EG UK Computer Graphics & Visual Computing (2019)
G. K. L. Tam and J. C. Roberts (Editors)

Short Paper

Sampling with Pinwheel Tiles

Extended Abstract

Abdalla G. M. Ahmed†

Abstract

We outline an adaptive sampling framework based on Conway's pinwheel tiles. It offers a unique feature that the generated tiles have infinite number of orientation angles, minimizing the strong auto-correlations found on all other lookup-sampling frameworks. We describe a novel string-based approach to enumerate distinct neighbor configurations in pinwheel tiles.

1. Introduction

Optimal point sampling is one of the longest-standing problems in computer graphics. The primary application for point sampling is Monte Carlo integration for rendering, which typically involves very high dimensions, making it infeasible to sample scene functions faithfully. Instead, many sampling strategies were developed based on tweaks, intuitions, and skills of the rendering practitioners.

Even with today's fast machines, only deterministic techniques can cope with the high sampling rates required by rendering applications [LD08]. The poor frequency spectrum of low-discrepancy sequences lead many researchers to look for sophisticated lookup techniques that offer the favorable blue-noise spectrum [CSHD03, ODJ04, LD06, KCODL06, Ost07, WPC*14, AHD15, ANHD17]. All these approaches, however, share one fundamental problem, that they exhibit strong auto-correlations (translational symmetries) in the neighborhoods of similarly-indexed tiles. The consequence of this is that, thanks to Wiener–Khinchin Theorem [LD08], the frequency spectrum of the generated point set deteriorates when the number of generated points grows significantly relative to the size of the lookup tables.

Among the many self-similar tiling sets, Conway's pinwheel tilings [Rad94] are one of the rare tessellations that offer a twirls effect of the tilings, producing infinitely many orientations. In the rest of this article we briefly outline our proposal of using these tiles for adaptive sampling.

2. Pinwheel Tilings

The base tile in pinwheel tiling is a right triangle with side length 1, 2, and $\sqrt{5}$. As illustrated in Figure 1, a recursive tiling is obtained by subdividing this triangle into 5 similar triangles. The simplest way for using such tiles for sampling is to place a sample point

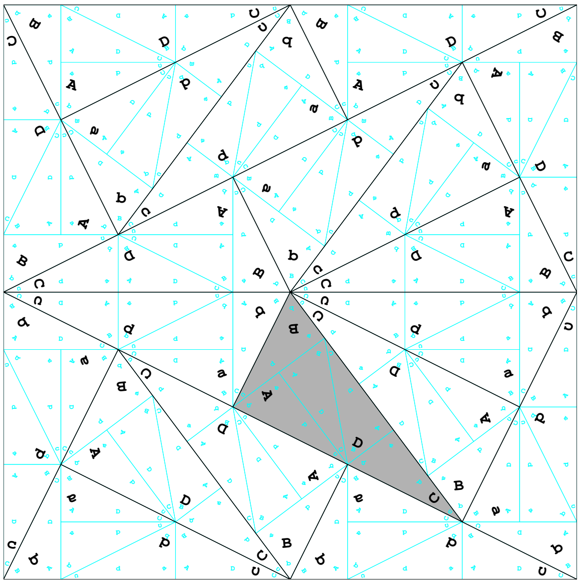

Figure 1: *Recursive subdivision of pinwheel tiles.*

at a fixed location inside each tile. Figure 2 shows three examples using different sample locations. As can be seen, without any optimization the tile set offers a rich frequency spectrum thanks the the continuously-rotating sub-tiles.

3. Enumerating the Tiles

Like the mentioned tiling techniques, the quality of sampling sets could be improved by positioning the sample points in optimized locations inside different tiles, in coordination with the sample locations inside adjacent tiles. Towards that end, it is desirable to identify geometrically identical configurations around each tile.

Following Ostromoukhov [Ost07], we start by assigning sym-

† Email:abdalla_gafar@hotmail.com

DOI: 10.2312/cgvc.20191271

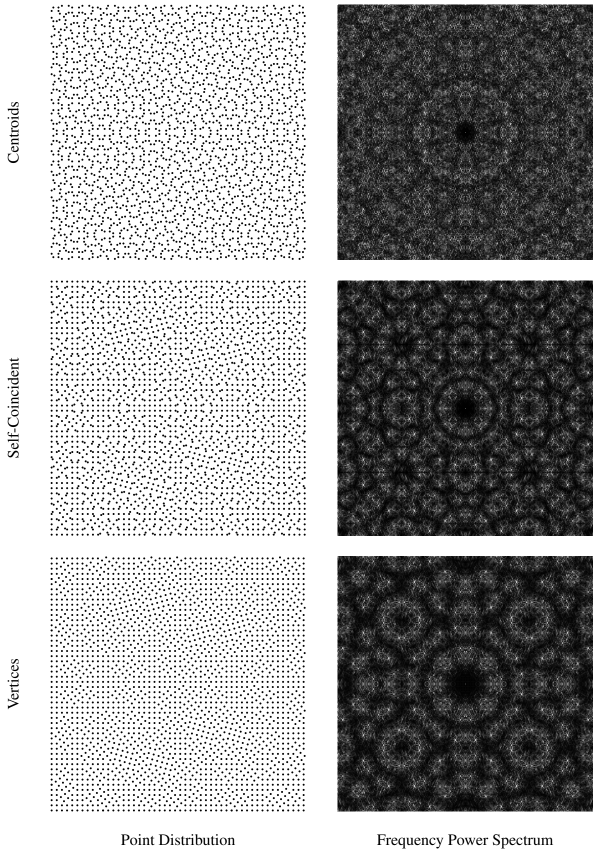

Figure 2: *Point distributions and frequency power spectra obtained from pinwheel tilings by placing sample points at centroids, self-coincident points with child tiles, and vertices.*

