

Fast Hierarchical Importance Sampling with Blue Noise Properties

Victor Ostromoukhov*
University of Montreal

Charles Donohue †
University of Montreal

Pierre-Marc Jodoin ‡
University of Montreal




Figure 1: A high dynamic range 1024×512 environment map [Debevec 98] sampled with 3000 point lights. In this image, importance density is represented by the lightness of the background. It took 0.064 seconds on a 2.6 GHz P4 to generate this point set. Similar results using a hardware accelerated Lloyd relaxation [Hoff et al. 1999] required 1 second, while *Structured Importance Sampling* [Agarwal et al. 2003] took 1393 seconds.

Abstract

This paper presents a novel method for efficiently generating a good sampling pattern given an importance density over a 2D domain. A Penrose tiling is hierarchically subdivided creating a sufficiently large number of sample points. These points are numbered using the Fibonacci number system, and these numbers are used to threshold the samples against the local value of the importance density. Pre-computed correction vectors, obtained using relaxation, are used to improve the spectral characteristics of the sampling pattern. The technique is deterministic and very fast; the sampling time grows linearly with the required number of samples. We illustrate our technique with importance-based environment mapping, but the technique is versatile enough to be used in a large variety of computer graphics applications, such as light transport calculations, digital halftoning, geometry processing, and various rendering techniques.

CR Categories: I.3.3 [Picture/Image Generation]: Anti-aliasing; I.3.m [Miscellaneous]: Sampling.

Keywords: Rendering, Importance Sampling, Deterministic Sampling, Hierarchical Representation, Environment Mapping, Digital Halftoning, Blue Noise, Lookup Table-based Techniques, Penrose Tiling, Fibonacci Number System.

*e-mail: ostrom@iro.umontreal.ca

†e-mail: donohuec@iro.umontreal.ca

‡e-mail: jodoinp@iro.umontreal.ca

1 Introduction

Sampling is ubiquitous in computer graphics. Many researchers have studied how the properties of sampling may affect the quality of the achieved results in applications such as ray tracing, Monte Carlo path tracing, motion blur, geometry processing, digital halftoning, etc. Nowadays, it is generally accepted that isotropic two-dimensional sampling with blue noise Fourier spectrum is well suited for a large range of applications – see [Cook 1986], [Ulichney 1988], [Shirley 1991], [Mitchell 1991], [McCool and Fiume 1992], [Glassner 1995], [Hiller et al. 2001], [Kollig and Keller 2002], [Kollig and Keller 2003].

Often, these graphics applications need distributions of samples proportional to an importance that results from a prior treatment (e.g., BRDF of a surface, distribution of light energy, and geometrical properties). The problem of 2D importance sampling with blue noise can be stated as follows:

- **Given** the importance density I on a domain D , as an analytical function or in the form of an array of discrete values. Without loss of generality, I can be normalized in such a way that $0 \leq I(x, y) \leq 1 \quad \forall (x, y) \in D$.
- **Find** a set of discrete samples, whose local density of samples (the number of samples per unit area, calculated locally) is proportional to the importance density I , and whose Fourier spectrum exhibits the following properties: (a) low angular anisotropy, and (b) characteristic blue noise profile of the radial component, i.e., a low-magnitude disk around the DC term, a high-magnitude annulus that corresponds to the mean distance between the samples, and a surrounding medium-magnitude background exterior to the annulus (see more details in [Ulichney 1987], [Hiller et al. 2001]).

Many different techniques have been developed in order to solve this problem. Some of them, known as relaxation techniques, can produce solutions of remarkable quality. In particular, Lloyd's relaxation [Lloyd 1983] and its variants lead to centroidal Voronoi tessellations [Du et al. 1999]. Unfortunately, the price paid for this

quality is high: relaxation techniques are fundamentally slow because they have to solve, often iteratively, the problem of neighborhood determination of each point with respect to all others. Even the most advanced and optimized implementations remain slow. The hardware-assisted implementation of Lloyd’s relaxation [Hoff et al. 1999] is faster but is limited by the resolution of the frame buffer. Some techniques use a form of stochastic sampling (dart throwing), such as the method proposed in [McCool and Fiume 1992], where random points are added or rejected according to the proximity to previous points. Due to the low convergence rate of these methods, their running times are at best in the same order as Lloyd’s. These methods are strictly descending and can be very sensitive to the initial point set.

