

Image Sampling with Quasicrystals

Mark GRUNDLAND [†], Jiří PATERA[‡], Zuzana MASÁKOVÁ* and Neil A. DODGSON [†]

[†] *Computer Laboratory, University of Cambridge, UK*

E-mail: *Mark@Eyemaginary.com, Neil.Dodgson@cl.cam.ac.uk*

[‡] *Centre de Recherches Mathématiques, Université de Montréal, Canada*

E-mail: *patera@crm.umontreal.ca*

* *Department of Mathematics FNSPE, Czech Technical University in Prague, Czech Republic*

E-mail: *zuzana.masakova@fjfi.cvut.cz*

Received December 15, 2008, in final form July 06, 2009; Published online July 20, 2009

doi:10.3842/SIGMA.2009.075

Abstract. We investigate the use of quasicrystals in image sampling. Quasicrystals produce space-filling, non-periodic point sets that are uniformly discrete and relatively dense, thereby ensuring the sample sites are evenly spread out throughout the sampled image. Their self-similar structure can be attractive for creating sampling patterns endowed with a decorative symmetry. We present a brief general overview of the algebraic theory of cut-and-project quasicrystals based on the geometry of the golden ratio. To assess the practical utility of quasicrystal sampling, we evaluate the visual effects of a variety of non-adaptive image sampling strategies on photorealistic image reconstruction and non-photorealistic image rendering used in multiresolution image representations. For computer visualization of point sets used in image sampling, we introduce a mosaic rendering technique.

Key words: Computer graphics; image sampling; image representation; cut-and-project quasicrystal; non-periodic tiling; golden ratio; mosaic rendering.

2000 Mathematics Subject Classification: 20H15; 52C23; 68U99; 82D25

1 Introduction

Non-periodic tilings have emerged as an important mathematical tool in a variety of computer graphics applications [27]. They have proven especially useful in the design of sampling algorithms, where they serve to direct the spatial distribution of rendering primitives by enforcing spatial uniformity while precluding regular repetition. Recently, Wang tilings [3, 19, 23, 24, 25], Penrose tilings [37], Socolar tilings [38] and polyominoes [39] have been used to generate point sets for non-periodic sampling. In one of the earliest applications of non-periodic tilings in computer graphics, Penrose tilings [12, 14, 44] were employed by Rangel-Mondragon and Abas [42] in the design of decorative patterns inspired by Islamic art. They had effectively reinvented the medieval trade secrets of the craftsmen of fifteenth century Islamic mosques [28] who created by hand highly intricate mosaics closely resembling quasicrystal tilings only discovered by modern science in the late twentieth century. Wang tilings [12, 14] were first introduced by Stam [48] in order to enable wave texture patches to cover water surfaces of arbitrary size without the appearance of regularly repeating artifacts. Further computer graphics applications of non-periodic tilings include texture mapping and synthesis [3, 23, 24, 25, 48, 49], photorealistic rendering using environmental maps [24, 37], and non-photorealistic rendering using stippling [23, 24]. For computer graphics applications, non-periodic tilings have usually been generated by geometric constructs, such as matching rules and hierarchical substitution [11, 14]. In this work, we present the cut-and-project method of generating quasicrystals as an alternative algebraic approach to the production of non-periodic tilings and point sets (Figure 1). This algebraic ap-




Figure 1. Quasicrystal tilings produced using spatial proximity graphs. In these visualizations, a non-periodic, rotationally symmetric point set (top left) is depicted as a planar tiling induced by a Voronoi diagram (top right), a Delaunay graph (bottom left), and a Delaunay triangulation (bottom right). This set of 1035 points comprises a cut-and-project quasicrystal derived from the standard root lattice of the non-crystallographic Coxeter group H_2 . Its viewing window is a square centered at the origin with radius 1, while its acceptance window is a decagon centered at the origin with radius $\tau^5 + \tau^3$, where τ is the golden ratio. The visualization of the quasicrystal tilings reveals some remarkable properties. It is well known that quasicrystals can exhibit five and ten fold rotational symmetry, an impossible feat for any periodic tiling. Recently, it has been shown analytically that a quasicrystal Delaunay graph can yield a non-periodic tiling with four distinct tile shapes [29]. As illustrated by this visualization, a Delaunay triangulation of a cut-and-project quasicrystal can yield a non-periodic tiling with just three distinct tile shapes.

proach has the advantages of being straightforward to implement, easy to calculate, and readily amenable to rigorous mathematical analysis. Moreover, it may be directly extended to higher dimensions as well as adaptive sampling applications, although this is outside the scope of our present work. We choose to base our method on the algebra of the golden ratio $\tau = \frac{1}{2}(1 + \sqrt{5})$, as its geometrical properties have been previously successfully exploited in computer graphics in the context of spatial sampling [37] and color quantization [33] techniques that rely on the Fibonacci number system. For an introduction to the theory of quasicrystals, consult Senechal’s comprehensive textbook [44], while a more advanced treatment of the cut-and-project method may be found in surveys by Patera [40] as well as Chen, Moody, and Patera [2].


In the evaluation of the effectiveness of quasicrystals as a non-adaptive image sampling strategy, our work is motivated by the use of image sampling in multiresolution image representation and progressive image rendering. In particular, we base our experimental investigations on our experience with the development of a point-based rendering approach to multiresolution image representation for digital photography [15, 16] based on scattered data interpolation techniques [1], which has been shown to support a secure and compact image encoding suitable for both photorealistic image reconstruction and non-photorealistic image rendering. A thorough discussion of the standard image sampling strategies can be found in Glassner’s textbook [10]. Their effectiveness has been extensively investigated for use in photorealistic computer graphics applications [47], such as Monte Carlo integration in 3D ray tracing. For non-adaptive sampling, the key trade-off is between aliasing and noise, as exemplified by the regular structure of periodic sampling using a square grid and the irregular clustering of random sampling using a uniform distribution. The classic compromise strategies are jittered sampling [4, 6, 22], which disrupts the regularity of periodic sampling by randomly perturbing the sample sites, and quasirandom sampling [41], which avoids the irregularity of random sampling by ensuring a consistent density of sampling is maintained. In our experiments, we demonstrate that quasicrystal sampling can permit more accurate photorealistic image reconstruction than either standard jittered sampling or standard quasirandom sampling. For photorealistic image reconstruction [10], the ideal strategy is generally considered to be the Poisson disk distribution [4, 6], random point sets conditioned on a minimum distance between the points, while for non-photorealistic image rendering [5, 17, 18, 20, 43], a popular strategy relies on centroidal Voronoi diagrams [7], optimized point sets with every point placed at the centroid of its Voronoi polygon. As both of these sampling strategies prove time consuming to compute exactly, a variety of approximation techniques have been proposed to produce sampling patterns that have similar properties in the frequency domain, in particular those that exhibit a blue noise Fourier power spectrum characteristic of a Poisson disk distribution. Historically, blue noise sampling strategies relied on slow, trial and error, stochastic procedures involving dart throwing algorithms that approximate the Poisson disk distribution by rejecting prospective locations for new sample sites whenever they are deemed to be too close to the preceding samples [4, 6, 30, 32]. Improved performance of blue noise sampling can be obtained through the use of efficient geometric data structures [8, 9, 21, 51] and parallel processing GPU hardware [50]. Alternatively, one can readily generate a blue noise sampling pattern using a non-periodic tiling composed of a suitable set of tiles, where each tile contains a precomputed optimal arrangement of sample sites [3, 19, 23, 24, 25, 37, 39]. A detailed evaluation of the spectral properties of various blue noise sampling algorithms can be found in the survey by Lagae and Dutre [26]. In practical applications, it can often be quite difficult to visually distinguish between renditions produced by blue noise sampling patterns generated using different algorithms. Hence, for the purpose of our evaluation, we relied on farthest point sampling [9] since in previous work we have shown this blue noise sampling technique to be highly suitable for multiresolution image representation [16]. In general, blue noise sampling algorithms tend to have higher requirements for either computational processing, data storage, or implementation complexity than simpler sampling strategies, such as quasicrystal

