

Final Radiosity Gather Step using a Monte Carlo Technique with Optimal Importance Sampling

Philippe Bekaert

Philip Dutré

Yves D. Willems

Report CW 275, January 1996



Katholieke Universiteit Leuven
Department of Computer Science

Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

Final Radiosity Gather Step using a Monte Carlo Technique with Optimal Importance Sampling

*Philippe Bekaert**

Philip Dutré

Yves D. Willems

Report CW 275, January 1996

*Research performed with financial support by a grant from the "Flemish Institute for the Promotion of Scientific-Technological Research in Industry" (IWT#941120).

Department of Computer Science, K.U.Leuven

Abstract

In this paper, an improved Monte Carlo based technique is presented to perform a per-pixel final gather step, also referred to as a local pass in radiosity computations. Using importance sampling based on the results from the finite-element radiosity computation, we achieve a very high image quality at a much lower cost compared to other common techniques. The theory in this paper exemplifies how finite element and Monte Carlo techniques for rendering can be combined elegantly.

Keywords : Global Illumination, Radiosity, Raytracing, Monte Carlo-techniques

CR Subject Classification : I.3.7

1 Introduction

First, we briefly present the main ideas behind radiosity methods, radiance reconstruction and Monte Carlo integration, introducing terminology. Some previous per-pixel reconstruction techniques are presented in §2. Our new improved technique will be outlined in section §3. Finally the results are discussed in §4.

1.1 Radiosity methods and radiance reconstruction

Light transport in an environment consisting of diffuse surfaces only and with absence of participating media is described by the equation

$$\begin{aligned} L(x) &= L_e(x) + L_{refl}(x) \\ &= L_e(x) + f_r(x) \int_{\Omega} L(rt(x, \Theta_x)) \cos \theta_x d\omega_{\Theta_x} \end{aligned} \quad (1)$$

where

- $L(x)$ is the total exitant radiance [$W/m^2 sr$] at point x on a surface. Although proportional to radiosity $B(x) = \pi L(x)$ [W/m^2] in a diffuse environment, it is a more fundamental quantity to describe light transport;
- $L_e(x)$ and $L_{refl}(x)$ are the self-emitted and reflected radiance at point x respectively;
- $f_r(x)$ is the bi-directional reflection distribution function (BRDF) [$1/sr$] at point x . The BRDF does not depend on incoming or outgoing directions and is proportional to the reflectivity $\rho(x) = \pi f_r(x)$ for diffuse reflectors;
- Ω the hemisphere of all directions at point x ;
- $d\omega_{\Theta_x}$ a differential solid angle around direction $\Theta_x = (\phi_x, \theta_x)$ at point x .
- $rt(x, \Theta_x)$ the nearest visible point from x in the direction Θ_x (ray-trace function);

It is often more appropriate to express (1) in terms of points x and y on surfaces only:

$$L(x) = L_e(x) + f_r(x) \int_A L(y) G(x, y) dA_y \quad (2)$$

with

$$G(x, y) = \frac{\cos \theta_x \cos \theta_y}{r_{xy}^2} vis(x, y)$$

the radiosity kernel, and

- A is the set of all points on surfaces in the scene;
- dA_y a differential area at point y ;

- θ_x and θ_y the angle which the line connecting the points x and y makes with the normal on the surface at respectively point x and y ;
- r_{xy} is the distance between the points x and y ;
- $vis(x, y)$ is 1 if both points are mutually visible and equals 0 otherwise.

The most common approach to compute light transport in diffuse environments is to approximate the surfaces in the scene by planar triangles and/or convex quadrilaterals called patches and to approximate the radiance on each patch i as a linear combination of basis functions $\phi_{i,\alpha}(x)$:

$$\begin{aligned}\tilde{L}(x) &= \sum_{\alpha} L_{i,\alpha} \phi_{i,\alpha}(x) \approx L(x) \\ \tilde{L}_e(x) &= \sum_{\alpha} E_{i,\alpha} \phi_{i,\alpha}(x) \approx L_e(x)\end{aligned}$$

