Interner Bericht

Stratification by Rank-1-Lattices

A. Keller

322/02

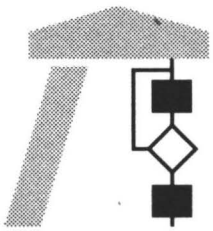

FACHBEREICH
INFORMATIK

UNIVERSITÄT
KAISERSLAUTERN

Postfach 3049 · D-67653 Kaiserslautern

Stratification by Rank-1-Lattices

A. Keller

322/02

Universität Kaiserslautern
AG Numerische Algorithmen
Postfach 30 49
67653 Kaiserslautern
Germany

Juli 2002

Herausgeber: AG Numerische Algorithmen
Leiter: Professor Dr. S. Heinrich

Stratification by Rank-1-Lattices

Alexander Keller

Kaiserslautern University

Abstract

Many rendering problems can only be solved using Monte Carlo integration. The noise and variance inherent with the statistical method efficiently can be reduced by stratification. So far only uncorrelated stratification methods were used that in addition depend on the dimension of the integration domain. Based on rank-1-lattices we present a new stratification technique that removes this dependency on dimension, is much more efficient by correlation, is trivial to implement, and robust to use. The superiority of the new scheme is demonstrated for standard rendering algorithms.

1 Introduction

Whenever the Monte Carlo method is used for numerical integration, it is known that increasing the uniformity of the random samples by stratification can only increase the convergence (see e.g. [Sobol' 1994] or any other textbook on Monte Carlo integration). In the context of computer graphics Mitchell [Mitchell 1996] investigated the benefits of jittered sampling [Cook et al. 1984] and demonstrated that stratification strictly improves accuracy. The experimentation was undertaken in two and four dimensions for pixel anti-aliasing and form factor estimation at $n = m^2$ and $n = m^4$ samples, respectively, for $m \in \mathbb{N}$.

Opposite to the core idea of Monte Carlo integration to get rid of the curse of dimension, the known constructions for jittered samples force sampling rates $n = \mathcal{O}(2^s)$ that exponentially increase with the dimension s of the integration domain, since the integration domain is partitioned along each dimension separately. In addition the independent stratified samples can happen to lie arbitrarily close along strata boundaries. A minimum distance property, however, is favorable as explored by blue noise sampling [McCool and Fiume 1992]. Even worse the independence prevents to exploit continuity when evaluating a family of integrals with smoothly varying parameters, as encountered for e.g. neighboring pixels.

The basic idea is illustrated in figure 1: Assuming a sampling rate of e.g. $n = 34$ samples in two dimensions, the standard decomposition would allow for 2 by 17 axis aligned strata, where each stratum is uniformly sampled once. Obviously sampling is not as efficient as it could be since samples can lie arbitrarily close together and huge areas remain without samples. On the contrary our new approach of correlated sampling using stratification by rank-1-lattices guarantees a minimum distance of the samples and a much better uniformity of the sample set. The resulting integration scheme shows much less noise converges much faster for the identical amount of work.

Based on rank-1-lattices that are introduced in section 2, we will present our new sampling schemes in section 3. These schemes allow to efficiently stratify any number n of samples independent of dimension s (section 3.1). By switching to correlated sampling in section 3.2 randomness and such noise are decreased while efficiency is increased. This is especially useful for trajectory splitting (section 3.3). The power of our new methods is demonstrated by application examples in section 4 before concluding the paper.

2 Rank-1-Lattices

A discrete set

$$L := P_n + \mathbb{Z}^s \subset \mathbb{R}^s, \text{ where } P_n := \{z_0, \ldots, z_{n-1}\},$$

that is closed under addition and subtraction is called a *lattice*. Simple examples for lattices are axis aligned regular grids. It is interesting to consider the *rank* of a lattice, which is the minimum number of basis vectors needed to generate all points of the lattice. For the case of axis aligned regular grids the dimension s is the number of basis vectors that in linear combinations yield the lattice points. However it is also possible to generate s-dimensional lattices by

$$z_j := \frac{j}{n} \cdot \mathbf{g} = (g_1, \ldots, g_s) \tag{1}$$

using only one generator vector $\mathbf{g} \in \mathbb{N}^s$. In consequence these lattices are called *rank-1-lattices*.

If for the components of $\mathbf{g}$ we have $\gcd(g_i, n) = 1$, the resulting rank-1-lattice is an instance of a Latin hypercube sample, i.e. provides perfect low dimensional projections (each axis is equidistantly stratified, see figure 1c) and a trivial lower distance bound of $\frac{1}{n}$. Depending on the choice of the generator vector $\mathbf{g}$ the point set even can be of low discrepancy [Niederreiter 1992]; then it is called *good lattice*. It has been shown that such constructions [Sloan and Joe 1994] exist for any number of points n independent of the dimension s. Thus it is possible to have good lattices in s dimensions, where however the number of lattice points does not exponentially dependent on dimension!

There is a well established theory on rank-1-lattices and their use for the integration of periodic functions from very restricted classes [Sloan and Joe 1994]. Besides the existence theorems for good lattices only a very small number of explicit constructions are known. The construction of the 2-dimensional *Fibonacci lattices* (section 2.1) will proof to be extremely useful for integrating the class of square-integrable functions as encountered in photorealistic rendering. For general dimension rank-1-lattices often are tabulated in *Korobov form* (section 2.2).

2.1 A useful Example: The Fibonacci Lattices

For the Fibonacci sequence $F_k := F_{k-1} + F_{k-2}$ with $F_2 = F_1 = 1$ and $k \in \mathbb{N}$, the Fibonacci lattice at $n = F_k$, $k \geq 3$, points has the generator vector

$$\mathbf{g} = (1, F_{k-1}).$$

In connection with equation (1) the Fibonacci lattice in the unit square is trivial to generate in a `for`-loop: The first component is $j\frac{1}{F_k}$, while the second component is computed by $j\frac{F_{k-1}}{F_k} \bmod 1$ (see also table 1). Figure 1b shows the Fibonacci lattice at $n = 34$ points. Obviously the unit square is much more uniformly partitioned and sampled by allowing for the use of rotated and sheared [Bouville et al. 1991] lattices instead of axis aligned constructions. The Fibonacci lattices are of low discrepancy and rectangular for odd k.

