing Society.

[SPNP96] Mateu Sbert, Xavier Pueyo, Laszlo Neumann, and Werner Purgathofer. Global multipath monte carlo algorithms for radiosity. *The Visual Computer*, 12(2):47–61, 1996.

[SPP95] Mateu Sbert, Frederic Perez, and Xavier Pueyo. Global monte carlo: A progressive solution. In *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, pages 231–239, New York, 1995. Springer-Verlag.

[SF97] Laszlo Szirmay-Kalos and Tibor Foris. Radiosity algorithms running in sub-quadratic time. In *Proceedings of WSCC'97*, Plzen, February 1997.

[SFNC97] Laszlo Szirmay-Kalos, Tibor Foris, Laszlo Neumann and Balasz Csebfalvi. An analysis of quasi-Monte Carlo integration applied to the transillumination radiosity method. In *Proceedings of Eurographics'97*, Budapest, September 1997.

[VG95] Eric Veach and Leonidas J. Guibas. Optimally combining sampling techniques for monte carlo rendering. In *Computer Graphics Proceedings, Annual Conference Series, 1995 (ACM SIGGRAPH '95 Proceedings)*, pages 419–428, 1995.