Other approaches employed in digital halftoning, known as error-diffusion techniques (see [Ulichney 1987], [Ostromoukhov 2001], [Zhou and Fang 2003]), are considerably faster because only a very limited neighborhood of each point is examined. An example of efficient usage of error-diffusion in geometry processing has been exploited in [Alliez et al. 2002]. The main drawback of error-diffusion is the discrete nature of the elements on which it operates: they must be rectangular tiles with fixed spatial resolution. This limits considerably the use of error-diffusion as a general-purpose sampling technique for computer graphics, where multi-resolution sampling is often needed. This drawback has been explicitly mentioned in [Surazhsky et al. 2003] where Lloyd’s relaxation was preferred to error-diffusion.

Another fast sampling technique that could compare to ours in terms of running times would be to use a cumulative density function (CDF), generated from the probability density, and to sample it with a stratified Monte-Carlo technique. Although such an approach can generate points that reflect the required local density, they do not follow the desired blue-noise distribution, as shown in Figure 13. Recently, [Secord et al. 2002] have used a similar approach, using some well-known low-discrepancy sequences such as Halton and Sobol sequences (see [Niederreiter 1992] combined with CDF, in order to distribute graphics primitives at interactive rates, in an NPR context. Although this deterministic approach is very promising, the convincing multi-purpose results have yet to be seen (see Figure 13).

In this paper, we introduce a novel Penrose tiling-based importance sampling technique that presents certain advantages over existing techniques. It belongs to the family of point sampling, that is, each point is processed independently of other points. The treatment for each point is simple and computationally inexpensive, which guarantees the very high speed of our algorithm. Moreover, thanks to off-line optimization and to a specially designed lookup table, the quality of the sampling is high, approaching the quality of centroidal Voronoi tessellations. The size of the lookup table is reasonably small (typically, less than 1K of data). No data-dependent pre-calculation is needed. Our technique is multi-resolution and can be successfully applied on high dynamic range images (we illustrate this in Section 5).

The rest of the paper is organized as follows. In Section 2, we recall some historical facts and properties of Penrose tiling. In Section 3 we describe the core of our sampling system. In Section 4, we enrich the basic technique with an advanced relaxation that produces an almost perfect blue noise Fourier spectrum at all importance levels. In Section 5, we apply, as a case study, the proposed technique to importance-based environment mapping. Finally, in Sections 6 and 7, we discuss future work and draw some conclusions.

2 Penrose Tiling

The history of Penrose tiling is fascinating. It goes back to the work of Johannes Kepler, a 17th-century astronomer and mathematician.




Figure 2: Left: Kepler’s drawing from *Harmonice Mundi* published in 1619. This tiling inspired Roger Penrose to discover his aperiodic tiling composed of two marked rhombs (top right). Tiles from his original 1979 article contain arrows as matching rules that force the aperiodicity of the tiling. Bottom right: Optical Fourier transform of the vertices of Penrose tiling obtained in 1982 by Alan Mackay.




Figure 3: Left: Original Penrose tiling with two kinds of rhombs. Right: An alternative representation of the same tiling, where the rhombs are split in two halves, and the pentagons of two kinds are placed at the vertices of the original tiling.

In his book *Harmonice Mundi*, he published an atlas of various tilings with regular polygons. One of them, shown in Figure 2 (left), excited the imagination of many mathematicians over a long period. Is it possible to tile the plane only with regular pentagons, decagons, and five-pointed stars? According to Kepler’s drawing, it was possible if one permitted also strange peanut-shaped figures (“monsters”), such as the one visible underneath the label “Aa”.