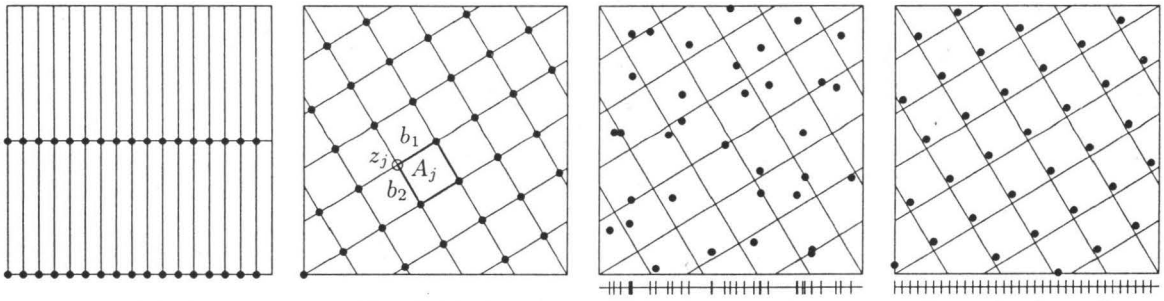

a) rank-2-lattice b) rank-1-lattice c) jittered sampling d) correlated sampling

Figure 1: Comparison of an axis aligned rank-2-lattice at $n = 2 \times 17$ points in a) and a Fibonacci rank-1-lattice at the same $n = 34$ lattice points in b). The basis vectors b_1 and b_2 are uniquely identified by the Delauney-tetrahedronization and the induced stratification is depicted. For the rank-1-lattice in b) the parallelepiped A_j offset by the lattice point z_j is outlined. Obviously the unit square is sampled and stratified much more uniformly by the rank-1-lattice. c) Each stratum contains an independent random sample, whereas in d) the same random instance is used in each stratum. The one-dimensional projections of the samples onto the x-axis underneath the unit squares exposes the much higher regularity of the correlated samples including a minimum distance property that cannot be guaranteed for uncorrelated sampling.

2.2 Lattices in Korobov Form

Lattices in Korobov form have a generator vector $\mathbf{g} = (1, a, a^2, \ldots, a^{s-1})$ and consequently can be specified by a pair (n, a), where n is the number of points. Besides various tables of pairs [Sloan and Joe 1994; Haber 1983], two useful pairs [Hickernell et al. 2001] are $(2^m, 17797)$ and $(2^m, 203)$ that produce a rank-1-lattice for any number of points $n = 2^m$, $m \in \mathbb{N}$, for any dimension s.

2.3 Domain Discretization by Rank-1-Lattices

The unit cell of a rank-1-lattice cannot be uniquely specified [Sloan and Joe 1994]. However specifying the unit cell in the sense of Delaunay, i.e. such that the volume of an inscribed sphere is maximized, almost makes the choice unique: The resulting unit cells are equivalent in the sense that they can be mapped onto each other by means of rigid body transformations.

The unit cells in the sense of Delaunay are identified by the Voronoi diagram of the rank-1-lattice, whose nerve is constructed out of the s shortest vectors $v_1, \ldots, v_s$ from the origin to points of the lattice, where the first component $v_{j,1} > 0$ (this implies $v_j \neq 0$ and $n > 1$). These vectors span parallelepipeds

$$\mathbf{B} := (v_1 \cdots v_s)$$

of volume $|\mathbf{B}| := |\det \mathbf{B}| = \frac{1}{n}$. Obviously $n\mathbf{B}$ is an integer matrix and so can be efficiently stored along with the generator vector $\mathbf{g}$.

Each of the n parallelepipeds contained in a rank-1-lattice restricted to the unit cube can be regarded as a stratum A_j induced by the bijections

$$\begin{aligned} R_j : [0,1)^s &\rightarrow A_j \\ x &\mapsto \left(\frac{j}{n} \cdot \mathbf{g} + \mathbf{B}x\right) \bmod [0,1)^s, \end{aligned} \qquad (2)$$

where the j-th lattice point $z_j = \frac{j}{n} \cdot \mathbf{g}$ is used as offset for the unit cube $[0,1)^s$ under the linear transformation $\mathbf{B}$ representing a rotation and shear (see figure 1b). We will refer to this mapping as *reduced Cranley-Patterson rotation* throughout the paper.

For the Fibonacci lattice at $n = F_k > 1$ points the unit cell in the sense of Delaunay is given by

$$\mathbf{B} = \frac{1}{F_k} \begin{pmatrix} j_1 & j_2 \\ j_1 F_{k-1} & -j_2 F_{k-1} \end{pmatrix}$$

where $j_1 = F_{2 \cdot \lfloor \frac{k-1}{4} \rfloor + 1}$ and $j_2 = F_{2 \cdot \lfloor \frac{k+1}{4} \rfloor}$ (see table 1).

$n = F_k$	$\mathbf{g}$	j_1	j_2
$2 = F_3$	(1,1)	$F_1 = 1$	$F_2 = 1$
$3 = F_4$	(1,2)	$F_1 = 1$	$F_2 = 1$
$5 = F_5$	(1,3)	$F_3 = 2$	$F_2 = 1$
$8 = F_6$	(1,5)	$F_3 = 2$	$F_2 = 1$
$13 = F_7$	(1,8)	$F_3 = 2$	$F_4 = 3$
$21 = F_8$	(1,13)	$F_3 = 2$	$F_4 = 3$
$34 = F_9$	(1,21)	$F_5 = 5$	$F_4 = 3$
$55 = F_{10}$	(1,34)	$F_5 = 5$	$F_4 = 3$
$89 = F_{11}$	(1,55)	$F_5 = 5$	$F_6 = 8$
$144 = F_{12}$	(1,89)	$F_5 = 5$	$F_6 = 8$
$233 = F_{13}$	(1,144)	$F_7 = 13$	$F_6 = 8$
$377 = F_{14}$	(1,233)	$F_7 = 13$	$F_6 = 8$
$610 = F_{15}$	(1,377)	$F_7 = 13$	$F_8 = 21$
$987 = F_{16}$	(1,610)	$F_7 = 13$	$F_8 = 21$

Table 1: The generator vectors $\mathbf{g} = (1, F_{k-1})$ for the first Fibonacci lattices at n points. The indices j_1 and j_2 identify the lattice points that form the unit cell in the sense of Delaunay (see section 2.3).

