

Efficient Bidirectional Path Tracing by Randomized Quasi-Monte Carlo Integration

Thomas Kollig and Alexander Keller

kollig@informatik.uni-kl.de and keller@informatik.uni-kl.de
Computer Science Dept., University of Kaiserslautern
Postfach 3049, D-67653 Kaiserslautern, Germany

Abstract. As opposed to Monte Carlo integration the quasi-Monte Carlo method does not allow for an error estimate from the samples used for the integral approximation and the deterministic error bound is not accessible in the setting of computer graphics, since usually the integrands are of unbounded variation. We investigate the application of randomized quasi-Monte Carlo integration to bidirectional path tracing yielding much more efficient algorithms that exploit low-discrepancy sampling and at the same time allow for variance estimation.

1 Introduction

The global illumination problem consists in rendering photorealistic images of a virtual scene and camera descriptions (for a detailed introduction see [Gla95]). A very basic algorithm for the solution of this light transport problem is the bidirectional path tracing algorithm [LW93, VG94], which in the context of the quasi-Monte Carlo method has been investigated in [Kel98b].

We first generalize this work by introducing multiple importance sampling using the balance heuristic for the quasi-Monte Carlo method, which is superior to the previous approaches of bidirectional path tracing. By sacrificing only little performance randomized quasi-Monte Carlo algorithms allow for integration error estimation. By numerical experiments we compare the efficiency of different randomized quasi-Monte Carlo approaches and illustrate how they smoothly blend between the pure Monte Carlo and the quasi-Monte Carlo case.

As a result the new scheme of *padded replications sampling* yields a bidirectional path tracing algorithm that is highly efficient, allows for an error estimate and is very simple to implement.

2 Bidirectional Path Tracing

We briefly recall the path integral formulation of the global illumination problem [Vea97], which in combination with multiple importance sampling yields the bidirectional path tracing algorithm. Furthermore we define the problem of insufficient techniques that is inherent with multiple importance sampling.

2.1 The Global Illumination Problem in Path Integral Form

A light path $\bar{x} = x_0x_1 \dots x_k$ of length k is characterized by its interaction points x_i with the scene surface S . The union of all path spaces

$$\mathcal{P}_k := \{\bar{x} = x_0x_1 \dots x_k \mid x_i \in S \text{ for } 0 \leq i \leq k\}$$

of a specific light path length k forms the path space $\mathcal{P} := \bigcup_{k=1}^{\infty} \mathcal{P}_k$. For Lebesgue measurable subsets $D_0, D_1, \dots, D_k \subseteq S$ we define the measure

$$\mu_k(D_0 \times D_1 \times \dots \times D_k) := A(D_0) \cdot A(D_1) \cdot \dots \cdot A(D_k) ,$$

where A is the area measure, and $\mu(D) := \sum_{k=1}^{\infty} \mu_k(D \cap \mathcal{P}_k)$ for $D \subseteq \mathcal{P}$. For a path $\bar{x} \in \mathcal{P}_k$ the measurement contribution function is

$$f_j(\bar{x}) := L_e(x_0 \rightarrow x_1) G(x_0 \leftrightarrow x_1) \cdot \left(\prod_{i=1}^{k-1} f_s(x_{i-1} \rightarrow x_i \rightarrow x_{i+1}) G(x_i \leftrightarrow x_{i+1}) \right) \cdot W_e^{(j)}(x_{k-1} \rightarrow x_k) ,$$

where the light sources are determined by the emittance L_e and $W_e^{(j)}$ are the detector functionals, which formalize the camera description. The bidirectional scattering distribution function f_s describes the surface properties.

$$G(x \leftrightarrow y) := V(x \leftrightarrow y) \frac{|\cos \theta_x| |\cos \theta_y|}{|x - y|^2}$$

is the geometry term, where θ_x is the angle between the surface normal in x and the vector between x and y ; θ_y is defined analogously. The visibility function $V(x \leftrightarrow y)$ is 1 if x and y are mutually visible and 0 otherwise. Then the global illumination problem consists in computing detector values I_j by the path integral

$$I_j = \sum_{k=1}^{\infty} \int_{\mathcal{P}_k} f_j(\bar{x}) d\mu_k(\bar{x}) = \int_{\mathcal{P}} f_j(\bar{x}) d\mu(\bar{x}) . \quad (1)$$

2.2 Multiple Importance Sampling

The problem of importance sampling [Sob94] is to find an efficient probability density function p . However, often it is possible to specify a whole set p_1, p_2, \dots, p_N of probability density functions instead of just one single p . While each probability density function of the set may reduce the variance of the importance sampling estimator only in a possibly unknown subdomain of D , multiple importance sampling, a variance reduction technique introduced by [VG94] and analyzed in [OZ99], allows for the combination of samples which are distributed according to different probability density functions.