By substituting these expressions for $L(x)$ and $L_e(x)$ in (2) and after projection on the basis functions $\phi_{i,\alpha}$ a set of linear equations with unknowns $L_{i,\alpha}$ results. Solving this set yields $\tilde{L}(x)$ for all points on the patches. A popular choice for $\phi_{i,\alpha}$ is one constant basis function per patch:

$$\phi_i(x) = 1 \text{ if } x \in \text{patch } i \text{ and } 0 \text{ if } x \notin \text{patch } i.$$

If the reflectivity $\rho_i = \pi f_r(x)$ is constant over the patch i , this choice leads to the classical radiosity equations [4]:

$$L_i = E_i + \rho_i \sum_j F_{ij} L_j$$

with form factors

$$F_{ij} = \frac{1}{A_i} \int_{A_i} \int_{A_j} \frac{G(x, y)}{\pi} dA_x dA_y.$$

Although it is not restricted to the constant radiance approximation, we will outline our technique for this case only.

Some defects in the computed radiance solution, using the constant radiance approximation, are:

- artificial discontinuities in the computed radiance function at element boundaries. With non-constant radiance approximations, the derivatives of the radiance function will be discontinuous, resulting in disturbing Mach-band effects;
- radiance is sometimes poorly reproduced inside the elements as well. Incorrectly reproduced shadow boundaries and light- or shadow leaks are the most noticeable symptoms of this problem. When using non-constant approximations, there is no guarantee that the computed radiance will be positive on all points of a single patch. Regions where the radiance has been clipped to zero are very disturbing as well.

A number of techniques to partially solve these problems are described in [3]. A very common solution is to compute radiances at the vertices of elements based on the computed radiance of the elements sharing the vertex, and then to apply Gouraud interpolation to reconstruct the radiance inside the elements. This remedies for discontinuities in radiance value at element boundaries, but shadow boundaries are still incorrectly reproduced and light- and shadow leaks still occur. First-order discontinuity meshing [3] remedies for incorrectly reproduced shadow boundaries and light- and shadow leaks. Disturbing Mach-band effects at element boundaries due to discontinuities in the first derivative of the reconstructed radiance function can be remedied by using appropriate interpolation techniques.

Alternatively, when a very high image quality is needed, one will evaluate

$$L(x) = L_e(x) + f_r(x) \int_A \tilde{L}(y)G(x, y)dA_y \quad (3)$$

at one or more points visible in every pixel of the image. The reconstruction technique presented in this paper belongs to this class and a short survey of such techniques will be presented in §2.

1.2 Monte Carlo integration

Our reconstruction technique is based on Monte Carlo integration using importance sampling. In this section we briefly present the main ideas of this integration technique [5].

The basic idea of Monte Carlo integration is to estimate the value of an integral

$$I = \int_0^1 f(x)dx$$

by taking a uniformly sampled random point \bar{x} in the domain (which is here for simplicity the interval $[0, 1]$) and using $\langle I \rangle_1 = f(\bar{x})$ as estimator. The expected value of the estimator $\langle I \rangle_1$ is equal to the integral I :

$$E[\langle I \rangle_1] = \int_0^1 f(x)dx = I.$$

However, if f is not a constant function, $f(\bar{x})$ will rarely yield the exact integral I . The variance

$$V[\langle I \rangle_1] = \int_0^1 (f(x) - I)^2 dx = \int_0^1 f^2(x)dx - I^2$$

is a measure for the reliability of $f(\bar{x})$ as an estimate for the integral I . Variance is visible as noise in images rendered with Monte Carlo techniques.

Variance can be reduced in several ways:

- By taking more samples: if N samples \bar{x}_i are taken according to a uniform probability distribution, the expected value $E(\langle I \rangle_N)$ of the estimator $\langle I \rangle_N = 1/N \sum_i f(\bar{x}_i)$ equals the integral I , but the variance

$$V(\langle I \rangle_N) = \frac{1}{N} \int_0^1 f^2(x) dx - \frac{1}{N} I^2$$

has been reduced by a factor N compared to $V(\langle I \rangle_1)$; Theoretically, the variance only becomes 0 if we take an infinite amount of samples. However, a “large” number of samples can also yield relatively reliable results. An image computed with a Monte Carlo algorithm, using a huge number of samples, has the advantage of yielding a very small variance and thus producing an almost exact result, since the expected value still is the value we want to compute. Such an image can therefore be used as a reference image to compare with other techniques;

