



Cascaded Sobol' Sampling

Lois Paulin, David Coeurjolly, Jean-Claude Iehl, Nicolas Bonneel, Alexander Keller, Victor Ostromoukhov

► To cite this version:

Lois Paulin, David Coeurjolly, Jean-Claude Iehl, Nicolas Bonneel, Alexander Keller, et al.. Cascaded Sobol' Sampling. ACM Transactions on Graphics, 2021, Proceedings of SIGGRAPH ASIA 2021, 40 (6), pp.274:1–274:13. 10.1145/3478513.3480482 . hal-03358957

HAL Id: hal-03358957

<https://hal.science/hal-03358957>

Submitted on 29 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Cascaded Sobol' Sampling

LOÏS PAULIN, Université de Lyon, CNRS, LIRIS, France
 DAVID COEURJOLLY, Université de Lyon, CNRS, LIRIS, France
 JEAN-CLAUDE IEHL, Université de Lyon, CNRS, LIRIS, France
 NICOLAS BONNEEL, Université de Lyon, CNRS, LIRIS, France
 ALEXANDER KELLER, NVIDIA, Germany
 VICTOR OSTROMOUKHOV, Université de Lyon, CNRS, LIRIS, France

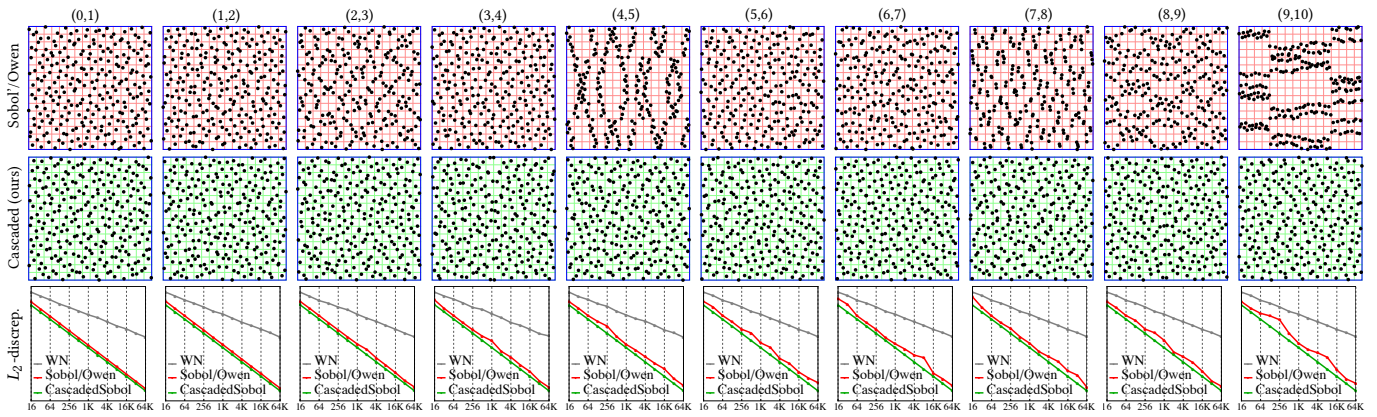


Fig. 1. **Cascaded Sobol' Point Sets.** For quasi-Monte Carlo integration problems, low discrepancy samplers, such as the Sobol' sequence with Owen scrambling [Owen 1998; Sobol' 1967], are widely used thanks to their ease of generating high-dimensional point sets. While being low discrepancy in high dimension, some projections may exhibit strong uniformity defects (illustrated here by consecutive 2-dimensional projections of an 11-dimensional point set in the first row). We propose a sampling strategy with perfect $(0, m, 2)$ -net properties (second row) for consecutive pairs of dimensions and optimized low discrepancy in high dimension, reducing errors in Monte Carlo rendering. The third row shows L_2 -discrepancies of first 10 consecutive 2-dimensional projections; see Figure 7 for s -dimensional discrepancies.

Rendering quality is largely influenced by the samplers used in Monte Carlo integration. Important factors include sample uniformity (e.g., low discrepancy) in the high-dimensional integration domain, sample uniformity in lower-dimensional projections, and lack of dominant structures that could result in aliasing artifacts. A widely used and successful construction is the Sobol' sequence that guarantees good high-dimensional uniformity and consequently results in faster convergence of quasi-Monte Carlo integration. We show that this sequence exhibits low uniformity and dominant

structures in low-dimensional projections. These structures impair quality in the context of rendering, as they precisely occur in the 2-dimensional projections used for sampling light sources, reflectance functions, or the camera lens or sensor. We propose a new cascaded construction, which, despite dropping the sequential aspect of Sobol' samples, produces point sets exhibiting provably perfect dyadic partitioning (and therefore, excellent uniformity) in consecutive 2-dimensional projections, while preserving good high-dimensional uniformity. By optimizing the initialization parameters and performing Owen scrambling at finer levels of binary representations, we further improve over Sobol's integration convergence rate. Our method does not incur any overhead as compared to the generation of the Sobol' sequence, is compatible with Owen scrambling and can be used in rendering applications.

CCS Concepts: • **Computing methodologies** → **Rendering**.

Additional Key Words and Phrases: Sobol' sequence, quasi-Monte Carlo integration, low-discrepancy sequences, Owen scrambling, path tracing

ACM Reference Format:

Loïs Paulin, David Coeurjolly, Jean-Claude Iehl, Nicolas Bonneel, Alexander Keller, and Victor Ostromoukhov. 2021. Cascaded Sobol' Sampling. *ACM Trans. Graph.* 40, 6 (September 2021), 13 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

Authors' addresses: Loïs Paulin, Université de Lyon, CNRS, LIRIS, France, lois.paulin@ens-lyon.fr; David Coeurjolly, Université de Lyon, CNRS, LIRIS, France, david.coeurjolly@liris.cnrs.fr; Jean-Claude Iehl, Université de Lyon, CNRS, LIRIS, France, jean-claude.iehl@univ-lyon1.fr; Nicolas Bonneel, Université de Lyon, CNRS, LIRIS, France, nicolas.bonneel@liris.cnrs.fr; Alexander Keller, NVIDIA, Germany, akeller@nvidia.com; Victor Ostromoukhov, Université de Lyon, CNRS, LIRIS, France, victor.ostromoukhov@liris.cnrs.fr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2021/9-ART \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

The numerical evaluation of integrals is a core computer graphics research problem, notably for rendering realistic images of 3D scenes. Many physically-based rendering techniques rely on the random sampling of an integrand, a process called Monte Carlo integration. For path tracing, it typically consists of following paths from the camera to light sources by randomly bouncing rays in the scene. In many cases, obtaining noise-free images requires computing hundreds to thousands of such paths. Improving the convergence rate of this integral estimator can be achieved by replacing these random values by samples that are particularly well distributed over the integration domain in a highly uniform fashion. Intuitively, correlating samples to avoid holes and clusters on the domain makes the estimate more efficient. This can be formally assessed by various uniformity measures and associated variance reduction theorems, such as the discrepancy and the Koksma-Hlawka inequality [Hlawka 1961].

Several techniques exist to obtain *low discrepancy samples*, either relying on sequences of values uniformly covering the domain by construction [Lemieux 2009] or finely optimizing point sets of fixed cardinality by minimizing well chosen energies [Keller 2013]. Among these options, the Sobol' sequence has gained significant popularity for rendering since it is fast and produces very well distributed samples that effectively reduce noise in rendered images.

However, it has been suggested that uniformity over the integration domain is not sufficient, and that uniformity over the two-dimensional projections used for sampling reflectance, light sources, camera lenses or sensors, is also important to improve image quality [Ahmed and Wonka 2020; Paulin et al. 2020; Perrier et al. 2018; Reinert et al. 2016]. In this context, we show that the popular Sobol' sequence does not always satisfy this requirement, and that consecutive pairs of dimensions can produce very poor distributions in 2-dimensional projections. They may exhibit dominant structures that may result in aliasing artifacts (see Figs. 9 and 10).

We propose to alleviate this problem by introducing a new sampler based on consecutive calls to Sobol' functions (a construction we call *cascaded Sobol' sampling*), which we prove to provide well distributed low discrepancy point sets in consecutive pairs of dimensions. For now, this comes at a cost: we cannot preserve the sequential aspect of the original Sobol' algorithm, and hence need to fix the number of samples in advance. Still, in addition to uniformity over 2-dimensional projections, we preserve uniformity over the high-dimensional integration domain by optimizing initialization tables over a range of sample cardinalities and dimensions useful for computer graphics applications [Joe and Kuo 2008].

We also propose a technique that improves uniformity, and hence the convergence rate of Monte Carlo rendering. While the deterministic Sobol' sequence is often accompanied with a randomization strategy – for example, Owen scrambling [Owen 1998], see Section 3.2 – we show that increasing the bit depth of this randomization technique allows one to more uniformly distribute samples in the domain of integration and to optimize the generation.

Aside from the predetermined number of samples, our approach can act as a drop-in replacement of Sobol' samplers in existing rendering engines to effectively obtain faster convergence. Our

sampler does not incur significant overhead as compared to standard Sobol' samplers and is simple to implement.

We now summarize our contributions. First, we propose a *cascaded Sobol' sampling* construction that provably yields sample uniformity in consecutive pairs of dimensions contrary to the original Sobol' sequence. Second, we compute and provide optimized initialization tables that ensure uniformity in high dimension as well.

We evaluate our method on characteristic integration and rendering examples, showing competitive convergence rates as compared to Sobol' and other state-of-the-art samplers.

