A Theoretical Comparison of Monte Carlo Radiosity Algorithms

Ph. Bekaert, H.-P. Seidel

Max Planck Institut für Informatik
Im Stadtwald Geb. 46.1, 66123 Saarbrücken, Germany
Email: Philippe.Bekaert@mpi-sb.mpg.de

Abstract

In this paper, we present a theoretical analysis of the error with three basic Monte Carlo radiosity algorithms, based on continuous collision shooting random walks, discrete collision shooting random walks and stochastic Jacobi iterations. We show that the variance of these three algorithms is in practice to very good approximation identical, explaining empirical results obtained before. The derivation in this paper can easily be adapted in order to calculate the variance of other Monte Carlo rendering algorithms, such as stochastic ray tracing and density estimation as well.

1 Introduction

This paper deals with algorithms for computing the illumination in a pure diffuse scene. It is well known that the illumination in such a scene is described by a integral equation similar to the rendering equation [5]:

$$B(x) = E(x) + \rho(x) \int_S G(x,y) B(y) dA_y \quad (1)$$

with kernel

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

$B(x)$ and $E(x)$ denote the total and self-emitted radiosity at a point x on the surfaces S of the scene. $\rho(x)$ is the reflectivity. r_{xy} is the distance between a pair of surface points x and y. $\cos\theta_x$ and $\cos\theta_y$ are the cosines of the angles between the surface normal at x and y and a line connecting both points. $\mathrm{vis}(x,y)$ is the usual visibility predicate. The algorithms we will compare here compute the average radiosity B_i on patches i in which the surfaces of the scene are discretised:

$$B_i = \frac{1}{A_i} \int_{S_i} B(x) dA_x. \quad (3)$$

The first algorithm we will study is the straightforward simulation of light particle paths with bucketing according to (1) and (3) as proposed by Pattanaik et al. in 1992 [9]. This algorithm will be referred to below as "continuous collision shooting random walks", for reasons explained below.

Traditionally, the integral equation (1) is reduced to a system of linear equations using Galerkins projection method with piecewise constant basis functions:

$$B'_i = E_i + \rho_i \sum_j F_{ij} B'_j. \quad (4)$$

The factors F_{ij} are called *patch-to-patch form factors*:

$$F_{ij} = \frac{1}{A_i} \int_{S_i} \int_{S_j} G(x,y) dA_y dA_x. \quad (5)$$

The coefficients B'_i that result after solving the system of linear equations (4) are only an *approximation* for the average radiosities (3) above. The difference is that discretisation errors get propagated, resulting in diffuse *reflections of* for instance light leaks. It is possible to construct scenes in which the difference is visible, but such cases are very rare in practice. We will denote both the average radiosity (3) and the radiosity coefficients in (4) by B_i in the remainder of this text.

The major problem with the system of linear equations (4) concerns the computation and storage of the form factors F_{ij}. Not only is the number of form factors huge (10,000,000,000 for a scene with 100,000 patches), but the form factor integral is moreover non-trivial to compute due to potential

discontinuities (changing visibility) and singularities (abutting patches). They are however all positive and their sum $\sum_j F_{ij} \leq 1$, so that they form a probability distribution function. This probability distribution function can be sampled efficiently by tracing rays with uniformly chosen origin on i and cosine distributed directions. Such sampling can be used in so called discrete random walks, introduced by Sbert et al. [10], or in stochastic adaptations of Jacobi's iterative method for linear systems, as proposed by Neumann et al. [8].

All three strategies mentioned above share the property that the numerical value of form factors is never needed, so that explicit form factor computation and storage is avoided. This leads to more reliable radiosity algorithms that require less storage. Moreover, it turns out the the resulting algorithms are easier to implement and to use and that they can be much more rapid than other radiosity algorithms. An overview with plenty of references can be found in [2, 1].

In previous work [3], we have performed an empirical comparison of the three algorithms mentioned above. Our experiments suggest that the accuracy of the three algorithms is in practice identical, for a fixed number of shot rays (ray shooting is the most expensive operation in the algorithms). An intuitive argument was used to explain this result. In this paper, we provide a thorough theoretical explication, by calculating the *variance* of the underlying Monte Carlo estimators.