In the early 1970s, a modern physicist and mathematician, Roger Penrose, was mesmerized by Kepler’s drawing. He modified it in such a way that he was able to tile the plane non-periodically with a similar set of tiles. And he did much more: he found that introducing special matching rules such as marks on the edges of the tiles will preclude any periodic arrangements of the tiles. Still, the tiling shows a clearly identifiable local order. This tiling belongs to the family of *aperiodic structures*, i.e., structures whose non-periodicity is forced by the matching rules. Penrose published a first account of his discovery in [Penrose 1974]. Later, Penrose published a paper where he presented three different but tightly related aperiodic tiling systems with matching rules [Penrose 1979]. One of them, shown in Figure 2 (top right), has only two extremely simple shapes, two different rhombs with matching rules. In 1977, Martin Gardner published in his column in *Scientific American*, an enthusiastic account of Penrose’s discovery [Gardner 1977]. After Gardner’s publication, Penrose tiling became well-known to a large number of mathematicians, physicists, and chemists.

The Penrose tiling with rhombs shown in Figure 2 (top right) and in Figure 3 (left) has attracted special attention, due to its simplicity. In the early 1980s, Alan Mackay generalized the Penrose tiling to

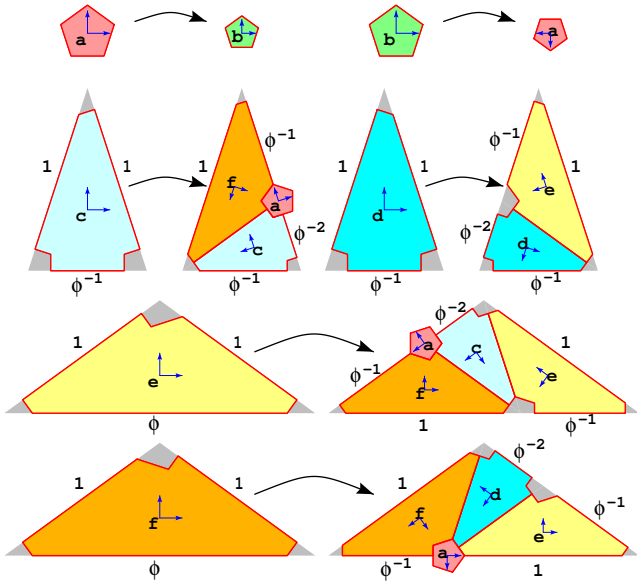


Figure 4: Subdivision rules for modified Penrose tiling, as shown in Figure 3 (right). The Golden Ratio $\phi = \frac{1+\sqrt{5}}{2} \approx 1.61803$. Pairs of orthogonal vectors form the basis for each tile.

three dimensions. He also performed an optical Fourier transform of the pattern of holes perforated at the vertices of Penrose tiling (see Figure 2 (bottom right), reproduced from [Mackay 1982]). This picture caused a sensation. In fact, it looks like a Bragg diffraction pattern, but it clearly violates a well-established principle, the so-called crystallographic restriction, which states that a diffraction pattern of a crystal may have only a two-, three-, four-, or six-fold rotational symmetry. The pattern obtained by Mackay in fact had ten-fold rotational symmetry! A new science of quasicrystallography was born. In 1984, Shechtman and his colleagues synthesized a new matter, a quasicrystal, that had a diffraction pattern close to that predicted by Mackay. It was the first time in human history that a new kind of matter was predicted and analyzed before being physically synthesized in the laboratory. Since then, hundreds of papers on quasicrystals have been published. A corpus of fundamental papers on the physics and mathematics of quasicrystals can be found in [Steinhardt and Ostlund 1987]. Grünbaum and Shephard, in their theory of tiling bible [Grünbaum and Shephard 1986], devote a whole section to Penrose tiling, where they provide a detailed analysis.

Alternatively, Penrose tiling with rhombs can be represented as a tiling with six shapes: both “fat” and “thin” rhombs are split into two triangular halves, as shown in Figure 3 (right). In addition, regular pentagons of two kinds, which we also call “sampling tiles” are placed at each vertex of the original Penrose rhombs. The pentagons play the role of matching rules that enforce aperiodicity. The tiling can be achieved by applying on each tile the subdivision process as shown in Figure 4 (Grünbaum and Shephard call this subdivision “inflations”). Geometrical proportions for all sides during the subdivision process are shown in Figure 4. Note that the size of the pentagons with respect to the triangles does not matter. Without loss of generality the pentagons can be taken to be infinitesimal, and the half-rhombs are triangular. The positions and orientations of tiles, schematically represented in Figure 4 by pairs of orthogonal vectors that form the basis for each tile, are important for our construction.

