

ANALYSIS OF THE QUASI-MONTE CARLO INTEGRATION OF THE RENDERING EQUATION

László Szirmay-Kalos, Werner Purgathofer

Department of Control Engineering and Information Technology, TU of Budapest
Budapest, Műegyetem rkp. 11, H-1111, HUNGARY
szirmay@fsz.bme.hu

ABSTRACT

Quasi-Monte Carlo integration is said to be better than Monte-Carlo integration since its error bound can be in the order of $O(N^{-(1-\epsilon)})$ instead of the $O(N^{-0.5})$ probabilistic bound of classical Monte-Carlo integration if the integrand has finite variation. However, since in computer graphics the integrand of the rendering equation is usually discontinuous and thus has infinite variation, the superiority of quasi-Monte Carlo integration has not been theoretically justified. This paper examines the integration of discontinuous functions using both theoretical arguments and simulations and explains what kind of improvements can be expected from the quasi-Monte Carlo techniques in computer graphics.

Keywords: Rendering equation, quasi-Monte Carlo quadrature, Hardy-Krause variation

1 Introduction

The fundamental task of computer graphics is to solve a Fredholm type integral equation describing the light transport. This equation is called *the rendering equation* and has the following form:

$$L(\vec{x}, \omega) = L^e(\vec{x}, \omega) + (\mathcal{T}L)(\vec{x}, \omega) \quad (1)$$

where $(\mathcal{T}L)(\vec{x}, \omega)$ is

$$\int_{\Omega'} L(h(\vec{x}, -\omega'), \omega') \cdot \cos \theta' \cdot f_r(\vec{x}, \omega', \omega) d\omega',$$

$L(\vec{x}, \omega)$ is the radiance of the surface in point \vec{x} at direction ω , $h(\vec{x}, \omega')$ is the visibility function defining the point that is visible from point \vec{x} at direction ω' , θ' is the angle between the surface normal and direction ω' , and $f_r(\vec{x}, \omega', \omega)$ is the bi-directional reflection/refraction function.

In equation (1) the unknown function L appears on both sides. Recursively substituting

the right side's L by $L^e + \mathcal{T}L$, which is obviously L according to the equation, we get the following Neumann series:

$$L = L^e + \mathcal{T}L = L^e + \mathcal{T}(L^e + \mathcal{T}L) = \\ L^e + \mathcal{T}L^e + \mathcal{T}^2L = \sum_{i=0}^{\infty} \mathcal{T}^i L^e \quad (2)$$

if integral operator \mathcal{T} is contractive.

The terms of this infinite Neumann series have intuitive meaning as well: $\mathcal{T}^0 L^e = L^e$ comes from the emission, $\mathcal{T}^1 L^e$ comes from a single reflection, $\mathcal{T}^2 L^e$ from two reflections, etc.

In order to compute $\mathcal{T}^n L^e(\vec{x}, \omega)$, the following function should be integrated: A ray is emanated recursively from \vec{x} at direction $-\omega'_1$, then from the found surface at $-\omega'_2$, etc. until $-\omega'_n$. The emission intensity at the end of the walk is read and multiplied by the BRDFs and the cosine terms of the stages of the walk.

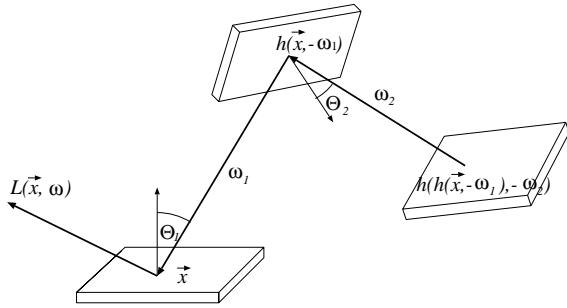


Figure 1: The integrand of $\mathcal{T}^2 L^e$ is a two-step walk

Since classical quadrature techniques based on uniform grids, such as trapezoidal rule, may have integration error of order $O(1/N^{1/s})$ where N is the number of dot points and s is the dimension of the domain, in order to guarantee a predefined error, the number of dot points should be in the order of $O(N^s)$, which is prohibitive for high-dimensional integrals. This phenomenon is called as the *dimensional explosion*. The dimensional explosion can be avoided by Monte-Carlo or Quasi-Monte Carlo integration methods. Monte-Carlo methods trace back the estimation of an integral to the calculation of an expected value, which is estimated by averaging random samples. Quasi-Monte Carlo techniques, on the other hand, use deterministic samples that are uniformly distributed in the integration domain.