Since good rank-1-lattices exist for any dimension and number of lattice points, the induced domain discretization has obvious applications for e.g. solving high dimensional partial differential equations. For the scope of this paper we will focus on stratification for variance reduction.

3 Stratified Sampling

The integrands in photorealistic image synthesis are square integrable but may contain unknown discontinuities and often are of high dimension. The Monte Carlo methods uses a random point set $P_r = \{\xi_1, \ldots, \xi_r\} \subset [0,1)^s$ of r independently realized uniformly distributed points ξ_i to estimate the integral of a square integrable function f yielding the probabilistic error bound

$$\text{Prob}\left(\left|\int_{[0,1)^s} f(x)dx - \frac{1}{r}\sum_{k=1}^{r} f(\xi_k)\right| \leq \frac{3\sigma(f)}{\sqrt{r}}\right) \approx 0.997.$$

Obviously the rate of convergence of $\mathcal{O}\left(\frac{1}{\sqrt{r}}\right)$ is independent of the smoothness and dimension of the integrand f.

Partitioning the integration domain $[0,1)^s = \cup_{k=1}^{K} A_k$ into K disjoint strata A_j of equal area and separately evaluating the result-

$\mathbb{N}$	the natural numbers excluding zero
$\mathbb{Z}$	the integers
$\mathbb{R}$	the reals
n	number of lattice points
s	dimension
$\mathbf{g}$	generator vector
z_j	j-th lattice points
v_i	i-th unit cell basis vector
$\mathbf{B}$	basis of a unit cell
R_j	reduced Cranley-Patterson rotation
R_j^{CP}	Cranley-Patterson rotation

Figure 2: Selected symbols.

ing integrals

$$\int_{[0,1)^s} f(x)dx = \sum_{j=1}^{K} \int_{A_j} f(x)dx$$

by a one-sample Monte Carlo scheme can only improve the convergence [Niederreiter 1992].

3.1 Jittered Sampling by Rank-1-Lattices

Jittered sampling has been introduced to computer graphics by Cook [Cook et al. 1984] et al. and its benefits have been illustrated by Mitchell [Mitchell 1996]. However the used partitions by axis-aligned grids of rank s suffer an exponential growth of strata with dimension s. On the contrary rank-1-lattices imply much more regular strata $(A_j)_{j=1}^{n}$ independent of dimension and for any number n of strata (compare figures 1a and 1b).

Using the strata induced by (2) yields

$$\begin{aligned}\int_{[0,1)^s} f(x)dx &= \sum_{j=0}^{n-1} \int_{A_j} f(x)dx \\ &= \frac{1}{n} \sum_{j=0}^{n-1} \int_{[0,1)^s} f(R_j(x))dx \\ &\approx \frac{1}{n} \sum_{j=0}^{n-1} f(R_j(\xi_j)), \end{aligned} \tag{3}$$

where the determinant $|\det \mathbf{B}| = \frac{1}{n} = |A_j|$ has been used for the substitution. The estimate uses independent uniform random samples $\xi_j \in [0,1)^s$; one for each stratum A_j as depicted in figure 1c.

Using rank-1-lattices the s-dimensional unit cube easily can be stratified into any number n of strata independent of dimension. The estimator is unbiased and applicable to any square-integrable function. This is a much larger class of functions than the one of bounded variation in the sense of Hardy and Krause [Niederreiter 1992] as used for quasi-Monte Carlo integration or periodic functions with rapidly decaying Fourier coefficients [Sloan and Joe 1994]. In addition no care of periodization transformations has to be taken [Sloan and Joe 1994]. This has not been recognized before.

3.2 Correlated Sampling

Correlated sampling is the core of the most powerful variance reduction method of separating the main part [Kalos and Whitlock 1986]. The same principle is applied here: Instead of evaluating each integral by an independent set of samples (see figure 1c), correlated sampling uses one set of samples for all integrals (see figure 1d). This is easily derived from (3) by just exchanging the finite sum and the integral yielding

$$\begin{aligned}\frac{1}{n} \sum_{j=0}^{n-1} \int_{[0,1)^s} f(R_j(x))dx &= \frac{1}{n} \int_{[0,1)^s} \sum_{j=0}^{n-1} f(R_j(x))dx \\ &\approx \frac{1}{n} \sum_{j=0}^{n-1} f(R_j(\xi)). \end{aligned} \tag{4}$$

Since stratified samples are independent, there is no correlation between them and especially no minimum distance property. Therefore no variance reduction can be guaranteed. Minimum distance constraints can be imposed by e.g. applying Lloyd's relaxation algorithm [McCool and Fiume 1992; Haeberli and Akeley 1990]. Combined with modified sample weights by post-stratification [Glassner 1993] the convergence is very much improved. Rank-1-lattices by nature are invariant under Lloyd's relaxation method and Hickernell [Hickernell 1998] proved that uniform weights are optimal. For Riemann-integrable functions the unbiased estimator (4) has been shown [Tuffin 1996] to have a reduced variance of $\mathcal{O}\left(n^{-2} \ln^{2s} n\right)$, if the sequence of points $(R_j(\xi))_{j=0}^{n-1}$ is of low discrepancy [Niederreiter 1992].