A probability density function p can be used as a *technique*, if we are able to generate p -distributed samples and to evaluate¹ $p(x)$ for a given $x \in D$. Assuming that we have N techniques with their associated probability density functions

$$p_1, p_2, \dots, p_N : D \rightarrow \mathbb{R}_0^+ ,$$

a so-called *heuristic* consists of N corresponding weight functions

$$w_1, w_2, \dots, w_N : D \rightarrow \mathbb{R}_0^+ ,$$

such that

1. $\sum_{i=1}^N w_i(x) = 1$ for all $x \in D$ with $f(x) \neq 0$ and
2. $w_i(x) = 0$ for all $x \in D$ with $p_i(x) = 0$ holds.

Note that these conditions imply that each $x \in \text{supp } f$ can be generated by at least one² technique p_i . Then the multiple importance sampling estimator

$$\int_D f(x) dx \approx \frac{1}{n} \sum_{j=1}^n \sum_{i=1}^N w_i(x_{i,j}) \frac{f(x_{i,j})}{p_i(x_{i,j})} \quad (2)$$

is unbiased, where the $x_{i,j}$ are p_i -distributed for $1 \leq i \leq N$ and $1 \leq j \leq n$. We use the so called *balance heuristic* which has the weight functions

$$w_i(x) := \frac{p_i(x)}{\sum_{\ell=1}^N p_\ell(x)} . \quad (3)$$

The behaviour of the estimator (2) using (3) is comparable to importance sampling with $p \equiv \frac{1}{N} \sum_{\ell=1}^N p_\ell$.

2.3 The Bidirectional Path Tracing Algorithm

The path space samples $\bar{x} \in \mathcal{P}$ are generated in three steps (see Fig. 1 (a)):

1. generate a light subpath by a random walk starting on a light source,
2. generate an eye subpath by a random walk starting on a detector, and
3. connect both subpaths deterministically.

Since the ray casting function is very expensive, we use all possible connections to form additional path space samples as illustrated in Fig. 1 (a). The resulting associated probability density functions are denoted by $p_{k,i}$, where k is the path length and i the number of points of the light subpath. Figure 1 (b) shows all possible techniques with their associated probability density functions for path length $k = 3$.

¹ At least we must be able to decide whether $x \in \text{supp } p$ holds for a given $x \in D$.
² This can happen due to disjoint $\text{supp } p_i$ or *the problem of insufficient techniques* as addressed in Sec. 2.4.




Fig. 1. (a) Generation of path space samples and (b) techniques with their associated probability density functions for path length $k = 3$, where each $p_{3,i}$ is positioned at its deterministic connection. A pinhole camera model is assumed where the eye subpaths originate from the pinhole through the image plane.

Applying multiple importance sampling (2) with the balance heuristic (3) to the path integral formulation of the global illumination problem (1) yields the bidirectional path tracing estimator

$$I_j = \sum_{k=1}^{\infty} \int_{\mathcal{P}_k} f_j(\bar{x}) d\mu_k(\bar{x}) \approx \sum_{k=1}^{\infty} \frac{1}{n} \sum_{j=1}^n \sum_{i=0}^k \frac{f_j(\bar{x}_{k,i,j})}{\sum_{\ell=0}^k p_{k,\ell}(\bar{x}_{k,i,j})}, \quad (4)$$

where n is the number of samples per technique and for $1 \leq j \leq n$ the $\bar{x}_{k,i,j}$ are generated according to $p_{k,i}$. In order to handle the infinite sum we can use absorbing Markov chains for subpath generation. A biased alternative is to compute the approximation up to a maximum path length k_{\max} . For example $k_{\max} = 2$ implies that instead of the full solution only direct illumination is calculated.

The possibility to achieve a valid path using the *eye connection techniques* $p_{k,k}$, where the end of a light subpath is connected with the eye point, is very small and most samples are of zero contribution, if the image is computed pixel by pixel. Therefore we allow samples of these techniques to contribute directly to any pixel of the image.