sampling. Naturally, simpler sampling strategies cannot replicate all of the desirable qualities of blue noise sampling. Nevertheless, as demonstrated by our evaluation, quasicrystal sampling is shown to be proficient at supporting a uniform sampling density, centroidal Voronoi regions, accurate image reconstruction, and progressive image rendering, despite having only a small number of local sampling configurations arranged in an anisotropic manner incompatible with a blue noise Fourier power spectrum. Hence, as a potential alternative to periodic sampling in image representation, sampling using cut-and-project quasicrystals can deterministically guarantee minimum and maximum distances between nearest neighbors in a uniformly space-filling sampling pattern without the overhead of geometric data structures or tiling lookup tables for tracking their placement.

2 Method


While a periodic point set is characterized by its translational symmetries, a non-periodic point set admits no translational symmetries. For use in image sampling, we focus on non-periodic point sets that are determined by their inflation symmetries. In such a non-periodic point set, a fixed configuration of sample sites can be repeated at different scales to generate a self-similar pattern. The simplest way of producing non-periodic point sets is to use hierarchical substitution tilings [13, 14]. For instance, hierarchical substitution can be readily applied to the famous Penrose tiling [12, 14, 37, 42, 44]. The strategy starts with a small set of polygonal tiles. The tile set is carefully designed so that each tile can be decomposed by geometric subdivision into smaller instances of itself and the other tiles. A hierarchy is formed whereby an existing tile becomes the parent of new child tiles. Starting with an initial configuration of the tiles scaled to cover the image plane, the tiling is refined through an iterative process of deflation and substitution. The tile vertices or centroids are used to derive a point set from the tiling. The choice of the initial configuration appears mirrored in the global structure and symmetry properties of the tiling and the resulting point set.

We focus on a more general class of non-periodic point sets corresponding to cut-and-project quasicrystals [2, 40, 44]. The cut-and-project method was originally introduced by Meyer [31] in the context of harmonic analysis and it was later adapted for generating quasicrystals by Moody and Patera [34]. The Fibonacci chain and Penrose tilings can be regarded as special cases of such quasicrystals. In our work, we employ the standard root lattice of the non-crystallographic Coxeter group H_2 , a group of reflections taken from Lie algebra theory. This approach to quasicrystals can be used to relate discrete, non-periodic point sets and tilings (Figure 1) with the level sets of continuous, non-periodic functions (Figure 2). To produce a 2D cut-and-project quasicrystal, a 4D periodic lattice is projected on a suitable 2D plane that is irrationally oriented with respect to the lattice. Using this method, we obtain a dense subspace consisting of integer coefficient linear combinations of the vertices of a regular decagon centered on the origin, which are the roots of the non-crystallographic Coxeter group H_2 . This construction ensures that the coordinates of all its elements can be expressed using only integers and an irrational number, the golden ratio. To obtain a finite 2D quasicrystal, we select only those elements of the dense subspace contained in a specified bounded region, called the viewing window, that are mapped by an everywhere discontinuous algebraic transformation, called the star map, to another specified bounded region, called the acceptance window. Through the gradual expansion of a rotationally symmetric acceptance window, a quasicrystal sequence can be uniquely ordered by radial distance and angle from then center of the acceptance window in order to produce a uniformly space-filling point set in the viewing window. As an important practical consequence, this property directly enables progressive sampling. It could also potentially enable adaptive sampling by varying the radius of the acceptance window according to an application dependent importance map defined for the viewing window. A 2D quasicrystal can also be



Phase function contour map

Phase function level set



Quasicrystal phase function: $f(z) = \sum_{j=0}^9 \exp(2\pi i(\zeta^j \cdot (2\tau^4 z)))$ defined for regular decagon vertices $\zeta^j = \exp(\frac{2\pi i}{10} j)$ and golden ratio $\tau = \frac{1}{2}(1 + \sqrt{5})$ with dot product $u \cdot v = \frac{1}{2}(u\bar{v} + \bar{u}v)$ and complex conjugate $\overline{u_1 + u_2 i} = u_1 - u_2 i$.

Figure 2. Quasicrystal construction using a continuous quasicrystal phase function. Observe that a quasicrystal point set is contained within a level set of a continuous phase function $f : \mathbb{C} \rightarrow \mathbb{R}$ defined in the complex plane [35]. This continuous phase function is formulated using the roots of the non-crystallographic Coxeter group H_2 , which comprise the vertices ζ^j of a regular decagon. To generate a quasicrystal point set, start by placing the first point at the origin. For each newly placed point $x \in \mathbb{C}$, consider the candidate points $z = x + \zeta^j$ with $j \in \{0, \dots, 9\}$, which are the vertices of a regular decagon centered at x , and only accept the candidates for which the phase function $f(z) \geq T$ exceeds a desired threshold T that controls the density of the resulting point set.

expressed as a subset of a Cartesian product of two 1D quasicrystals, in accordance with the fact that the points that lie on any straight line through a 2D quasicrystal correspond to some linearly transformed 1D quasicrystal. Therefore, in practice, a 2D quasicrystal (Figure 4) can be generated from a 2D lattice of 1D quasicrystals. Meanwhile, a 1D quasicrystal (Figure 3) is produced by taking a strip of a 2D periodic lattice, having finite width and irrational slope, and orthogonally projecting its points onto a line of the same slope. The resulting 1D quasicrystal is composed of at most three distinct tiles. It is easy to generate 1D quasicrystal points using an iterative numerical algorithm. Alternatively, it is possible to exploit the self-similar structure of a 1D quasicrystal, viewing it as the fixed point of a set of substitution rules that act recursively on a finite alphabet of possible tile arrangements.


Based on the geometry and algebra of the golden ratio, these quasicrystal point sets exhibit some remarkable properties. They display pentagonal and decagonal rotational symmetries, which cannot occur in any periodic point set. Originally, the theory of quasicrystals was motivated by solid state physics as a model of the non-periodic geometric structures that describe the symmetries of a new kind of long-range atomic order discovered in certain metallic alloys [46]. While translational symmetries define periodic crystals, inflation symmetries can be used to describe quasicrystals based on algebraic irrational numbers, such as the golden ratio. When the

Cut-and-project principle:

To generate a non-periodic point set in an n -dimensional space, take a region of a periodic point set in a $2n$ -dimensional space and orthogonally project it onto an irrationally oriented n -dimensional subspace. For the resulting point set in n -dimensional space to be discrete rather than dense, the region of $2n$ -dimensional space undergoing projection, typically a cylinder, must be bounded along the directions of projection.

Ruler and graph paper construction:

Start with the 2D integer lattice. Consider three parallel lines with slope τ , the inner one passing through the origin $(0, 0)$. The outer lines serve to cut out a strip of the integer lattice, while the inner line provides the central axis. The width of the strip defines the acceptance window $\Omega = [c, d]$, while the visible length of the strip defines the viewing window $V = [l, r]$. Observe that when the lattice points contained in the strip are orthogonally projected onto the central axis, the resulting 1D sequence Σ_Ω is non-periodic if and only if the slope τ of the strip is irrational.