- By stratified sampling: if each of the samples x_i is selected according to a uniform probability distribution in separate intervals X_i such that $\cup_i X_i = [0, 1]$, the variance is reduced to

$$V(\langle I \rangle) = \frac{1}{N} \int_0^1 f^2(x) dx - \sum_{i=1}^N I_i^2$$

where $I_i = \int_{X_i} f(x) dx$. We assumed equal-sized intervals X_i in this formula;

- By importance sampling: when a sampling point \bar{x} is chosen according to a non-uniform probability density function (pdf) $p(x)$, the expected value $E(\langle I \rangle)$ of the estimator $\langle I \rangle = f(\bar{x})/p(\bar{x})$ is equal to the integral to be computed:

$$E(\langle I \rangle) = \int_0^1 \frac{f(x)}{p(x)} p(x) dx = \int_0^1 f(x) dx = I$$

The variance $V(\langle I \rangle)$ of this estimator

$$V(\langle I \rangle) = \int_0^1 \left(\frac{f(x)}{p(x)} \right)^2 p(x) dx - I^2 = \int_0^1 \frac{f^2(x)}{p(x)} dx - I^2$$

will be zero if $p(x)$ is chosen proportional to $f(x)$: $p(x) = f(x)/I$. In this case one sample would suffice to know the integral exactly (the sample yields the proportionality factor of $f(x)$ and $p(x)$). Constructing such a pdf however requires knowing I , which is exactly the integral we want to compute. One may expect however that fewer samples x_i will be needed to estimate the integral with equal accuracy if a pdf $p(x)$ is chosen with shape that matches the shape of $f(x)$ rather well.

Our radiance reconstruction technique, outlined in §3, makes use of importance sampling using the form factors and approximate radiance solutions obtained during the finite-element computations.

2 Per-pixel radiance reconstruction techniques

In this section a short survey of per-pixel radiance reconstruction techniques is given.

The basic idea is to compute the average radiance L_{pix} received through each pixel pix at the eye point eye :

$$L_{pix} = \frac{1}{\Omega_{pix}} \int_{\Omega_{pix}} L(rt(eye, \Theta_{eye})) d\omega_{\Theta_{eye}} \quad (4)$$

where Ω_{pix} is the projection of the pixel pix onto the unit sphere around the eye point eye . The integral in (4) is estimated using a Monte Carlo technique: one or more random directions $\Theta_{eye} \in \Omega_{pix}$ to the pixel are chosen. For each Θ_{eye} , $L(x)$ with $x = rt(eye, \Theta_{eye})$ is computed using (3):

$$L(x) = L_e(x) + f_r(x) \int_A \tilde{L}(y) G(x, y) dA_y.$$

The average of the $L(x)$ obtained this way is an estimate for L_{pix} .

The per-pixel reconstruction methods described in the literature differ in the way (3) is evaluated:

2.1 Form factors per pixel

Equation (3) is equivalent to

$$L(x) = L_e(x) + f_r(x) \sum_j \int_{A_j} \tilde{L}(y) G(x, y) dA_y.$$

All terms

$$\int_{A_j} \tilde{L}(y) G(x, y) dA_y = L_j \int_{A_j} G(x, y) dA_y = L_j \pi F_{dA_x, j},$$

where $F_{dA_x, j}$ is the point- x to patch- j form factor, are evaluated with equal precision regardless of the contribution of element j to the illumination at x [8]. When the point-to-patch form factors are computed by integration over the surface of the patches j , some artifacts may occur near discontinuities in the radiosity kernel $G(x, y)$. These can be remedied by evaluating the point-to-patch form factors using integration over the spherical triangle or quadrilateral Ω_j subtended by element j on the hemisphere at point x instead of integrating over the surface of j . Other artifacts are due to the constant radiance assumption, and disappear when higher-order bases are used. This technique results in images that are indistinguishable from the reference image in figure 4.