2 RELATED WORK

Monte Carlo Integration. The beauty of Monte Carlo integration is its simplicity yielding a convergent estimate by just averaging evaluations of the integrand at independent random points in the integration domain. Yet, introducing some correlations within the samples may improve the convergence rate (see for instance [Singh et al. 2019]). Enhancing the uniformity of the samples can be done in many ways. One can stratify the domain while keeping the stochastic nature of the process (for example, jittered sampling or [Christensen et al. 2018]), we can optimize a point set minimizing some objective functions [Balzer et al. 2009; Bridson 2007; Fattal 2011; Heck et al. 2013; Ostromoukhov et al. 2004; Paulin et al. 2020; Zhou et al. 2012], or rely on arithmetic and algebraic properties of lattices to generate samples [Grünschloß et al. 2008; L'Ecuyer and Munger 2016; Liu et al. 2021] or using low discrepancy sequences (for example, [Halton 1964; Niederreiter 1992; Sobol' 1967]). Our proposal belongs to the latter category since low discrepancy samplers imply fast sample generation in high dimensions and result in the fastest convergence speed thanks to the Koksma-Hlawka inequality [Hlawka 1961] that bounds the integration error by the product of the discrepancy of the point set and the variation of the integrand.

Low Discrepancy Sequences and Generator Matrices. Low discrepancy sequences often rely on a combinatorial approach to sets and permutations, and Galois field arithmetic [Dick and Pillichshammer 2010; Lemieux 2009; Niederreiter 1992]. Many of the constructions generate points as follows: to retrieve the i -th sample in the s -dimensional unit cube $[0, 1]^s$ of a given point set, we express i as a vector in some integer base b and compute the j -th component of the sample by multiplying the vector by a so-called *generator matrix* using operations on finite fields. Uniformity properties of the point set are obtained by structural properties of the generator matrices used for all dimensions. There exist many algebraic constructions of such matrices [Halton 1964; Niederreiter 1992; Sobol' 1967]. Besides construction, there is a number of efforts to identify good generator matrices by optimization [Ahmed et al. 2016; Grünschloß et al. 2008; Perrier et al. 2018]. We propose a new construction for generator matrices resulting in high uniformity in low-dimensional projections and good discrepancy in high dimension, outperforming classical approaches.

Monte Carlo Rendering. In the specific case of Monte Carlo rendering additional properties besides the uniformity of distribution need

to be considered. There are two main ways to sample the image plane. One can globally sample the image and attribute samples to the pixel they fall in [Grünschloß et al. 2012] or one can sample each pixel individually. The second method requires to take special care to avoid similarity in the samples used across pixels. One thus may randomize deterministic samplers, for example by techniques [Cranley and Patterson 1976; Kuipers and Niederreiter 2012; Owen 1998] that preserve the uniformity of distribution. Among these strategies, we elaborate on Owen scrambling [Owen 1998] as it preserves the low discrepancy properties of certain point sets while improving on practical aspects of Monte Carlo rendering [Burley 2020; Perrier et al. 2018].

In rendering, the function to be integrated exposes a low-dimensional structure implied by how paths are traced through a scene for light transport simulation. If samples are not highly uniformly distributed with respect to such low dimensional projections, a decreased rendering quality is the consequence [Paulin et al. 2020; Perrier et al. 2018; Reinert et al. 2016]. This leads to approaches disregarding the requirement for high-dimensional uniformity and using independent samples for each bounce [Ahmed and Wonka 2020] or the ZeroTwo sampler [Pharr et al. 2016]. We show that this results in a loss of integration quality and that high-dimensional uniformity is beneficial in addition.

More recent complementary research tackled the issue of distributing the integration error across the picture in a way that makes it less perceptible to the human eye. These methods worked either by optimizing sample set distributions across the pixel grid [Heitz and Belcour 2019; Heitz et al. 2019] or by using algebraic properties to reorder a global sampler [Ahmed and Wonka 2020].

In this paper we will introduce a new way to construct generator matrices for low discrepancy point sets with highly uniform low-dimensional projections across consecutive dimensions that are also compatible with state of the art error diffusion methods.

3 PRELIMINARIES

For the purpose of the paper, we assume familiarity with the basics of quasi-Monte Carlo integration and low discrepancy sequences, especially with the concepts of (t, s) -sequences and (t, m, s) -nets that are well established in graphics [Pharr et al. 2016] and refer to the textbooks by Niederreiter [1992], Lemieux [2009], or Dick and Pillichshammer [2010]. In what follows, we recall the details of the Sobol' low discrepancy sequence and Owen scrambling that are required to establish our contribution.

3.1 The Sobol' Sequence

The Sobol' sequence [Sobol' 1967] is one of the most popular high-dimensional low discrepancy sequences used for quasi-Monte Carlo integration. The j -th component of the i -th point x_i is computed by

$$x_{i,j} := \sum_{k=0}^{m-1} b_k \cdot 2^{k-m} \in [0, 1), \text{ where} \quad (1)$$

$$(b_{m-1}, \dots, b_0) := (a_0, \dots, a_{m-1}) \cdot C_j^T. \quad (2)$$

(a_0, \dots, a_{m-1}) denotes the (row) vector representation of i in base 2 (a_0 being the least-significant digit). The matrix multiplication is performed in the Galois field \mathbb{F}_2 (or $\text{GF}(2)$). Vector operations on

\mathbb{F}_2^m can be efficiently implemented as bit-vector XOR and AND operations for addition and multiplication, respectively.

The so-called *generator matrix* C_j is determined by the j -th primitive polynomial over \mathbb{F}_2 . The sequence of these primitive polynomials is enumerated as increasing integers¹. As the number-theoretic construction of the generator matrices is beyond the scope of our article (see the aforementioned standard textbooks), we focus on the properties and algorithmic aspects of the Sobol' sequence.

Using the generator matrices C_j as depicted in Figure 2, we exemplify the generation process for $m = 4$ and $i = 13_{10} = 1101_2$ that is the 14-th point of the Sobol' sequence: the computation of the first component yields $(b_3, b_2, b_1, b_0) = (1, 1, 0, 1) \cdot C_0^T = (1, 0, 1, 1)$, resulting in $x_{13,0} = 0.6875$. The second and third component are $(b_3, b_2, b_1, b_0) = (1, 1, 0, 1) \cdot C_1^T = (1, 1, 0, 1) \Rightarrow x_{13,1} = 0.8125$ and $(b_3, b_2, b_1, b_0) = (1, 1, 0, 1) \cdot C_2^T = (0, 1, 1, 1) \Rightarrow x_{13,2} = 0.4375$, respectively.

The Sobol' sequence is a (t, s) -sequence in base 2 [Niederreiter 1992], where s denotes the dimension and t is a non-negative integer determining the quality of the points. The smaller t , the more uniformly distributed is the point set. For Sobol's construction [Sobol' 1967], the parameter t is the sum of the degrees (minus one) of the primitive polynomials used for the generator matrices C_0, \dots, C_{s-1} . As a consequence, the higher j , the higher t , limiting the uniformity, especially of low-dimensional projections. Hence, constructing a (t, s) -sequence or a (t, m, s) -net with minimal t has high practical impact and is the major challenge in the quasi-Monte Carlo community.

By construction, the generator matrices C_j of the Sobol' sequence are full rank. In fact, they are infinite-dimensional and become finite only by selecting a maximum number m of bits to operate on. Joe and Kuo [2008] published optimized generator matrices up to dimension $s = 21201$ that are widely used in finance and graphics (see Figure 2).

As the generator matrices C_j are full rank, they are bijections. Hence setting $N := 2^m$, we can define σ_j^N as the permutation on $\{0, \dots, N-1\}$ that relates the integers i ($a_{m-1} \dots a_0$ in base 2) and b (resp. $b_{m-1} \dots b_0$) by

$$(b_0, \dots, b_{m-1}) = (a_0, \dots, a_{m-1}) \cdot C_j^T \cdot J_m, \quad (3)$$

where J_m is the antidiagonal $m \times m$ unit matrix. This allows us to rewrite the first N points of the Sobol' sequence as

$$\mathbf{x}_i := \frac{1}{N} \begin{pmatrix} \sigma_0^N(i) \\ \sigma_1^N(i) \\ \vdots \\ \sigma_{s-1}^N(i) \end{pmatrix}^T. \quad (4)$$

3.2 Owen Scrambling

Owen scrambling [Owen 1998] is a method to randomize low discrepancy sequences while preserving their low discrepancy properties. It consists of constructing independent trees Π_j of random permutations, one per dimension j , and applying these permutations to the digits of the j -th component of the points. More precisely, if

¹ A058947 sequence in Sloane's *On-Line Encyclopedia of Integer Sequences* [Sloane 2017]

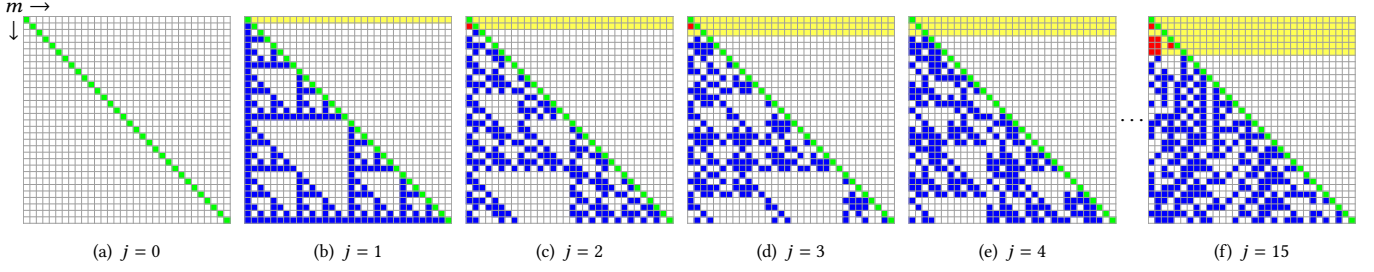


Fig. 2. **Sobol' generator matrices:** the triangular binary matrices C_j^T introduced by Joe and Kuo [2008] (dimensions $j = 0, 1, 2, 3, 4$, and 15). Yellow lines represent the so-called *direction vectors*, i.e. digits that may be initialized or optimized freely. The number of yellow lines is equal to the degree n of the primitive polynomial used (e.g., $x^6 + x^5 + 1$ of degree $n = 6$ for dimension $j = 15$). Red squares correspond to the '1's in the direction vectors as proposed by Joe and Kuo [2008]. Green squares correspond to mandatory '1's in the matrix that cannot be changed, according to Sobol's construction. Blue squares correspond to the '1's induced by the corresponding primitive polynomials within Sobol's construction. Shown matrices are of size 32×32 , which can be used for generating up to 2^{32} points. The construction is suitable for matrices of arbitrary size.

