EG UK Computer Graphics & Visual Computing (2026)
Y. Sheng and M. Elshehaly (Editors)

# Generalized ART-Owen Scrambling

Abdalla G. M. Ahmed[1] and Cheng Wang[2]

[1]Khartoum, Sudan
[2]University of East Anglia, School of Computing Sciences, UK

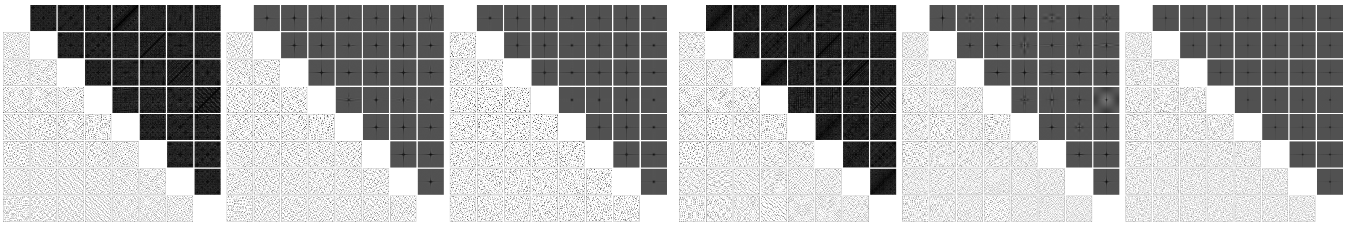

(a) 256 point, Unscrambled (b) Binary ART-Owen (c) Base-16 ART-Owen (d) 512 point, Unscrambled (e) Binary ART-Owen (f) Base-16 ART-Owen

**Figure 1:** *Pairwise spatial and spectral plots of the first eight dimensions of the universal SZ sequence [APOH25]. We show unscrambled samples with (a) 256 and (d) 512 points, and compare (b, e) binary scrambling using the original ART-Owen scrambling [APW23] to (c, f) base-16 scrambling using our proposed generalized ART-Owen scrambling. Base-16 scrambling breaks the nested substructure of the sequence, reducing pairwise correlations across dimensions, leading to similar plots over all pairs.*

**Abstract**
*Owen scrambling is a widely used randomization technique for sampling distributions in quasi-Monte Carlo applications such as rendering. ART-Owen scrambling is a flexible and scalable algorithm for implementing binary Owen scrambling, leveraging the binary structure of Adaptive Regular Tiles (ART). In this formulation, scrambling data is dynamically composed over a static ART tree. Q-ART Owen scrambling extends the approach to base 4 by factoring the* 4! *possible permutations into affine transformations that are easier to compose over the ART tree. We further extend the model to higher power-of-2 bases using a similar factorization into two components, and discuss a few implementation alternatives.*

## 1. Introduction

After three decades of competition with explicitly controlled semi-stochastic distributions such as blue noise, low-discrepancy (LD) constructions—specifically the so-called $(t,s)$-sequences—have become a dominant choice for sample generation in rendering [PJH23]. Indeed, research interest has recently shifted toward imposing spectral control over existing constructions [BCIO25, AW21]. This has, in turn, led to growing interest in Owen scrambling as a principled approach for introducing stochasticity into otherwise deterministic distributions.

Unlike many other Monte Carlo applications that use the same LD sequences, rendering ultimately produces a raster of pixels that must be computed independently and, preferably, in parallel. This setting calls for random-access Owen-scrambling data. Ahmed et al. [APW23] introduced ART-Owen scrambling, a versatile implementation that enables parallel processing while scaling smoothly with available memory. Ahmed [Ahm25] later extended the algorithm to the base-4 case. For a long time, the Sobol sequence has been the primary $(t,s)$-sequence used in graphics applications, and its binary nature makes binary Owen scrambling sufficient in most cases, with a base-4 variant used for scrambling pixel indices in Z-curve sampling [AW20, PJH23].

Recently, Ahmed et al. [APOH25] introduced a new family of $(t,s)$-sequences, SZ, with nested structures of $(0,2^q)$-sequences. For example, in the *universal* variant, SZU, each consecutive pair of dimensions forms a binary $(0,2)$-sequence, each consecutive group of four dimensions forms a $(0,4)$-sequence in base 4, each consecutive group of sixteen dimensions forms a $(0,16)$-sequence in base 16, and so on. Although these sequences are constructed in base 2 and can therefore be scrambled using binary Owen scrambling, the design space can be significantly enriched by matching the scrambling base to that of the sequence. Such base-$2^q$ methods are—to our knowledge—not readily available while also meeting the requirements of random access.



