

Reuse of Paths in Final Gathering Step with Moving Light Sources

Mateu Sbert, Francesc Castro

*Institut d'Informàtica i Aplicacions, Universitat de Girona.
Edifici P4. Campus Montilivi. 17071 Girona, Spain.
mateu|castro@ima.udg.es*

Abstract. In this paper we extend our previous result on reusing shooting paths when dealing with moving light sources in radiosity to the final gathering step algorithm, which combines a shooting solution and a gathering step. Paths are reused for each frame in both shooting and gathering steps, obtaining a noticeable reduction of cost in the computation of the solutions for every frame. Moreover, the correlation between consecutive frames happens to be higher than in the case of independent solutions, resulting in a clear reduction of the temporal aliasing effects in the animation.

Keywords Radiosity, Random Walk, Animation

1 Introduction

The radiosity solution [3] can be obtained by combining shooting and gathering random walks [1]. One of these hybrid algorithms is the *final gathering step* [5]. From a radiosity solution obtained by shooting, a simple gathering step allows to obtain a final solution for a finer level of subdivision.

We present in this paper an acceleration strategy applied to this final gathering step algorithm and valid for animations involving moving light sources. Our strategy is based on the reuse of paths [2], which is done for both shooting and gathering ones. The reuse of shooting paths has been already presented in [6], and it will be described in next section. Here we incorporate the reuse of gathering paths, based on the use of each path for all the frames (that is, for all positions of the light source). The reuse of both shooting and gathering paths permits to calculate the solution for all frames with a noticeable reduction of cost regarding to the independent computation of each solution (classic approach).

Moreover, the reuse of the paths (in both shooting and gathering) produces a correlation between the solutions of consecutive frames. This fact clearly reduces the flickering effect observed in the video sequences. This is not applicable when computing each frame independently, as done in the classic approach.

This paper is organized as follows. In next section we will refer to previous work, including the reuse of shooting paths. The description of our contribution appears in section 3. Section 4 presents the results, with both error graphs and images. Finally, in last section we present the conclusions and future work.

2 Previous work

2.1 Final Gathering Step

Given a radiosity solution at a coarse level, say patches, the final gathering step is meant to solve at a finer level, say elements or pixels, for the interesting (or important) part of the scene. Rays can be cast from the eye to get a view dependent solution (usually to obtain a glossy solution enhancing a first radiosity solution) for the pixels in the screen plane as in the photon-map technique [7], or from the patches of interest to get a new radiosity solution for the elements. On the other hand the first radiosity solution can be obtained by a deterministic method (for instance classic hierarchical radiosity) or a Monte Carlo method (random walk, hierarchical Monte Carlo) (see [5] for a short survey). When gathering from a complete radiosity solution, the gathering path can be limited to length one. The foundation in this case is the following one. Consider the radiosity equation:

$$B_i = R_i \sum_{j=1}^{n_p} F_{ij} B_j + E_j \quad (1)$$

where B_i , E_i , R_i are respectively the radiosity, emissivity and reflectance value for patch (element) i , and F_{ij} the form factor from patch (element) i to patch (element) j . We can solve with Monte Carlo equation (1), by choosing N patches j with a given probability vector (or pdf in the continuous case). Taking as probabilities the form factors F_{ij} , and considering a first step radiosity approximation B_j^* , the final gathering estimator for B_i results in:

$$\hat{B}_i = \frac{1}{N} R_i \sum_{j=1}^N B_j^* + E_j \quad (2)$$

2.2 Reuse of shooting paths in light source animation

Shooting paths can be reused in the context of radiosity when dealing with moving light sources [6]. The main idea is that a path exiting from a point of the light source can be used to distribute power from every light source position, or in other words, to calculate the illumination for every frame. This means that each path will be used to build n paths, n being the number of light source positions (see Fig. 1). This involves a noticeable reduction of cost.