Penrose tiling has attracted our attention for several reasons. First, it is obvious from a glance at Mackay’s optical Fourier transform

shown in Figure 2 (bottom right) that it is surprisingly close to the blue noise spectrum, a goal we fixed for our sampling system. In fact, annuli of spectral peaks around the DC term are clearly visible. Second, it can be easily observed that the pentagons obtained at one level of subdivision are enriched by the pentagons of the next level, which are placed between the pentagons of the previous level (see Figure 5).

Although another famous aperiodic tiling, Wang tiling, has recently been successfully exploited in computer graphics for the generation of Poisson distributions of points [Hiller et al. 2001], [Cohen et al. 2003], Penrose tiling has been used only in the context of visualization of decorative properties of the tiling (see for example [Glassner 1998]).




Figure 5: Two consecutive subdivision levels of Penrose tiling. Red dots mark sampling tiles of level n , blue dots – that of level $n + 1$.

3 Two-dimensional Penrose-based Hierarchical Importance Sampling

Let us consider the Penrose subdivision process shown in Figure 4 as a recursive subdivision process. A special binary code called F-code is assigned to each tile. This subdivision process can be described by the following production rules:

$$\mathcal{P}_{Penrose} := \begin{cases} a_* \mapsto \{b_{00*}\} \\ b_* \mapsto \{a_{00*}\} \\ c_* \mapsto \{f_{00*}, c_{10*}, a_{10*}\} \\ d_* \mapsto \{e_{00*}, d_{10*}\} \\ e_* \mapsto \{f_{00*}, c_{10*}, e_{01*}, a_{10*}\} \\ f_* \mapsto \{e_{00*}, d_{10*}, f_{01*}, a_{01*}\}, \end{cases} \quad (1)$$

where x_y means a tile of type x having F-code y . The symbol ‘*’ replaces the F-code of a tile before subdivision. Each subdivision left-concatenates two symbols to the current F-code. Thus, after n subdivisions, the F-code will have the length of $2n$ symbols. F-codes can be interpreted as integer numbers in the Fibonacci number system as described in [Knuth 1997] and [Graham et al. 1994]. Appendix B provides some basic facts about the Fibonacci number system, together with the pseudo-code of the routine FIBOTODECIMAL that converts F-codes to the conventional decimal representation.

Figure 6 shows the first three subdivisions applied to a pair of tiles of type ‘e’ and ‘f’ (top left).

Three important observations can be made:

- Decimal numbers that correspond to F-codes assigned to pentagonal sampling tiles of type ‘a’ and ‘b’ are all in the range $[1..(F_{2(n+1)} - 1)]$, where n is the subdivision level of the initial tiles, and F_i are the Fibonacci numbers.
- Successive subdivisions enrich sampling points obtained with the previous subdivisions, putting new sampling points in between.

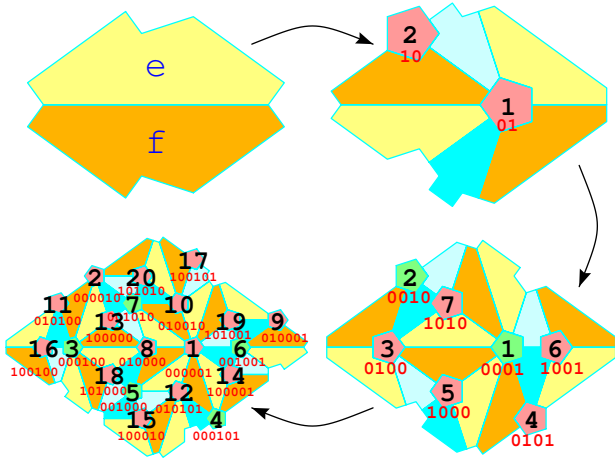


Figure 6: Three subdivisions according to the production rules (1). Only the F-codes of pentagonal sampling tiles are shown (in red), together with corresponding decimal values (in black).