DOI: 10.2312/aeta,0.04/./5

In this paper, we generalize the base-4 extension of ART-Owen scrambling to higher power-of-2 bases while preserving the key properties of ART, including random access and scalability. We discuss the structure of the resulting transformations and their practical implications, and evaluate their effect on inter-dimensional correlations.

## 2. Technical Background and Related Work

Niederreiter [Nie92] introduced the concept of a "$(t,m,s)$-net in base $b$" to describe a stratified distribution of $b^m$ points in the $s$-dimensional unit hypercube, such that exactly $b^t$ points fall in each congruent hyperrectangular stratum of volume $b^{t-m}$. He also introduced the notion of a "$(t,s)$-sequence in base $b$" where, for all $m$, every consecutive block of $b^m$ points is a $(t,m,s)$-net in base $b$.

Owen [Owe95] showed that *nested scrambling* of the digits of the base-$b$ fractional representation of point coordinates preserves the $(t,m,s)$-net and $(t,s)$-sequence properties. Thus, Owen scrambling maps a given net or sequence into another valid one in the same base. Implementing Owen scrambling amounts to constructing, for each dimension, a scrambling tree with branching factor $b$, where each node indexed by a digit prefix $.x_1x_2\cdots x_m$ defines a permutation $\pi_{.x_1x_2\cdots x_m}$ that remaps the next digit in the base-$b$ expansion of the coordinates.

### 2.1. ART

Adaptive Regular Tiles [ANHD17] (ART) is a scalable hierarchical indexing scheme for tree structures. It comprises an alphabet $\Sigma$ of symbols and a grammar that maps each symbol to $b$ symbols from the same alphabet, where $b$ is the branching factor of the target tree. The original ART model is binary, with its symbols and grammar derived from the Thue–Morse word [Lot02], leveraging its repetition-avoidance properties.

### 2.2. ART-Owen Scrambling

ART-Owen scrambling implements binary Owen scrambling as follows. An ART alphabet and associated grammar are constructed at the prescribed size. Each symbol is equipped with a binary XOR scrambling vector, with one entry per level of the tree. At runtime, a root symbol is selected, and a tree of symbols is generated according to the production rules. The scrambling value at each node is obtained by XOR-composing the values inherited from its ancestors with the contribution of the symbol at that node. Although the same tree structure is generated from each node due to the context-free grammar, the resulting scrambling data varies through inherited values.

### 2.3. Q-ART Owen Scrambling

Ahmed [Ahm25] extended the approach to base 4 by factoring the $4! = 24$ possible permutations into affine transformations of the form

$$X \mapsto AX + B, \quad X, B \in \left\{ \begin{smallmatrix}0\\0\end{smallmatrix}, \begin{smallmatrix}1\\0\end{smallmatrix}, \begin{smallmatrix}0\\1\end{smallmatrix}, \begin{smallmatrix}1\\1\end{smallmatrix} \right\}, \; A \in \left\{ \begin{smallmatrix}1&0\\0&1\end{smallmatrix}, \begin{smallmatrix}1&1\\0&1\end{smallmatrix}, \begin{smallmatrix}1&0\\1&1\end{smallmatrix}, \begin{smallmatrix}0&1\\1&0\end{smallmatrix}, \begin{smallmatrix}1&1\\1&0\end{smallmatrix}, \begin{smallmatrix}0&1\\1&1\end{smallmatrix} \right\} \tag{1}$$

over the binary Galois field $\mathrm{GF}(2)$. A key property of this decomposition is that both components admit efficient composition along the tree, unlike the original permutation representation. According to Ahmed [Ahm25], no noticeable difference is observed when composing the linear component in practice, and an efficient implementation can therefore be obtained by composing only the translation component.