A less accurate but generally faster alternative is to compute all the required point-to-patch form factors using the hemicube method [2]. An exact solution, by projecting the scene onto the hemisphere at point x and then evaluating the form factors to each visible part by contour integration is described in [10].

In general however, most of the elements j contribute very little to the illumination at point x . Evaluating the point-to-patch form factors $F_{dA_x, j}$ with equal precision regardless of the contribution of j to the illumination at x is highly suboptimal. The new technique we propose in §3 does an optimal division of work over the linked elements j .

2.2 Monte Carlo integration

Equation (3) is also equivalent to

$$L(x) = L_e(x) + f_r(x) \int_{\Omega} \tilde{L}(rt(x, \Theta_x)) \cos \theta_x d\omega_{\Theta_x}. \quad (5)$$

The integral in the right hand side can be estimated by Monte Carlo techniques. One easy way is to use stratified sampling of the hemisphere at point x according to a $\cos \theta_x$ probability distribution.

As is evident in figure 1 this technique may result in very poor sampling of the direct illumination at point x . This is due to the fact that light sources or other very influential elements may subtend only a small solid angle on the hemisphere at most points x . Spikes from direct illumination in the incoming radiance distribution can be handled efficiently by a technique often referred to as *next event estimation*: the integral in (5) is split in two parts

$$\begin{aligned} \int_{\Omega} \tilde{L}(rt(x, \Theta_x)) \cos \theta_x d\omega_{\Theta_x} &= \int_{\Omega} \tilde{L}_e(rt(x, \Theta_x)) \cos \theta_x d\omega_{\Theta_x} \\ &+ \int_{\Omega} \tilde{L}_{refl}(rt(x, \Theta_x)) \cos \theta_x d\omega_{\Theta_x}. \end{aligned}$$

The indirect-illumination integral (over $\tilde{L}_{refl} = \tilde{L} - \tilde{L}_e$) is estimated using Monte Carlo techniques while the direct-illumination integral, a sum of integrals over each light source directly contributing light to x can be handled similarly to the per-pixel form factors technique of §2.1:

$$\begin{aligned} \int_{\Omega} \tilde{L}_e(rt(x, \Theta_x)) \cos \theta_x d\omega_{\Theta_x} &= \sum_l \int_{A_l} \tilde{L}_e(y) G(x, y) dA_y \\ &= \sum_l E_l \pi F_{dA_x, l}. \end{aligned}$$

With next event estimation, the direct illumination at point x is computed accurately (see figure 2) but variance on the indirect illumination estimation is visible as noise. Noise is the main drawback of this technique. To reduce noise by a factor S , S^2 more samples need to be taken. In the technique we propose in §3 noise is reduced by exploiting as much as possible the knowledge about the scene that is available from the finite element computations.

In [9] only the direct-illumination integral (over \tilde{L}_e) is evaluated, while indirect illumination is obtained by interpolating the radiances from the radiosity mesh. The distinction between primary and secondary light-sources is however rather arbitrary. In [7] an algorithm is presented in which the illumination from the most important secondary light sources in the scene is re-evaluated as well. The technique we propose in §3 does an unbiased selection of sources however.

3 Monte Carlo per-pixel radiance reconstruction using optimal importance sampling

3.1 The basic idea

Optimal Monte Carlo estimation by means of importance sampling of

$$\int_A \tilde{L}(y)G(x, y)dA_y = \sum_j \int_{A_j} \tilde{L}(y)G(x, y)dA_y \quad (6)$$

will be obtained if we can construct a probability density function which closely resembles the integrand $\tilde{L}(y)G(x, y)$. From the finite-element computations, the following data are available that can be used to construct a good pdf:

- Approximations for the average radiance on each patch j :

$$L_j = \frac{1}{A_j} \int_{A_j} \tilde{L}(y)dA_y;$$

- The average point- x -to-patch- j form factor $F_{dA_x, j}$ for points x on patch p , visible through the pixel being considered, to each linked element j of p :

$$\pi F_{pj} = \frac{1}{A_p} \int_{A_p} \int_{A_j} G(x, y)dA_x dA_y \approx \pi F_{dA_x, j} = \int_{A_j} G(x, y)dA_y;$$

We propose to perform the choice of a sample to evaluate (6) in two steps:

1. The choice of an element j to sample, or in other words, the choice of a term in the sum of (6): the probability of choosing patch j should be proportional to

$$\int_{A_j} \tilde{L}(y)G(x, y)dA_y \approx \pi F_{pj}L_j$$

We thus assign each patch j a probability

$$P_j = \frac{F_{pj}L_j}{\sum_j F_{pj}L_j} \quad (7)$$

of being sampled.