Before doing so, we remind the reader that the basic idea of Monte Carlo methods is to formulate a quantity to be computed as the *expected value* of a *random variable*. The mean of independent samples of the random variable yields an estimate for its expected value and thus for the quantity to be computed. A random variable is a set of possible *outcomes*, say a_i, with associated probabilities p_i that indicate the chance that the outcome will show up in a random trial. The outcomes can form a discrete (countable) or a continuous set. The expected value of a discrete random variable $\hat{A} = (a_i, p_i), i = 1 \ldots n$ (n is the number of potential outcomes), is defined as

$$E[\hat{A}] = \sum_{i=1}^{n} a_i p_i. \tag{6}$$

The variance is the mean square deviation of the outcomes from the expected value:

$$V[\hat{A}] = \sum_{i=1}^{n} \left(a_i - E[\hat{A}]\right)^2 p_i = \sum_{i=1}^{n} a_i^2 p_i - E[\hat{A}]^2. \tag{7}$$

The expected value and variance of a *continuous* random variable are obtained by replacing the sums by an integral over the set of outcomes. It can be shown that an estimate with N trials will be off by less than one *standard error* $\sqrt{V[\hat{S}]/N}$ 68.3% of the time. It will be off by less than twice the standard error 95.4% of the time. The probability that the estimate is off by less than three times the standard error is 99.7%. If the variance $V[\hat{S}]$ is large, more samples (larger N) will be required in order to obtain estimates which are within a fixed threshold from the true sum with given confidence. An excellent introduction to Monte Carlo methods can be found in [6].

2 Continuous collision shooting random walks

The random walk algorithm proposed by Pattanaik et al. [9] proceeds as follows: imaginary particles are created on light sources with *birth density* $\pi(x) = E(x)/\Phi_T$ proportional to self-emitted radiosity. Φ_T denotes the total self-emitted power. The propagation of these particles throughout the environment is sampled in a two stages: first, a *survival* test is done, in which the probability of survival $\sigma(x) = \rho(x)$ is taken equal to the reflectivity. If the particle survives the test, a ray originating at x with cosine distributed direction Θ_x w.r.t. the surface normal at x is shot. The nearest point $y = h(x, \Theta_x)$ along the ray on a surface becomes the new position of the particle. It can be shown that the new position y is chosen with *transition density* $p(x, y) = \rho(x)G(x, y)$. If the particle does not survive, it is absorbed. The absorption probability is $\alpha(x) = 1 - \sigma(x)$. No survival testing is performed at the origin (on the light sources) of the particles. An estimate for the average non-self-emitted radiosity $b_k = B_k - E_k$ on each patch k is obtained by counting the number C_k of particle visits to the patch. The origin of the particles does not count (source term estimation suppression). With

N random walks:

$$\tilde{b}_k = \frac{\Phi_T}{N}\frac{\rho_k}{A_k}C_k \approx b_k. \tag{8}$$

The resulting random walks are called *continuous* because they are simulated in a continuous state space (points x on surfaces S of the scene). The estimator above is called a *collision* estimator, because the number of collisions of particles is counted. Alternatively, one could only count particles when they are absorbed, or only when they survive a collision. The algorithm has the physical interpretation of particles being shot from light sources, hence the description *continuous collision shooting random walk*. In the remainder of this section, we will show that the variance $V[\hat{b}_k^{\mathrm{CRW}}]$ of the corresponding random walk estimator is to good approximation given by:

$$V[\hat{b}_k^{\mathrm{CRW}}] \approx \frac{\rho_k}{A_k}\Phi_T b_k. \tag{9}$$

We will present a general derivation, which is also valid for calculating the variance of other global illumination algorithms based on continuous random walks such as density estimation or stochastic ray tracing.