**Golden ratio:**

$\tau = \frac{1}{2}(1 + \sqrt{5}) \approx 1.618$ and its conjugate $\tau' = \frac{1}{2}(1 - \sqrt{5}) \approx -0.618$ are the solutions of $x^2 = x + 1$.

Golden integers:

$\mathbb{Z}[\tau] = \{a + b\tau \mid a, b \in \mathbb{Z}\}$ is an Euclidean domain that is dense in \mathbb{R} .

Conjugate golden integers:

$\mathbb{Z}[\tau'] = \{a + b\tau' \mid a, b \in \mathbb{Z}\}$ is the set of conjugates $(a + b\tau)' = a + b\tau' = a - b\tau^{-1}$, where $\tau' + \tau = 1$ and $\tau'\tau = -1$.

Cut-and-project 1D quasicrystals:

$\Sigma_\Omega = \{a + b\tau \in \mathbb{Z}[\tau] \mid a + b\tau' \in \Omega \cap \mathbb{Z}[\tau']\}$ quasicrystal is specified by a bounded acceptance window $\Omega = [c, d]$.

Duality of 1D quasicrystals:

$x_k \in (\Sigma_\Omega \cap V) \subset \mathbb{Z}[\tau]$ restricted to the bounded viewing interval $V = [l, r]$ implies a dual quasicrystal $x'_k \in (\Sigma'_V \cap \Omega) \subset \mathbb{Z}[\tau']$ contained in the bounded acceptance interval $\Omega = [c, d]$.


Translation and scaling of 1D quasicrystals:

$\Sigma_{[c,d]} + \lambda = \Sigma_{[c+\lambda, d+\lambda]}$ for $\lambda \in \mathbb{Z}[\tau]$ and $\xi'\Sigma_{[c,d]} = \Sigma_{[\xi c, \xi d]}$ for $\xi = \tau^k$ and $k \in \mathbb{Z}$.

Stepping function for 1D quasicrystals:

Assume a standard acceptance window $\Omega = [c, d]$ such that $0 \in \Omega$ and $d - c \in [1, \tau)$ and observe that a 1D quasicrystal is an arrangement of just three possible tiles.

$$x'_{k+1} = \begin{cases} x'_k + 1 & \text{if } x'_k \in [c, d-1) & \Rightarrow x_{k+1} - x_k = 1 \text{ (small tile)} \\ x'_k + 1 + \tau' & \text{if } x'_k \in [d-1, c-\tau') & \Rightarrow x_{k+1} - x_k = 1 + \tau \text{ (large tile)} \\ x'_k + \tau' & \text{if } x'_k \in [c-\tau', d) & \Rightarrow x_{k+1} - x_k = \tau \text{ (medium tile)} \end{cases}$$

**Iterative construction of 1D quasicrystals:**

For a standard acceptance window $\Omega = [c, d]$, containing the origin $c \leq 0 < d$ inside an interval of suitable width $1 \leq d - c < \tau$, a quasicrystal sequence $x_k \in \Sigma_\Omega$ is generated by successively applying the stepping function $\sigma : \Omega \rightarrow \Omega$ to obtain the conjugate of the right neighbor $x'_{k+1} = \sigma(x'_k)$ or the inverse stepping function $\sigma^{-1} : \Omega \rightarrow \Omega$ to obtain the conjugate of the left neighbor $x'_{k-1} = \sigma^{-1}(x'_k)$.

Stepping algorithm for 1D quasicrystals:

1. Translate and scale the desired acceptance Ω and viewing V intervals to make the acceptance window Ω into a standard acceptance window.
2. Starting from $x_0 = 0$, apply the stepping function $x'_{k+1} = \sigma(x'_k)$ and its inverse $x'_{k-1} = \sigma^{-1}(x'_k)$ to generate the consecutive quasicrystal points until both edges of the translated and scaled viewing window are reached.
3. Reverse the translation and scaling to obtain the desired quasicrystal Σ_Ω .

Figure 3. Algorithm for 1D cut-and-project quasicrystals.

Quasilattice principle:


For a quasicrystal point set in n -dimensional space, the points that lie on any straight line can be mapped by an affine transformation to a valid quasicrystal sequence in 1-dimensional space.

Star map:

$(xe_1 + ye_2)^* = x'e_1^* + y'e_2^*$ semilinear map acts on golden integers $x, y \in \mathbb{Z}[\tau]$ and basis vectors $e_1, e_2 \in \mathbb{R}^2$.

Star basis:

$e_1 = (0, 1)$ and $e_2 = (-\frac{1}{2}\tau, \frac{1}{2}\sqrt{3-\tau})$ are mapped by the star map to $e_1^* = (0, 1)$ and $e_2^* = (\frac{1}{2}(\tau-1), -\frac{1}{2}\sqrt{2+\tau})$, as defined by the standard basis of the non-crystallographic Coxeter group H_2 .

**Golden integer lattice:**

$M = \mathbb{Z}[\tau]e_1 + \mathbb{Z}[\tau]e_2 = \left\{ (a_1 + b_1\tau)e_1 + (a_2 + b_2\tau)e_2 \mid a_1, b_1, a_2, b_2 \in \mathbb{Z} \right\}$ is a $\mathbb{Z}[\tau]$ -module that is dense in \mathbb{R}^2 .

Star mapped golden integer lattice:

$M^* = \mathbb{Z}[\tau']e_1^* + \mathbb{Z}[\tau']e_2^* = \left\{ a_1 + b_1\tau'e_1^* + (a_2 + b_2\tau')e_2^* \mid a_1, b_1, a_2, b_2 \in \mathbb{Z} \right\}$ is determined by the star map.

Cut-and-project 2D quasicrystals:

$\Sigma_\Omega = \left\{ xe_1 + ye_2 \in M \mid x'e_1^* + y'e_2^* \in \Omega \cap M^* \right\}$ quasicrystal is specified by a bounded acceptance window Ω .

Duality of 2D quasicrystals:

$xe_1 + ye_2 \in (\Sigma_\Omega \cap V) \subset M$ restricted to the bounded viewing window region V implies a dual quasicrystal $x'e_1^* + y'e_2^* \in (\Sigma_V^* \cap \Omega) \subset M^*$ contained in the bounded acceptance window region Ω .

Set laws of 2D quasicrystals:

$\Sigma_{\Omega_1 \cap \Omega_2} = \Sigma_{\Omega_1} \cap \Sigma_{\Omega_2}$ and $\Sigma_{\Omega_1 \cup \Omega_2} = \Sigma_{\Omega_1} \cup \Sigma_{\Omega_2}$ as well as $\Sigma_{\Omega_1} \subseteq \Sigma_{\Omega_2}$ whenever $\Omega_1 \subseteq \Omega_2$.

Translation and scaling of 2D quasicrystals:

$\Sigma_\Omega + \lambda = \Sigma_{\Omega + \lambda^*}$ for $\lambda \in M$ and $\xi'\Sigma_\Omega = \Sigma_{\xi\Omega}$ for $\xi = \tau^k$ and any $k \in \mathbb{Z}$.

Inflation symmetry of 2D quasicrystals:

$\tau^2(\Sigma_\Omega - z) + z \subset \Sigma_\Omega$ implies that every point $z \in \Sigma_\Omega$ is a center of inflation symmetry, if Ω is a convex set.