2.4 The Problem of Insufficient Techniques

Multiple importance sampling tries to hide the weaknesses of single probability density functions, but nevertheless can fail. In order to illustrate the limits of the estimator (2) suppose we have a subdomain $G \subseteq D$ for which only one technique is accessible. Then multiple importance sampling degenerates to standard importance sampling on G due to an *insufficient set of techniques*.

For bidirectional path tracing the problem of insufficient techniques arises for singular surface properties, e.g. mirrors that usually are modeled by a Dirac delta distribution in the bidirectional scattering distribution function f_s . In Fig. 2 (a) such a difficult path is sketched: It can only be generated by the technique $p_{k,0}$ that uses no light subpath. The resulting high variance is clearly visible in Fig. 2 (b) and is perceived as white dots on the mirror.

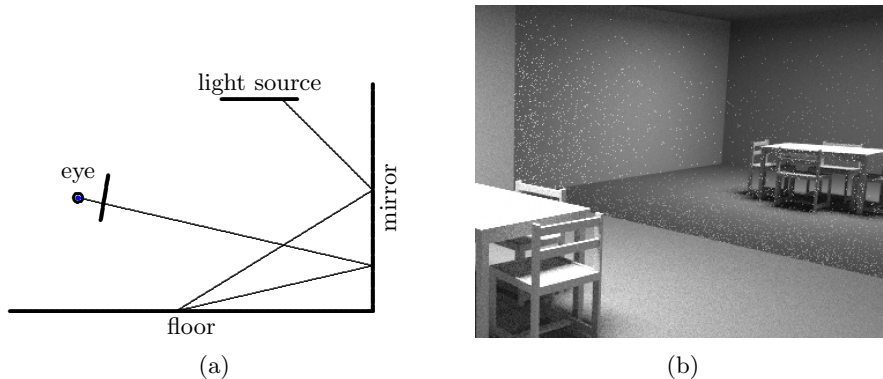


Fig. 2. The problem of insufficient techniques: (a) sketch of difficult path and (b) resulting high variance on the mirror.

3 Quasi-Monte Carlo Bidirectional Path Tracing

The Koksma-Hlawka inequality predicts that quasi-Monte Carlo integration performs superior to Monte Carlo integration for integrands of bounded variation in the sense of Hardy and Krause [Nie92]. For integrands with unknown discontinuities like the measurement contribution function in (1) only pessimistic error bounds are available [Hla71] due to unbounded variation. Nevertheless numerical experiments [Kel98a] reveal that even for these functions low-discrepancy sampling performs better than random sampling.

3.1 Multiple Importance Sampling for Quasi-Monte Carlo

In order to apply quasi-Monte Carlo integration to bidirectional path tracing the subpath generation has to be done using high-dimensional low-discrepancy points, where the dimension depends on the length of the subpaths. Due to the transport operator points at the beginning of a subpath affect the integration error more than points at the end of a subpath. In accordance the lower dimensions of low-discrepancy points often are better equidistributed than their higher dimensions. Therefore the first four dimensions are used to determine the first point of each subpath, the next four for the first scattering events and so on. This interleaving scheme is similar to [Kel98b], however, now we use the multiple importance sampling estimator (2) with the balance heuristic (3) and deterministic low-discrepancy sampling. In order to avoid aliasing different light subpaths have to be used for the estimation of each pixel functional. This is particularly important for the eye connection techniques $p_{k,k}$ (see also Fig. (5)) and is achieved by using consecutive subsequences of a low-discrepancy point sequence instead of a repeated finite point set.

4 Randomized Quasi-Monte Carlo Bidirectional Path Tracing

In [Owe98b] Owen surveys randomization techniques for quasi-Monte Carlo integration. Randomized quasi-Monte Carlo integration exploits the benefits of low-discrepancy sampling and at the same time allows for an efficient error estimate, which is not accessible for quasi-Monte Carlo integration.

From an initial low-discrepancy point set $P := \{a_1, a_2, \dots, a_m\} \subset I^s$ we generate r randomized replications $X_j := \{x_{1,j}, x_{2,j}, \dots, x_{m,j}\} \subset I^s$ with $1 \leq j \leq r$ such that

1. each replication X_j preserves the low-discrepancy properties of the initial point set P and
2. the replications $x_{i,1}, x_{i,2}, \dots, x_{i,r}$ of each point $a_i \in P$ are independent and uniformly distributed on I^s .