Proof: Consider a second kind Fredholm equation

$$u(x) = e(x) + \int_S k(x,y)u(y)dy$$

with positive and contractive kernel $\int_S k(x,y)dy < 1, \forall x \in S$. Scalar products

$$W = \int_S w(x)u(x)dx \tag{10}$$

can then be estimated by means of random walks $J = x_0, x_1, \ldots, x_\tau$ as follows: x_0 is sampled with birth density $\pi(x_0)$ with cannot vanish if $w(x_0) \neq 0$. Subsequent collisions $x_k, k = 1, \ldots, \tau$ are obtained from previous collision x_{k-1} with transition density $p(x_{k-1}, x_k)$ using a survival/propagation two-step sampling procedure similar to the one outlined above, the survival probabilities being $\sigma(x_{k-1})$ and the transition probabilities shall not vanish when the kernel $k(x_{k-1}, x_k)$ doesn't. The probability associated with such a random walk is $p(J) = \pi(x_0)p(x_0,x_1)\cdots p(x_{\tau-1},x_\tau)\alpha(x_\tau)$. Associate with such a random walk a score

$$\begin{aligned} s(J) &= \frac{w(x_0)}{\pi(x_0)}e(x_0) \\ &+ \frac{w(x_0)}{\pi(x_0)}\frac{k(x_0,x_1)}{p(x_0,x_1)}e(x_1) + \cdots \\ &+ \frac{w(x_0)}{\pi(x_0)}\frac{k(x_0,x_1)}{p(x_0,x_1)}\cdots\frac{k(x_{\tau-1},x_\tau)}{p(x_{\tau-1},x_\tau)}e(x_\tau). \end{aligned}$$

The interpretation of this random walk estimator is that of gathering: the random walks originate at a region of interest (determined by $w(x)$) and each time they visit a point x with non-zero source term $e(x) \neq 0$, a score is contributed. Using straightforward calculation according to the definition (6), it can be shown that the expected value of the scores $s(J)$ above for all possible random walks J is $E[s] = W$.

Now consider the following scores of random walks with fixed origin $x_0 = z$:

$$\begin{aligned} s_z(J|_{x_0=z}) &= e(z) + \frac{k(z,x_1)}{p(z,x_1)} \\ &\times \left(e(x_1) + \frac{k(x_1,x_2)}{p(x_1,x_2)}\left(e(x_2) + \cdots\right)\right). \end{aligned}$$

By construction, the expected value of these scores is $E[s_z] = u(z)$ (take $w(x) = \pi(x) = \delta(x-z)$ with $\delta(x-z)$ Diracs delta function). We can thus re-write (10) like

$$E[s] = \int_S \pi(x)\frac{w(x)}{\pi(x)}E[s_x]dx = E[E[\frac{w}{\pi}s_x]].$$

An easy to prove and well-known property of such multivariate random variables states that then:

$$V[s] = V[E[\frac{w}{\pi}s_x]] + E[V[\frac{w}{\pi}s_x]]$$

leading to:

$$V[s] = \int_S \frac{w^2(x)}{\pi(x)}\left(V[s_x] + u^2(x)\right)dx - W^2.$$

On the other hand, a random walk originating at $x_0 = z$ is either directly absorbed, with probability $\alpha(x_0)$, or propagated to a different point x_1 with probability $p(x_0,x_1)$. In the former case, the score is $e(x_0)$, in the latter case, the expected score with the first two collisions x_0 and x_1 fixed, is $e(x_0) + \frac{k(x_0,x_1)}{p(x_0,x_1)}E[s_{x_1}]$. By explicitly introducing a *death state* Δ, with properties $k(x,\Delta) = k(\Delta,x) = 0$, $p(x,\Delta) = \alpha(x)$, $p(\Delta,x) = 0$, $p(\Delta,\Delta) = 1$, the expected value $E[s_{x_0}]$ can again be written as an expected value, over all x_1 and Δ of the expected value of a second random variable related with the score of random walks originating at x_1 [7]:

$$E[s_{x_0}] = \int_{S\cup\Delta} p(x_0,x_1)\left(e(x_0) + \frac{k(x_0,x_1)}{p(x_0,x_1)}E[s_{x_1}]\right)dx_1.$$

The same relation $E[X] = E[E[Y]] \Longrightarrow V[X] = V[E[Y]] + E[V[Y]]$ then yields:

$$v(x) = \mu(x) + \int_S \frac{k^2(x,y)}{p(x,y)}v(y)dy \tag{11}$$

$$v(x) = V[s_x] + u^2(x) \tag{12}$$

$$\mu(x) = e(x)\left[e(x) + 2(u(x) - e(x))\right]. \tag{13}$$

Now consider Greens function $\mathcal{G}^*(x,z)$ for the equation (11):

$$\begin{aligned} \mathcal{G}^*(x,z) &= \delta(z-x) + \frac{k^2(x,z)}{p(x,z)} \\ &+ \int_S \frac{k^2(x,y_1)}{p(x,y_1)}\frac{k^2(y_1,z)}{p(y_1,z)}dy_1 + \cdots \end{aligned}$$