bols $\{A, B, C\}$ to the vertices of each tile, and assign one more symbol, D, to the midpoint between A and C, where neighbor vertices may meet; see Figure 1. We use lower-case symbols for the mirror-reflected triangles. For each triangle in a tiling, then, each vertex is identified by the symbols from all the triangles that meet there, starting from the one belonging to the inspected triangle, and reading down clockwise. Each tile is then identified by a string of identifications of its vertices A to D. For example, the gray tile in Figure 1 is identified by the string ADa-BbBbcCcC-CBad-DaA.

The challenge now is to exhaustively enumerate all distinct tiles. For Ostromoukhov's polyominos, it is considerably easier thanks to the rectilinear nature of them that makes it easy to locate adjacent units. It does not seem equally feasible in our case to trace adjacency relations geometrically. We consider the main contribution of this work so far is the novel string-based approach we developed to solve this problem. The starting point is to note that, upon subdivision, each symbol maps into a fixed sequence:

$$A \rightarrow cb; \quad B \rightarrow b; \quad \cdots.$$

Thus, given the string that identifies a tile, we may use this mapping to enumerate the new identification of each of its vertices, and use this, in turn, to read out the identification string of its child tiles; without having to know the full identity of the neighbor child tiles. We insert the distinct newly-found identification strings in a list, and recursively process them, until we find all distinct strings, which comprise the whole set. In the supplementary materials we included a program to implement this search procedure, given the full identification string of a distinct tile. The whole pinwheel set amounts to 108 distinct geometrical configurations of neighbors.

Once the set of tiles is indexed, optimized locations for the sample points could be computed for each of the tiles using, for example, blue noise through optimal transport (BNOT) [dGBOD12], and applying similar techniques to those described in [ODJ04, KCODL06, Ost07, WPC*14, ANHD17]. We are yet to investigate on this.

4. Conclusion

In this brief discussion we shed the light on a potentially-interesting tile-based technique for sampling. Our research is still at an early stage, but the initial results are quite promising. Pinwheel tiles magically offer a combination of the rich frequency spectrum of Polyhexes [WPC*14] and the tight lookup table size of Penrose tiles [ODJ04], using only 108 entries. In addition, they offer two unique features not found in other existing alternatives: the infinite orientations of the tiles, and the simple triangular shape of these tiles that could be useful in applications like sampling over meshes.

References

[AHD15] Ahmed A. G. M., Huang H., Deussen O.: AA Patterns for Point Sets with Controlled Spectral Properties. *ACM Trans. Graph. 34*, 6 (Oct. 2015), 212:1–212:8. 1

[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. 1, 2

[CSHD03] Cohen M. F., Shade J., Hiller S., Deussen O.: Wang Tiles for Image and Texture Generation. *ACM Trans. Graph. 22*, 3 (July 2003), 287–294. 1

[dGBOD12] de Goes F., Breeden K., Ostromoukhov V., Desbrun M.: Blue Noise through Optimal Transport. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 171:1–171:11. 2

[KCODL06] Kopf J., Cohen-Or D., Deussen O., Lischinski D.: Recursive Wang Tiles for Real-Time Blue Noise. *ACM Trans. Graph. 25*, 3 (July 2006), 509–518. 1, 2

[LD06] Lagae A., Dutré P.: An Alternative for Wang Tiles: Colored Edges Versus Colored Corners. *ACM Trans. Graph. 25*, 4 (Oct. 2006), 1442–1459. 1

[LD08] Lagae A., Dutré P.: A Comparison of Methods for Generating Poisson-Disk Distributions. In *Computer Graphics Forum* (2008), vol. 27, Wiley Online Library, pp. 114–129. 1

[ODJ04] Ostromoukhov V., Donohue C., Jodoin P.-M.: Fast Hierarchical Importance Sampling with Blue-Noise Properties. In *ACM SIGGRAPH 2004 Papers* (New York, NY, USA, 2004), SIGGRAPH '04, ACM, pp. 488–495. 1, 2

[Ost07] Ostromoukhov V.: Sampling with Polyominoes. *ACM Trans. Graph. 26*, 3 (July 2007). 1, 2

[Rad94] Radin C.: The pinwheel tilings of the plane. *Annals of Mathematics 139*, 3 (1994), 661–702. 1

[WPC*14] Wachtel F., Pilleboue A., Coeurjolly D., Breeden K., Singh G., Cathelin G., de Goes F., Desbrun M., Ostromoukhov V.: Fast Tile-Based Adaptive Sampling with User-Specified Fourier Spectra. *ACM Trans. Graph. 33*, 4 (July 2014), 56:1–56:11. 1, 2