Then the randomized quasi-Monte Carlo estimator with a total of $n = rm$ samples

$$\int_{I^s} f(x) dx \approx \frac{1}{r} \sum_{j=1}^r \frac{1}{m} \sum_{i=1}^m f(x_{i,j}) \quad (5)$$

is unbiased. The expected error is bounded by the square root of the variance σ^2 of the above estimator and can be estimated in an unbiased way using the samples of (5):

$$\sigma^2 \approx \frac{1}{r(r-1)} \sum_{k=1}^r \left(\frac{1}{m} \sum_{i=1}^m f(x_{i,k}) - \frac{1}{r} \sum_{j=1}^r \frac{1}{m} \sum_{i=1}^m f(x_{i,j}) \right)^2. \quad (6)$$

Choosing the number r of replications just large enough to obtain a good variance estimate very little performance of the low-discrepancy quadrature is sacrificed and adaptive sampling controlled by error estimation yields much more efficient rendering algorithms.

As long as the replications $x_{i,1}, x_{i,2}, \dots, x_{i,r}$ are independent and uniformly distributed on I^s the estimator (5) is unbiased and the variance estimator (6) remains valid. Thus the initial point set P and its replications X_j do not need to be of low-discrepancy, however, their choice affects variance and therefore error.

For bidirectional path tracing the initial point set P with m points and the replication scheme have to be selected. Then each pixel functional is estimated by r independent random replications of P . The dimensions are assigned identically to the quasi-Monte Carlo setting in Sect. 3.1.

4.1 Cranley-Patterson Rotations

Cranley and Patterson [CP76] suggested the following form of randomization: For a replication X_j they added a random shift ξ_j to each point a_i of the

initial point set P . Thus we have $x_{i,j} = (a_i + \xi_j) \bmod 1$ with independent realizations $\xi_j \sim U(I^s)$ for $1 \leq i \leq m$ and $1 \leq j \leq r$.

Most low-discrepancy constructions are designed to minimize the star-discrepancy in the sense of (t, m, s) -nets or (t, s) -sequences. By randomly shifting a point set P this discrepancy of a replication X_j can be different and especially worse than the original discrepancy [Tuf96], since the (t, m, s) -net property is not shift invariant. Points designed to also minimize the torus discrepancy [BC87] are better suited for Cranley-Patterson rotations. The equidistribution properties of good lattice points remain almost unaffected when being shifted [CP76,SJ94].

4.2 Owen Scrambling

In [Owe95] Owen presented a randomization scheme for (t, m, s) -nets and (t, s) -sequences in base b . Starting with $H = I^s$ the following algorithm is applied to each coordinate:

1. Slice H into b equal volumes H_1, H_2, \dots, H_b along the coordinate.
2. Randomly permute these volumes in an independent way.
3. For each volume H_h repeat the procedure with $H = H_h$.

Due to the finite precision of computer arithmetic the infinite scheme in fact becomes a finite algorithm [FK00].

4.3 Padded Replications Sampling

For the light transport problem the benefits of quasi-Monte Carlo integration diminish in high dimensions [Kel98a]. So instead of using a computationally expensive high-dimensional low-discrepancy point set as initial point set P (see Fig. 4 (a)), the structure underneath the transport problem can be exploited: The light and eye subpaths are generated by area sampling and scattering, which both are two-dimensional problems. Thus the idea of padded replications sampling is to use a random replicate of one two-dimensional basis pattern for each two-dimensional subproblem as illustrated in Fig. 3. This in fact only changes the initial point set P (see Fig. 4 (b)) of the replication scheme. Of course P no longer is of low-discrepancy, but is much cheaper to generate and performs at least as good as shown in the experimental section.

5 Latin Supercube Sampling

Similar to padded replications sampling Latin supercube sampling [Owe98a] is a method to expand low-dimensional samples to high dimensions: The low-dimensional point sets are randomly permuted before padding. Suppose that

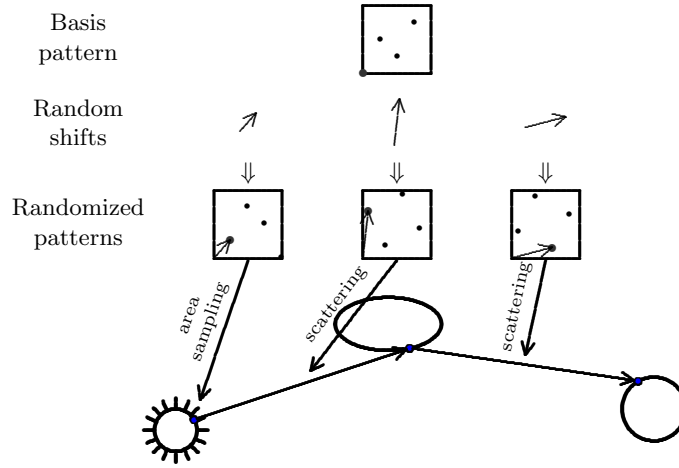


Fig. 3. Subpath generation by padded replications sampling using Cranley-Patterson rotations for e.g. a light subpath. The basis pattern size is $m = 4$ and the dimension of the padded points is $s = 6$.