The functions $v(x)$ can then be written as

$$v(x) = \int_S \mathcal{G}^*(x,z)\mu(z)dz$$

and the variance of the random walk estimator for W becomes:

$$V[s] = \int_S \mu(z) \int_S \mathcal{G}^*(x_0,z)\frac{w^2(x_0)}{\pi(x_0)}dx_0 dz - W^2 \tag{14}$$

With source term estimation suppression, the scores $s(J)$ are modified to reflect suppression of absorption at the origin by multiplying with the survival probability, and the fact that no counting happens at the origin:

$$\tilde{s}(J) = \frac{w(x_0)}{\pi(x_0)}\sigma(x_0)\frac{k(x_0,x_1)}{p(x_0,x_1)} \times \left(e(x_1) + \frac{k(x_1,x_2)}{p(x_1,x_2)}\left(e(x_2)+\cdots\right)\right).$$

The expected value is

$$\tilde{W} = \int_S w(x)\left(u(x) - e(x)\right)dx.$$

Using the same reasoning as before, the variance can be shown to be

$$V[\tilde{s}] = \int_S \frac{w^2(x_0)}{\pi(x_0)}\sigma(x_0)\int_S \frac{k^2(x_0,x_1)}{p(x_0,x_1)}v(x_1)dx_1 dx_0 - \tilde{W}^2.$$

Since

$$\int_S \frac{k^2(x_0,x_1)}{p(x_0,x_1)}\mathcal{G}^*(x_1,z)dx_1 = \mathcal{G}^*(x_0,z) - \delta(z-x_0),$$

the variance $V[\tilde{s}]$ can also be written as

$$V[\tilde{s}] = \int_S \mu(z)\int_S \left(\mathcal{G}^*(x_0,z) - \delta(z-x_0)\right) \frac{w^2(x_0)}{\pi(x_0)}\sigma(x_0)dx_0 dz - \tilde{W}^2. \tag{15}$$

In the important special case that the transition density $p(x,y) = k(x,y)$ corresponds to the kernel, $\mathcal{G}^*(x,z) = \mathcal{G}(x,z)$ is Greens function for the integral equation to be solved itself. According to its definition,

$$u(x) = \int_S \mathcal{G}(x,z)e(z)dz. \tag{16}$$

In the case of radiosity, the kernel $k(x,y) = \rho(x)G(x,y)$. It can be shown easily that Greens function then also fulfills

$$\mathcal{G}(x,z)\rho(z) = \mathcal{G}(z,x)\rho(x). \tag{17}$$

The relations (16) and (17) will be helpful in order to get rid of Greens function in the resulting variance expression.

The random walk estimator discussed so far in the proof is a gathering random walk estimator. A shooting random walk estimator however corresponds with a gathering random walk estimator for an adjoint integral equation. First, it is clear that the average radiosity (3) is a scalar product like (10):

$$B_k = \int_S V(x)B(x)dA_x$$

with $V(x) = \chi_k(x)/A_k$, $\chi_k(x)$ being the characteristic function of patch k ($\chi_k(x) = 1$ if $x \in S_k$ and 0 otherwise). Such scalar products can also be obtained as

$$B_k = \int_S I(x)E(x)dA_x$$

where $I(x)$ is the solution of the *adjoint* radiosity integral equation

$$I(x) = V(x) + \int_S I(y)\rho(y)G(y,x)dA_y.$$

Or alternatively:

$$\begin{aligned} B_k &= \int_S I(x)\rho(x)\frac{E(x)}{\rho(x)}dA_x \\ I(x)\rho(x) &= V(x)\rho(x) + \int_S \rho(x)G(x,y)I(y)\rho(y)dA_y. \end{aligned}$$