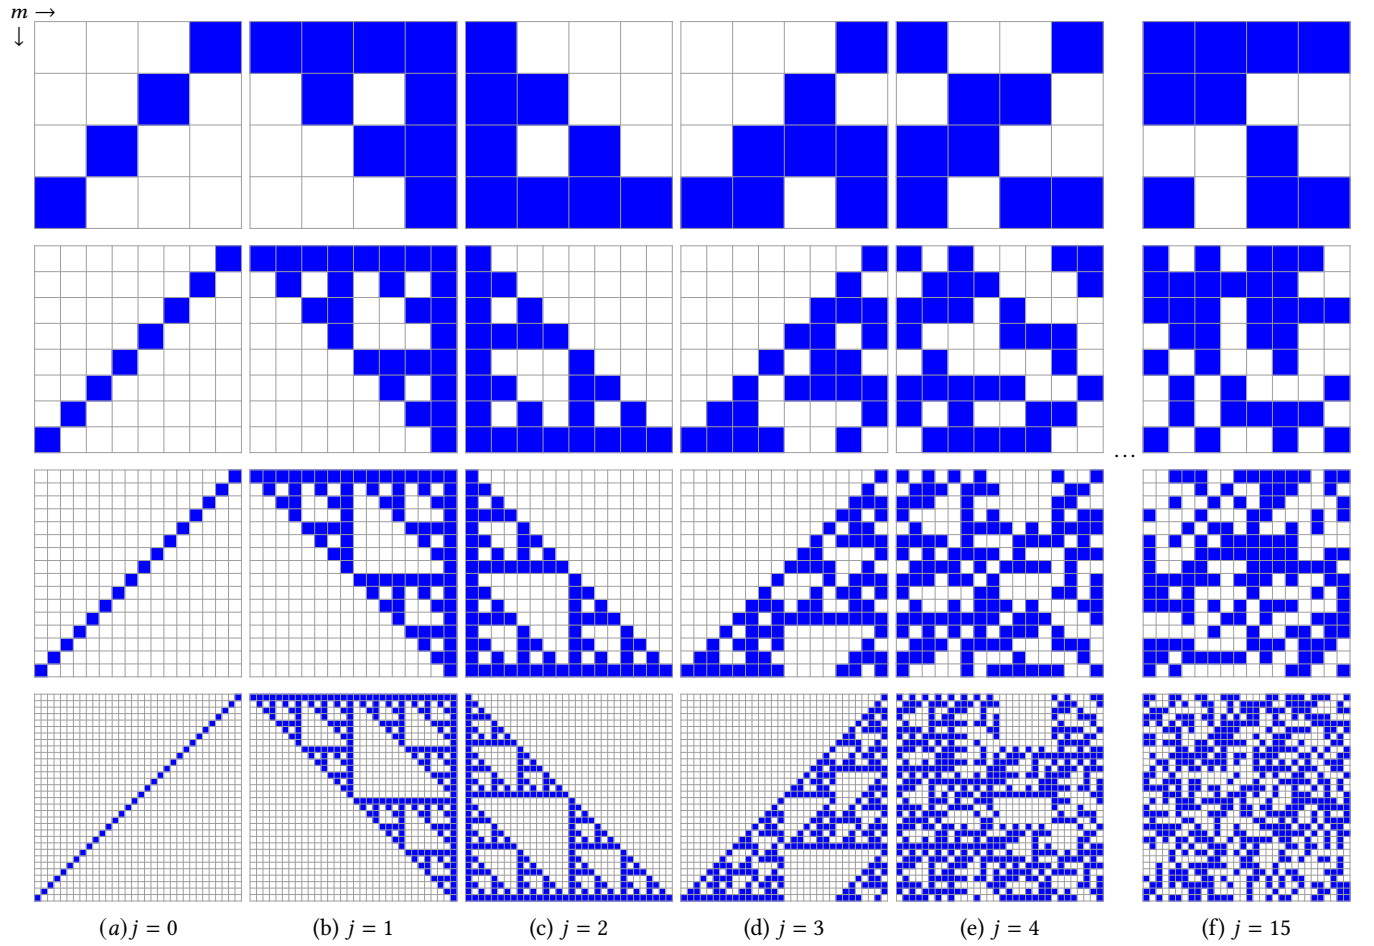


Fig. 3. **Optimized cascaded Sobol' generator matrices M_j** for generating 2^4 , 2^8 , 2^{16} , and 2^{32} points, respectively. Other than Sobol's construction, the generator matrices are not necessarily lower triangular matrices.

$a_0 \cdots a_{m-1}$ are the first m bits of the binary representation of the component j , the Owen scrambling $c_0 \cdots c_{m-1}$ of $a_0 \cdots a_{m-1}$ is determined by a tree of permutations Π_j of depth m where the boolean value, denoted $\pi(\cdot)$, at each node is XORed with the corresponding digit:

$$\begin{aligned} c_0 &= \pi \oplus a_0 \\ c_1 &= \pi_{a_0} \oplus a_1 \\ c_2 &= \pi_{a_0, a_1} \oplus a_2 \\ &\vdots \\ c_{m-1} &= \pi_{a_0, a_1, \dots, a_{m-2}} \oplus a_{m-1} \end{aligned}$$

Note that the permutation π applied to the k -th bit depends on the leading $k-1$ bits of $a_0 \cdots a_{m-1}$.

An important fact to keep in mind is that Owen scrambling applied to (t, s) -sequences or (t, m, s) -nets does not change the quality parameter t . In many cases Owen scrambling can improve the uniformity of a point set and hence can be used to optimize point sets.

4 CASCADED SOBOLO' SAMPLING

We propose the cascaded Sobol' sampler that combines the advantages of optimized Sobol' sequences, Owen scrambling, and optimization of the correlation between dimensions.

4.1 New Construction

Using the permutation property of the components of the Sobol' sequence established in Equation 3, our new construction iteratively applies the j -th Sobol' permutation to the result of the $(j-1)$ -th permutation yielding the points

$$\mathbf{x}_i := \frac{1}{N} \begin{pmatrix} \sigma_0^N(i) \\ \sigma_1^N \circ \sigma_0^N(i) \\ \vdots \\ \sigma_{s-1}^N \circ \sigma_{s-2}^N \circ \dots \circ \sigma_0^N(i) \end{pmatrix}^T. \quad (5)$$

We then apply Owen scrambling in order to benefit from randomization without compromising the low-discrepancy properties. Denoting the application of Owen scrambling by composition, for example $\Pi_0 \circ \sigma_0^N(i)$, this leads to the central

Definition 4.1 (Cascaded Sobol' Sampling). For a fixed number of samples $N = 2^m$, given the sequence of Sobol's permutations σ_j^N and a sequence of permutations trees Π_j for Owen scrambling, the i -th point $\mathbf{x}_i \in [0, 1]^s$ is defined as:

$$\mathbf{x}_i := \frac{1}{N} \begin{pmatrix} \Pi_0 \circ \sigma_0^N(i) \\ \Pi_1 \circ \sigma_1^N \circ \sigma_0^N(i) \\ \vdots \\ \Pi_{s-1} \circ \sigma_{s-1}^N \circ \sigma_{s-2}^N \circ \dots \circ \sigma_0^N(i) \end{pmatrix}^T. \quad (6)$$

Note that for $i = a_{m-1} \cdots a_0$, the composition of Sobol' permutations $\sigma_j^N \circ \dots \circ \sigma_0^N(i)$ admits the linear algebraic formulation

$$(b_0, \dots, b_{m-1}) = (a_0 \dots a_{m-1}) \cdot M_j, \quad (7)$$

where

$$M_j := C_0^T \cdot J_m \cdot C_1^T \cdot J_m \cdot \dots \cdot C_j^T \cdot J_m. \quad (8)$$

Unlike the original generator matrices, the cascaded matrices M_j are not necessarily triangular, as can be seen in Figure 3. Yet, the matrices are still full rank matrices and expose a very strong property for consecutive dimensions as analyzed in the next section.

4.2 Perfect Consecutive Nets

A key property of the cascaded Sobol' sampling is that any pair of two consecutive dimensions results in a highly uniform point set. More formally, in Appendix A we prove

THEOREM 4.2. *For any $j \in \{0, \dots, s-2\}$, the points $(x_{i,j}, x_{i,j+1})$ form a $(0, m, 2)$ -net.*

Being a (t, m, s) -net in dimension $s = 2$ with $t = 0$ has important consequences for quasi-Monte Carlo integration: each two-dimensional elementary interval of the form $[\alpha/2^{m-1}, \beta/2^{m-1}) \times [\gamma/2^{m-1}, \delta/2^{m-1}) \in [0, 1]^2$ with $\alpha, \beta, \gamma, \delta \in \{0, \dots, 2^{m-1} - 1\}$ and area 2^{-m} contains exactly one sample point. Note that the elementary intervals include the strata regularly used for stratified jittered sampling. $(0, m, 2)$ -nets hence expose a stratification superior to regular stratified sampling as illustrated in Figure 5.