Let Φ be the power of a source taking n positions. M points are considered on the source. Let x be a point on the source repeated in each different position x_1, x_2, \dots, x_n . Let j be the position taken into account and $F(x_j, y)$ the form factor from x_j to a point y (y being any point to receive direct illumination from x_j). The unbiased estimator for the power arriving to point y from x_j is (see [6])

$$\frac{F(x_j, y)}{F(x_1, y) + F(x_2, y) + \dots + F(x_n, y)} \times \frac{\Phi}{M} \quad (3)$$

The analysis of the cost shows that the theoretical acceleration factor due to the reuse of the shooting paths is bounded by l , l being the average length of

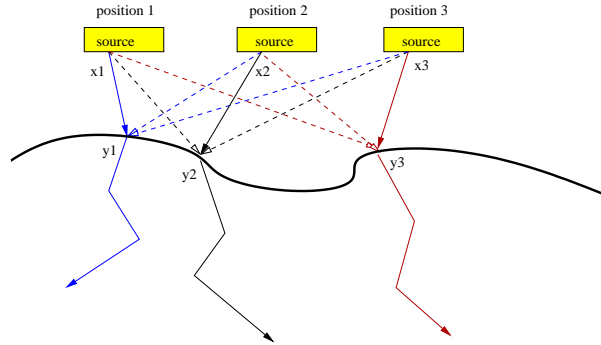


Fig. 1. 2D representation of a shooting random walk with reuse of the paths. Path from light source in position 1 (blue) is reused to distribute power from positions 2 and 3 of the light source, and the same with paths from position 2 (black) and 3 (red).

a path, when n grows. On the other hand, the expected error using the same number of paths per position will be the same in the best case. This corresponds to the limiting case in which all the light source positions were the same. In general, the more distant the positions, the bigger the error, due to the fact that when positions get distant, the sampling pdf for y goes away from the importance sampling function.

3 Reuse of gathering paths in light source animation

The new contribution presented in this paper consists in the reuse of paths in the final gathering step algorithm (reviewed in section 2.1) for the case of moving light sources. We will reuse both shooting and gathering paths. The reuse of shooting paths is done in the same way as explained in section 2.2. The reuse of gathering paths is based on the fact that each gathering path can be used to gather the radiosity for each of the n frames.

3.1 Dealing with direct illumination

The final gathering step algorithm presents a problem when dealing with small light sources. This problem concerns the distribution of the direct illumination. Such illumination is dealt with when a gathering ray hits a light source. But the usual small size of the light source makes the number of gathering rays that hit the source from a given patch to be small, and thus the error for the received direct illumination is big. This results in an unacceptable aliasing in the final image, only avoided if the number of gathering rays is very large.

This problem can be solved by calculating direct illumination in the shooting step instead of in the gathering one. Thus, the gathering step will only deal with the indirect illumination. Final radiosity values will be obtained as the sum of

direct and indirect values. In this approach, the shooting step will work with two different levels of subdivision. The finest level will be used to store direct illumination, while the coarse one will be used for the complete shooting solution, which will be gathered in the final step. From here on, we will refer to the finest level subdivision units as *elements*, and to the roughest level ones as *patches*.

Since direct illumination solution will be calculated at the element level, and a complete shooting solution will be calculated at the patch level, the number of paths needed to get a determined accuracy is not the same in each case. The number of shooting paths used to calculate direct illumination must be higher than the one for the complete shooting solution. A compromise consists of eliminating part of the paths after the first reflection, using a parameter λ ($0, 1]$ as the probability for a path to continue after the first hit. This means that each shooting path will have a double function. On the one hand, it will be used to distribute primary power from the sources to the elements, computing in this way the direct illumination. On the other hand, it will be used to distribute power (not only primary) from the sources to the patches in the scene. After the first hit, the continuation of paths will be decided according to λ . If a path survives, its carried power has to be multiplied by $1/\lambda$ to compensate for the reduction of the number of paths.

3.2 Reusing gathering paths

In the gathering step one-length paths are cast from each element i in the scene, and the nearest intersected patch j is obtained for each path. Then, radiosity of patch j is gathered by i . The reusing strategy proposed here consists of using every path to gather at once the radiosity for each of the n solutions corresponding to the n frames (n positions of the light source). Note that each element and each patch in the scene store an array of n radiosity accumulators (1 per frame). All n gathering solutions are computed at the cost of a single one. As the expected error for our solutions is exactly the same as from n results obtained in n independent executions, the expected speed-up factor of our strategy for the gathering step will be n .