2. Evaluation of

$$\frac{1}{P_j} \int_{A_j} \tilde{L}(y)G(x, y)dA_y \quad (8)$$

for the element j selected in step 1. Several methods for evaluating this integral are possible. We have implemented and compared the following sampling schemes:

- (a) Method A: Choosing a sample point y with uniform probability $\frac{1}{A_j}$ on patch j . An estimate for the integral given in (8) is

$$A_j L_j \frac{\cos \theta_x \cos \theta_y}{r_{xy}^2} vis(x, y)$$

- (b) Method B: In [1] an algorithm is presented to sample a random direction with uniform probability $\frac{1}{\Omega_j}$ on a spherical triangle Ω_j . This can be used to sample the equivalent integral

$$\int_{\Omega_j} \tilde{L}(y) vis(x, y) \cos \theta_x d\omega_{\Theta_x}$$

where y is the point on patch j visible along the direction $\Theta_x \in \Omega_j$. An estimate for the integral in (8) is then

$$\Omega_j L_j vis(x, y) \cos \theta_x$$

- (c) Method C: The $\cos \theta_x$ can be taken into account as well when choosing a sample using rejection sampling on a bounding region $\bar{\Omega}_j$ for Ω_j on the hemisphere. The bounding region is defined by intervals $[\phi_1, \phi_2]$ and $[\theta_1, \theta_2]$ bounding all $(\phi, \theta) \in \Omega_j$. The integral in (8) is then equivalent with

$$L_j \int_{\bar{\Omega}_j} \cos \theta vis(x, y) R(y) d\omega_{\Theta} = L_j \int_{\phi_1}^{\phi_2} \int_{\theta_1}^{\theta_2} \cos \theta vis(x, y) R(y) d\phi \sin \theta d\theta$$

where

- y is the intersection of a line through x along direction $\Theta = (\phi, \theta)$ with the plane of patch j ;
- $R(y)$, the rejection function, is 1 if y is inside patch j and 0 otherwise, equivalent with whether $\Theta = (\phi, \theta) \in \Omega_j$.

Two uniformly distributed random numbers ξ_1 and ξ_2 define a sample $(\phi, \theta) \in \bar{\Omega}_j$ according to a $\cos \theta$ distribution with

$$\begin{aligned} \phi &= \phi_1 + \xi_1(\phi_2 - \phi_1) \\ \theta &= \text{asin} \sqrt{\sin^2 \theta_1 + \xi_2(\sin^2 \theta_2 - \sin^2 \theta_1)}. \end{aligned}$$

An estimate for the integral is then

$$L_j vis(x, y) R(y) \int_{\bar{\Omega}_j} \cos \theta d\omega_{\Theta}$$

Alternatively one can generate samples until a non-rejected one ($R(y) = 1$) is encountered. In this case,

$$L_j vis(x, y) \frac{\Omega_j}{\bar{\Omega}_j} \int_{\bar{\Omega}_j} \cos \theta d\omega_{\Theta}$$

is an estimate for the integral. However, if Ω_j is malformed ($\Omega_j \ll \bar{\Omega}_j$), it may take an arbitrary long time before a sample inside Ω_j is generated. In practice, one will switch to method A or B after a given number of rejections. We have used this strategy in our implementation.

- (d) Method D: when patch j is fully visible from point x , the point-to-patch formfactor $F_{dA_x,j}$ can be computed analytically [3]. This will e.g. be the case if there is full visibility between the patches p and j . However, if this information is not available from the finite element computations, one of the previous methods should be used due to the cost of reliable visibility determination.