The estimator (4) is trivial to realize by drawing one random vector $\xi \in [0,1)^s$ and replicating it by (2). The resulting correlated samples (see figure 1d) for a suited rank-1-lattice expose Latin hypercube structure, guaranteed minimum distance, and very good uniformity properties for suited choices of the generator vector $\mathbf{g}$. If x is uniformly distributed, in fact the transformation $\mathbf{B}$ in (2) can be omitted (see figures 3a) yielding

$$\begin{aligned} R_j^{\mathrm{CP}} : [0,1)^s &\to [0,1)^s \\ x &\mapsto \left(\frac{j}{n} \cdot \mathbf{g} + x\right) \bmod [0,1)^s, \end{aligned} \tag{5}$$

which is the so-called Cranley-Patterson rotation [Cranley and Patterson 1976].

3.3 Correlated Trajectory Splitting

Trajectory splitting [Sobol' 1994] can improve the efficiency depending on the correlation of the integrand with respect to its lower dimensional projections. A beneficial application example from computer graphics is tracing multiple shadow rays to an area light source for only one ray sent from the eye. Here the variance of the integrand with respect to the pixel area usually is much smaller than with respect to the the support of the light source.

Contrary to current practice, the samples split off one instance do not need to be independent. Thinking of x as the position in the pixel and y a location on the light source, we can formalize:

$$\begin{aligned} &\int_{[0,1)^{s_1}} \int_{[0,1)^{s_2}} f(x,y)\, dydx \\ &= \int_{[0,1)^{s_1}} \sum_{j=0}^{n-1} \int_{A_j} f(x,y)\, dydx \\ &= \int_{[0,1)^{s_1}} \frac{1}{n} \int_{[0,1)^{s_2}} \sum_{j=0}^{n-1} f(x, R_j(y))\, dydx \\ &\approx \frac{1}{r} \sum_{i=0}^{r-1} \frac{1}{n} \sum_{j=0}^{n-1} f(\xi_i, R_j(\zeta_i)). \end{aligned} \tag{6}$$

If now Monte Carlo integration is applied for integral estimation, by exchanging the sum and the integral the n split samples for one random realization are correlated due to the mapping R_j. In fact this

Cranley Patterson rotation

$$R_j^{CP}(\zeta_i) = \left(\frac{j}{n}\cdot\mathbf{g}+\zeta_i\right) \bmod [0,1)^s$$

reduced
Cranley Patterson rotation

$$R_j(\zeta_i) = \left(\frac{j}{n}\cdot\mathbf{g}+\mathbf{B}\zeta_i\right) \bmod [0,1)^s$$

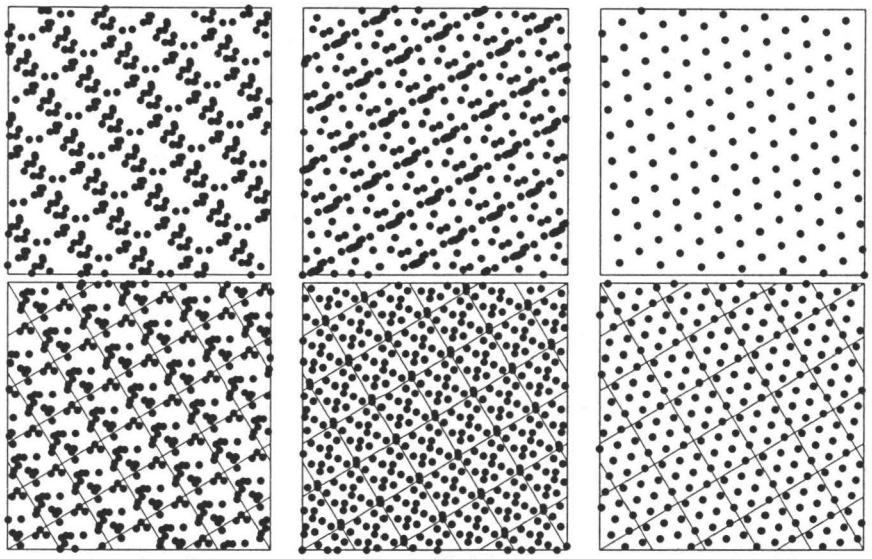

a) random ζ_i b) stratified ζ_i c) correlated ζ_i

Figure 3: Comparison of different replications of the $n = 34$ point Fibonacci lattice with $\mathbf{g} = (1, 21)$ (see sections 3.2 and 3.3). The top row shows the replications of the points $\zeta_i \in [0,1)^s$ by Cranley Patterson rotations R_j^{CP}, whereas the bottom row shows the reduced Cranley Patterson rotations R_j. For a) 15 random points ζ_i there is no difference, whereas the structure of 3×5 stratified samples ζ_i becomes resolved by the Cranley Patterson rotation in b). The character of the reduced Cranley Patterson rotation becomes even more obvious for correlated samples: In c) the third and fourth dimension of a (shifted) $(8, 3)$ Korobov lattice are replicated. Using the Cranley Patterson rotation replicated points coincide, since 8 and 34 are not relatively prime (instead of $8 \times 34 = 272$ only 136 different points are generated). The reduced Cranley Patterson rotation avoids the coincidence of points by stratification and preserves the structure of the ζ_i.

sampling scheme then corresponds to subdividing the light source into smaller ones, i.e. stratification by rank-1 lattices. However in each stratum the same sample position is used and thus the samples of one instance i preserve their good properties of uniform distribution although being independent from instance to instance.

In figure 3 the application of Cranley Patterson rotations and reduced Cranley Patterson rotations is illustrated. While the standard Cranley-Patterson rotation [Cranley and Patterson 1976] (5) can destroy the structure of the samples ζ_j, replicating the samples using reduced Cranley Patterson rotations (2) preserves their structure.

Due to its tensor product structure trajectory splitting should be applied only once along a trajectory. Otherwise the recursive application of the mapping (2) yields exponential complexity.

4 Applications

The new unbiased Monte Carlo estimators have many applications in image synthesis. We illustrate the principles from the previous sections for some representative examples.

4.1 Photon Map Generation

The photon map method [Jensen 2001] simulates random walks of portions of light starting from the sources through a synthetic scene. The photons, i.e. packages of energy, are stored in their location of incidence. Then elements of density estimation are applied for irradiance estimation.