Dimensions	1,2	3,4	5,6	7,8
(a) Scrambled Hammersley				
(b) Padded Hammersley				
(c) Latin supercube				

Fig. 4. Illustration of the initial point sets P used by the different approaches for randomized quasi-Monte Carlo integration: (a) high-dimensional low-discrepancy point set, (b) padded replications sampling, and (c) decorrelated padded replications sampling by index permutations.

$Q_i := \{a_{i,1}, a_{i,2}, \dots, a_{i,m}\} \subset I^{s_i}$ for $1 \leq i \leq q$ with $\sum_{i=1}^q s_i = s$ are (randomized) quasi-Monte Carlo point sets. Then the Latin supercube samples are

$$x_j := (a_{1,\pi_1(j)}, a_{2,\pi_2(j)}, \dots, a_{q,\pi_q(j)}) \subset I^s$$

for $1 \leq j \leq m$, where the π_i are independent uniform random permutations over $\{1, 2, \dots, m\}$.

In computer graphics Latin supercube sampling has been applied by Cook [Coo86] (later formalized by Shirley [Shi90]) for distribution ray tracing, which is not a consistent algorithm in the sense of (1) since it uses only a subset of the techniques $p_{k,0}$ and $p_{k,1}$, where the end of an eye subpath has to hit a light source or is connected with a point on a light source. However, Cook and Shirley did not use (randomized) quasi-Monte Carlo point sets but stratified random point sets for Latin supercube sampling.

5.1 Latin Supercube Samples from Deterministic Low-Discrepancy Points

Using large two-dimensional quasi-Monte Carlo point sets for Latin supercube sampling is prohibitive due to the considerable amount of permutation memory of order $\mathcal{O}(qm) = \mathcal{O}(k_{\max}m)$, if a finite maximum path length k_{\max} is used. On the other hand the number of different light subpaths that can be generated by Latin supercube sampling is limited by $m^{k_{\max}}$. Therefore for small values of m Latin supercube sampling is only practicable if the eye connection techniques $p_{k,k}$ are not used, otherwise severely disturbing aliasing artifacts will be visible (see Fig. (5)) that only can be avoided by using huge values of m that are of the order of pixels in the image.




Fig. 5. Aliasing caused by Latin supercube samples from deterministic points with eye connection techniques.

5.2 Latin Supercube Samples from Randomized Low-Discrepancy Points

For padded replications sampling the same basis pattern is padded together to form the initial point set P . The resulting correlation between the dimensions (see Fig. 4 (b)) can cause an increased variance for a larger number m of points in the basis pattern. Latin supercube resolves this correlation when being applied to the initial point set P before replication (see Fig. 4 (c)).

6 Numerical Experiments

For the numerical experiments we chose the GLASS SPHERE and OFFICE test scenes. Figure 6 shows the master images, which have been rendered with the original bidirectional path tracing algorithm using more than a thousand samples per pixel and technique. In our experiments the error of an image is approximated by its L_2 -distance to these master images. Instead of using




Fig. 6. The GLASS SPHERE and OFFICE scene are used as test scenes.

absorbing Markov chains for subpath generation we restricted the maximum path length to $k_{\max} = 6$ for the GLASS SPHERE and to $k_{\max} = 3$ for the OFFICE scene.

The difficulties of the GLASS SPHERE scene are the caustic on the floor and the light, which is reflected by the glass sphere onto the ceiling. The OFFICE scene has only diffuse surface properties. Besides the two big luminaries at the ceiling the small spherical light source of the table lamp makes the rendering complicated. Since here we have no singular surface properties we omit the eye connection techniques $p_{k,k}$, which are only useful in order to render directly seen caustics.