3.2 Handling colour

In §3.1 we have ignored the fact that radiance $L(x, \lambda)$ is a function of wavelength λ . Most rendering codes represent radiance only at a limited number of wavelengths (typical three channels L^r, L^g, L^b perceived as red, green and blue). This implies that there are potentially very different probabilities (7), corresponding to each colour channel, to choose terms in the sum (6), e.g.:

$$P_j^r = \frac{F_{pj}L_j^r}{\sum_j F_{pj}L_j^r}, \quad P_j^g = \frac{F_{pj}L_j^g}{\sum_j F_{pj}L_j^g}, \quad P_j^b = \frac{F_{pj}L_j^b}{\sum_j F_{pj}L_j^b}$$

The question is how these pdf's should be used in order to generate optimal results and without losing the advantages of our "good" sampling strategy.

One simple solution is to construct one combined pdf P_j by using the sum of the radiance components:

$$P_j = \frac{F_{pj}(L_j^r + L_j^g + L_j^b)}{\sum_j F_{pj}(L_j^r + L_j^g + L_j^b)}.$$

However, [11] presents a number of more efficient ways to combine different sampling methods for the same process. We have used the *balance heuristic* described there to combine samples generated according to each pdf P_j^α , $\alpha = r, g, b$. The number of samples n_α generated according to each pdf was chosen proportional to $P^\alpha = \sum_j P_j^\alpha$:

$$n_\alpha = N \times \frac{P^\alpha}{\sum_\alpha P^\alpha}$$

where N is the total number of samples to be generated. The weight $1/P_j$ in (8) was replaced by $1/\bar{P}_j$ with combined probability

$$\bar{P}_j = \sum_\alpha \frac{n_\alpha}{N} P_j^\alpha.$$

4 Results and discussion

The figures show the results of the various techniques mentioned in this paper for a simple test-scene consisting of 361 initial patches. The radiosity solution was obtained using a hierarchical radiosity method [6]. After convergence, the 361 initial patches were refined into 6437 elements, with 74464 links. The cost of tracing a ray through the scene is determined by the number of initial patches, which is quite low in this example scene.

Figure 3 shows the results of our new technique, handling colour using the balance heuristic and sampling patches using method B. As evidenced from this figure, there is no need for next-event estimation, unlike previously proposed Monte Carlo techniques (figure 1). Figure 1, 2 and 3 were generated using the same number of rays per pixel. The reduction of variance due to the importance sampling is very clear. Remaining noise is to be attributed to discrepancies between the proposed pdf and the integrand in (6). These are due to

- inaccuracies in the evaluation of the form factors F_{pj} during the finite-element computations:

$$\int_{A_p} \int_{A_j} G(x, y) dA_y dA_x \neq \pi A_p F_{pj}$$

- variation of $\tilde{L}(y) \neq L_j$ over the elements j with non-constant radiance approximations;
- variation of $F_{dA_x, j} \neq F_{pj}$ over p . This variation is due to
 - changing visibility $vis(x, y)$ for $x \in A_p, y \in A_j$. This effect occurs near shadow boundaries;
 - the projection Ω_j of j onto the hemisphere at point x changes over p .
- suboptimal sampling to estimate (8): with method B $\cos \theta_x$ is not taken into account when choosing a sample. With method A $\cos \theta_y / r_{xy}^2$ is not taken into account either.

The cost of selecting a term in (6) is proportional to the number of terms, i.e. the number of elements that contribute to the illumination at the point x under consideration. This extra cost in our new technique can be quite large but is independent of the number of rays per pixel. If only few rays per pixel are needed in order to obtain a satisfactory image, traditional Monte Carlo techniques will be preferred. For a sufficiently large number of rays per pixel, when the cost of selecting a patch to sample in (6) becomes a smaller fraction of the total computing time, the advantage of much lower variance for equal number of rays per pixel will be appreciated. This is particularly true when noise-less images are to be generated.

However, the cost per element that may contribute to the illumination at x is much lower than with the per-pixel-form-factors techniques, where a point-to-patch form factor needs to be computed for each such element. Also the fact that with our technique more rays are cast to important elements and none to unimportant ones, while the contribution of unsampled elements is brought into account implicitly in an unbiased way, is definitely an advantage over per-pixel-form-factors techniques.