Monte-Carlo quadrature considers the integral as an expected value which can be estimated by averaging random points:

$$\int_V f(\mathbf{z}) d\mathbf{z} = V \cdot \int_V f(\mathbf{z}) \cdot \frac{1}{V} d\mathbf{z} = V \cdot E[f(\mathbf{z})] \quad (3)$$

since $1/V$ can be considered as the probability density of a random variable uniformly distributed in V .

If N number of points are used to estimate an integral of f in the domain of size V , the Monte-Carlo quadrature will have the following (stochastic) error bound:

$$\int_V f(\mathbf{z}) d\mathbf{z} \approx \frac{V}{N} \cdot \sum_{i=1}^N f(\mathbf{z}_i) \pm V \cdot \frac{\sigma}{\sqrt{N}}. \quad (4)$$

If $f(\mathbf{z})$ is square-integrable function, then

variance σ is finite, thus the stochastic error bound will be in the order of $O(N^{-0.5})$.

Quasi-Monte Carlo integration, on the other hand, uses deterministic point-sets that contain very uniformly distributed points in the integration domain. For the normalized, s -dimensional integration domain $[0, 1]^s$, the quasi-Monte Carlo approximation is:

$$\int_{[0,1]^s} f(\mathbf{x}) d\mathbf{x} \approx \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_i). \quad (5)$$

Sample points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ should be selected to minimize the error of the integral quadrature. Intuitively this error must depend on two independent factors. On the one hand, if the distribution of the sample points is not uniform, then there are large regions where there are no sample point at all, which increases the error. From mathematical point of view, the uniformity is characterized by the discrepancy of a point set. The star-discrepancy is defined by

$$D^*(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) = \sup_A \left| \frac{m(A)}{N} - V(A) \right| \quad (6)$$

where A is an s -dimensional sub-cube parallel to the coordinate axes and originating at the center, $V(A)$ is its volume, and $m(A)$ is the number of sample points inside this subcube. For carefully selected sample points, called *low-discrepancy sequences* [Niede92], the discrepancy can be in the order of $O(\log^{s-1} N/N)$.

On the other hand, the error of quadrature also depends on how quickly the function changes between the sample points. If the function can change significantly in small domain, then the error can be quite large. However, if the slope of the function is small, then nothing dramatic happens between the sample points, thus the error will be small.

Measures describing how rapidly a function can change are called **variations**. For a 1-dimensional function, the *variation in the sense of Vitali* is defined as:

$$\mathcal{V}_V(f(x)) = \limsup \sum_{i=1}^n |f(x_{i+1}) - f(x_i)|. \quad (7)$$

For higher dimensions, the variation of Vitali does not provide a good measure: if the function is constant in x , for instance, then the variation is zero, regardless how the function changes depending on y . Thus, it is worth using a somehow more stronger variation type, called the *Hardy-Krause variation*. The variation in the sense of Hardy and Krause is the sum of the variations of the function and its restrictions to the end of the domain.

From these two factors — the discrepancy of the sample locations and the variation of the function — an upper-bound can be established for the error of the quadrature. If the integrand f has finite variation in the sense of Hardy and Krause, then the error of the quasi-Monte Carlo quadrature can be bounded as stated by the Koksma-Hlawka inequality:

$$\left| \int_{[0,1]^s} f(\mathbf{x}) d\mathbf{x} - \frac{1}{N} \sum_{i=1}^N f(\mathbf{x}_i) \right| \leq \mathcal{V}_{HK} \cdot D^*(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N), \quad (8)$$

where \mathcal{V}_{HK} is the variation in the sense of Hardy and Krause and $D^*(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N)$ is the star-discrepancy of the sample points [Niede92]. As mentioned, for low-discrepancy point sets the discrepancy can be in the order of $O(\log^{s-1} N/N)$, which can provide much better error bounds than the $O(1/N^{1/s})$ bound of classical quadrature rules or the $O(1/\sqrt{N})$ probabilistic error bound of Monte Carlo methods. Moreover, quasi-Monte Carlo methods guarantee this accuracy in a deterministic way, unlike Monte Carlo methods where the error bound is also probabilistic.