Quasilattice for 2D quasicrystals:

$\Sigma_{\Gamma_1 e_1^* + \Gamma_2 e_2^*} = \Sigma_{\Gamma_1} e_1 + \Sigma_{\Gamma_2} e_2$ is a lattice of 1D quasicrystals Σ_{Γ_1} and Σ_{Γ_2} with acceptance windows Γ_1 and Γ_2 .


Quasilattice algorithm for 2D quasicrystals:

1. Find a quasilattice viewing window $W = W_1 e_1 + W_2 e_2$ that contains the desired quasicrystal viewing window $V \subseteq W$. The quasilattice viewing window W is a parallelogram where the viewing intervals satisfy $x \in W_1$ and $y \in W_2$ for all $xe_1 + ye_2 \in V \cap M$.
2. Find a quasilattice acceptance window $\Gamma = \Gamma_1 e_1^* + \Gamma_2 e_2^*$ that contains the desired quasicrystal acceptance window $\Omega \subseteq \Gamma$. The quasilattice acceptance window Γ is a parallelogram where the acceptance intervals satisfy $x' \in \Gamma_1$ and $y' \in \Gamma_2$ for all $x'e_1^* + y'e_2^* \in \Omega \cap M^*$.
3. Generate the 1D quasicrystals $\Sigma_{\Gamma_1} \cap W_1$ and $\Sigma_{\Gamma_2} \cap W_2$ according to viewing intervals W_1 and W_2 as well as the acceptance intervals Γ_1 and Γ_2 .
4. Generate the 2D quasilattice $\Sigma_\Gamma \cap W = \Sigma_{\Gamma_1 e_1^* + \Gamma_2 e_2^*} \cap W = (\Sigma_{\Gamma_1} \cap W_1) e_1 + (\Sigma_{\Gamma_2} \cap W_2) e_2$ according to quasilattice viewing window W and quasilattice acceptance window Γ from the coordinates supplied by the 1D quasicrystals $\Sigma_{\Gamma_1} \cap W_1$ and $\Sigma_{\Gamma_2} \cap W_2$.
5. Discard all quasilattice points $xe_1 + ye_2 \in \Sigma_\Gamma \cap W$ that do not belong to the desired quasicrystal $xe_1 + ye_2 \notin \Sigma_\Omega \cap V$ because either they do not belong to the desired quasicrystal viewing window $xe_1 + ye_2 \notin V$ or acceptance window $x'e_1^* + y'e_2^* \notin \Omega$.


Figure 4. Algorithm for 2D cut-and-project quasicrystals.



Evolution of a 1D quasicrystal point set Σ_Ω from $\Omega = [0, 1)$ to $\Omega = [0, \tau)$.



Evolution of a Delaunay graph of a 2D quasicrystal, where 10 new sites are added at each generation.



Evolution of a Delaunay triangulation of a 2D quasicrystal, where the acceptance window is a decagon centered at the origin that is expanded by a factor of τ , from radius $\tau^4 + \tau^2$ on the left and to radius $\tau^5 + \tau^3$ on the right.

Figure 5. Evolution of cut-and-project quasicrystals reveals their self-similar structure.

acceptance window is a convex region, every point of a quasicrystal can be viewed as a center of inflation symmetry. A quasicrystal can have no translational symmetries and no periodic subsets. Moreover, it can be partitioned into subsets such that each subset forms a valid quasicrystal. Consider the local configurations of tiles in an infinite quasicrystal mosaic formed by a Voronoi diagram or a Delaunay triangulation. According to the repetitivity property implied by suitable regularity conditions, each fragment is repeated infinitely many times in the mosaic. Yet no single fragment is ever sufficient to determine the structure of the whole mosaic because every finite fragment, no matter its size, occurs in an uncountable infinity of nonequivalent mosaics. Furthermore, according to the Delaunay point set property, quasicrystals are both uniformly discrete and relatively dense, creating space-filling tilings.

As Delaunay point sets, quasicrystals are particularly well suited to image sampling. They enforce both a minimal and a maximal distance between each sample site and its closest neighboring site. In quasicrystal sampling, we rely on the golden ratio to ensure symmetry and self-similarity, which are generally absent when other than algebraic irrational numbers are used with the cut-and-project method to produce non-periodic point sets. By taking this approach, we can endow a rendition with a decorative symmetry, which viewers may find attractive in the context of non-photorealistic rendering. Compared with the regular grids of periodic point sets, the self-similar, space-filling structure of non-periodic point sets (Figure 5) appears less monotonous, especially during progressive image rendering. In effect, the geometric structure of quasicrystal sampling eliminates the possibility of aliasing artifacts regularly repeating in the rendition. However, in quasicrystal sampling, fixed local configurations of sample sites can be repeated at multiple places and orientations, albeit not at regular intervals, with the potential to yield some recurring, anisotropic aliasing artifacts. Although we did do so in this work, for photorealistic image reconstruction, it is preferable to avoid inducing global rotational symmetry in the sampling pattern, which is done by ensuring the viewing window does not contain the origin when the acceptance window is symmetric with respect to the origin.

3 Evaluation

We now present a qualitative evaluation of quasicrystal sampling in the context of non-adaptive sampling strategies for use in image representation. In this application, non-adaptive sampling strategies serve as building blocks for interactive sampling, adaptive sampling, and importance sampling techniques. In effect, they dictate the placement of sample sites in image regions sampled at a uniform resolution. For this purpose, a non-adaptive sampling strategy should satisfy several image representation objectives. The sample sites should be distributed in a manner that fairly and accurately represents the image. The sampling pattern should be evenly space-filling in order to enable progressive image rendering. Without any preconceptions about the distribution of visually salient features in the image, the same amount of information should be devoted to capturing each part of the image. Hence, the number of sample sites placed in any region of the image should be proportional to its area, so the sampling density remains the same throughout the image. Both globally and locally, the placement of sample sites should be uniform and isotropic while still allowing for a variety of different sample site configurations. To prevent aliasing, the sample sites should not be arranged into fixed configurations that visibly repeat locally or globally. To prevent clustering, a minimum distance between sample sites should be maintained throughout the image. Assuming that the correlation between pixels decreases with distance, a sample site should be placed close to the centroid of its Voronoi polygon to enable its sampled color to be most representative of its region of influence when the image is reconstructed using a local interpolation technique. Naturally, the relative importance of these considerations depends on the requirements of a particular application. We have formulated these priorities based on our previous work on a multiresolution image representation [16] designed to simultane-

ously support both photorealistic image reconstruction and non-photorealistic image rendering. While previous qualitative surveys [10] and quantitative comparisons [47] of image sampling have focused on applications in photorealistic image reconstruction, they did not cover quasicrystal sampling and farthest point sampling. Furthermore, their tests were not carried out on an image representation of digital photographs and they did not specifically address the needs of non-photorealistic image rendering.