Filling in $k(x,y) = p(x,y) = \rho(x)G(x,y)$, $\sigma(x) = \rho(x)$, $\pi(x) = E(x)/\Phi_T$, $w(x) = E(x)/\rho(x)$, $e(x) = V(x)\rho(x)$, $u(x) = I(x)\rho(x)$, and using (17) and (16) to get rid of Greens function, equation (15) immediately leads to the result:

$$V[\hat{b}_k^{\mathrm{CRW}}] = \Phi_T \int_S \left\{\rho(z)\left[V(z) + 2\left(I(z) - V(z)\right)\right] V(z)b(z)\right\}dA_z - \left(\int_S V(z)b(z)dA_z\right)^2. \tag{18}$$

Filling in $V(z) = \chi_k(z)/A_k$ and considering that $I(z) - V(z) \ll V(z)$ and $A_k b_k \ll \Phi_T$, then yields (9).

The variance of other Monte Carlo rendering algorithms exploiting random walks with brdf times cosine transition sampling fulfills the same relation. The difference between algorithms such as light path tracing and density estimation, the photon map, random walk algorithms for higher order approximations, and instant-radiosity like final gathering algorithms is mainly in the definition of the direct potential term $V(z)$ [1].

3 Discrete collision shooting random walks

Just like a continuous random walk can be used in order to solve a second kind Fredholm integral equation with contractive kernel, a discrete random walk can be used in order to solve a system of linear equations with contractive matrix. The states visited by the random walk now correspond with

components of the solution, the patches i of the scene, rather than surface points x. Sbert et al. [10] has proposed the following algorithm for computing the power $P_k = A_k B_k$ emitted by a patch i. Random walks are generated by sampling their origin j_0 corresponding to a light source patch, with birth density $\pi_i = \Phi_i/\Phi_T$, $\Phi_i = A_i E_i$ the self-emitted power. Next, transitions are sampled by a two-step procedure similar to continuous random walks: first, a survival test is performed with survival probability $\sigma_i = \rho_i$ equal to the reflectivity. If the random walk survives, a transition is sampled to a patch j with probability corresponding to the form factor F_{ij}. The combined transition density is $p_{ij} = \rho_i F_{ij}$. Sampling transitions according to the form factor is done by choosing a uniform random ray origin x on the source patch i, and a cosine distributed direction Θ_x w.r.t. the surface normal at x. It can be shown easily that the patch containing the first hit point of this ray with the surfaces in the scene is selected with probability equal to the form factor. Basically by counting the number of visits also here, an estimate for the non-self-emitted radiosity on each patch is obtained:

$$\rho_k \frac{\Phi_T}{N} C_k \approx A_k b_k.$$

The difference between the continuous and discrete random walk algorithms is very small in practice: in the discrete random walk algorithm, a particle hitting a patch i is reflected from a uniform random other location x on the patch i, whereas in the continuous random walk algorithm, it is reflected from the point of incidence.

The variance of this discrete random walk estimator has been calculated before by Sbert et al. [10], or can be derived using the same reasoning as in the previous section (as in [2]). The result is:

$$V[\hat{b}_k^{\text{DRW}}] = \frac{\rho_k}{A_k}\Phi_T(1+2\zeta_k)b_k - b_k^2 \quad (19)$$

$$\approx \frac{\rho_k}{A_k}\Phi_T b_k \quad (20)$$

In this equation, ζ_k is the recurrent potential at k: consider k the only source of potential: $V_i = \delta_{ki}$. Then $\zeta_k = I_k - V_k$. In practice, $\zeta_k \ll 1$ and $A_k b_k \ll \Phi_T$. These approximations lead to (20), which is identical as for a continuous collision shooting random walk (9).