## 3. Generalized ART Owen Scrambling

The Q-ART algorithm is motivated by the observation that the set of affine transformations over $\mathrm{GF}(2)$ acting on 2-bit vectors has the same cardinality as the $4! = 24$ permutations of base-4 digits. While this correspondence is specific to base 4, as noted by Ahmed [Ahm25], the underlying model can still be extended to higher power-of-two bases. Inspecting the Q-ART mechanism, and noting that the translation component is reported to be sufficient to induce context awareness, we consider a more general decomposition

$$\pi(X) \mapsto \psi(X) + B \tag{2}$$

of a permutation $\pi$ over digits in base $2^q$ into a bijection $\psi$ and a translation component $B$. The bijection $\psi$ encodes a context-free permutation for each symbol, while the translation component is composed along the tree to induce context awareness across nodes.

Note that $\psi$ can be chosen to fix one digit (e.g., 0) while permuting the remaining digits, so that Eq. (2) effectively reduces a permutation over a $2^q$ base to permuting over $2^q - 1$ elements. This is primarily useful for analysis, since handling a power-of-2 base is typically more efficient. In particular, the binary and base-4 cases arise as special cases of this construction: $(2-1)! = 1$ eliminates the $\psi$ component in the binary case, while $(4-1)! = 6$ corresponds to the six linear transformations $\{A\}$ in Eq. (1).

## 4. Implementation

The translation vector $B$ can always be encoded as a 32- or 64-bit word, similarly to the binary case, and used in the same way. In contrast, there is considerable flexibility in modeling and encoding the bijection $\psi$, allowing different trade-offs between computational speed, memory footprint, and the level of randomization.

A full permutation of eight digits can be encoded compactly as a packed array in 24 bits, which can share the same 32-bit word as the translation vector by reusing the permuted digits as translations for deeper descendant nodes in the tree. Thus, a full-permutation base-8 ART-Owen scrambling can be encoded with essentially the same table size as base-4 Q-ART. In contrast, a full permutation of base-16 digits fits naturally in a 64-bit word. Combined with an additional 32-bit translation vector, this results in roughly three times the memory footprint of Q-ART. It is also slightly less GPU-friendly due to the use of 64-bit words, but remains compact and efficient.

For larger power-of-2 bases, and even for the cases above, the bijection $\psi$ can be modeled as a random invertible $\mathrm{GF}(2)$ matrix, which is always more compact than storing explicit permutations. For example, an $8\times 8$ matrix over $\mathrm{GF}(2)$ can be stored in a 64-bit word, compared to 256 bytes for storing a permutation over 256 symbols. It is important to note that matrix-based permutations are


Proceedings published by Eurographics - The European Association for Computer Graphics.

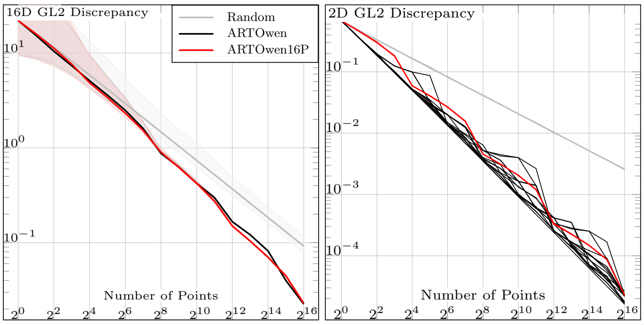


(a) Global 16D Discrepancy (b) Pairwise 2D Discrepancy

**Figure 2:** *Average discrepancy plots comparing binary to base-16 ART-Owen scrambling, showing (a) joint 16D and (b) 2D pairings of the first with the other 15 dimensions. The gray Random and red ARTOwen16P in (b) are actually 15 overlapped curves each.*

not uniformly random in a statistical sense. That said, they are commonly used in practical random vector generators [Nie92], and may therefore still be suitable for some applications.

## 5. Results and Applications

The goal of using base-$2^q$ Owen scrambling with $2^q$-dimensional SZ sequences is to break the inherent substructure of these sequences, thereby reducing inter-dimensional dependencies. This may be beneficial in rendering applications; for example, making Russian roulette decisions independent of other sampled dimensions. A full evaluation of this effect in rendering systems is nontrivial, as it would require redesigning samplers to take advantage of the newly available scrambling capabilities. For this introductory paper, we therefore focus on analysis in sample space. To this end, we have implemented a baseline full-tree base-16 Owen scrambling and observed no noticeable difference compared to our base-16 ART-Owen scrambling, allowing us to treat the latter as an equivalent and efficient realization in the following comparisons.