We compared quasicrystal sampling to a number of standard non-adaptive image sampling strategies [10], reflecting different approaches to the inherent trade-off between aliasing and noise. For our evaluation, we chose approaches that exemplify divergent aims in sampling. For each approach, we selected a representative implementation. As noted below, alternative implementations are certainly possible but they are likely to produce similar qualitative results. From the deterministic to the stochastic, we tested a range of sampling strategies (Figure 7):

1. **Periodic sampling** [36] aims for global regularity. Our implementation relies on a square lattice refined in scan line order. An alternative implementation could use a hexagonal lattice, the densest periodic lattice in the plane.
2. **Quasicrystal sampling** [44] aims for local regularity. Our implementation relies on the cut-and-project method applied using the golden ratio. An alternative implementation could use a Penrose tiling produced using a hierarchical substitution algorithm.
3. **Farthest point sampling** [9] aims for spatial uniformity. Our implementation relies on the principle of progressively sampling at the point of least information, placing each new sample site at the point farthest from any preceding sample site, which is necessarily a vertex of the Voronoi diagram of the preceding sample sites. An alternative implementation could position sample sites to conform to a centroidal Voronoi diagram so that each sample site is placed at the centroid of its Voronoi polygon.
4. **Jittered sampling** [22] aims for local variability. Our implementation relies on a full random displacement of a square lattice refined in scan line order. An alternative implementation could use a partial random displacement of a hexagonal lattice.
5. **Quasirandom sampling** [41] aims for low discrepancy. Our implementation relies on the Halton sequence. An alternative implementation could use a Sobol sequence.
6. **Random sampling** [10] aims for global variability. Our implementation relies on a uniform distribution. An alternative implementation could use a random walk on a unit square with toroidal boundary conditions.

To perform a qualitative evaluation of the image sampling strategies, we applied a number of computer visualization techniques. For each non-adaptive sampling strategy, we visualize its sample sites (Figure 7) in the spatial domain using a Voronoi diagram (Figure 8) and in the frequency domain using a Fourier power spectrum (Figure 9). We examined the visual effects of applying the image sampling strategies in the context of various image rendering techniques that are used in multiresolution image representations. For photorealistic image reconstruction [1], we tested the accuracy of Shepard’s interpolation (Figure 13), an inverse distance weighted interpolation method that applies the Voronoi diagram to determine the color of each pixel based on its four nearest sample sites, as well as Gouraud shading (Figure 14), a piecewise linear interpolation method that applies the Delaunay triangulation to determine the color of each pixel based on its three surrounding sample sites. To quantitatively assess the results of these widely used local interpolation techniques, we relied on the peak signal-to-noise ratio (PSNR). This standard image fidelity metric [45] estimates the accuracy of the




Figure 6. Techniques for visualizing the spatial configuration of sample sites: Voronoi diagram (far left), Delaunay triangulation (center left), Delaunay star shapes (center right), and our mosaic rendering technique (far right).

rendition according to the negative logarithm of the mean squared error between the rendered and actual RGB color values of the pixels. Hence, higher peak signal-to-noise ratio scores are considered better. For non-photorealistic image rendering [16], we experimented with a simple “paint strokes” rendering style (Figure 12), which applies geometric subdivision to the Delaunay triangulation of the sample sites and then uses linear and nonlinear interpolation to emphasize the transitions between the sampled colors.

Our mosaic rendering style (Figure 6) offers a new computer visualization tool for evaluating the spatial properties of point sets, such as the sample sites produced by the various image sampling strategies. To produce a mosaic rendering, we apply geometric subdivision to the Delaunay triangulation of the sample sites. The midpoints of the edges of each Delaunay triangle are joined to form three outer triangles and one inner triangle. Each outer triangle is rendered with the color sampled at its vertex of the original Delaunay triangle, while the central triangle is colored black. As each sample site is represented by a star-shaped polygonal tile, the resulting mosaic appears packed as tightly as possible, with the black central triangles serving as grout between the tiles. As a sample site’s local neighborhood (Figure 6, far left) comprises the surrounding sample sites connected to it by edges in the Delaunay triangulation (Figure 6, center left), the sample site’s mosaic tile (Figure 6, far right) is shaped to reflect the star of its surrounding Delaunay triangles (Figure 6, center right). For instance, a sample site’s mosaic tile is a convex or concave polygon according to whether its neighboring sites are arranged in a convex or concave configuration. As a visualization tool, the advantage of mosaic rendering is that the layout of the star-shaped mosaic tiles makes the spatial properties of a sampling strategy easier to see at a glance than the triangles of a Delaunay triangulation or the convex polygons of a Voronoi diagram. The sizes of the mosaic tiles are indicative of the uniformity of sampling, as coarsely sampled regions give rise to large tiles and finely sampled regions give rise to small tiles, making common defects such as clustering, undersampling, and oversampling easy to detect. The orientations of the mosaic tiles are indicative of the isotropy of sampling, as the preferred directions of the sampling are revealed in the preferred rotations of the tiles, making global or local grid structures easy to detect. The shapes of the mosaic tiles are indicative of the heterogeneity of sampling, as the local configurations of neighboring sites uniquely determine the tile polygons, making repetitive patterns easy to detect. For instance, farthest point sampling produces tiles of uniform size and similar shape to create the appearance of a pebble mosaic, while quasicrystal sampling yields a decorative tiling with just a small set of possible tile shapes.

Our qualitative evaluation of image sampling strategies uses seven criteria (Figure 10) known to affect the visual quality of photorealistic image reconstruction and non-photorealistic image rendering [16]:

1. **Accurate reconstruction** requires the rendition to faithfully represent the likeness of the original image. It is a necessary but not sufficient condition of success in both photorealistic and non-photorealistic image rendering. This objective appears to be closely associated with uniform coverage and centroidal regions. It is assessed by measuring the peak signal-to-noise ratio for the results of photorealistic image reconstruction (Figure 13 and Figure 14). Its most pronounced effects can also be observed in the results of non-photorealistic image rendering (Figure 12). When the resolution of sampling is uniform, periodic sampling yields the most accurate image interpolation (in Figure 14, for the regular square grid, this takes place when there are $33^2 = 1089$, $65^2 = 4225$, and $129^2 = 16641$ samples). However, when regions of varying resolution arise during progressive refinement, the accuracy of periodic sampling can substantially deteriorate. Similar behavior is observed in jittered sampling since it applies random perturbations to a periodic point set. Farthest point sampling produces image reconstructions that are nearly as accurate as periodic sampling, but its performance does not diminish during progressive refinement. Intermediate accuracy is offered by quasicrystal sampling, which appears to be slightly more accurate than quasirandom sampling. The least accurate reconstructions are produced by jittered and random sampling. Given its popularity in computer graphics implementations, the poor performance of jittered sampling is rather disappointing. Of course, the accuracy of jittered sampling can always be made closer to that of periodic sampling by reducing the amount of random displacement, which risks reintroducing the aliasing artifacts of periodic sampling. In general, accuracy is improved by uniformity and reduced by randomness, an effect that can be readily seen as producing tight or loose image stylization.

2. **Progressive refinement** requires the sample sites to smoothly fill the available space, avoiding abrupt changes in appearance as new sample sites are sequentially added to the rendition. This objective serves to enable a multiresolution image representation to support progressive rendering of compressed images based on an incremental sampling of the image data. It is assessed by examining the spatial layout of the sequence of sample sites (Figure 7). Under ideal circumstances, progressive refinement should yield a smooth curve for the peak signal-to-noise ratio (Figure 14). The best progressive refinement results are produced by farthest point sampling and quasicrystal sampling, as these methods maintain a uniform sampling density by ensuring that new sample sites are placed in the largest empty spaces between the existing sample sites. Quasirandom sampling proves slightly less proficient, as it places some sample sites very close together while keeping others far apart. Random sampling is even less effective due to its tendency to locally cluster sample sites. The regular grids used in periodic and jittered sampling are not suitable for smooth progressive refinement, especially when they are refined in scan line order. While other refinement schemes can be applied to regular grids, such as refinement in random order, their intrinsic symmetry makes it difficult to smoothly increase the sampling density throughout the image.