4.3 Generator Matrix Optimization

In Sobol's construction, a generator matrix C_j is determined by two entities: an primitive polynomial of a certain degree n , and the first n rows of the matrix C_j^T , which are called *direction vectors* (see Figure 2). While consecutive dimensions of the cascaded Sobol' sampling are guaranteed to have excellent stratification properties, this may not be the case for other projections as illustrated in Figure 4. Optimizing the direction vectors allows one to alleviate the issue.

The goal of our offline optimization is to provide the direction vectors that insure the best low discrepancy in high dimensions. This process resembles the seminal work by Joe and Kuo for the Sobol'/Owen construction [Joe and Kuo 2008]. Once it is done, the direction vectors are tabulated (for instance, Joe and Kuo's direction vectors are hardcoded in the PBRT Sobol' implementation). The same is true for our construction: once such an optimization is performed, and the resulting initialization table is shared, the user can efficiently generate low discrepancy point sets in high dimensions.

Our optimization minimizes the generalized L_2 -discrepancy [Hickernell 1998] for any dimension up to a certain bound and for any number $N = 2^m$ of points. Applying a greedy search and starting from dimension 2, we go up to dimension 100. For each dimension, we test all primitive polynomials up to degree 10, that is 161 polynomials². For each polynomial, we assign n random direction vectors (according to the degree of the polynomial). For this particular assignment we then calculate the resulting L_2 -discrepancy for all point sets of cardinality 2^0 to 2^{18} . Among the billions of tested configurations, we accept the one with the minimal L_2 -discrepancy, for all N , and for all tested polynomials among 161 available.

²Sequence A000020 in Sloane's *On-Line Encyclopedia of Integer Sequences* [Sloane 2017] determines the number of primitive polynomials of degree n over \mathbb{F}_2 , which is 2, 1, 2, 2, 6, 6, 18, 16, 48, 60, ...

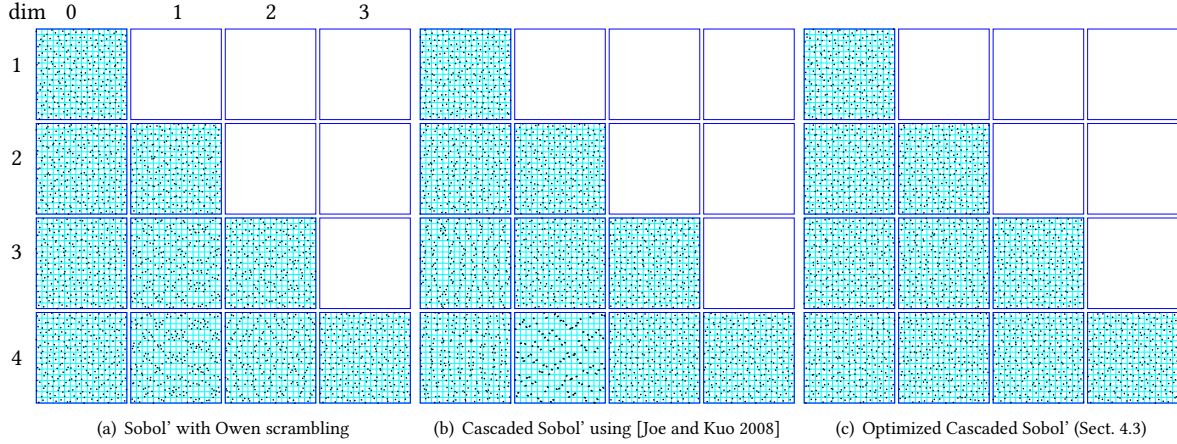


Fig. 4. **Uniformity of projections.** We show the 2-dimensional projections of 256 5-dimensional points when using the Sobol' construction with Owen scrambling (a), our construction according to Definition 4.1 using the original generator matrices by Joe and Kuo [2008] (b), and our cascaded Sobol' sampler with optimized matrices. Both (b) and (c) highlight perfect $(0, m, 2)$ -net properties for consecutive pairs of dimensions (e.g. $(0,1)$, $(1,2)$, $(2,3)$, and $(3,4)$). Optimizing the generator matrices allows us to fix potential uniformity issues in other projections (for example, $(1,4)$ in (b)).

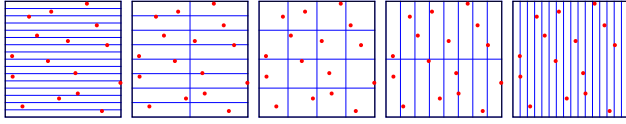


Fig. 5. **$(0, m, 2)$ -nets in base 2** are important particular cases of (t, m, s) -nets where 2-dimensional point sets of cardinality $N = 2^m$ are organized in such a way that all elementary dyadic partitions of size $1/N$ of the $[0, 1]^s$ domain contain exactly $2^t = 2^0 = 1$ sampling point, as illustrated here for 16 points. For a more detailed explanation, see Niederreiter [1992], Lemieux [2009], or Dick and Pillichshammer [2010].

Computational resources needed for optimization depend on the quality criterion, which is the ratio between the L_2 -discrepancy obtained in our construction and that of the reference Sobol'/Owen one, for all tested cardinality numbers. A typical timing for finding a good solution may vary from a few seconds ("loose" quality criterion) to many hours (tight quality criterion) per dimension. The results presented in this publication for 100 dimensions have been obtained with extremely tight quality criterion; it required about one week of computation using a farm of 300 modern 16-core computers or approximatively 50,000 CPU hours.

Figure 4 illustrates the outcome of this optimization procedure: using the generator matrices given by [Joe and Kuo 2008] (Figure 4-(b)), consecutive pairs of dimensions define $(0, m, 2)$ -nets, but with uniformity defects in other projections (for example, the pair $(1,4)$). The optimized matrices, denoted by C_j^* hereafter, lead to high quality samples for all consecutive pairs of projections (Figure 4-(c)). Figure 7 (right) presents quantitative results in high dimensions.

It is worth mentioning that the L_2 -discrepancy of all consecutive 2-dimensional projections of our point sets is clearly superior to that of Sobol' s construction with Owen scrambling, thanks to our unique property of $(0, m, 2)$ -nets in all consecutive 2-dimensional projections, as derived in the previous section.

In the supplementary material (<https://projet.liris.cnrs.fr/cascaded>), we provide our optimized initialization table for the first 100 dimensions in the file format introduced by Joe and Kuo [2008]. Each entry of this initialization table contains, one dimension per line, the index of the dimension, the degree of the primitive polynomial associated with this dimension, the primitive polynomial itself, followed by a list of direction vectors. The table contains all necessary information for generating the C_j^* .

4.4 Implementation of the Cascaded Sobol' Sampler

Algorithm 1 shows the Cascaded Sobol' Sampler that combines all previous ingredients: for a given number N of samples, the iterative construction of Definition 4.1 is used with the optimized generator matrices C_j^* as described above.

ALGORITHM 1: Generation of the sample $\mathbf{x}_i \in [0, 1]^s$, where $0 \leq i < 2^m$.

Data: Optimized generator matrices C_0^*, \dots, C_{s-1}^* , 32-bit Owen scrambling permutation trees Π_0, \dots, Π_{s-1}

Result: (x_0, \dots, x_{s-1})

$\mathbf{a} = (a_{m-1}, \dots, a_0) \leftarrow \text{BASE2INTTOVEC32}(i);$

$\mathbf{b} \leftarrow \mathbf{a} \cdot J_m; \quad // \text{ product in } \mathbb{F}_2$

for $j \leftarrow 0$ **to** $s - 1$ **do**

$\mathbf{b} \leftarrow \mathbf{b} \cdot C_j^{*T} \cdot J_m; \quad // \text{ products in } \mathbb{F}_2$

$q \leftarrow \text{BASE2VECTOINT32}(\mathbf{b});$

$p \leftarrow \Pi_j(q);$

$x_j \leftarrow \frac{1}{2^m} p;$

end

By considering Owen scrambling not only for the first $\log_2 N$ digits of the sample coordinates (see Sect. 3.2), but for all digits, we can further improve the distribution of the sample points. The interest is twofold: first, it allows us to use fixed precision operands for fast bitwise operations when performing computations on \mathbb{F}_2 .

Second, applying Owen scrambling to the least significant bits (digits between $m = \log_2 N$ and 32) performs an additional random micro-jittering across the strata of volume $1/N$ as illustrated in Figure 6. Technically speaking, this is equivalent to generating random bits for the $(32 - m)$ least significant bits. This allows for an optimized implementation of Owen scrambling that does neither generate nor store permutation beyond the m most significant bits.

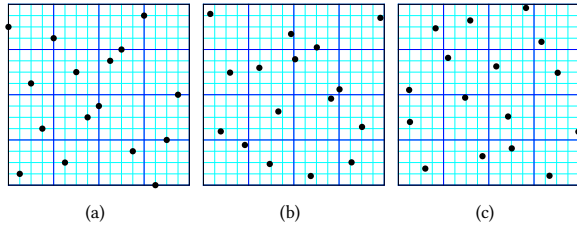


Fig. 6. **Micro-jittering by 32-bit Owen scrambling.** 16 samples of a $(0, 4, 2)$ -net with Owen scrambling applied to the $m = 4$ first bits (a). Scrambling 32 bits, the samples become jittered in their respective strata (b). The properties transfer to the cascaded Sobol' sampler (c).