4 Stochastic Jacobi iterations

Consider the power system of equations, obtained by multiplying both sides of (4) with the patch area A_i:

$$P_i = \Phi_i + \sum_j P_j F_{ji} \rho_i.$$

The basic idea of the Jacobi iterative method for solving this linear system is to construct a sequence of approximations $P_i^{(k)}$ for P_i, eventually leading to the right solution. A convenient starting guess is to take $P_i^{(0)} = \Phi_i$. A subsequent approximation $P^{(k+1)}$ is then obtained from a previous approximation $P^{(k)}$ by filling $P^{(k)}$ into the right hand side of the equations above. Each resulting iteration can also be described as a double sum, by introducing Kroneckers delta function $\delta_{li} = 1$ if $l = i$ and 0 if $l \neq i$:

$$P_i^{(k+1)} = \Phi_i + \sum_{j,l} P_j^{(k)} F_{jl} \rho_l \delta_{li}. \quad (21)$$

These double sums can be estimated using Monte Carlo as follows [8, 2]:

1. Select terms (j, l) (pairs of patches) as follows: first select a patch j with probability p_j proportional to $P_j^{(k)}$. Next select a patch l conditional on j with probability $p_{l|j} = F_{jl}$ equal to the form factor. The latter can be done easily again by tracing a ray with uniform random chosen origin on j and cosine distributed direction. The combined probability of sampling the pair (j, l) is

$$p_{jl} = \frac{P_j^{(k)}}{P_T^{(k)}} F_{jl} \quad \text{with:} \quad P_T^{(k)} = \sum_j P_j^{(k)};$$

2. The score associated with each sampled term is

$$\tilde{P}_i^{(k+1)} = \frac{P_j^{(k)} F_{jl} \rho_l \delta_{li}}{p_{jl}} = \rho_l P_T^{(k)} \delta_{li}$$

Using N rays, a new power estimate $P_i^{(k+1)}$ is obtained simultaneously on all patches i by basically counting the fraction N_i/N of the rays that lands on i.

Unlike for random walk estimators, the variance of the above estimator is extremely easy to calculate

according to the definition of variance, formula (7). It yields:

$$V[\hat{b}_i^{(k+1)}] = \frac{\rho_i}{A_i} P_T^{(k)} b_i^{(k+1)} - \left(b_i^{(k+1)}\right)^2. \quad (22)$$

Now suppose the input radiosities to the iteration $B_i^{(k)} \approx B_i$ are close to the true solution already, then the output radiosities will also be close to the solution. The variance on the output radiosities then is approximately equal to

$$V[\hat{b}_i^{SJ}] \approx \frac{\rho_i}{A_i} P_T b_i. \quad (23)$$

Proof of (22): According to the definition, the variance $V[\hat{P}_i^{(k+1)} - \Phi_i]$ is

$$V = \sum_{j,l} \frac{\left(P_j^{(k)} F_{jl} \rho_l \delta_{li}\right)^2}{P_j^{(k)} F_{jl} / P_T^{(k)}} - \left(P_i^{(k+1)} - \Phi_i\right)^2$$

The double sum immediately simplifies to

$$\rho_i P_T^{(k)} \sum_j P_j^{(k)} F_{ji} \rho_i = \rho_i P_T^{(k)} \left(P_i^{(k+1)} - \Phi_i\right).$$

Division by A_i^2 yields the required result.

In practice, such a stochastic Jacobi iteration transforms an approximation for the power distribution in a scene into a new, almost entirely independent, one. Rather than replacing the previous approximation with the new one, the approximations are averaged. In this way, variance is progressively reduced. A first approximation, sufficiently close to the solution in order to make such averaging work correctly, can be obtained by a sequence of similar stochastic Jacobi iterations in which *unshot* power is propagated rather than total power, similar to in the progressive refinement radiosity method [4]. The result of these iterations is added together, rather than averaged, until the unshot power drops below a small threshold. In this way, a first complete, but possibly very noisy radiosity solution is obtained. The iterations steps as described above are then used in order to progressively reduce variance. It is very easy to prove that, when the number of samples is chosen proportional to the unshot power to be propagated in each incremental iteration, the variance on the first complete radiosity solution will be identical to the variance of a regular iteration with same total number of samples [2, 1].

5 Discussion and Validation

According to (9) and (20), the variance with N^{RW} random walks is approximately equal to

$$\frac{V^{RW}}{N^{RW}} \approx \frac{\Phi_T}{N^{RW}} \frac{\rho_k}{A_k} b_k.$$

The variance of a stochastic Jacobi iteration with N^{SJ} rays is, according to (22)

$$\frac{V^{SJ}}{N^{SJ}} \approx \frac{P_T}{N^{SJ}} \frac{\rho_k}{A_k} b_k.$$

It can be shown easily that the mean number of rays traced for simulating N random walks, continuous or discrete, is NP_T/Φ_T. For an equal number of rays, choosing $N^{SJ} = N^{RW} P_T/\Phi_T$, the variance with collision shooting random walks, both discrete or continuous, and stochastic Jacobi iterations will therefore be equal to good approximation.