2 Generation of Low-Discrepancy Sequences

As a conclusion of the error analysis we can state that we need very uniform sequences for quasi-Monte Carlo quadrature. The discrepancy of the best sequences known is in the order of $O(\log^s N/N)$ or even in the order of $O(\log^{s-1} N/N)$ if N is known before starting the sequence. There are many sequences published in the literature [Niede92,

Warno95, Deák89, Knuth81]. The most famous one is probably the *Halton sequence*.

The element i of the one-dimensional Halton sequence of base b is defined as the radical inverse of the expansion of i in base b . This means that number i is expanded in radix b , then the number is mirrored onto the “radix” point.

Why is this sequence uniform? Note that the construction algorithm generates as binary form of i all binary combinations of length k before producing a combination of length $k+1$. This means that after the radical inverse the sequence H_i will visit all intervals of length 2^{-k} before putting a new point in an interval already visited.

On the other hand, as k increases, the algorithm produces a single point in each interval of length $1/2$, then in each interval of length $1/4$, etc. thus the sequence is not only asymptotically uniform, but also the first N points are fairly uniformly distributed (this is guaranteed by the property that the radical inverse makes the most significant bit the most rapidly changing bit).

Note that the same construction can be used for any base, thus in this way infinitely many different Halton sequences can be constructed. For higher dimensional Halton sequences, the interdependence of different coordinates should be as little as possible. Thus an appropriate way of the construction is selecting relative primes (say b_1, b_2, \dots, b_s) as bases for the different coordinates.

3 Integrating functions of unbounded variation

As mentioned, the integrand of the original form of the rendering equation is discontinuous where the discontinuity is not aligned with the coordinate axes, thus its variation is infinite. These discontinuities are usually produced by the projected object boundaries. This property makes the Koksma-Hlawka inequality not appropriate for the error analysis

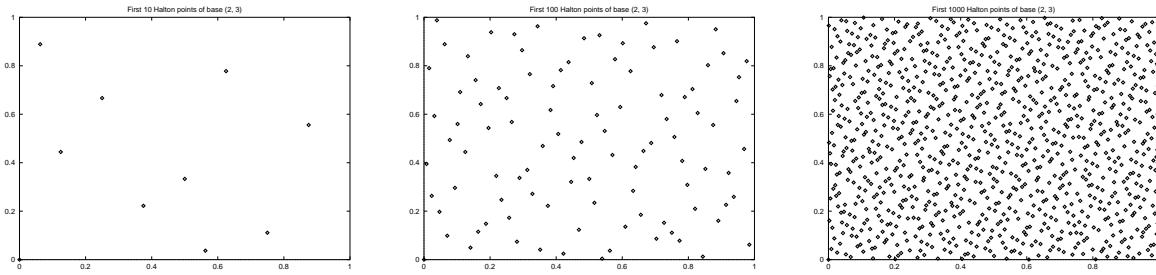


Figure 2: The distribution of the first 10,100 and 1000 Halton points in 2 dimensions

of the solution of the rendering equation and for the prediction of the convergence rates.

In this section the convergence speed is examined for functions which are generally smooth but have general discontinuities of finite “length”. First the domain of the integration is assumed to be 2-dimensional, then the results will be generalized to arbitrary dimension.