3. **Uniform coverage** requires the sample sites to be evenly distributed regardless of position, avoiding configurations that place sample sites too close or too far from their nearest neighbors. This objective is assessed using Voronoi diagrams (Figure 8) as well as mosaic renderings (Figure 11). Its effect determines the sizes of brush strokes in non-photorealistic image rendering (Figure 12). Uniform coverage is associated with a Fourier power spectrum (Figure 9) that displays an empty ring around the central spike, as low frequencies are attenuated in favor of a threshold frequency corresponding to the most commonly observed nearest neighbor distance between sample sites. Although a blue noise spectrum can ensure uniform coverage, it is not a necessary condition. When the resolution of sam-




Figure 7. Non-adaptive sampling strategies: periodic sampling (top left), quasicrystal sampling (top center), farthest point sampling (top right), jittered sampling (bottom left), quasirandom sampling (bottom center), and random sampling (bottom right). Sampling starts with the dark blue sites and finishes with the light green sites.

pling is uniform, periodic sampling generates uniform coverage, as its mosaic tiles are all exactly the same size. However, periodic sampling cannot sustain uniform coverage during progressive refinement. By design, farthest point sampling maintains uniform coverage at all times, as its mosaic tiles are all approximately the same size. Quasicrystal sampling maintains nearly as uniform coverage, as its mosaic tiles are limited to just a few comparable sizes. While quasirandom and jittered sampling strive to uphold a uniform density of sampling, they nevertheless are less effective at providing uniform coverage, as their mosaic tiles come in many sizes. In the case of jittered sampling, uniform coverage can be improved by reducing the amount of random displacement. Random sampling does not give uniform coverage, as its mosaic tiles exhibit the greatest range of different sizes.

4. **Isotropic distribution** requires the sample sites to be evenly distributed regardless of orientation, avoiding configurations that align sample sites along globally or locally preferred directions. This objective is assessed using Voronoi diagrams (Figure 8) as well as mosaic renderings (Figure 11). Its effect is to determine the orientations of brush strokes in non-photorealistic image rendering (Figure 12). An isotropic distribution produces a Fourier power spectrum (Figure 9) that displays a rotational symmetry around the central spike, as the power at each frequency does not depend on its orientation. Although a blue noise spectrum can ensure isotropic distribution, it is not a necessary condition. Random sampling is the most isotropic, as its sample sites are both locally and globally uncorrelated. Farthest point and jittered sampling are nearly as isotropic, as their sample sites can exhibit slight local alignment. Farthest point samples can appear to be placed in roughly hexagonal local configurations. Jittered samples can appear to retain some of the




Figure 8. Voronoi diagrams of image sampling strategies applied to a color spiral test image: periodic sampling (top left), quasicrystal sampling (top center), farthest point sampling (top right), jittered sampling (bottom left), quasirandom sampling (bottom center), and random sampling (bottom right).

structure of the underlying square grid, as the isotropy of jittered sampling reflects the amount of random displacement used to generate the sampling. Quasirandom sampling has intermediate isotropy, as its sample sites can exhibit slight global alignment, which can be verified in the lack of radial symmetry in its Fourier power spectrum. Periodic sampling and quasicrystal sampling do not have isotropic distribution since their sample sites are globally aligned along predetermined axes.

5. **Blue noise spectrum** requires the sample sites to be distributed similarly to a Poisson disk distribution, a random point field conditioned on a minimum distance between the points. According to this objective, for an image sampling strategy to provide effective antialiasing for image rendering, it should attempt to mimic the idealized distribution of photoreceptors in the human eye. Usually implying both uniform coverage and isotropic distribution, a blue noise spectrum is highly desirable in many computer graphics applications, particularly photorealistic image reconstruction. It is assessed by examining the Fourier power spectrums of the sampling strategies (Figure 9) for a radially symmetric profile that concentrates noise in the high frequencies while attenuating the power of the low frequencies, thereby eliminating the aliasing artifacts associated with low frequency patterns that can appear distracting to the eye. In effect, a blue noise spectrum exhibits a disk of low power around the origin, surrounded by roughly constant power at the higher frequencies. Its effects can be judged according to the amount of aliasing present in photorealistic image reconstruction (Figure 13) and non-photorealistic image rendering (Figure 12). Farthest point sampling has a Fourier power spectrum that is closest to a blue noise spectrum. Jittered sampling attempts to replicate the blue noise spectrum, but




Figure 9. Fourier power spectrums of image sampling strategies: periodic sampling (top left), quasicrystal sampling (top center), farthest point sampling (top right), jittered sampling (bottom left), quasirandom sampling (bottom center), and random sampling (bottom right).

it does not clearly exhibit the threshold frequency ripple around the central spike. Quasirandom is even less successful because its Fourier power spectrum lacks radial symmetry. Periodic, quasicrystal, and random sampling have Fourier power spectrums that do not resemble the blue noise spectrum. The Fourier power spectrums of periodic and quasicrystal sampling reflect the spatial structures and directional symmetries of the lattices used to place the sample sites. On the other hand, the white noise spectrum of random sampling assigns roughly the same power to all frequencies.

6. **Centroidal regions** require sample sites to be well centered with respect to their Voronoi polygons, approximating a centroidal Voronoi diagram. Typically associated with uniform coverage and accurate reconstruction, this objective is popular in non-photorealistic image rendering. Sampling strategies, such as periodic sampling, that produce centroidal regions can still be prone to aliasing artifacts since centroidal regions do not guarantee a blue noise spectrum. Centroidal regions can be readily assessed using the Voronoi diagrams (Figure 8). The effects can also be observed in the shapes of tiles in mosaic rendering (Figure 11) and brush strokes in non-photorealistic image rendering (Figure 12). Periodic sampling, placing each sample site at the same distance from all of its nearest neighbors, generates exact centroidal Voronoi regions. Quasicrystal and farthest point sampling produce approximately centroidal Voronoi regions. To place sample sites close to the center of their Voronoi polygons, quasicrystal sampling relies on local symmetries while farthest point sampling relies on nearest neighbor distance. Jittered sampling and quasirandom sampling have difficulty ensuring centroidal regions because they are less effective at maintaining a minimal nearest neighbor distance. Finally, random sampling can only produce centroidal regions by chance.

Sampling Strategies	Accurate Reconstruction	Progressive Refinement	Uniform Coverage	Isotropic Distribution	Blue Noise Spectrum	Centroidal Regions	Heterogeneous Configurations
Periodic	★★★★	★	★★★★	★	★	★★★★	★
Quasicrystal	★★★	★★★★	★★★	★	★	★★★	★★
Farthest Point	★★★★	★★★★	★★★★	★★★	★★★★	★★★	★★
Jittered	★	★	★★	★★★	★★★	★★	★★★
Quasirandom	★★	★★★	★★	★★	★★	★★	★★★
Random	★	★★	★	★★★★	★	★	★★★★

★★★★ Superior ★★★ Good ★★ Fair ★ Poor

Figure 10. Qualitative evaluation of image sampling strategies.