There are however minor differences: first, we already explained in the introduction that a continuous and discrete random walk do not solve the same result. The difference in the variance of a continuous or discrete random walk is of the same order: rarely noticeable in practice because propagation of discretisation error is negligible.

A discrete random walk, and the stochastic Jacobi method, both do solve for the same result. There is a clear, and more interesting, difference between the variance of these methods in scenes with high recurrent potential ζ_k. The exact expressions for the variances are:

$$\begin{aligned} \frac{V^{RW}}{N^{RW}} &= \frac{1}{N^{RW}} \left(\frac{\rho_k}{A_k} \Phi_T (1 + 2\zeta_k) b_k - b_k^2\right) \\ \frac{V^{SJ}}{N^{SJ}} &= \frac{1}{N^{SJ}} \left(\frac{\rho_k}{A_k} P_T b_k - b_k^2\right). \end{aligned}$$

We have validated these expressions by measuring the mean square error in a simple cubical scene with unit square faces. The reflectivities and emissivities are chosen constant so that $\rho + E = 1$. In such a scene, the radiosity $B = 1$ can be computed analytically, and so are the variance expressions above:

- $A_k = 1$, the total area $A_T = 6$;
- $\Phi_T = 6E = 6(1 - \rho)$;
- $b_k = \rho$;
- $P_T = \Phi_T/(1 - \rho)$, so that N^{RW} must be chosen $N^{SJ}(1 - \rho)$.

ρ	1/10	1/3	1/2	2/3	9/10	19/20
ζ	0.00218	0.03125	0.0909	0.2353	1.3729	3.0336
RW (theory)	0.04915	0.54167	1.2727	2.588	10.076	20.218
RW (experiment)	0.0477	0.532	1.236	2.639	10.55	20.088
SJ (theory)	0.05	0.5555	1.25	2.2222	4.05	4.5125
SJ (experiment)	0.0500	0.554	1.198	2.180	3.887	4.323

Table 1: Observed mean square error (MSE) per ray for a selected patch in a homogeneous cube with unit sides and $\rho + E = 1$. There is a good correspondence between the empirical and the predicted values for both the stochastic Jacobi method and the discrete collision shooting random walk. RW=Random Walk, SJ=Stochastic Jacobi.

With these choices, the variance expressions become:

$$\frac{V^{RW}}{N^{RW}} = \frac{1}{N^{SJ}}\left[6\rho^2(1+2\zeta) - \frac{\rho^2}{1-\rho}\right]$$

$$\frac{V^{SJ}}{N^{SJ}} = \frac{1}{N^{SJ}}5\rho^2$$

$\zeta = I_k - V_k$ can be determined by solving

$$I_i = V_i + \sum_j F_{ij}\rho_j I_j$$

analytically with $V_i = \delta_{ik}$, taking values 0.2 for the form factors between different patches. The true values of the form factors are 0.200043 for abutting patches and 0.1998 for parallel patches in the cube. This yields the following solution for ζ:

$$\zeta = \frac{0.2\rho^2}{1 - 0.2\rho(4+\rho)}.$$

Filling in these values for ζ leads to the theoretical values shown in table 1.

The reported empirical results are the average mean square error (MSE) per ray observed on a selected patch. They are the average MSE, after a sufficiently high number of runs (more than 20,000 in all cases) with $N^{SJ} = 10000$ rays per run, multiplied with N^{SJ}. The result for the stochastic Jacobi iterative method have been obtained with the analytical solution $B = 1$ as input distribution. The experiment has also been repeated with the (inexact) result of a previous regular stochastic Jacobi iteration as the input for the next iteration. The observed mean square errors were not significantly different.