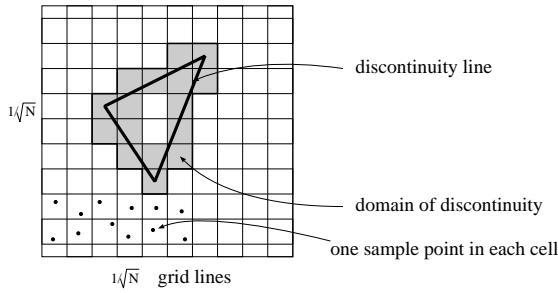


Figure 3: A typical integrand of the rendering equation

Suppose that N number of samples has been generated to estimate the integral of a function such as in figure 3 using a low-discrepancy sequence. In order to overcome the difficulty that the integrand f has infinite variation, the function is decomposed into two functions, one is smooth \tilde{f} having continuous mixed derivatives and the other \hat{f} inherits the discontinuity of f (figure 4). Low-discrepancy sequences generate points in a cellular grid in a way that the difference of the number of points in two cells is at most 1. If there are already N number of points, the size of a cell on the finest, filled level is approximately $1/\sqrt{N} \times 1/\sqrt{N}$. Let us define the domain of \hat{f} as the set of those cells that are intersected by the discontinuity. This do-

main is called the **domain of discontinuity**. The number of such cells is in the order of the “digital length” of the discontinuity curve, which is the product of the maximum extent l and the resolution of the grid \sqrt{N} . Since each cell has at least 1 (and at most 2) points, the number of points in this domain is at least $l\sqrt{N}$.

The error of quadrature is as follows:

$$\begin{aligned}
 & \left| \int_{\mathbf{z} \in [0,1]^2} f(\mathbf{z}) \, d\mathbf{z} - \frac{1}{N} \sum_{i=1}^N f(\mathbf{z}_i) \right| \leq \\
 & \left| \int_{\mathbf{z} \in [0,1]^2} \tilde{f}(\mathbf{z}) \, d\mathbf{z} - \frac{1}{N} \sum_{i=1}^N \tilde{f}(\mathbf{z}_i) \right| + \\
 & \left| \int_{\mathbf{z} \in [0,1]^2} \hat{f}(\mathbf{z}) \, d\mathbf{z} - \frac{1}{N} \sum_{i=1}^N \hat{f}(\mathbf{z}_i) \right|. \quad (9)
 \end{aligned}$$

Since \tilde{f} has finite variation, the first term in the error is bounded by

$$\mathcal{V}_{HK}(\tilde{f}) \cdot D^*(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N).$$

Concerning the second term, the integration of \hat{f} is estimated taking $l\sqrt{N}$ samples and averaging the result. Since the samples and the discontinuity are not related in any way, we can suppose that this is a normal Monte-Carlo integration [Press92]. The uniform property of low-discrepancy sequence guarantees that this pseudo-random set can be assumed to have uniform distribution.

If Δf is the difference between the maximum and minimum values in the domain of discontinuity, then $\sigma^2 \leq (\Delta f)^2$. In our case

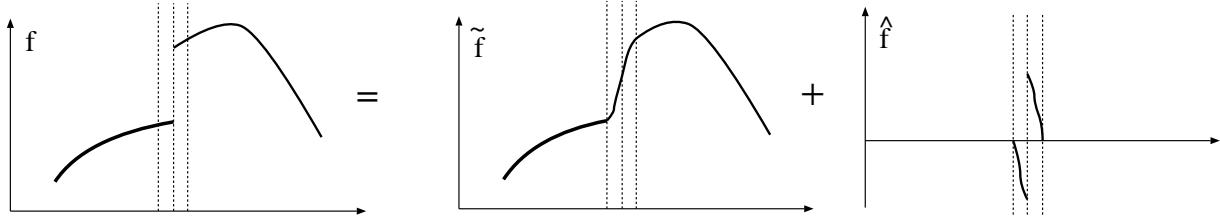


Figure 4: Decomposition of f into a smooth (\tilde{f}) and a discontinuous (\hat{f}) function

the number of sample points M is $l\sqrt{N}$ and $V = l/\sqrt{N}$, thus we obtain:

$$\begin{aligned} \int_V \hat{f}(\mathbf{z}) d\mathbf{z} &= \frac{V}{M} \cdot \sum_{i=1}^M \hat{f}(\mathbf{z}_i) \pm V \cdot \frac{\Delta f}{M} = \\ &= \frac{1}{N} \cdot \sum_{i=1}^M \hat{f}(\mathbf{z}_i) \pm \Delta f \cdot \sqrt{l} \cdot N^{-3/4}. \end{aligned} \quad (10)$$