Also, the reuse of the paths will provide us with an important additional gain: the correlation between consecutive frames, due to the reuse of the paths, makes the transition between frames smoother, avoiding the annoying flickering effect observed with independent solutions. This is clearly visible in the videos accompanying the paper.

3.3 Comparing the expected cost

Next we compare the cost of a radiosity solution S for n frames obtained by reusing both shooting and gathering paths against another radiosity solution S' obtained without reusing any path (that is, from independent executions for each frame in both shooting and gathering). Both solutions S and S' have been obtained using the same number of shooting and gathering paths per frame.

Let l be the average length of a path. Let N be the total number of shooting paths per frame, and let N_g be the total number of gathering paths per frame, being $N_g = kN$ for a positive value k (note that finding an optimal value of k is not an objective of this paper). We have to remark that, regarding to the number of shooting paths per frame, their reuse permits to reduce their number by a factor of n ; that is, to obtain N effective shooting paths per frame, only $M = N/n$ actual paths have to be cast from each position of the source.

For the sake of simplicity in the following formulas, the cost has been expressed as the number of nearest intersection computations (actually, most of the real cost is due to such computations).

The results in [6] can be now generalized by considering the parameter λ (see section 3.1). A simplification has been done by considering the cost of a visibility computation (used in the point-to-point form factor computation) to be the same as the cost of computing the nearest intersection (actually, and depending on the implementation, the cost of a visibility computation can be lower or, in the worst case, equal to the nearest intersection one). The cost of the shooting step with reusing will be $N(n + \lambda(l - 1))$, while the cost without reusing will be $Nn(1 + \lambda(l - 1))$.

The total cost of the gathering step when computing independently the n frames (classic approach) is given by nkN , since we have to cast kN paths for each frame. Conversely, if we reuse the gathering paths (new approach) we just have to cast kN one-length paths in total, so this is the cost in this case.

The total expected acceleration factor, that considers the addition of both shooting and gathering cost, is given by

$$\frac{nN(1 + \lambda(l - 1)) + nkN}{N(n + \lambda(l - 1)) + kN} = \frac{n(1 + \lambda(l - 1)) + kn}{n + \lambda(l - 1) + k} \quad (4)$$

When the number of light source positions (frames) grows to infinity, we have

$$\lim_{n \rightarrow \infty} \frac{n(1 + \lambda(l - 1)) + kn}{n + \lambda(l - 1) + k} = k + 1 + \lambda(l - 1) \quad (5)$$

Observe incidentally that the acceleration factor for indirect illumination can be obtained by putting $\lambda = 1$ in (5).

3.4 Expected error

In the best case, using the same number of paths per position in both classical and new approach will produce the same error. This best case corresponds to the limiting case in which all the light source positions were the same. In general, the more distant the positions, the bigger the error. This is due to the fact that when positions get distant, $p(y)$ goes away from the importance sampling function corresponding to each source position j .

This means that the new algorithm presents a theoretical speed-up factor bounded by the expression (5).

4 Implementation and results

We have used in our implementation a collision random walk estimator (see [1]). Parameter λ (see section 3.1) has been set to 0.25 in all our tests. The reason for this value is that we have subdivided each patch in 4 elements, so that the area of an element is approximately 1/4 of the area of the patch, and the variance of the radiosity random walks is known to be proportional to the area of the elements.

Regarding to the gathering step, the value of k (see section 3.3) that indicates the relation between the number of shooting and gathering paths, has been set to 1 in all our tests.

The new algorithm implies some additional storage: for each patch, the accumulated values for each frame must be stored, and also for each element in which patches are subdivided. This corresponds to store an array of n triplets (r,g,b) per patch and element, that is, a total of $n(P+4P) = 5nP$ triplets, where P is the number of patches. Note that using the classical approach we just have to store $5P$ triplets (one per patch and element) for each of the n executions. However, no storage for the shooting or gathering paths is required, due to the immediate update of values for each position when dealing with a path.

The tested scene (see Fig. 3) has been discretized in 2686 patches and 10744 elements. We have considered 30 very close light source positions that are tracing a diagonal route near the ceiling.