For Monte Carlo rendering, this micro-jittering also has an important impact as detailed in Section 5.2. Note that fixing the precision is not a limitation of the approach since for sample counts greater than 2^{32} , we can fall back to the classical Owen scrambling.

We implemented Owen scrambling [Owen 1995] using one seed for a pseudo-number random generator to generate the tree of random permutations. An alternative implementation could consider hash-based scrambling as proposed by Burley [2020]. Compared to the classic Sobol' sequence, generating 16 (4 bits) Owen scrambled samples in $s = 6$ dimensions is 1.5 slower on the average. Our current implementation of the Owen scrambled cascaded Sobol' sampling (32 bits) is 2 times slower than the original Sobol' sequence (but generates still more than 7M samples per second on a single core AMD Ryzen 5800X). There is room for further optimizations, such as caching the generator matrices M_j and skipping the explicit matrix products of Eq. (8).

The code for the cascaded Sobol' sampler and the optimized matrices up to dimension $s = 100$ are provided (<https://projet.liris.cnrs.fr/cascaded>). For dimensions greater than 100, we use the Sobol' matrices from [Joe and Kuo 2008].

5 NUMERICAL EXPERIMENTS AND DISCUSSION

We provide numerical evidence for the practical advantages of the cascaded Sobol' sampler. We therefore measure the uniformity of distribution by integrating smooth and discontinuous test functions and computing the generalized L_2 -discrepancy. Then, we illustrate the advantages of randomization across pixels by full precision scrambling. Finally, we demonstrate the benefits of the excellent low-dimensional projections with overall high-dimensional low discrepancy that the cascaded Sobol' sampler provides by construction.

5.1 Integrating Test Functions and Discrepancy

In Figure 7, we assess the performance of the new sampler by integrating test functions and measuring the uniformity of distribution. For comparison, we consider other state-of-the-art and popular samplers: random uniform sampling, Poisson disk sampling by dart throwing, samplers relying on stratification (Jittered sampling, PMJ02 [Christensen et al. 2018], OrthogonalArrays [Jarosz et al. 2019]), samplers relying on a low discrepancy sequence (in particular, the Sobol' sequence with Owen scrambling, rank-1 lattices [Keller 2004] (with optimized vectors from [L'Ecuyer and Munger 2016]), PMJ02, ZeroTwo [Pharr et al. 2016], and our cascaded Sobol' sampler), and optimized point sets such as Sliced Optimal Transport Sampling (SOT) [Paulin et al. 2020]. Following the strategy proposed by Paulin et al. [2020], random multivariate Gaussians of the form $g(x) = \exp(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu))$ with bounded covariance matrix eigenvalues represent the class of smooth functions, while random Heaviside functions represent the class of discontinuous functions. The experiments are conducted across a range of dimensions and the uniformity of the point sets is measured by the generalized L_2 -discrepancy.

With respect to the integration of test functions, the cascaded Sobol' sampler is among the best sampling strategies while outperforming previous samplers in terms of generalized L_2 -discrepancy. For the integration tests, SOT is the most competitive set of samples ($s \in \{2, 6, 20\}$ for smooth integrands), but requires a costly optimization process that does not scale to high-dimensional sampling, which may limit its practicality in production rendering. From that perspective, cascaded Sobol' sampling has the same level of code complexity and is as fast to evaluate as the classic Sobol' with Owen scrambling.

5.2 Full-precision Scrambling

For path-tracing based Monte Carlo rendering, we use cascaded Sobol' sampling on a per-pixel basis, i.e. we generate N samples of dimension s for each pixel. As reviewed in Section 2, several strategies exist to randomize point sets such that they are uncorrelated across pixels in order to reduce visual artifacts in rendering. These strategies can be optimized by searching for shifts or pseudo-random number seeds such that the error is distributed across the image plane in a much more visually agreeable way [Ahmed and Wonka 2020; Georgiev and Fajardo 2016; Heitz and Belcour 2019; Heitz et al. 2019].

Similar to previous work, the cascaded Sobol' sampler relies on Owen scrambling to create uncorrelated point sets. As Owen scrambling preserves the low discrepancy properties, it does not alter the per-pixel convergence of Monte Carlo estimators. However, limited precision scrambling may be prone to aliasing artifacts, because the number of possible coordinates per dimension is limited to N different values. This is illustrated in Figures 6 and 9. We resolve this issue by considering a 32-bit Owen scrambling as discussed in Section 4.4. This remedies the artifacts problem as illustrated in Figure 8 as now 2^{32} positions are possible in each dimension (before floating point conversion, see Equation 1). Furthermore, the uniformity of distribution is improved almost surely.