Taking into account that \hat{f} is zero outside the domain of discontinuity, equality (9) can be expressed as:

$$\begin{aligned} \left| \int_{\mathbf{z} \in [0,1]^2} f(\mathbf{z}) d\mathbf{z} - \frac{1}{N} \sum_{i=1}^N f(\mathbf{z}_i) \right| &\leq \\ \mathcal{V}_{HK}(\tilde{f}) \cdot D^*(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N) + \Delta f \cdot \sqrt{l} \cdot N^{-3/4}. & \end{aligned} \quad (11)$$

For large N values the second term will be dominant, which results in $O(N^{-3/4})$ error bound. This is poorer than the $O(\log^2 N/N)$ bound suggested by the Koksmo-Hlawka inequality. Note that the point from where the second term dominates the first one depends on “intensity” Δf and size of the discontinuity l .

The same analysis can be carried out in higher dimensions as well. In D dimensions a discontinuity of size l would intersect $V = l \cdot N^{-1/D}$ volume of cells which would contain $M = l \cdot N^{(D-1)/D}$ sample points. Thus the general error bound is:

$$\begin{aligned} \left| \int_{\mathbf{z} \in [0,1]^D} f(\mathbf{z}) d\mathbf{z} - \frac{1}{N} \sum_{i=1}^N f(\mathbf{z}_i) \right| &\leq \\ \mathcal{V}_{HK}(\tilde{f}) \cdot D^*(\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N) + \Delta f \cdot \sqrt{l} \cdot N^{-\frac{(D+1)}{2D}}. & \end{aligned} \quad (12)$$

Thus, for quasi-Monte Carlo integration of discontinuous functions, the order of the error bound will be in between the $O(N^{-(1-\epsilon)})$

bound of finite variation functions and $O(N^{-0.5})$ stochastic bound of Monte-Carlo quadrature. The higher the dimension of the integral, the closer the two methods get in this order. However, the proportionality ratio of the bounds may still be very different. Thus it is still worth using quasi-Monte Carlo quadrature if the size of the discontinuity l is not very large, since in this case the error could be significantly less then for Monte-Carlo quadrature.

When quasi-Monte Carlo quadrature is used for the solution of the rendering equation, it should be taken into account that the Neumann series contains a series of integrals of increasing dimension. If shooting type walks are used, then the effect of higher dimensional integrals gets lower due to the dissipation of energy. Thus the degradation of the convergence for higher dimensional integrals is partly compensated by their reduced importance.

Unfortunately this does not happen for gathering type walk, since here the final power may result from any step.

4 Numerical evidence for simple functions

In order to demonstrate the previous results, the convergences of a 2-dimensional and a 3-dimensional functions are examined, that are simple enough to analytically compute their integrals.

The 2-dimensional function is:

$$f_2(x, y) = \begin{cases} (x + y - 2) \cdot a + 1 & \text{if } x + y > 1, \\ (x + y) \cdot a & \text{otherwise,} \end{cases} \quad (13)$$

where a is a free parameter in the range of $[0, 0.5]$. Note that by setting a appropriately, the intensity of the discontinuity can be controlled without altering either the value of the integral or the variation of the continuous part. If $a = 0.5$, then the function has finite variation, otherwise it has infinite variation.

The results of the simulation are shown in figure 5. This figure shows the maximum error after a given number of samples.

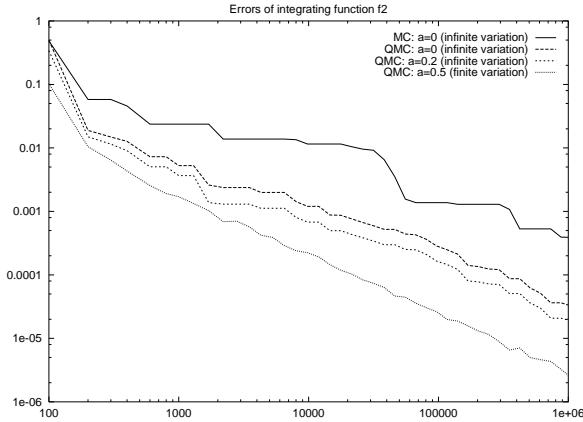


Figure 5: Error of integrating f_2

The 3-dimensional function is:

$$f_3(x, y, z) = \begin{cases} (x + y + z - 1.8) \cdot a + 0.6 & \text{if } x + y + z > 1, \\ (x + y + z) \cdot a \text{ otherwise,} \end{cases} \quad (14)$$

where a is a free parameter in the range of $[0, 1/3]$. If $a = 1/3$, then f_3 has finite variation, otherwise it has not. The error of integration of f_3 is summarized in figure 6.

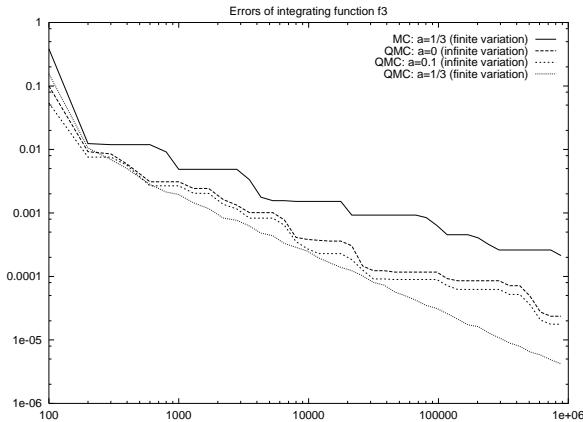


Figure 6: Error of integrating f_3

5 Numerical evidence for rendering

In order to obtain accurate error measurements for rendering algorithms, we need scenes for which the exact solution is known. Generally, we have two opportunities. Either, we use simple scenes for which analytical solution is possible, or numerical methods are applied using so many samples that the approximate solution can be accepted as a reference.

One of the famous analytically solvable case is the homogeneous environment, where the geometry is arbitrary, but all surfaces have the same diffuse reflection and emission [Shir91, Kelle95]. In this scene, however, the emission distribution is very homogeneous, thus quasi-Monte Carlo methods would be given an unfair handicap.

Thus we shall use another reference scene proposed by [Hyben98], which is an inner surface of a sphere, where the reflection is diffuse and homogeneous, but only a part of the surface has non-zero emission. Due to the not axis-aligned discontinuity of the lightsource the integrand of the rendering equation has infinite variation.

Here we derive the analytical solution somehow differently than in [Hyben98], since we also show that not only the complete solution can be obtained but also the partial solutions of different bounces. The knowledge of the contribution of different bounces allows us to use the reference solution to evaluate biased, finite random walk algorithms.

5.1 Analytical solution of the reference scene

Let the scene be an inner surface S of a sphere of radius R , the BRDF be constant $f_r = \varrho/\pi$ and the emission be diffuse and constant L^e in a subset A_e of the spherical surface and zero otherwise.

Using the

$$d\omega' = \frac{\cos \theta_{\vec{y}} \cdot d\vec{y}}{\|\vec{y} - \vec{x}\|^2}$$

substitution for the solid angle, we obtain the following form of the one bounce:

$$(\mathcal{T}L^e) = \int_S f_r \cdot \cos \theta_{\vec{x}} \cdot L^e(\vec{y}) \cdot \frac{\cos \theta_{\vec{y}}}{\|\vec{y} - \vec{x}\|^2} d\vec{y} \quad (15)$$

Since $L^e(\vec{y})$ is constant and non-zero only inside A_e we get:

$$(\mathcal{T}L^e) = f_r \cdot L^e \cdot \int_{A_e} \frac{\cos \theta_{\vec{x}} \cdot \cos \theta_{\vec{y}}}{\|\vec{y} - \vec{x}\|^2} d\vec{y}. \quad (16)$$