We have obtained an acceleration factor of about 6, much higher than the theoretical prevision (about 2.8 according to (4)) due to the lower cost of the visibility computations. On the other hand, we have noted that the Mean Square Error (MSE) using the new method is lower for frames placed in the middle of the frame sequence than for frames placed in the extremes. This behavior can be theoretically justified by the fact that central positions have an importance sampling function closer to the pdf that we have used than the extreme positions.

We present in Fig. 2 the time vs. MSE graph corresponding to frame 17, including also the results of reusing only gathering paths (from independent shooting solutions).

Figure 3 shows two frames in the animation. The one on the left has been obtained using the classic algorithm, in which no reuse of paths is done. The one on the right has been obtained using the new algorithm, involving reuse of both shooting and gathering paths. The number of shooting and gathering paths is the same in both classic and new approach, but in the new approach the number of shooting paths includes the reuse. No remarkable difference is observed between the quality of left and right images, but the cost is much lower using the new algorithm.

Animation sequences accompanying the paper ¹ show the reduction of the flickering effect when reusing the paths (see section 3.2). Animation 1 involves no reuse of paths (classic approach). Each image has been obtained in about 200 seconds. That means a total of $200 \times 30 = 6000$ sec. = 100 min. Animation 2

¹ ima.udg.es/~castro/videoCGGM/anim1.gif|anim2.gif|anim3.gif




Fig. 2. Graph of time in seconds (horizontal axis) vs. MSE (vertical axis). This graph corresponds to frame 17. Note the clear advantage of new method in front of the classical one, and also in front of reusing only gathering paths.



Fig. 3. (left) Frame 17 without reusing paths. 600.000 paths. Time= 70 sec. (right) Frame 17 reusing shooting and gathering paths. 600.000 paths. Time= 11.7 sec. MSE is a bit higher in the second image. Speed-up factor about 4.4.

involves reuse of the gathering paths, but not of the shooting ones. Each shooting has spent about 114 sec., and the combined gathering has spent about 86 sec. This is a total of $114 \times 30 + 86 = 3506$ sec. = 58.5 min. Animation 3 has been obtained by reusing both shooting and gathering paths. The combined shooting has spent 1017 sec., and the combined gathering about 86 sec. This means a total of $1017 + 86 = 1103$ sec. = 18.4 min.

5 Conclusions and future work

We have extended in this paper our previous results on reusing shooting paths to the final gathering step algorithm, which combines a shooting solution and a gathering step. It has been applied in the context of radiosity to the case of moving light sources. Both shooting and gathering paths have been reused for each frame, so that different frames in an animation (corresponding to different positions of the light source) have been obtained with a considerable gain respect to the classical algorithm.

The gain will come not only from the reusing strategy, but also from the inner-frame coherence obtained by eliminating the annoying flickering due to temporal aliasing.

The scope of the idea presented in this paper is not limited to radiosity. It could be applied to different techniques in global illumination, like virtual light sources [4], and photon maps [7].

6 Acknowledgments

This project has been funded in part with a Catalan-Flemish Joint Action n. ACI2002-52 from the Catalan Government, and with grant number TIC2001-2416-C03-01 from the Spanish Government.

References

1. P.Bekaert *Hierarchical and Stochastic Algorithms for Radiosity*. Ph.D. thesis. Katholic Univ. of Leuven. 1999.
2. P.Bekaert, M.Sbert, and J.Halton *Accelerating Path Tracing by Re-Using Paths*. Proceedings of Workshop on Rendering 125-134. 2002.
3. M.Cohen, and J.Wallace *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Cambridge. 1993.
4. A.Keller *Instant Radiosity*. Computer Graphics Proceedings, Siggraph'97. 49-56. 1997.
5. F.Perez *Global Illumination Techniques for the Computation of High Quality Images in General Environments*. Ph.D. thesis. Universitat Politècnica de Catalunya. 2003.
6. M.Sbert, F.Castro, and J.Halton *Reuse of Paths in Light Source Animation*. Computer Graphics International Proceedings (to appear as short paper). 2004.
7. H.Wann Jensen *Global Illumination using Photon Maps*. Rendering Techniques. p. 21-30. 2004.