A random walk is determined by a high-dimensional vector $\xi \in [0,1)^s$ of random numbers (For the explicit details we refer to Jensen's book [Jensen 2001]). High-dimensional stratification has not been available so far [Jensen 2001]. However using rank-1-lattices provides a simple solution to that problem: The j-th random walk is determined by the stratified sample $R_j(\xi_j)$ as defined in (2).

In figure 5 the locations of incidence of the photons are displayed for an example scene. Using the lattice points $R_j(0)$ directly results in obvious correlation patterns. These patterns of course persist randomization by correlated sampling (4), i.e. $R_j(\xi)$, since locally close photon map queries still reveal the just shifted correlation patterns. However the correlations are completely resolved by the independent stratified samples $R_j(\xi_j)$.

The general proof for the rate of stratified sampling (3) only guarantees that the convergence rate cannot be worse. However hardly any improvement in the convergence rate can be observed for the general photon map method. This relates to observations from the domain of quasi-Monte Carlo integration [Keller 1996b; Keller 1996a], where a slightly faster but much smoother convergence has been observed. The reason is the small support of the estimation kernel that selects only a small fraction of the random walks and thus attenuates variance reduction. However benefits are to be expected in high density regions as caused by convergent beams of light, e.g. by caustics.

4.2 Anti-Aliasing

Since in most cases it is impossible to fulfil Shannon's sampling theorem, aliases are unavoidable. However they can be masked by noise from randomization. To achieve this the samples must be independent for different pixels, however can be correlated in one pixel. The estimator (4) perfectly matches these requirements: The samples in one pixel are a shifted lattice exposing very good uniformity of distribution and a maximized minimum distance guaranteeing for a fast convergence. For different pixels they are independent still pushing aliasing artifacts to noise.

In figure 6 jittered sampling is compared to Cranley-Patterson rotated Fibonacci lattices. The enlargements are taken from the scene completely shown in figure 4. Clearly the correlated sampling scheme outperforms jittered sampling: At a lower sampling rate aliasing and noise noticeably are reduced, which becomes even

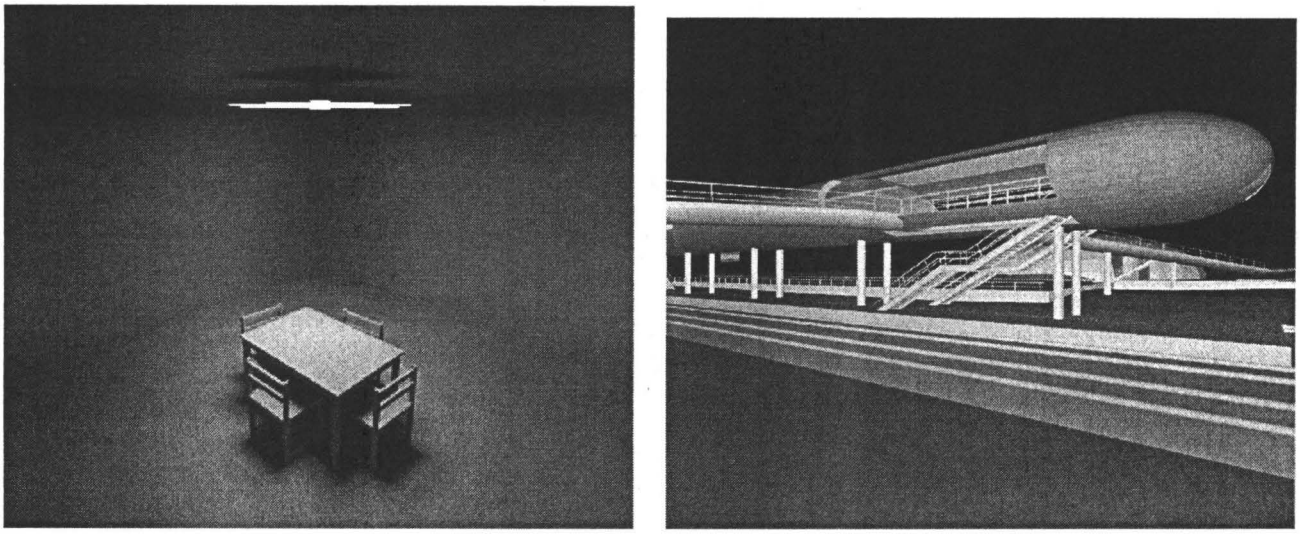

Figure 4: The test scenes. The office is used for photon map simulation in figure 5. The train station is used for the anti-aliasing, path tracing, and distribution ray tracing numerical experiments.

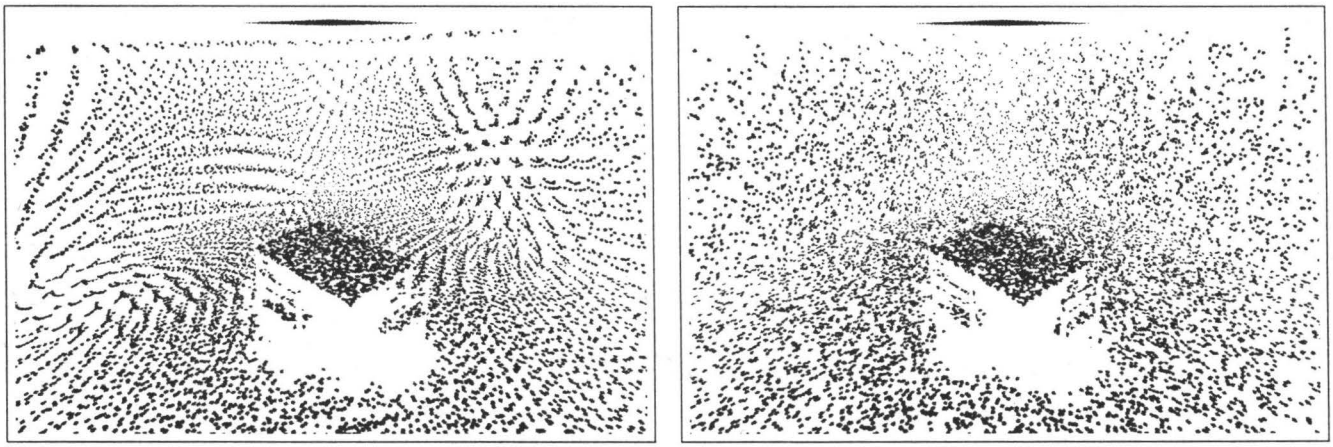

Figure 5: Direct visualization of a photon map: Using the lattice $(2^m, 17797)$ of Korobov form at $m = 14$ exposes the correlation of the lattice points as clearly visible in the patterns on the left. Jittering the lattice by randomly sampling each lattice cell once (right) completely resolves the structures yielding stratified random walks.