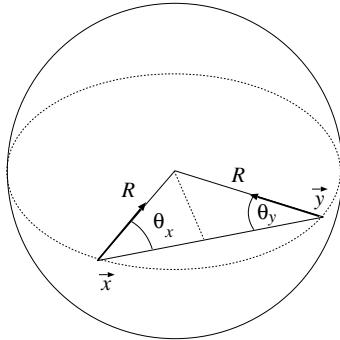


Figure 7: Geometry of the reference scene

Looking at figure 7, we can see that inside a sphere

$$\cos \theta_{\vec{x}} = \cos \theta_{\vec{y}} = \frac{\|\vec{y} - \vec{x}\|}{2R},$$

thus we can obtain the following final form for the one-bounce:

$$(\mathcal{T}L^e) = \frac{f_r \cdot L^e}{4R^2} \cdot A_e = \varrho \cdot \tilde{L}^e \quad (17)$$

where $\tilde{L}^e = L^e \cdot A_e / S$ is the average emission of the total spherical surface.

Summarizing, after the first bounce the radiance distribution is constant. Substituting this into L^e and replacing A_e by the total spherical surface S , we can derive the two bounce radiance transfer as well:

$$(\mathcal{T}^2 L^e)(\vec{x}, \omega) = \left(\frac{f_r}{4R^2} \right) \cdot S \cdot \varrho \cdot \tilde{L}^e = \varrho^2 \tilde{L}^e. \quad (18)$$

Similarly, the n bounce radiance distribution is:

$$(\mathcal{T}^n L^e)(\vec{x}, \omega) = \varrho^n \tilde{L}^e. \quad (19)$$

5.2 Error measurements

The efficiency of Monte-Carlo and quasi-Monte Carlo quadratures have been tested for the presented spherical scene assuming a single pixel camera. The error has been measured separately for the different bounces.

Looking at the error measurements of figure 8, we can see that even for integrands of infinite variation quasi-Monte Carlo methods are still better but they lose their advantage when computing higher bounces as predicted by the theoretical results. The other important problem in higher dimensions is that although a low-discrepancy series has almost linearly decreasing discrepancy in the asymptotic sense, this discrepancy can still be high for not very many points (in the solution of the rendering equation we rarely use more than 1000 samples for the estimation of a single pixel). In the case of the Halton series, for example, the *base* of the series strongly affects the initial behavior of the discrepancy. These base numbers are different prime numbers for different dimensions, thus for high-dimensional integrals the base numbers can be quite high, which results in degraded performance.

6 Conclusions

This paper reviewed the application of quasi-Monte Carlo integration to solve the rendering equation. It showed that although the integrand of the rendering equation is not of finite variation, an error analysis is still possible that combines the Koksma-Hlawka inequality with the probabilistic bounds of Monte-Carlo techniques. This error analysis leads us to the conclusion that the convergence rate of quasi-Monte Carlo integration can be better than Monte-Carlo quadrature even for the discontinuous rendering equation, especially when the required dimension — that is the average reflectivity of the scene — is not very high. The theoretical results have also been demonstrated by simulation.

7 Acknowledgments

This work has been supported by the National Scientific Research Fund (OTKA ref.No.: F015884, T029135) and the Austrian-Hungarian Action Fund (ref.No.: 2964 and 32609).