6.1 Quasi-Monte Carlo Bidirectional Path Tracing

In Fig. 7 (a) the performance of several deterministic low-discrepancy sequences is compared. The Niederreiter sequence [Nie92] in base 2 is slightly worse than the scrambled Halton sequence [Fau92] for a small number of samples per technique and pixel and almost as good for more than 64 samples. Optimizing the t parameter by increasing the construction base of the Niederreiter sequence leads to worse results. The Niederreiter-Xing sequence [NX96,Pir00] in base 2 is even far less efficient except for multiples of 64 samples. The reason why the simple scrambled Halton sequence performs best is that it better fits the structure of the global illumination problem: The initial segments of the subpaths contribute most and consequently the lower dimensions of the points are most important. In addition only a very small number of samples is used to estimate a pixel functional. Looking at the low-dimensional projections of short consecutive subsequences yields that the projections of the scrambled Halton sequence expose better discrepancy than the projections of the (t, s) -sequences, causing the superior performance of the scrambled Halton sequence.




Fig. 7. Convergence graphs for (a) quasi-Monte Carlo bidirectional path tracing using different low-discrepancy sequences and (b) randomized quasi-Monte Carlo bidirectional path tracing using different high-dimensional point sets as initial point set P and different randomization schemes.

6.2 Randomized Quasi-Monte Carlo Bidirectional Path Tracing

Now different approaches to randomized quasi-Monte Carlo bidirectional path tracing are compared, where high-dimensional low-discrepancy point sets are selected as initial point set P . While the scrambled Hammersley point set is chosen for Cranley-Patterson rotations (see 4.1), different (t, s) -nets are used with Owen scrambling (see 4.2). Here each pixel functional is estimated using only one independent replication. The resulting convergence graphs in Fig. 7 (b) show that the scrambled Hammersley version performs best. Using a Niederreiter sequence is slightly worse, where the increased construction base affects the error less than in the quasi-Monte Carlo setting. The Niederreiter-Xing sequence again performs even worse except for multiples of 64 samples.

Usually for randomized quasi-Monte Carlo integration the size m of the initial point set P is fixed and the desired sampling rate is obtained by increasing the number r of replications, yielding a convergence rate of $\mathcal{O}(r^{-\frac{1}{2}})$. Seen that way the points of the graphs in Fig. 7 (b) can be considered as starting points of the convergence graphs for fixed m and increasing r .

6.3 Blending between Monte Carlo and Quasi-Monte Carlo

So far we have analyzed quasi-Monte Carlo and randomized quasi-Monte Carlo separately. Now the best of the above sampling schemes are compared:

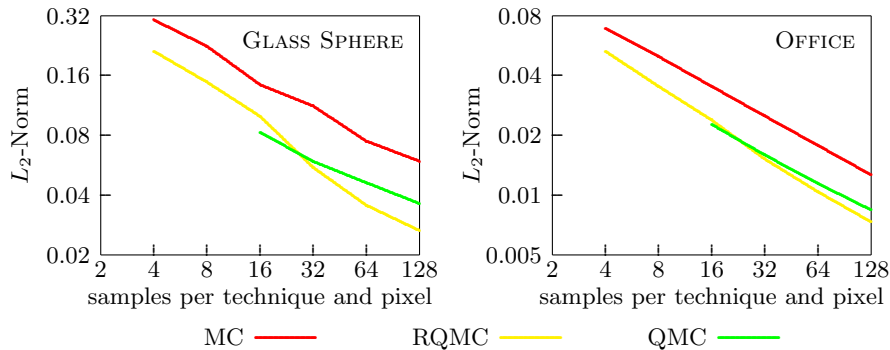


Fig. 8. Convergence graphs for MC, RQMC ($m = 16$), and QMC.

- **Monte Carlo (MC).** The original bidirectional path tracing algorithm uses pure random sampling. For comparison we also implemented a version using Latin hypercube sampling (LHS).
- **Randomized quasi-Monte Carlo (RQMC).** Besides using Cranley-Patterson rotations with the scrambled Hammersley point set also its padded replications sampling (see Sect. 4.3) counterpart using the two-dimensional Hammersley point set has been applied. In both approaches we can choose the fixed size m of the initial point set P .
- **Quasi-Monte Carlo (QMC)** using the scrambled Halton sequence.

In Fig. 8 convergence graphs are shown. As expected the convergence rates for the Monte Carlo and randomized quasi-Monte Carlo versions are $\mathcal{O}(n^{-1/2})$, where n denotes the total number of samples per technique and pixel. A slightly improved rate of $\mathcal{O}(n^{-1/2-\alpha})$ can be observed for the quasi-Monte Carlo approach, where $\alpha \in [0, \frac{1}{2}]$ decreases with the maximum path length k_{\max} used in the simulation due to the discontinuities in the measurement contribution function.