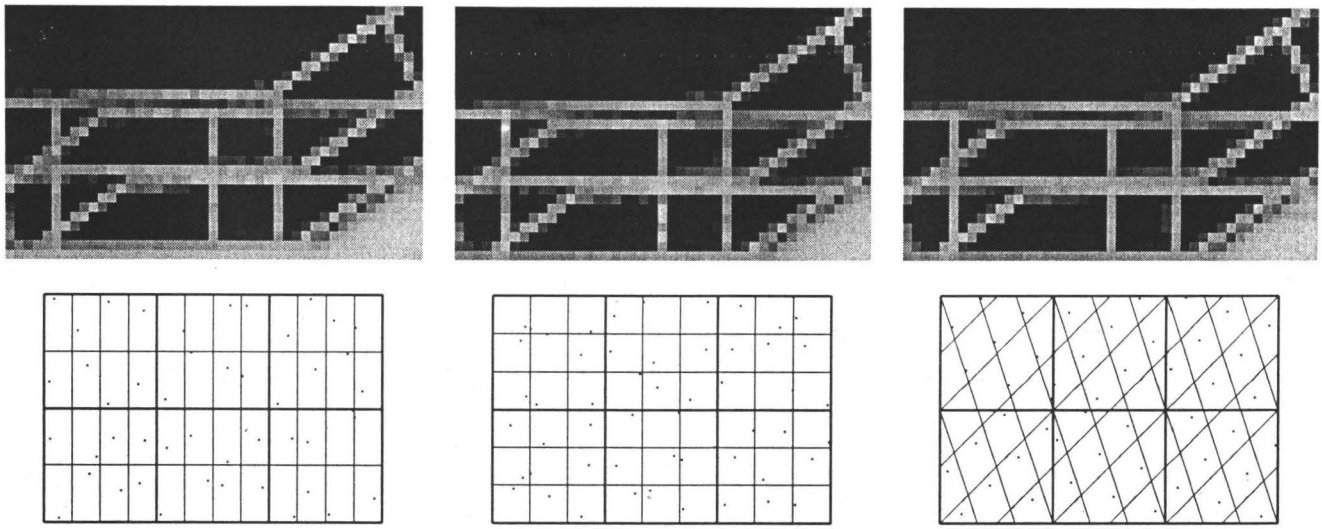

Figure 6: Pixel anti-aliasing by 4×2 jittered sampling (left), 3×3 jittered sampling (middle), and Cranley-Patterson rotated Fibonacci lattices at $n = F_6 = 8$ points (right). The latter correlated sampling scheme at only 8 samples outperforms jittered sampling at 9 samples by reproducing the fine structures of the handrail with less noise and aliasing. The resulting reduction of flicker is even more noticeable in animations.

more apparent in animations. As compared to less correlated sampling schemes, the faster convergence is obtained by a much more uniform coverage of the pixel. Sample generation effort is almost negligible and much less random numbers are required as compared to previous schemes.

4.3 Path Tracing

Contrary to the photon map generation path tracing uses all trajectories for integral computation. Consequently the efficiency of correlated sampling (4) can be exploited.

In figure 7 one randomly shifted lattice per pixel is used for anti-aliasing and illumination by a very long thin area light source. This is a difficult situation for blind stratification, since it can happen that the shorter side is stratified better than the longer (which of course is counterintuitive). The lattices however perfectly handle this case: Since Korobov lattices with $\gcd(a, n) = 1$ are instances of Latin hypercube samples, they have perfectly stratified projections and a minimum distance property (see also figure 1). The comparison shows that the Korobov lattice $(8, 3)$ in $s = 4$ performs as good as the best case of a 9 sample multidimensional stratified pattern. It such is more efficient, simpler to implement and more robust to use.

4.4 Distribution Ray Tracing

Distribution ray tracing can be regarded as extending path tracing by trajectory splitting. For one dimension correlated splitting coincides with random offset sampling [Pauly et al. 2000; Keller and Heidrich 2001] as introduced for volume rendering. The advantages of faster convergence at improved speed transfer to correlated trajectory splitting in higher dimensions as e.g. final gathering or illumination by area light sources.

In figure 8 we illustrate splitting by reduced Cranley-Patterson rotations. We use the same setting as in the previous section, i.e. the four dimensional samples (ξ_i, ζ_i) of a randomly shifted Korobov lattice $(8, 3)$ are replicated on the light source using correlated splitting (6) by the Fibonacci lattice $(F_9, F_8) = (34, 21)$. For one instance the resulting points $R_j(\zeta_i)$ on the light source are those shown in bottom row of figure 3c. Clearly the points are much more uniformly distributed than jittered samples yielding a much better convergence as can be seen in the images of figure 8. The implementation is almost trivial. Randomness is reduced to only one random shift and such aliases still are pushed to noise, however the noise level is much lower as compared to uncorrelated sampling.

5 Deterministic Application of Rank-1-Lattices

The ideas of strictly deterministic sampling schemes for image synthesis [Keller 2000] in a straightforward way apply to the path tracing and distribution ray tracing algorithm from the previous sections: Instead of the random shifts deterministic low discrepancy points sets are applied, i.e. the Monte Carlo quadratures are replaced by interleaved [Keller 2000] quasi-Monte Carlo quadrature rules.

Some of the Fibonacci lattices implicitly have been used in computer graphics before. We briefly point out how these applications fit into our general framework.