7. **Heterogeneous configurations** require sample sites to be placed in a variety of different local arrangements, avoiding regularly or randomly repeating the same sampling patterns. While this objective is not traditionally a concern in photorealistic image reconstruction, it helps to prevent non-photorealistic image rendering from appearing too perfect, seemingly mechanical and artificial. For instance, it helps to give a vibrant appearance to brush stroke rendering. Typically, when sampling strategies yield centroidal regions, they also tend to produce homogeneous configurations and vice versa, illustrating an apparent trade-off between these competing objectives. Heterogeneous configurations are assessed using the Voronoi diagrams (Figure 8) and mosaic renderings (Figure 11). Their effect is also visible in the arrangement of brush strokes in non-photorealistic image rendering (Figure 12). Random sampling produces the most heterogeneous local configurations, as its mosaic tiles exhibit the greatest variety of shapes. Jittered and quasirandom sampling are nearly as heterogeneous, since their mosaic tiles are almost as widely varied, though few of them are exceptionally large in size. Quasicrystal and farthest point sampling are far less heterogeneous. By upholding local symmetries, quasicrystal sampling causes sample sites to have only a few possible local configurations, resulting in mosaic tiles that have only a few possible shapes. By upholding nearest neighbor distance, farthest point sampling causes sample sites to have similar local configurations, resulting in mosaic tiles that look very much alike, mostly convex and rounded. Based on repetitions of a single local configuration, periodic sampling is entirely homogeneous.

Our qualitative analysis (Figure 10) indicates that quasicrystal sampling offers a useful compromise between the ordered behavior of standard periodic sampling using a regular square lattice and the disordered behavior of standard Monte Carlo sampling using jittered, quasirandom, or random sampling. Compared with periodic sampling, quasicrystal sampling displays a greater variety of local sample site configurations resulting in smoother progressive refinement, although its sampling patterns are somewhat less uniform, leading to lower accuracy of image reconstruction. Compared with jittered, quasirandom, and random sampling, quasicrystal sampling displays more uniform coverage resulting in better accuracy of image reconstruction, although its sampling patterns are anisotropic, exhibiting significantly less variety of local sample site configurations. By virtue of its deterministic construction, quasicrystal sampling does not suffer from the variability that can affect the results of random sampling, jittered sampling, and, to a much lesser degree, farthest point sampling. Nevertheless, its lack of a blue noise power spectrum renders it rather susceptible to aliasing artifacts. Research on quasicrystal sampling based on the Penrose tiling [37] suggests that it may be possible to partially alleviate this problem by taking advantage of the symmetries and the repetitions of the local sample site configurations in order to systematically displace the sample sites in a manner that improves

the spectral properties of the sampling pattern.

Based on our qualitative evaluation of the various non-adaptive sampling strategies (Figure 10), we recommend a blue noise sampling strategy, such as farthest point sampling, for general use in image representation. In particular, farthest point sampling does not perform poorly on any of our seven evaluation criteria. Overall, our qualitative evaluation of non-adaptive sampling strategies is in broad agreement with previous studies, which did not consider quasicrystal sampling. They emphasized the importance of Poisson disk distributions [10] and low discrepancy distributions [47], which are exemplified in our evaluation by farthest point sampling and quasirandom sampling respectively. Hence, the good overall performance of these two techniques should come as no surprise. Farthest point sampling performs better than quasirandom sampling on six out of the seven evaluation criteria. For the majority of our evaluation criteria, quasicrystal sampling performs no better than farthest point sampling and no worse than quasirandom sampling. Nevertheless, from a practical point of view, quasicrystal sampling is significantly simpler to implement and calculate than farthest point sampling, which relies on maintaining complex geometric data structures to keep track of the vertices of a Voronoi diagram. This could be an important consideration for imaging applications on mobile devices that have limited processing and storage capabilities. From a theoretical point of view, the deterministic algebraic construction of quasicrystals renders their sampling patterns particularly well suited to mathematical analysis. Presenting possibilities for future research, the cut-and-project method could be adapted for higher dimensional sampling or adaptive sampling applications.


In future work, it would also be interesting to explore the relationship between local symmetry and sampling quality. The cut-and-project method can be used to generate non-periodic point sets with different symmetries, not just the pentagonal and decagonal symmetries associated with the golden ratio, as shown in this work. Just as for periodic sampling it would be interesting to compare the image reconstruction accuracy of square and hexagonal grids, for non-periodic sampling it would be interesting to compare our decagonal quasicrystal tiling with the dodecagonal Socolar tiling, which was recently proposed for use in sampling applications [38].

4 Conclusion


Cut-and-project quasicrystals present new possibilities for image sampling in computer graphics. This non-periodic sampling approach deterministically generates uniformly space-filling point sets, ensuring that sample sites are evenly distributed throughout the image. It offers a useful compromise between predictability and randomness, between the standard periodic sampling and the standard Monte Carlo sampling methods. Although blue noise sampling can generate higher quality sampling patterns for photorealistic image reconstruction, quasicrystal sampling may prove much simpler to implement and calculate. In the context of non-photorealistic image rendering, quasicrystal sampling may prove attractive for its symmetry properties.

Acknowledgements


We are grateful to the referees for valuable remarks that led to improvements of the original manuscript. We also acknowledge the financial support of the Natural Sciences and Engineering Research Council of Canada, le Fonds Qubcois de la Recherche sur la Nature et les Technologies, as well as the grants MSM6840770039 and LC06002 of the Ministry of Education of the Czech Republic, and GA201/09/0584 of the Czech Science Foundation. We are also grateful for the support of the MIND Research Institute and the Merck Frosst Company. Mark Grundland further acknowledges the financial assistance of the Celanese Canada Internationalist Fellowship, the British Council, the Cambridge Commonwealth Trust, and the Overseas Research Student Award Scheme. The images were provided by FreeFoto.com and the Waterloo Brag Zone.




Periodic sampling




Quasicrystal sampling




Farthest point sampling



Jittered sampling



Quasirandom sampling




Random sampling


Figure 11. Image sampling strategies rendered using the mosaic style (4225 samples \approx 2.6%).




Periodic sampling



Quasicrystal sampling




Farthest point sampling



Jittered sampling




Quasirandom sampling




Random sampling


Figure 12. Image sampling strategies rendered using the “paint strokes” style (4225 samples \approx 2.6%).




Periodic sampling (PSNR = 18.58)




Quasicrystal sampling (PSNR = 18.52)




Farthest point sampling (PSNR = 18.55)



Jittered sampling (PSNR = 18.31)



Quasirandom sampling (PSNR = 18.24)



Random sampling (PSNR = 17.93)

Figure 13. Image sampling strategies rendered using Shepard interpolation (4225 samples \approx 2.6%).




Figure 14. Quantitative evaluation of image sampling strategies rendered using Gouraud shading.