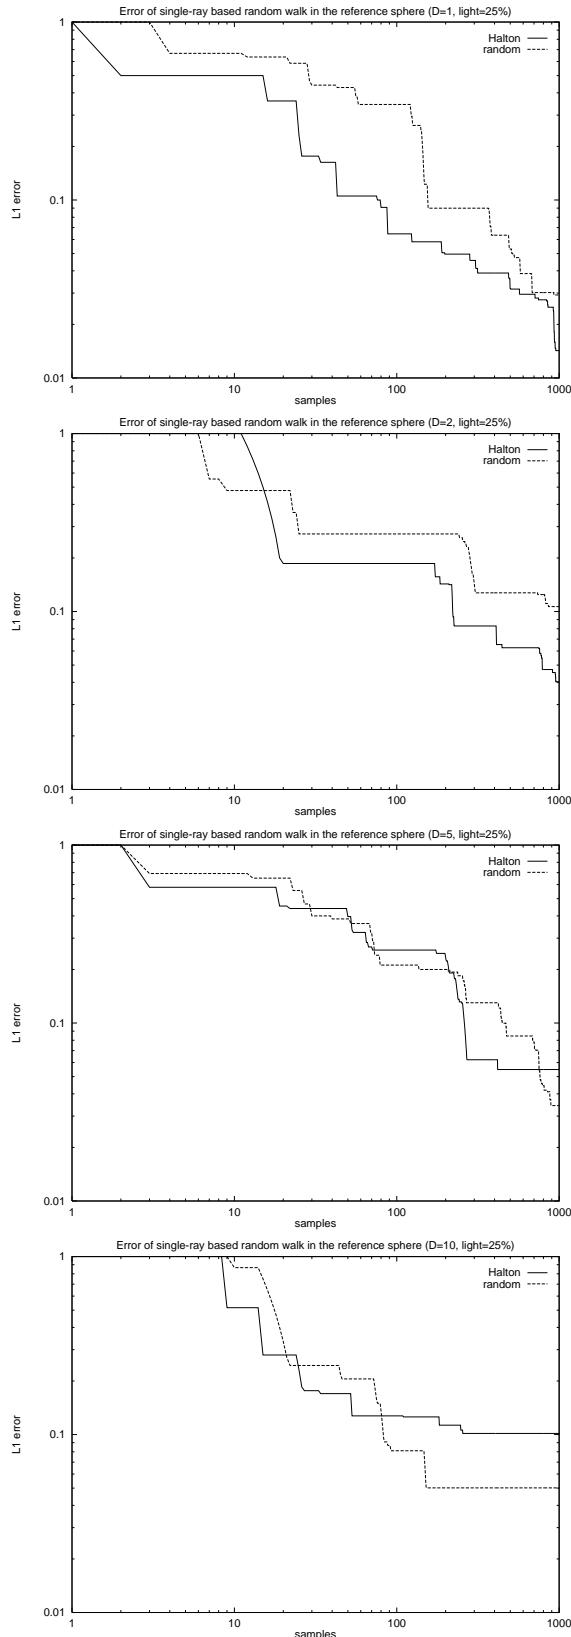


Figure 8: Error measurements for 1, 2, 5 and 10 bounces

REFERENCES

- [Castr98] F. Castro, R. Martinez, and M. Sbert. Quasi Monte-Carlo and extended first-shot improvements to the multi-path method. In *Spring Conference on Computer Graphics '98*, pages 91–102, 1998.
- [Deák89] I. Deák. *Random Number Generators and Simulation*. Akadémia Kiadó, Budapest, 1989.
- [Dobki96] D. P. Dobkin, D. Eppstein, and D. P. Mitchell. Computing the discrepancy with applications to supersampling patterns. *ACM Transactions on Graphics*, 15(4):354–376, 1996.
- [Hyben98] M. Hyben, I. Martisovits, and A. Ferko. Scene complexity for rendering in flatland. In L. Szirmay-Kalos, editor, *Spring Conference on Computer Graphics*, pages 112–120, 1998.
- [Kelle95] A. Keller. A quasi-Monte Carlo algorithm for the global illumination in the radiosity setting. In H. Niederreiter and P. Shiue, editors, *Monte-Carlo and Quasi-Monte Carlo Methods in Scientific Computing*, pages 239–251. Springer, 1995.
- [Knuth81] D.E. Knuth. *The art of computer programming. Volume 2 (Seminumerical algorithms)*. Addison-Wesley, Reading, Mass. USA, 1981.
- [Niede92] H. Niederreiter. *Random number generation and quasi-Monte Carlo methods*. SIAM, Pennsylvania, 1992.
- [Press92] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C (Second Edition)*. Cambridge University Press, Cambridge, USA, 1992.
- [Shirl91] P. Shirley. Time complexity of Monte-Carlo radiosity. In *Eurographics '91*, pages 459–466. Elsevier Science Publishers, 1991.
- [Warno95] T. Warnock. Computational investigations of low-discrepancy point sets. In H. Niederreiter and P. Shiue, editors, *Monte-Carlo and Quasi-Monte Carlo Methods in Scientific Computing*, pages 354–361. Springer, 1995.