Rasterization algorithms are optimized by exploiting the structure of regular grids, which correspond to the lattice $L = \mathbb{Z}^2$. The Quincunx [Bouville et al. 1991] sampling pattern as used for anti-aliasing by the nVidia graphics boards embodies the Fibonacci lattice with $n = F_3 = 2$ points. $\sqrt{5}$-sampling [Stamminger and Drettakis 2001] in fact is a special case of the recursive application of (2) using the Fibonacci-lattice at $n = F_5 = 5$ points.

6 Conclusion and Future Work

Using the classical rank-1-lattices in a completely novel way we introduced more efficient sampling schemes. Randomness has been reduced to the necessary minimum to still result in unbiased Monte Carlo estimates. The intrinsic correlation preserves the smoothness of continuously varying parameters when evaluating a family of integrals thus very much improving image quality and interpolation algorithms. The superior performance has been demonstrated by examples.

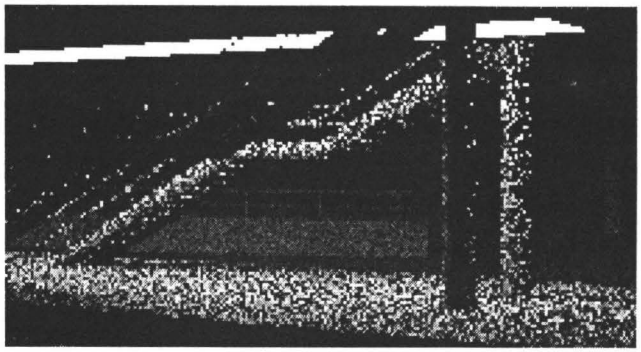

a) 1×9 jittered sampling

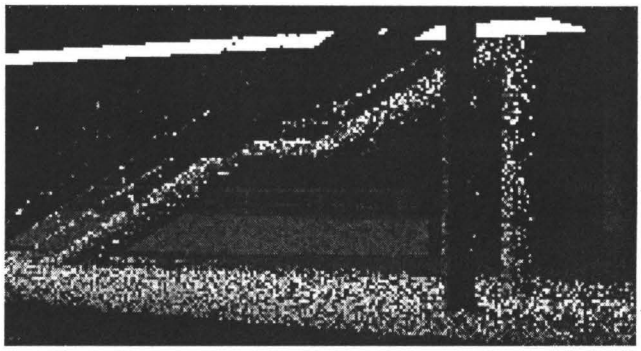

b) 3×3 jittered sampling

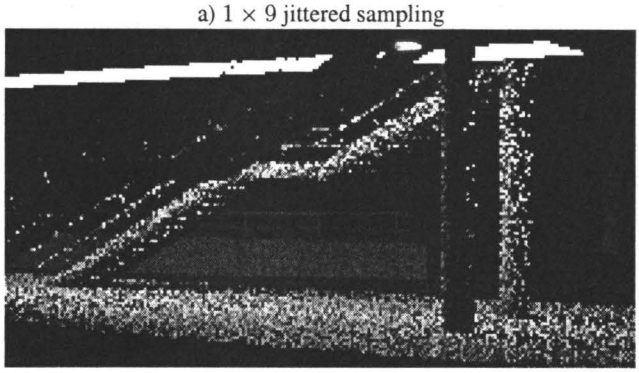

c) 9×1 jittered sampling

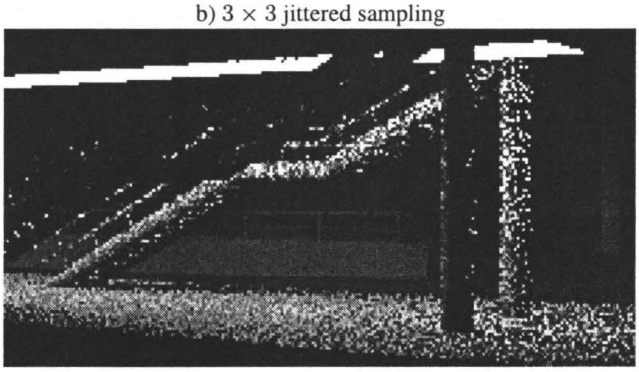

d) correlated sampling by an $(8,3)$ Korobov lattice

Figure 7: Comparison of multidimensional sampling. Pixel samples are 3×3 stratified in a) to c), the lattice samples in d) are those from figure 6. The light source samples in a) and b) are chosen not optimal, since they do not stratify the light source along the maximum side length. Choosing the stratification along the longest side in c) clearly reduces noise. The randomly shifted four-dimensinoal $(8,3)$ Korobov lattices in d) perform as good as the optimally stratified jittered sampling at only 8 samples. Due to their intrinsic correlation any low dimensional projection is good and no extra care has to be taken of the stratification.