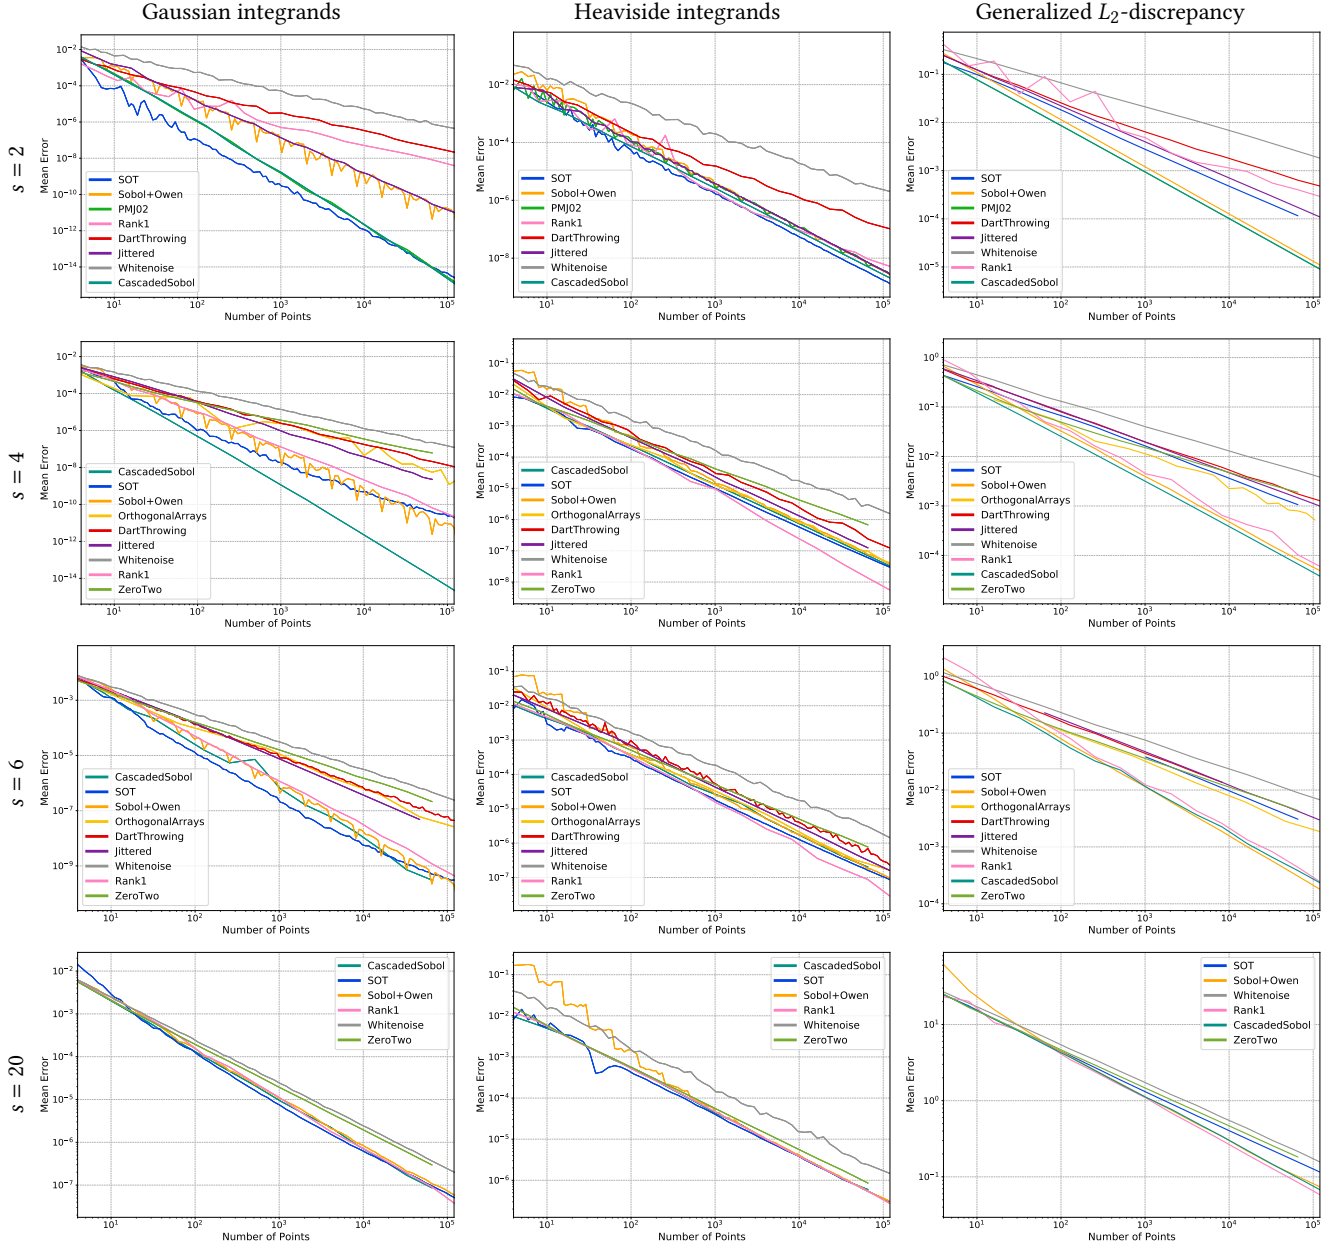


Fig. 7. Monte Carlo integration of canonical functions and discrepancy tests. We consider various integration tests in dimensions 2, 4, 6, and 20 to evaluate Monte Carlo integration error as a function of the number of sample points N . The Gaussian integrands column (left) depicts the error averaged over 1024 integrations of random multivariate Gaussian distributions in the $[0, 1]^s$ domain; each curve indicates the error of a different sampler over the 1024 integral evaluations. The Heaviside integrands column (middle) depicts the error averaged over 1024 integrations of random Heaviside functions going through the center of the $[0, 1]^s$ domain. Note that Orthogonal Arrays refers to the CMJND sampler of Jarosz et al. [2019], PMJ02 refers to the method of Christensen et al. [2018], and SOT refers to Paulin et al. [2020]. The ZeroTwo sampler corresponds to PBRT's *ZeroTwoSequenceSampler* which considers a (0,2)-sequence per pair of dimensions with a random pairing across the pair of dimensions [Pharr et al. 2016]. Rank1 uses rank-1 lattices [Keller 2004] generated via the implementation of [L'Ecuyer and Munger 2016]. The uniformity of distribution is measured by the generalized L_2 -discrepancy (right) as defined in [Hickernell 1998].

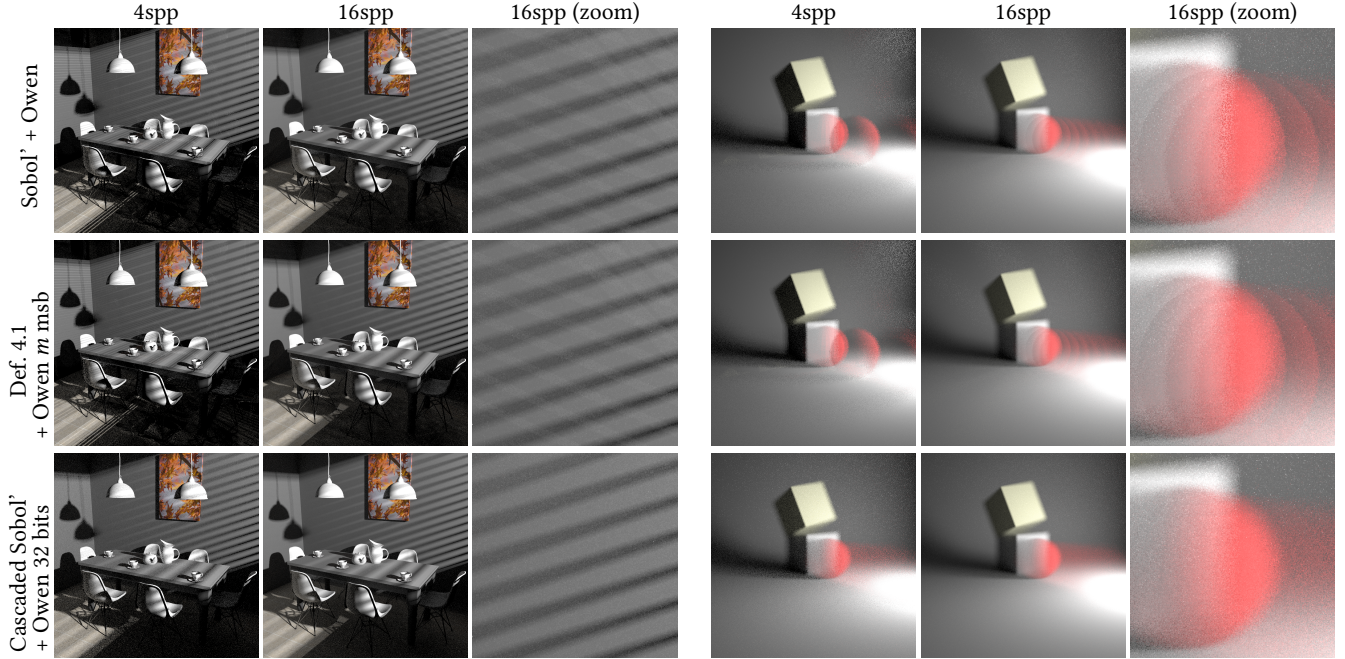


Fig. 8. **Per-pixel scrambling artifacts:** on two different scenes ($s = 6$ and $s = 10$) and low sample counts $N \in \{4, 16\}$, classic Owen scrambling applied to the $m = \log_2 N$ most significant bits (MSB) produces aliasing issues (first two rows) that no longer exist when using our $\max(32, \log_2 N)$ -bit scrambling. Note that even with the cascaded construction but only scrambling the $\log_2 N$ most significant bits, the artifacts are still visible (middle row). CC-BY "The breakfast room" scene by Wig42.

5.3 Importance of High-dimensional Low Discrepancy

We now discuss how the intrinsic structure of cascaded Sobol' helps path tracing in Monte Carlo rendering. By construction, our sampler is a low discrepancy point set with perfectly stratified pairs of consecutive projections. Implementing a path tracer, dimensions are consumed in a consecutive fashion, too, for example, consecutive scattering events along a path consume two consecutive dimensions at a time. The same goes for sampling area light sources or lenses.

For that reason, sampling strategies with good low-dimensional projections, especially for two dimensions, have attracted considerable interest [Ahmed and Wonka 2020; Joe and Kuo 2008; Perrier et al. 2018; Reinert et al. 2016]. An almost perfect sampler in that respect is the ZeroTwo sampler [Pharr et al. 2016] that has been combined with screen space properties [Ahmed and Wonka 2020]: each pair of components corresponds to a perfect $(0, 2)$ -sequence derived from the first two dimensions of the Sobol' sequence with a random pairing between pairs of dimensions (per-pixel random pairing). However, even with perfect low discrepancy of pairs of dimensions, Figure 7 shows that the ZeroTwo sampler has a bad uniformity across all dimensions – due to a lack of overall low discrepancy.

In Figure 9, we evaluate the impact of our sampler on rendering, and compare it to the Sobol' sequence and the cascaded Sobol' sampler applied to a Cornell box scene with high frequency content on the gray wall as suggested in [Burley 2020] with one-bounce

indirect lighting (thus requiring 6-dimensional samples). As discussed in Sect. 5.2, the Owen scrambled Sobol' sampling exhibits aliasing issues. ZeroTwo does not expose such artifacts, however the lack of high uniformity across all 6 dimensions causes disturbing high frequency artifacts visible on the back wall. The cascaded Sobol' sampler does not have such issues as it has the same quality as ZeroTwo with respect to low-dimensional projections (see Theorem 4.2) without sacrificing high-dimensional uniformity.

In Figure 10, we provide more qualitative and quantitative results of various scenes including longer paths and hence higher dimensions. For the qualitative analysis, we color each pixel according to the sampler achieving the minimal mean square error (MSE) at 256spp for the first three, and 4096spp for the fourth one, as compared to a ground truth image (black pixels correspond to equivalent MSE across all samplers). The graphs show MSE plots for various sample counts indicating that the cascaded Sobol' sampler outperforms its closest competitors. From the qualitative images and for low-dimensional rendering problems, we can relate the better performances of cascaded Sobol' sampler in both the projection and full space to pixels corresponding integration problems where the global uniformity matters. In higher dimensions, the cascaded Sobol' sampler still has lower MSE but it is more complex to relate this to some specific optical effects.

To summarize, the best results are achieved with perfect pairs of projections without sacrificing high-dimensional uniformity, which is exactly what the cascaded Sobol' sampler provides.




Fig. 9. **Cornell box with high frequency sub-pixel texture pattern**: highly discontinuous integrands can cause aliasing at low sample counts (the yellow rectangle highlights the region used for the zooms).

6 LIMITATIONS

First, our construction only produces point sets. However, thanks to the flexibility of our cascaded approach, there is room for more advanced constructions that unify both sequentiality and the $(0, m, 2)$ -net property in 2D projections.

As detailed in Section 4.3, our offline generator matrix optimization process has some internal parameters, which can be adapted to specific application areas. As long as direction vectors are generated, our sampler definition (Definition 4.1) and its fundamental properties (Theorem 4.2) remain the same.

Our optimization is slow when the uniformity criteria are tight. Speeding up this process has not been the focus of the new construction, but bears a lot of potential. In any case, this optimization needs to be done only once, and we provide a table of well-optimized parameters for the first 100 dimensions.

7 CONCLUSIONS

Inspired by the popular Sobol' sequence and Owen scrambling, we constructed a new sampling method. Unlike previous methods, our sampling technique offers several unique and advantageous features. First, the point sets generated by our construction have a proven $(0, m, 2)$ -net property for any consecutive pair of dimensions in a multi-dimensional setting. This property leads to unprecedented uniformity as shown in Section 5. Second, by optimizing the direction vectors as explained in Section 4.3, we are able to keep the L_2 -discrepancy at the level of the construction by Sobol' (for dimensions up to 5, our L_2 -discrepancy clearly surpasses that of the Sobol' sequence with Owen scrambling; for higher dimensions

the L_2 -discrepancy differs by a small amount). We demonstrate the benefits of the proposed sampler by examples in rendering using path tracing.

In future work, we hope to overcome the limitation of finite point sets by constructing infinite point sequences with the aforementioned advantages. Furthermore, we like to extend our sampler to incorporate screen space blue noise characteristics [Ahmed and Wonka 2020; Georgiev and Fajardo 2016; Heitz et al. 2019].