While patch sampling method B was only slightly more expensive than method A, the RMS difference of the resulting image compared to a reference image (figure 4) was about 9% lower. The reference image was obtained using method B with 1000 samples per pixel. Method C yielded almost no improvement over method B while it was about 3 times more expensive. Note that these results depend highly on the scene to be visualised. For other scenes these numbers were totally different and we found up to 20% improvement with method C over method B. Not using a different sampling method, 56% more samples would be required to achieve the same reduction of noise. Which method should be used depends on the scene complexity as well: for



Fig.1: Monte-Carlo without Next Event Estimation.



Fig.2: Monte-Carlo with Next Event Estimation.



Fig.3: "Optimal" Monte-Carlo.




Fig.4: Reference Solution.

complex scenes, where the cost of tracing a ray out-weighs the cost of choosing a ray direction, the more costly method C will be preferred.

Using the balance heuristic to combine samples chosen according to each pdf yielded a noise reduction of about 3% compared to the simple solution at negligible extra cost. The noise reduction will be more important in scenes with e.g. light-sources having different colours.

5 Conclusion and future directions

We have presented a new Monte Carlo based technique for performing a final gather step following finite-element radiance computations. In this new technique optimal advantage is taken from the knowledge available from the finite-element computations: geometric relations between elements (reflected in the form factors) and an approximate solution for the radiance distribution. With this new technique more work is spent in the computation of important contributions to the illumination than in less important contributions. Unsampled contributions are brought into account implicitly. There is an important reduction of noise using the same number of samples per pixel compared to traditional Monte Carlo techniques. The disadvantage is the potential high cost incurred in selecting a term to sample in (6). However, noiseless images, requiring a high number of samples per pixel, can be obtained at much lower cost with this new technique than with previously proposed techniques. Extension of the technique to radiance computations using non-constant bases on the elements is straightforward.

Two directions of future research that will be explored are:

- Integration of this technique in a final gather step accounting for specular reflection as well. We expect this to work well for scenes in which most of the reflection is purely diffuse, which is very often the case in practice;
- Adaption of the technique for use with finite-element algorithms to compute radiance in environments exhibiting glossy reflection.

References

- [1] James Arvo. Stratified sampling of spherical triangles. In *Computer Graphics Proceedings, SIGGRAPH Annual Conference Series*, pages 437–438, August 1995.
- [2] Michael F. Cohen and Donald P. Greenberg. The Hemi-Cube: A radiosity solution for complex environments. In *Computer Graphics (SIGGRAPH '85 Proceedings)*, volume 19, pages 31–40, August 1985.
- [3] Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, San Diego, CA, 1993.
- [4] Cindy M. Goral, Kenneth E. Torrance, Donald P. Greenberg, and Bennett Battaile. Modeling the interaction of light between diffuse surfaces. In *Computer Graphics (SIGGRAPH '84 Proceedings)*, volume 18, pages 212–22, July 1984.

- [5] J. M. Hammersley and D. C. Handscomb. *Monte Carlo Methods*. Chapman and Hall, London, UK, 1979.
- [6] Pat Hanrahan, David Salzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. In *Computer Graphics (SIGGRAPH '91 Proceedings)*, volume 25, pages 197–206, July 1991.
- [7] Arjan J. F. Kok and Frederik W. Jansen. Adaptive sampling of area light sources in ray tracing including diffuse interreflection. *Computer Graphics Forum (Eurographics '92)*, 11(3):289–298, September 1992.
- [8] Mark C. Reichert. A two-pass radiosity method driven by lights and viewer position. Master's thesis, Program of Computer Graphics, Cornell University, January 1992.
- [9] Peter Shirley. *Physically Based Lighting Calculations for Computer Graphics*. Ph.D. thesis, Dept. of Computer Science, U. of Illinois, Urbana-Champaign, November 1990.
- [10] Wolfgang Stürzlinger. Exact projections onto the hemisphere. In *Spring School on Computer Graphics, Bratislava, Slovakia*, June 1995.
- [11] Eric Veach and Leonidas Guibas. Optimally combining sampling techniques for monte carlo rendering. In *Computer Graphics Proceedings, SIGGRAPH Annual Conference Series*, pages 419–428, August 1995.