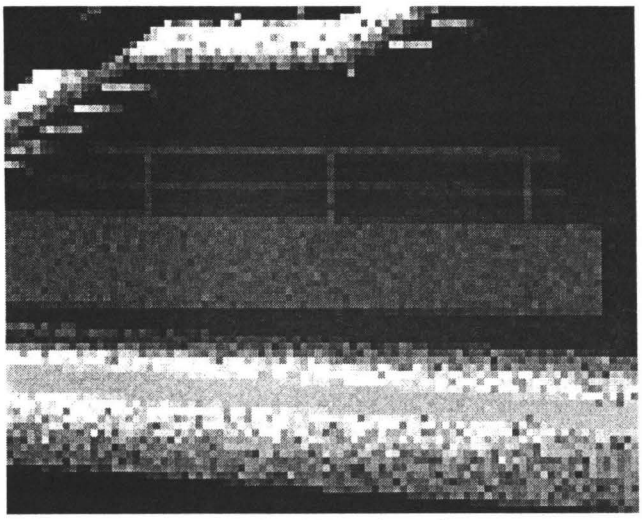

uncorrelated jittered sampling

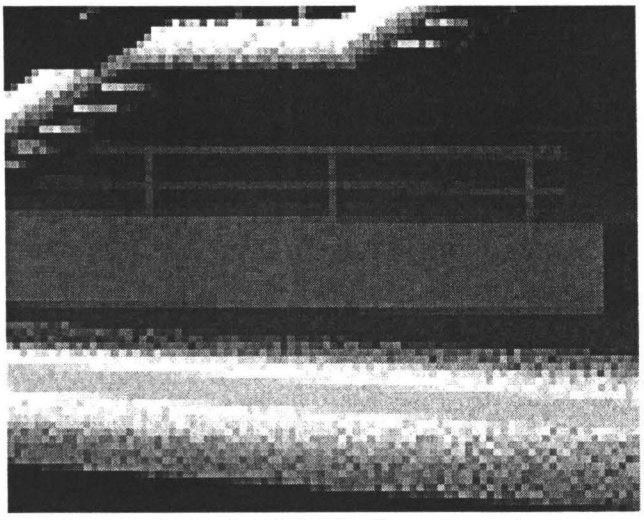

correlated reduced Cranley-Patterson rotation

Figure 8: Correlated sampling of a long thin light source in a distribution ray tracer: Jittered sampling (left) and samples from a rank-1-lattice replicated using reduced Cranley-Patterson rotations (right). At the same computational effort the correlated trajectory splitting noticeably reduces the noise level. Due to the projection regularity of the shifted lattice the long thin light source is stratified much better than jittered sampling can guarantee.

The repeated application of the new estimators allows for adaptive sampling and in addition for variance estimation, i.e. simple error control [Cranley and Patterson 1976]. However, further investigations will consider nested sequences of lattices [Hickernell et al. 2001]. Also the extension to bidirectional path tracing has to be investigated, where initial investigation have been done by Kollig et al. [Kollig and Keller 2002].

Since the number of strata determined by Voronoi tessellations of rank-1-lattices can be chosen independent of dimension, it is highly interesting to investigate the resulting domain discretization in connection with e.g. the solution of partial differential equations and the Fourier transform on rank-1-lattices [Keller 2001].

References

BOUVILLE, C., TELLIER, P., AND BOUATOUCH, K. 1991. Low Sampling Densities using a psychovisual approach. In *Eurographics '91*, Elsevier Science Publishers, Amsterdam, North-Holland, 167–182.

COOK, R., PORTER, T., AND CARPENTER, L. 1984. Distributed Ray Tracing. In *Computer Graphics (SIGGRAPH 84 Conference Proceedings)*, 137–145.

CRANLEY, R., AND PATTERSON, T. 1976. Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis 13*, 904–914.

GLASSNER, A. 1993. Dynamic Stratification. In *Proc. 4th Eurographics Workshop on Rendering*, 5–14.

HABER, S. 1983. Parameters for integrating periodic functions of several variables. *Math. of Computation 41*, 163, 115–129.

HAEBERLI, P., AND AKELEY, K. 1990. The Accumulation Buffer: Hardware Support for High-Quality Rendering. In *Computer Graphics (SIGGRAPH 90 Conference Proceedings)*, 309–318.

HICKERNELL, F., HONG, H., L'ÉCUYER, P., AND LEMIEUX, C. 2001. Extensible lattice sequences for quasi-Monte Carlo quadrature. *SIAM J. Sci. Comput. 22*, 1117–1138.

HICKERNELL, F. 1998. Lattice rules: How well do they measure up? In *Random and Quasi-Random Point Sets, volume 138 of Lecture Notes in Statistics*, P. Hellekalek and G. Larcher, Eds. Springer, 109–166.

JENSEN, H. 2001. *Realistic Image Synthesis Using Photon Mapping*. AK Peters.

KALOS, M., AND WHITLOCK, P. 1986. *Monte Carlo Methods, Volume I: Basics*. J. Wiley & Sons.

KELLER, A., AND HEIDRICH, W. 2001. Interleaved Sampling. In *Rendering Techniques 2001 (Proc. 12th Eurographics Workshop on Rendering)*, Springer, K. Myszkowski and S. Gortler, Eds., 269–276.

KELLER, A. 1996. Quasi-Monte Carlo Methods in Computer Graphics: The Global Illumination Problem. *Lectures in App. Math. 32*, 455–469.

KELLER, A. 1996. Quasi-Monte Carlo Radiosity. In *Rendering Techniques '96 (Proc. 7th Eurographics Workshop on Rendering)*, Springer, X. Pueyo and P. Schröder, Eds., 101–110.

KELLER, A. 2000. Strictly deterministic sampling methods in computer graphics. Tech. rep., mental images, Berlin, Germany.

KELLER, A. 2001. Random Fields on Rank-1-Lattices. Interner Bericht 307/01, University of Kaiserslautern.

KOLLIG, T., AND KELLER, A. 2002. Efficient bidirectional path tracing by randomized quasi-Monte Carlo integration. In *Monte Carlo and Quasi-Monte Carlo Methods 2000*, H. Niederreiter, K. Fang, and F. Hickernell, Eds. Springer, 290–305.

MCCOOL, M., AND FIUME, E. 1992. Hierarchical Poisson Disk Sampling Distributions. In *Proceedings of Graphics Interface '92*, 94–105.

MITCHELL, D. 1996. Consequences of Stratified Sampling in Graphics. In *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, 277–280.

NIEDERREITER, H. 1992. *Random Number Generation and Quasi-Monte Carlo Methods*. SIAM, Pennsylvania.

PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis Light Transport for Participating Media. In *Rendering Techniques 2000 (Proc. 11th Eurographics Workshop on Rendering)*, Springer, B. Péroche and H. Rushmeier, Eds., 11–22.

SLOAN, I., AND JOE, S. 1994. *Lattice Methods for Multiple Integration*. Clarendon Press, Oxford.

SOBOL', I. 1994. *A Primer for the Monte Carlo Method*. CRC Press.

STAMMINGER, M., AND DRETTAKIS, G. 2001. Interactive Sampling and Rendering for Complex and Procedural Geometry. In *Rendering Techniques 2001 (Proc. 12th Eurographics Workshop on Rendering)*, Springer, K. Myszkowski and S. Gortler, Eds., 151–162.

TUFFIN, B. 1996. On the Use of Low Discrepancy Sequences in Monte Carlo Methods. *Monte Carlo Methods and Applications 2*, 4, 295–320.