Our construction offers benefits that reach beyond computer graphics. We are confident that this work will inspire specialists in other domains which use quasi-Monte Carlo integration, e.g., in finance and physics.

ACKNOWLEDGMENTS

This work was partially funded by ANR-16-CE33-0026 (CALiTrOp). We thank Nikolaus Binder and the reviewers for their helpful comments. We are also thankful to the Computing Center of Institut National de Physique Nucléaire et de Physique des Particules (CC-IN2P3) for providing computational facilities, crucial for this project. We thank B. Bitterli [Bitterli 2016] for pre-processing and adapting the 3D scenes to common rendering tools.

A PROOF OF THEOREM 4.2

First, we show that any two consecutive dimensions have the same canonical structure:

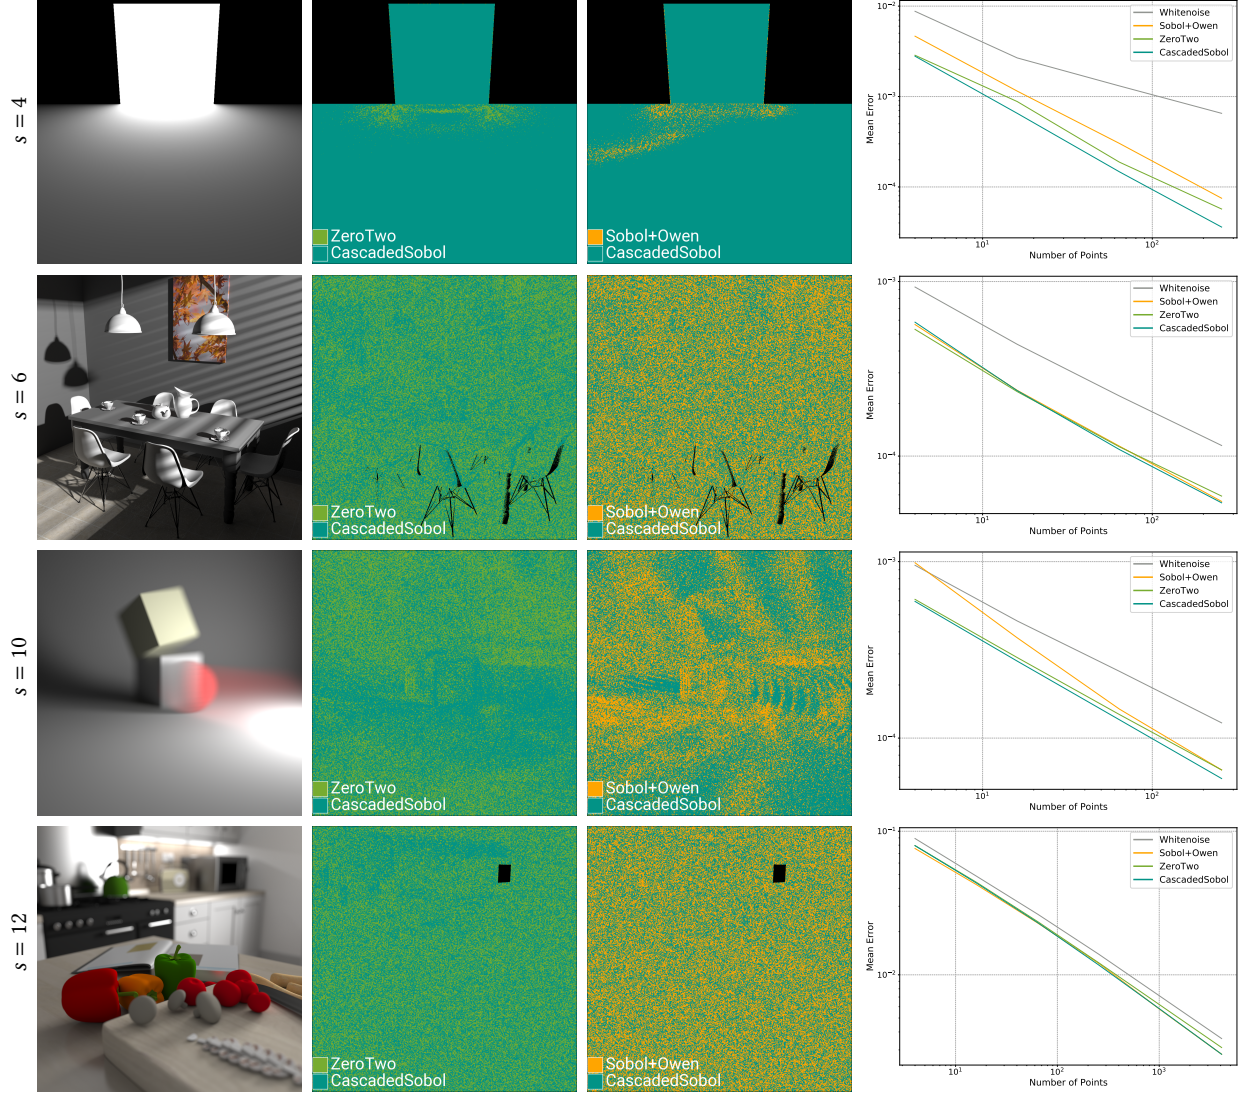


Fig. 10. **Importance of high-dimensional low discrepancy.** Each row shows path tracing results for different path lengths, i.e. different dimensions s . In the second and third column pixels are colored by the sampler with the lowest mean square error (MSE), while the last column graphs the MSE with respect to the number of samples per pixel. The MSE is computed with respect to a reference solution using 64k samples per pixel. The second column clearly indicates that low discrepancy across all dimensions in addition to low discrepancy in low-dimensional projections is clearly superior to only paying attention to the low dimensional projections. The third column shows that our cascaded Sobol' sampler is a safe bet across all dimensions, especially including the lower dimensional setting. The overall reliability and versatility of the cascaded Sobol' sampler is confirmed by the MSE plots in the last column. CC-BY "Country-Kitchen" scene by Jay-Artist. "The breakfast room" scene by Wig42.

LEMMA A.1. *Points generated by two consecutive dimensions j and $j + 1$ of our sampler are the same as those generated by*

$$\frac{1}{N} \left\{ \left(i, \sigma_{j+1}^N(i) \right) \mid i \in \{0, \dots, N-1\} \right\}.$$

PROOF. All σ_k^N are permutations on the set $\{0 \dots N-1\}$ and hence $\left\{ \sigma_k^N(i) \mid i \in \{0, \dots, N-1\} \right\} = \{0, \dots, N-1\}$. Consequently, the composition $\sigma_j^N \circ \dots \circ \sigma_0^N$ is a permutation on the same set

$\{0, \dots, N-1\}$. Therefore, the point set

$$\frac{1}{N} \left(\sigma_j^N \circ \dots \circ \sigma_0^N(i), \sigma_{j+1}^N \circ \dots \circ \sigma_0^N(i) \right)$$

for $i \in \{0, \dots, N-1\}$ is equivalent to the set

$$\left\{ \frac{1}{N} \left(i, \sigma_{j+1}^N(i) \right) \mid i \in \{0, \dots, N-1\} \right\},$$

which concludes the proof. \square

For the canonical point set of Lemma A.1 we now show

LEMMA A.2. $\left\{ \frac{1}{N} \left(i, \sigma_{j+1}^N(i) \right) \mid i \in \{0, \dots, N-1\} \right\}$ is a $(0, m, 2)$ -net.

PROOF. Following the formulation of Eq.(2), the point set

$$\left\{ \left(i, \sigma_{j+1}^N(i) \right) \mid i \in \{0, \dots, N-1\} \right\}$$

is generated by the two generator matrices

$$G_0 := \begin{pmatrix} 0 & \dots & 0 & 1 \\ \vdots & & & \vdots \\ 0 & \dots & 0 & 0 \\ 1 & 0 & \dots & 0 \end{pmatrix}$$

and

$$G_1 := C_{j+1} = \begin{pmatrix} 1 & \star & \dots & \star \\ 0 & \vdots & & \vdots \\ \vdots & & & \vdots \\ 0 & \dots & 0 & 1 \end{pmatrix}.$$

By definition, G_1 is an upper triangular matrix in \mathbb{F}_2 . Let $d = (k, m-k)$ and let M^d be the matrix composed of the k first lines of G_0 and $m-k$ first lines of G_1 , i.e.

$$M^d = \left(\begin{array}{c|c} \begin{matrix} 0 & \dots & 0 & 1 \\ \vdots & & & \vdots \\ 0 & \dots & 0 & 0 \\ 1 & 0 & \dots & 0 \end{matrix} & \begin{matrix} 0 & \dots & 0 & 1 \\ \vdots & & & \vdots \\ 0 & \dots & 0 & 0 \\ 1 & 0 & \dots & 0 \end{matrix} \\ \hline \begin{matrix} 1 & \star & \dots & \star \\ 0 & \vdots & & \vdots \\ \vdots & & & \vdots \\ 0 & \dots & 0 & 1 \end{matrix} & \begin{matrix} \star & \dots & \star \\ \vdots & & \vdots \\ \vdots & & \vdots \\ \vdots & & \vdots \end{matrix} \end{array} \right).$$

Obviously, $\det(M^d)$ is equivalent to the determinant of the lower left block of the matrix M^d that is a triangular matrix with 1 on the diagonal. Hence, $\det(M^d) = 1$ for any choice of d for $0 \leq k \leq m$.