Supplementary materials

For a full resolution version of this paper, along with supplementary materials, please visit:
<http://www.Eyemaginary.com/Portfolio/Publications.html>

References

- [1] Amidror I., Scattered data interpolation methods for electronic imaging systems: A survey, *Journal of Electronic Imaging* 11, 2 (2002), 157–176.
- [2] Chen L., Moody R. V., Patera J., Non-crystallographic root systems. In *Quasicrystals and discrete geometry*, Fields institute monograph, (1998), vol. 10, 135–178.
- [3] Cohen M. F., Shade J., Hiller S., Deussen O., Wang tiles for image and texture generation. In *Proceedings of SIGGRAPH*, (2003), 287–294.
- [4] Cook R. L., Stochastic sampling in computer graphics, *ACM Transactions on Graphics* 5, 1 (1986), 51–72.
- [5] Deussen O., Hiller S., van Overveld C., Strothotte T., Floating points: A method for computing stipple drawings. In *Proceedings of EUROGRAPHICS*, (2000), 41–50.
- [6] Dippe M. A. Z., Wold E. H., Antialiasing through stochastic sampling. In *Proceedings of SIGGRAPH*, (1985), 69–78.
- [7] Du Q., Faber V., Gunzburger M., Centroidal Voronoi tessellations: Applications and algorithms, *SIAM Review* 41, 4 (1999), 637–676.
- [8] Dunbar D., Humphreys G., A spatial data structure for fast Poisson-disk sample generation. In *Proceedings of SIGGRAPH*, (2006), 503–508.
- [9] Eldar Y., Lindenbaum M., Porat M., Zeevi Y. Y., The farthest point strategy for progressive image sampling, *IEEE Transactions on Image Processing* 6, 9 (1997), 1305–1315.
- [10] Glassner A., Principles of digital image synthesis, vol. 1, Morgan Kaufmann, 1995.
- [11] Glassner A., Aperiodic tiling, *IEEE Computer Graphics and Applications* 18, 3 (1998), 83–90.
- [12] Glassner A., Penrose tiling, *IEEE Computer Graphics and Applications* 18, 4 (1998), 78–86.
- [13] Goodman-Strauss C., Aperiodic hierarchical tilings. In *Foams, emulsions, and cellular materials*, Proceedings of NATO-ASI, (1999), vol. E 354, 481–496.
- [14] Grunbaum B., Shephard G. C., Tilings and patterns, WH Freeman, 1987.
- [15] Grundland M., Style and content in digital imaging: Reconciling aesthetics with efficiency in image representation, VDM, 2008.
- [16] Grundland M., Gibbs C., Dodgson N. A., Stylized multiresolution image representation, *Journal of Electronic Imaging* 17, 1 (2008), 013009: 1–17.
- [17] Hausner A., Simulating decorative mosaics. In *Proceedings of SIGGRAPH*, (2001), 573–580.
- [18] Hausner A., Pointillist halftoning. In *Proceedings of the International Conference on Computer Graphics and Imaging*, (2005), 134–139.
- [19] Hiller S., Deussen O., Keller A., Tiled blue noise samples. In *Proceedings of Vision, Modeling and Visualization*, (2001), 265–271.
- [20] Hiller S., Hellwig H., Deussen O., Beyond stippling – Methods for distributing objects on the plane. In *Proceedings of EUROGRAPHICS*, (2003), 515–522.
- [21] Jones T. R., Efficient generation of Poisson-disk sampling patterns, *Journal of Graphics Tools* 11, 2 (2006), 27–36.
- [22] Klassen R. V., Filtered jitter, *Computer Graphics Forum* 19, 4 (2000), 223–230.
- [23] Kopf J., Cohen-Or D., Deussen O., Lischinski D., Recursive Wang tiles for real-time blue noise. In *Proceedings of SIGGRAPH*, (2006), 509–518.
- [24] Lagae A., Dutre P., A procedural object distribution function, *ACM Transactions on Graphics* 24, 4 (2005), 1442–1461.
- [25] Lagae A., Dutre P., An alternative for Wang tiles: Colored edges versus colored corners, *ACM Transactions on Graphics* 25, 4 (2006), 1442–1459.

-
- [26] Lagae A., Dutre P., A comparison of methods for generating Poisson disk distributions, *Computer Graphics Forum* 27, 1 (2008), 114–129.
- [27] Lagae A., Kaplan C. S., Fu C.-W., Ostromoukhov V., Deussen O., Tile-based methods for interactive applications, *SIGGRAPH 2008 Class Notes* ACM, 2008.
- [28] Lu P. J., Steinhardt P. J., Decagonal and quasi-crystalline tilings in medieval Islamic architecture, *Science* 315, 5815 (2007), 1106–1110.
- [29] Maskov Z., Patera J., Zich J., Classification of Voronoi and Delone tiles of quasicrystals III: Decagonal acceptance window of any size, *Journal of Physics A* 38, 9 (2005), 1947–1960.
- [30] McCool M., Fiume E., Hierarchical Poisson disk sampling distributions. In *Proceedings of Graphics Interface*, (1992), 94–105.
- [31] Meyer Y., Algebraic numbers and harmonic analysis, North-Holland, 1972.
- [32] Mitchell D. P., Spectrally optimal sampling for distribution ray tracing. In *Proceedings of SIGGRAPH*, (1991), 157–164.
- [33] Mojsilovic A., Soljanin E., Color quantization and processing by Fibonacci lattices, *IEEE Transactions on Image Processing* 10, 11 (2001), 1712–1725.
- [34] Moody R. V., Patera J., Quasicrystals and icosians, *Journal of Physics A* 26, 12 (1993), 2829–2853.
- [35] Moody R. V., Patera J., Dynamical generation of quasicrystals, *Letters in Mathematical Physics* 36, 3 (1996), 291–300.
- [36] Ostromoukhov V., Mathematical tools for computer-generated ornamental patterns. In *Lecture Notes in Computer Science*, (1998), vol. 1375, 193–223.
- [37] Ostromoukhov V., Donohue C., Jodoin P. M., Fast hierarchical importance sampling with blue noise properties. In *Proceedings of SIGGRAPH*, (2004), 488–495.
- [38] Ostromoukhov V., Building 2D low-discrepancy sequences for hierarchical importance sampling using dodecagonal aperiodic tiling. In *Proceedings of GRAPHICON*, (2007), 139–142.
- [39] Ostromoukhov V., Sampling with polyominoes. In *Proceedings of SIGGRAPH*, (2007), 078: 1–6.
- [40] Patera J., Non-crystallographic root systems and quasicrystals. In *The mathematics of long-range aperiodic order*, Proceedings of NATO-ASI, (1997), vol. C 489, 443–465.
- [41] Press W., Teukolsky S. A., Vetterling W. T., Flannery B. P., Numerical recipes in C, 2 ed, Cambridge University Press, 1992.
- [42] Rangel-Mondragon J., Abas S. J., Computer generation of Penrose tilings, *Computer Graphics Forum* 7, 1 (1988), 29–37.
- [43] Secord A., Weighted Voronoi stippling. In *Proceedings of the Second International Symposium on Non-photorealistic Animation and Rendering*, (2002), 37–43.
- [44] Senechal M., Quasicrystals and geometry, Cambridge University Press, 1995.
- [45] Sharma G., Digital color imaging handbook, CRC Press, 2003.
- [46] Shechtman D., Blech I., Gratias D., Cahn J. W., Metallic phase with long-range orientational order and no translational symmetry, *Physical Review Letters* 53, 20 (1984), 1951–1953.
- [47] Shirley P., Discrepancy as a quality measure for sample distributions. In *Proceedings of EUROGRAPHICS*, (1991), 183–194.
- [48] Stam J., Aperiodic texture mapping. *European Research Consortium for Informatics and Mathematics Technical Report ERCIM-01/97-R046*, 1997.
- [49] Wei L.-Y., Tile-based texture mapping on graphics hardware. In *Proceedings of the ACM Conference on Graphics Hardware*, (2004), 55–63.
- [50] Wei L.-Y., Parallel Poisson disk sampling. In *Proceedings of SIGGRAPH*, (2008), 020: 1–10.
- [51] White K. B., Cline D., Egbert P. K., Poisson disk point sets by hierarchical dart throwing. In *Proceedings of the IEEE Symposium on Interactive Ray Tracing*, (2007), 129–132.