In Figure 1, we present spatial and spectral pairwise plots comparing binary and base-16 Owen scrambling applied to an ensembled 16D SZ sequence [APOH25]. Base-16 scrambling breaks inter-dimensional correlations, especially for non-power-of-16 sample counts, resulting in similar distributions across all dimension pairs. In contrast, binary Owen scrambling preserves all binary elementary intervals [Nie92], i.e., strata, and therefore retains pairwise correlations.

We further validate these observations by measuring the $L_2$ discrepancy, shown in Figure 2, which provides a quantitative measure of sample coverage across dimensions. Although the curves exhibit slightly different behavior, the 16D discrepancy remains of the same order for both binary and base-16 scrambling, suggesting that high-dimensional coverage is not significantly affected by the choice of scrambling. In contrast, all 2D discrepancy curves coincide under base-16 scrambling, similar to random points, whereas they are clearly distinct under binary scrambling, quantitatively confirming the visual observations in Figure 1. While individual pairs under base-16 scrambling may be less uniform than the best binary pairs, they remain more uniform than the worst binary pairs and significantly better than random sampling.

## 6. Conclusion

In this paper, we generalized ART-Owen scrambling to power-of-two bases by factoring permutations into a static component assigned to ART symbols and a dynamic translation component composed along the tree, analogous to the original binary construction. We demonstrated that aligning the scrambling base with the intrinsic base of the sequence effectively breaks its substructure, thereby reducing inter-dimensional correlation. We expect this to enable improved sampler designs that take advantage of this decorrelation facility. We will make our code publicly available to encourage further research and practical adoption of the proposed approach.

## References

[Ahm25] AHMED A. G. M.: Q-ART Owen Scrambling. In *Computer Graphics and Visual Computing (CGVC)* (2025), The Eurographics Association. doi:10.2312/cgvc.20251215. 1, 2

[ANHD17] AHMED A. G. M., NIESE T., HUANG H., DEUSSEN O.: An Adaptive Point Sampler on a Regular Lattice. *ACM Trans. Graph. 36*, 4 (July 2017), 138:1–138:13. doi:10.1145/3072959.3073588. 2

[APOH25] AHMED A. G. M., PHARR M., OSTROMOUKHOV V., HUANG H.: SZ Sequences: Binary-Constructed (0, 2q)-Sequences. *ACM Trans. Graph. 44*, 6 (Dec. 2025). doi:10.1145/3763272. 1, 3

[APW23] AHMED A. G. M., PHARR M., WONKA P.: ART-Owen Scrambling. *ACM Trans. Graph. 42*, 6 (Dec. 2023). doi:10.1145/3618307. 1

[AW20] AHMED A. G. M., WONKA P.: Screen-Space Blue-Noise Diffusion of Monte Carlo Sampling Error via Hierarchical Ordering of Pixels. *ACM Trans. Graph. 39*, 6 (Nov. 2020). doi:10.1145/3414685.3417881. 1

[AW21] AHMED A. G. M., WONKA P.: Optimizing Dyadic Nets. *ACM Trans. Graph. 40*, 4 (July 2021). URL: https://doi.org/10.1145/3450626.3459880, doi:10.1145/3450626.3459880. 1

[BCIO25] BONNEEL N., COEURJOLLY D., IEHL J.-C., OSTROMOUKHOV V.: Sobol' Sequences with Guaranteed-Quality 2D Projections. *ACM Trans. Graph. 44*, 4 (July 2025). URL: https://doi.org/10.1145/3730821, doi:10.1145/3730821. 1

[Lot02] LOTHAIRE M.: *Algebraic Combinatorics on Words*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2002. 2

[Nie92] NIEDERREITER H.: *Random Number Generation and Quasi-Monte Carlo Methods*. SIAM, 1992. 2, 3

[Owe95] OWEN A. B.: Randomly Permuted (t,m,s)-Nets and (t, s)-Sequences. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing* (1995), Springer, pp. 299–317. 2

[PJH23] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation*, 4th ed. MIT Press, 2023. 1


Proceedings published by Eurographics - The European Association for Computer Graphics.