- Decimal numbers already assigned at a subdivision level n will remain at the same positions during all further subdivisions. The principle of this distribution is similar to that of Bayer’s dispersed-dot dither [Bayer 1973], [Foley et al. 1990] or that of rotated dispersed-dot dither [Ostromoukhov et al. 1994].

This allows us to build an adaptive importance sampling system based on the Penrose subdivision system with the production rules (1). Our adaptive importance sampling system is simple. First, we cover the area of interest, where the importance is defined, with a pair of tiles of type ‘e’ and ‘f’, as shown in Figure 6 (top left). Then, we apply the recursive subdivision process according to the production rules (1). We stop subdividing when the required local subdivision level κ is reached. In this case, we output the center of the ‘a’ and ‘b’ type tiles, if the local importance is greater than the decimal value of the F-code of the current tile. Pseudo-code for this algorithm is shown in Appendix A.

Importance density may be scaled by a factor mag , constant for the entire importance density image, in order to obtain the desired number of points. This effect is illustrated in the companion video¹. The required local level of subdivision κ can be determined as

$$\kappa = \lceil \log_{\phi^2} \max_{tile} (mag \cdot I(x,y)) \rceil, \quad (2)$$

where $\lceil \cdot \rceil$ is the usual notation for ceiling, $I(x,y)$ is the importance value at position (x,y) , and $\phi = \frac{1+\sqrt{5}}{2}$ is the *Golden Ratio*. The factor \log_{ϕ^2} can be explained as the factor of self-similarity of Penrose tiling. In fact, from one level of subdivision to the next, the area of Penrose tiles diminishes by factor ϕ^2 . The value $\max_{tile}(\cdot)$ can be achieved with standard scan-conversion on the triangle, for tiles of type ‘c’, ‘d’, ‘e’, and ‘f’ (no scan-conversion is needed for tiles of type ‘a’ and ‘b’ that are supposed to be infinitesimal). This scan-conversion is opened to possible optimization. If less precision is required but speed is capital, the importance can be tested only at a few points within the tile.

4 Lookup Table-based Relaxation

To improve the spatial distribution of the sampling points, we create a table of corrective vectors, which is used at run-time to relocate the sampling points. These corrective vectors, expressed in terms of

¹Siggraph 2004 Full Conference DVD-ROM; also available on the web site of the first author: www.iro.umontreal.ca/~ostrom/ImportanceSampling




Figure 7: Lloyd’s relaxation applied on sampling point set of constant importance, produced with our Penrose tiling-based system. Small corrections are shown as yellow lines connecting the centers of uncorrected (blue) and corrected (red) sampling points.

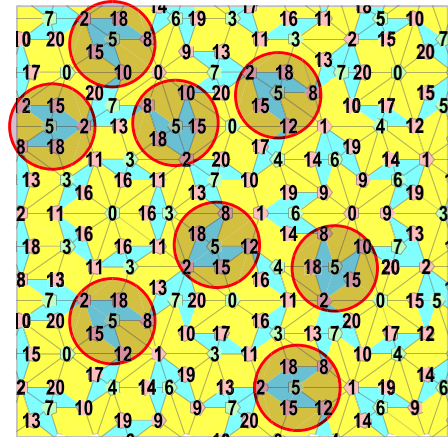


Figure 8: Structural indices i_s obtained by converting the 6 most significant bits of the F-code assigned to each tile to the conventional decimal representation. Only sampling (pentagonal) tiles are considered. In this figure, we use identical colors for tiles of type ‘b’ and ‘c’, ‘d’ and ‘e’, to better visually identify the structure. Notice how the local neighborhoods around tiles with the same label are similar, after rotation. Highlighted are the neighborhoods around pentagons labeled with $i_s = 5$.

orthogonal basis proper to each tile, will have the effect of “relaxing” the point distribution. See Figure 7 and the companion video¹.