For a more detailed comparison we measured the number of samples required to achieve a given error in relation to the number needed by the original bidirectional path tracing algorithm, i.e. the pure Monte Carlo algorithm. The results are shown in Fig. 9. Far more than half of the expensive samples can be saved by the quasi-Monte Carlo version. The randomized quasi-Monte Carlo approaches form a smooth transition between the pure random and the deterministic algorithm. With an increasing size m of the initial point set P the error decreases due to the better equidistribution of the samples. It is an interesting result that for bidirectional path tracing padded replications sampling performs at least as good as a high-dimensional low-discrepancy point set.

Along the lines of Cook [Coo86] and Shirley [Shi90] also Latin supercube sampling by deterministic low-discrepancy points (see Sect. 5.1) has been

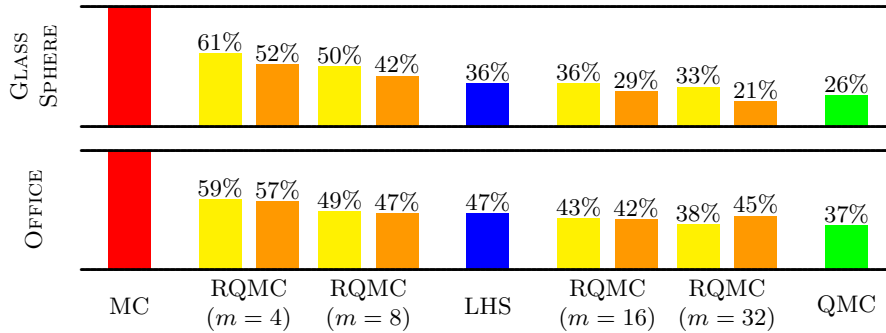


Fig. 9. Number of samples needed to achieve a given error in relation to MC. Padded replications sampling (right bars of RQMC) performs at least as good as the straight forward approach (left bars of RQMC).




Fig. 10. Using Latin supercube samples as input for Cranley-Patterson rotation reduces the effect of correlation caused by padded replications sampling. For this experiment Fibonacci lattice points have been used as basis pattern.

applied to render the OFFICE scene, since here the eye connection techniques are not required. For padding the two-dimensional Hammersley point set was chosen. In comparison to the pure Monte Carlo algorithm only 35% of the samples are needed to achieve the same error. Thus it performs similar to the purely deterministic quasi-Monte Carlo approach (37% of the samples, see Fig. 9).

6.4 Decorrelation of Padded Replications Sampling

Latin supercube sampling can reduce the correlation between the dimensions of the points used by padded replications sampling. An increased variance due to correlation becomes visible in Fig. 9 for the OFFICE scene when increasing the basis pattern size m from 16 (42%) to 32 (45%).

In Fig. 10 padded replications sampling with and without decorrelation using Latin supercube samples is compared, where we padded replications of the Fibonacci lattice points [SJ94]. The reduced correlation results in a better performance when using bigger basis pattern sizes m . In computer




Fig. 11. Image comparison. The close-ups in the upper row were rendered with pure Monte Carlo bidirectional path tracing and the close-ups in the lower row were rendered by the padded replications sampling approach using the Hammersley point set without decorrelation.

graphics, however, $rm = n \leq 128$ so that the effect of decorrelation by Latin supercube sampling is hardly perceivable.

6.5 Visual Comparison

For a visual comparison of images we rendered the OFFICE scene with the original bidirectional path tracing algorithm and with the padded replications sampling using 16 samples per technique and pixel. Since the padded replications sampling version needs fewer pseudo random numbers its rendering time was about 12% shorter. Figure 11 shows two close-ups of the images. The reduced error results in a less noisy image. Even in only indirectly illuminated regions (right column) there is less noise.

7 Conclusion

We investigated several new sampling approaches to bidirectional path tracing speeding up the original algorithm by a factor of 2 to 5. By numerical

evidence we showed that *padded replications sampling* is almost as efficient as the best quasi-Monte Carlo integration approach. However, *padded replications sampling* allows for variance estimation, is much simpler to implement as compared to the high-dimensional low-discrepancy constructions, requires much less random numbers than pure random sampling, and perfectly fits the intrinsic two-dimensional structure of the global illumination problem.