By Niederreiter [1992, Thm. 4.28] it follows that

$$\frac{1}{N} \left\{ \left(i, \sigma_{j+1}^N(i) \right) \mid i \in \{0, \dots, N-1\} \right\}$$

is a $(0, m, 2)$ -net. \square

Our proof of Lemma A.2 uses a fundamental result from Niederreiter [1992] and Gröschloß et al. [2008], which we recall here for the sake of completeness:

THEOREM A.3 ([NIEDERREITER 1992, THM. 4.28][GRÖSCHLOSS ET AL. 2008, THM. 1]). *Let P_N be a C_0, \dots, C_{s-1} -generated point set of cardinality $N = 2^m$ in base 2 and dimension s , then P_N is a $(0, m, s)$ -net in base 2 if for all $d = (d_1, \dots, d_s) \in \mathbb{N}_0^s$ with $\|d\|_1 = m$ the*

following holds:

$$\det(M^d) \neq 0, \text{ where } M^d = \begin{pmatrix} C_{1,1}^0 & \dots & C_{1,m}^0 \\ \vdots & & \vdots \\ C_{d_1,1}^0 & \dots & C_{d_1,m}^0 \\ \vdots & & \vdots \\ C_{1,1}^{s-1} & \dots & C_{1,m}^{s-1} \\ \vdots & & \vdots \\ C_{d_s,1}^{s-1} & \dots & C_{d_s,m}^{s-1} \end{pmatrix}$$

is an $m \times m$ matrix consisting of the first d_j rows of the generator matrix C_j for $0 \leq j < s$.

It is now straightforward to prove Theorem 4.2.

PROOF. Without loss of generality, selecting one $j \in \{0, \dots, s-2\}$, we have

$$\begin{aligned} (x_{i,j}, x_{i,j+1}) \\ = \frac{1}{N} \left(\Pi_j \circ \sigma_j^N \circ \dots \circ \sigma_0^N(i), \Pi_{j+1} \circ \sigma_{j+1}^N \circ \dots \circ \sigma_0^N(i) \right). \end{aligned}$$

By Lemma A.1 and Lemma A.2 the point set formed by the two consecutive dimensions forms a $(0, m, 2)$ -net in base $b = 2$. \square

REFERENCES

- Abdalla GM Ahmed, Hélène Perrier, David Coeurjolly, Victor Ostromoukhov, Jianwei Guo, Dong-Ming Yan, Hui Huang, and Oliver Deussen. 2016. Low-discrepancy blue noise sampling. *ACM Trans. Graph.* 35, 6 (2016), 247.
- Abdalla GM Ahmed and Peter Wonka. 2020. Screen-space blue-noise diffusion of Monte Carlo sampling error via hierarchical ordering of pixels. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–15.
- Michael Balzer, Thomas Schlömer, and Oliver Deussen. 2009. Capacity-constrained point distributions: a variant of Lloyd's method. *ACM Trans. Graph.* 28, 3 (2009).
- Benedikt Bitterli. 2016. Rendering resources. <https://benedikt-bitterli.me/resources/>.
- Robert Bridson. 2007. Fast Poisson Disk Sampling in Arbitrary Dimensions. In *ACM SIGGRAPH Sketches*. 22–23.
- Brent Burley. 2020. Practical Hash-based Owen Scrambling. *Journal of Computer Graphics Techniques (JCGT)* 10, 4 (2020), 1–20.
- Per Christensen, Andrew Kensler, and Charlie Kilpatrick. 2018. Progressive Multi-Jittered Sample Sequences. *Computer Graphics Forum* 37, 4 (2018), 21–33.
- Roy Cranley and Thomas NL Patterson. 1976. Randomization of number theoretic methods for multiple integration. *SIAM J. Numer. Anal.* 13, 6 (1976), 904–914.
- Josef Dick and Friedrich Pillichshammer. 2010. *Digital Nets and Sequences: Discrepancy Theory and Quasi-Monte Carlo Integration*. Cambridge University Press.
- Raanan Fattal. 2011. Blue-Noise Point Sampling Using Kernel Density Model. *ACM Trans. Graph.* 30 (2011), 48:1–48:12.
- Iliyan Georgiev and Marcos Fajardo. 2016. Blue-Noise Dithered Sampling. 35:1–35:1.
- Leonhard Gröschloß, Johannes Hanika, Ronnie Schwede, and Alexander Keller. 2008. (t, m, s) -Nets and Maximized Minimum Distance. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 397–412.
- Leonhard Gröschloß, Matthias Raab, and Alexander Keller. 2012. Enumerating Quasi-Monte Carlo Point Sequences in Elementary Intervals. In *Monte Carlo and Quasi-Monte Carlo Methods 2010*, Leszek Plaskota and Henryk Woźniakowski (Eds.). Springer, 399–408.
- John Halton. 1964. Radical-inverse quasi-random point sequence [G5]. *Comm. ACM* 7, 12 (1964), 701–702.
- Daniel Heck, Thomas Schlömer, and Oliver Deussen. 2013. Blue Noise Sampling with Controlled Aliasing. *ACM Trans. Graph.* 32, 3 (2013), 25:1–25:12.
- Eric Heitz and Laurent Belcour. 2019. Distributing Monte Carlo Errors as a Blue Noise in Screen Space by Permuting Pixel Seeds Between Frames. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 149–158.
- Eric Heitz, Laurent Belcour, Victor Ostromoukhov, David Coeurjolly, and Jean-Claude Iehl. 2019. A Low-Discrepancy Sampler that Distributes Monte Carlo Errors as a Blue Noise in Screen Space. In *ACM SIGGRAPH Talk*.
- Fred Hickernell. 1998. A generalized discrepancy and quadrature error bound. *Math. Comp.* 67, 221 (1998), 299–322.

- Edmund Hlawka. 1961. Funktionen von beschränkter Variation in der Theorie der Gleichverteilung. *Annali di Matematica Pura ed Applicata* 54, 1 (1961), 325–333.
- Wojciech Jarosz, Afnan Enayet, Andrew Kensler, Charlie Kilpatrick, and Per Christensen. 2019. Orthogonal Array Sampling for Monte Carlo Rendering. *Computer Graphics Forum* 38, 4 (2019), 135–147.
- Stephen Joe and Frances Kuo. 2008. Constructing Sobol' sequences with better two-dimensional projections. *SIAM J. on Scientific Computing* 30, 5 (2008), 2635–2654.
- Alexander Keller. 2004. Stratification by rank-1 lattices. In *Monte Carlo and Quasi-Monte Carlo Methods 2002*, Harald Niederreiter (Ed.). Springer, 299–313.
- Alexander Keller. 2013. Quasi-Monte Carlo Image Synthesis in a Nutshell. In *Monte Carlo and Quasi-Monte Carlo Methods 2012*, Joseph Dick, Frances Kuo, Gareth Peters, and Ian Sloan (Eds.). Springer, 203–238.
- Lauwerens Kuipers and Harald Niederreiter. 2012. *Uniform distribution of sequences*. Courier Corporation.
- Pierre L'Ecuyer and David Munger. 2016. Lattice Builder: A General Software Tool for Constructing Rank-1 Lattice Rules. <https://github.com/umontreal-simul/latnetbuilder>.
- Christiane Lemieux. 2009. *Monte Carlo and Quasi Monte Carlo Sampling*. Springer.
- Hongli Liu, Honglei Han, and Min Jiang. 2021. Rank-1 Lattices for Efficient Path Integral Estimation. *Computer Graphics Forum* (2021).
- Harald Niederreiter. 1992. *Random Number Generation and quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, USA.
- Victor Ostromoukhov, Charles Donohue, and Pierre-Marc Jodoin. 2004. Fast Hierarchical Importance Sampling with Blue Noise Properties. *ACM Transactions on Graphics* 23, 3 (2004), 488–495. Proc. SIGGRAPH 2004.
- Art Owen. 1995. Randomly Permuted (t, m, s) -Nets and (t, s) -Sequences. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing (Lecture Notes in Statistics, Vol. 106)*, Harald Niederreiter and Peter Shiue (Eds.). Springer, 299–315.
- Art Owen. 1998. Scrambling Sobol' and Niederreiter–Xing Points. *Journal of Complexity* 14, 4 (1998), 466–489.
- Lois Paulin, Nicolas Bonneel, David Coeurjolly, Jean-Claude Iehl, Antoine Webanck, Mathieu Desbrun, and Victor Ostromoukhov. 2020. Sliced optimal transport sampling. *ACM Trans. Graph* 39 (2020).
- Hélène Perrier, David Coeurjolly, Feng Xie, Matt Pharr, Pat Hanrahan, and Victor Ostromoukhov. 2018. Sequences with Low-Discrepancy Blue-Noise 2-D Projections. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 339–353.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory to Implementation* (3 ed.). Morgan-Kaufmann.
- Bernhard Reinert, Tobias Ritschel, Hans-Peter Seidel, and Iliyan Georgiev. 2016. Projective blue-noise sampling. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 285–295.
- Gurprit Singh, Cengiz Öztireli, Abdalla GM Ahmed, David Coeurjolly, Kartic Subr, Oliver Deussen, Victor Ostromoukhov, Ravi Ramamoorthi, and Wojciech Jarosz. 2019. Analysis of sample correlations for Monte Carlo rendering. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 473–491.
- Neil James Alexander Sloane. 2017. The On-Line Encyclopedia of Integer Sequences. <https://oeis.org/A058943> (2017).
- Il'ya Meerovich Sobol'. 1967. On the distribution of points in a cube and the approximate evaluation of integrals. *Zhurnal Vychislitel'noi Matematiki i Matematicheskoi Fiziki* 7, 4 (1967), 784–802.
- Yahan Zhou, Haibin Huang, Li-Yi Wei, and Rui Wang. 2012. Point sampling with general noise spectrum. *ACM Trans. Graph* 31, 4 (2012), 76.