The difference in variance between discrete random walk and stochastic Jacobi iterations is due to a subtle difference in survival decisions. Survival decisions are made independently for each random walk in the random walk method, while they are made "in group" in the stochastic Jacobi method. Consider that 10 particles have landed on a patch with reflectivity 0.45 in the stochastic Jacobi method. In a subsequent iteration, the number of particles that will be shot from this patch will be 4 or 5, with an average of 4.5. With random walks, the average number of surviving particles also will be 4.5 but it can be anything from 0 to 10 in an actual experiment. The difference is however only noticeable in scenes where particles have a high chance of visiting the same patch twice or more. In practice, this is very rarely the case.

6 Conclusion

In previous work [3], an empirical study has shown that the accuracy for fixed number of samples (rays) is identical with continuous collision shooting random walks, discrete collision shooting random walks, and with the stochastic Jacobi algorithm. Only an intuitive argument was given for this observation. In this paper, we have calculated the variance of the underlying Monte Carlo estimators in these algorithms, and we have shown that for the same number of ray, the expected square error will indeed be the same, to good approximation.

Our previous study also indicated that low-discrepancy sampling is more effective in the discrete algorithms than in the continuous random walk algorithm. A theoretical explication for this observation has not been found. To the contrary: the theory of low-discrepancy sampling rather indicates that the lower-dimensional sampling in a continuous random walk algorithm should be more efficient. The theory however only provides very conservative error bounds that are valid asymptotically,

for very large number of samples. We suspect that the number of samples required in radiosity is not high enough to exhibit the asymptotic properties.

The discussion in this paper also made clear that the analysis of the stochastic Jacobi iterative method is considerably easier than that of random walk methods. Compare the length of the proofs in §2 and §4! Not only is the analysis of the stochastic Jacobi iterative method itself easier, but also the design and analysis of variance reduction techniques, such as view-importance sampling, correlated sampling, the combination of shooting and gathering, and weighted importance sampling appears to be far more easy, and more effective than for random walks.

Stochastic Jacobi iterations have been studied extensively for diffuse reflection. An overview with pointers to literature can be found in [1]. With the exception of the work by Szirmay-Kalos, they have to date hardly been investigated for general light emission and scattering. With an appropriate solution for the problem of storing non-diffuse illumination information however, we believe that the advantages of stochastic Jacobi iterations will carry over to the general case as well.

References

[1] Ph. Bekaert. Stochastic radiosity: doing radiosity without form factors. SIGGRAPH'2001 course notes, Course 20, "Advanced Global Illumination", edited by Ph. Dutré and K. Bala, Los Angeles, August 2001, 52 pages.

[2] Ph. Bekaert. *Hierarchical and Stochastic Algorithms for Radiosity*. PhD thesis, K. U. Leuven, department of computer science, December 1999.

[3] Ph. Bekaert, R. Cools, and Y. D. Willems. An empirical comparison of Monte Carlo radiosity algorithms. In *7th International Conference in Central Europe on Computer Graphics, Visualization and Digital Interactive Media, Plzen, Czech Republic*, pages 9–16, February 1999.

[4] M. F. Cohen, S. E. Chen, J. R. Wallace, and D. P. Greenberg. A progressive refinement approach to fast radiosity image generation. In *Computer Graphics (SIGGRAPH '88 Proceedings)*, volume 22, pages 75–84, August 1988.

[5] J. T. Kajiya. The rendering equation. *Computer Graphics (SIGGRAPH '86 Proceedings)*, 20(4):143–150, August 1986.

[6] M. H. Kalos and P. Whitlock. *The Monte Carlo method*. J. Wiley and sons, 1986.

[7] C. Kollman, K. Baggerley, D. Cox, and R. Picard. Adaptive importance sampling on discrete markov chains. Technical Report LA-UR-96-3998, Los Alamos National Lab, November 1996.

[8] L. Neumann, W. Purgathofer, R. Tobler, A. Neumann, P. Elias, M. Feda, and X. Pueyo. The stochastic ray method for radiosity. In P. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95 (Proceedings of the Sixth Eurographics Workshop on Rendering)*, July 1995.

[9] S. N. Pattanaik and S. P. Mudur. Computation of global illumination by Monte Carlo simulation of the particle model of light. *Third Eurographics Workshop on Rendering*, pages 71–83, May 1992.

[10] M. Sbert. Error and complexity of random walk Monte Carlo radiosity. *IEEE Transactions on Visualization and Computer Graphics*, 3(1):23–38, March 1997.