Acknowledgement

The authors would like to thank Fred Hickernell for his comments.

References

- BC87. J. Beck and W. Chen, *Irregularities of Distribution*, Cambridge University Press, 1987. [4.1](#)
- Coo86. R. Cook, *Stochastic Sampling in Computer Graphics*, ACM Transactions on Graphics **5** (1986), no. 1, 51–72. [5](#), [6.3](#)
- CP76. R. Cranley and T. Patterson, *Randomization of Number Theoretic Methods for Multiple Integration*, SIAM Journal on Numerical Analysis **13** (1976), 904–914. [4.1](#)
- Fau92. H. Faure, *Good Permutations for Extreme Discrepancy*, J. Number Theory **42** (1992), 47–56. [6.1](#)
- FK00. I. Friedel and A. Keller, *Fast Generation of Randomized Low Discrepancy Point Sets*, 2000, this volume. [4.2](#)
- Gla95. A. Glassner, *Principles of Digital Image Synthesis*, Morgan Kaufmann, 1995. [1](#)
- Hla71. E. Hlawka, *Discrepancy and Riemann Integration*, Studies in Pure Mathematics (L. Mirsky, ed.), Academic Press, New York, 1971, pp. 121–129. [3](#)
- Kel98a. A. Keller, *Quasi-Monte Carlo Methods for Photorealistic Image Synthesis*, Ph.D. thesis, Shaker, Aachen, 1998. [3](#), [4.3](#)
- Kel98b. ———, *The Quasi-Random Walk*, Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing 1996 (H. Niederreiter, P. Hellekalek, G. Larcher, and P. Zinterhof, eds.), Lecture Notes in Statistics, vol. 127, Springer, 1998, pp. 277–291. [1](#), [3.1](#)
- LW93. E. Lafortune and Y. Willems, *Bidirectional Path Tracing*, Proc. 3rd International Conference on Computational Graphics and Visualization Techniques (Compugraphics), 1993, pp. 145–153. [1](#)
- Nie92. H. Niederreiter, *Random Number Generation and Quasi-Monte Carlo Methods*, SIAM, Pennsylvania, 1992. [3](#), [6.1](#)
- NX96. H. Niederreiter and C. Xing, *Quasirandom Points and Global Function Fields*, Finite Fields and Applications (S. Cohen and H. Niederreiter, eds.), Cambridge University Press, 1996, 269–296. [6.1](#)
- Owe95. A. Owen, *Randomly Permuted (t, m, s) -Nets and (t, s) -Sequences*, Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing (H. Niederreiter and P. Shiue, eds.), Lecture Notes in Statistics, vol. 106, Springer, 1995, pp. 299–317. [4.2](#)

- Owe98a. ———, *Latin Supercube Sampling for Very High Dimensional Simulations*, ACM Transactions on Modeling and Computer Simulation **8** (1998), 71–102. [5](#)
- Owe98b. ———, *Monte Carlo Extension of Quasi-Monte Carlo*, Winter Simulation Conference, IEEE Press, 1998, pp. 571–577. [4](#)
- OZ99. A. Owen and Y. Zhou, *Safe and Effective Importance Sampling*, Tech. report, Stanford University, Goldman-Sachs, 1999. [2.2](#)
- Pir00. G. Pirsic, *A Software Implementation of Niederreiter-Xing Sequences*, 2000, this volume. [6.1](#)
- Shi90. P. Shirley, *Physically Based Lighting Calculations for Computer Graphics*, Ph.D. thesis, University of Illinois, Urbana-Champaign, 1990. [5](#), [6.3](#)
- SJ94. I. Sloan and S. Joe, *Lattice Methods for Multiple Integration*, Clarendon Press, Oxford, 1994. [4.1](#), [6.4](#)
- Sob94. I. Sobol, *A Primer for the Monte Carlo Method*, CRC Press, 1994. [2.2](#)
- Tuf96. B. Tuffin, *On the Use of Low Discrepancy Sequences in Monte Carlo Methods*, Monte Carlo Methods and Applications **2** (1996), no. 4, 295–320. [4.1](#)
- Vea97. E. Veach, *Robust Monte Carlo Methods for Light Transport Simulation*, Ph.D. thesis, Stanford University, December 1997. [2](#)
- VG94. E. Veach and L. Guibas, *Bidirectional Estimators for Light Transport*, Proc. 5th Eurographics Workshop on Rendering (Darmstadt, Germany), June 1994, pp. 147–161. [1](#), [2.2](#)