Unfortunately, even though there are only two different sampling tiles, the fact that Penrose tiles fill the plane aperiodically makes it impossible to account for every possible correction vector. Nevertheless, the self-similar nature of the tiling can be harnessed to obtain a limited number of corrective vectors. To accomplish this, we relabel the sampling tiles with what we call a “structural index”, i_s , which is calculated from the first six bits of their F-code (see Figure 8). This gives a total of 21 different labels (the maximum value encoded with the F-system over 6 bits). Each of these 21 labels has a corresponding corrective vector. This six-bit structural indexing has been found experimentally.

Because these corrective vectors have to be representative of any importance density function, we must optimize them with regards to different importance values. We chose to optimize the vectors over n importance values that are represented by what we call the “importance index”, i_v , which is calculated as follows:

$$i_v = \lfloor n \cdot \psi(mag \cdot I(x,y)) \rfloor, \quad (3)$$

where $\lfloor \cdot \rfloor$ is the usual notation for the floor, $I(x,y)$ is the importance value at position (x,y) and $\psi(\cdot)$ maps a real number onto the interval $[0..1]$. We found that with $n = 8$ we get rather smooth gradations across importance values.

The resulting corrective vectors optimized over each importance

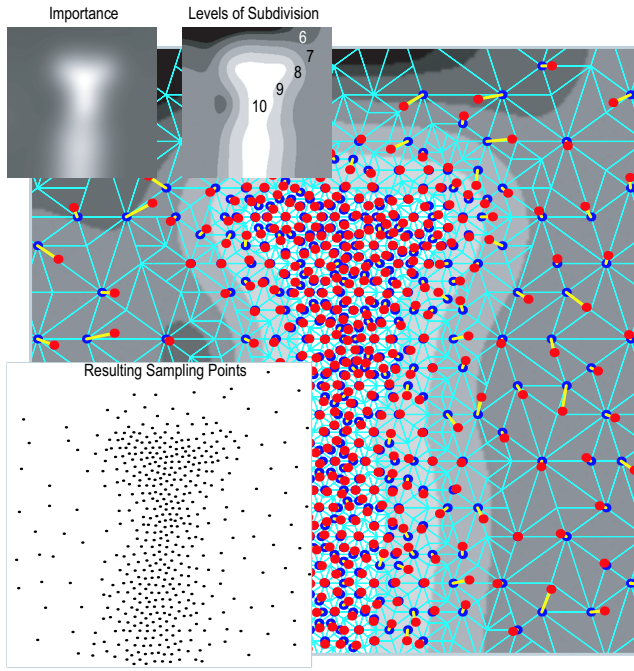


Figure 9: A typical importance map sampled with our system. Blue dots indicate the centers of the sampling tiles. Yellow lines show lookup table-based corrections applied on sampling points. Red dots show the corrected sampling points. Please note seamless transitions between zones of different levels of subdivision.

value are stored in a 8×21 lookup table where the first index is addressed by the importance index i_v and the second dimension is addressed by the structural index i_s .

The procedure used to generate the lookup table is the following.

First, we initialize the lookup table with zero vectors. Then, for each importance index, i_v , we apply the following process.

1. With our sampling system, create a large patch of sampling points that correspond to the current importance, using the latest version of the lookup table for corrections as will be explained below.
2. Apply Lloyd's relaxation on this set. Be careful to make peripheral points immovable. This is needed because the patch is finite. See Figure 7.
3. For each sampling point, calculate the difference between the uncorrected position and the relaxed position.
4. Calculate the mean value of correction vectors, for all sampling tiles with the same structural index i_s . Store the results in the table at position (i_v, i_s) .
5. Repeat steps (1)-(4) until convergence is attained (typically 5 to 10 iterations).

Considering the lookup table as a vector field, a low-pass filter is then applied to the vectors across importance indices. Then, the whole process is repeated until convergence is attained (typically 5 to 10 iterations). Thus, relaxation and low-pass filtering are applied alternatively.

At the end of the optimization, the corrected points closely match those obtained by true Lloyd relaxation, for all importance levels. The low-pass filtering between each iteration ensures that the points will be distributed adequately over gradients in non-constant importance density functions.

Figures 9, 13 (top), 10, and 11 (top right) illustrate the results achieved with our system. Please note how our technique captures nuances of importance in all subranges of the dynamic range.

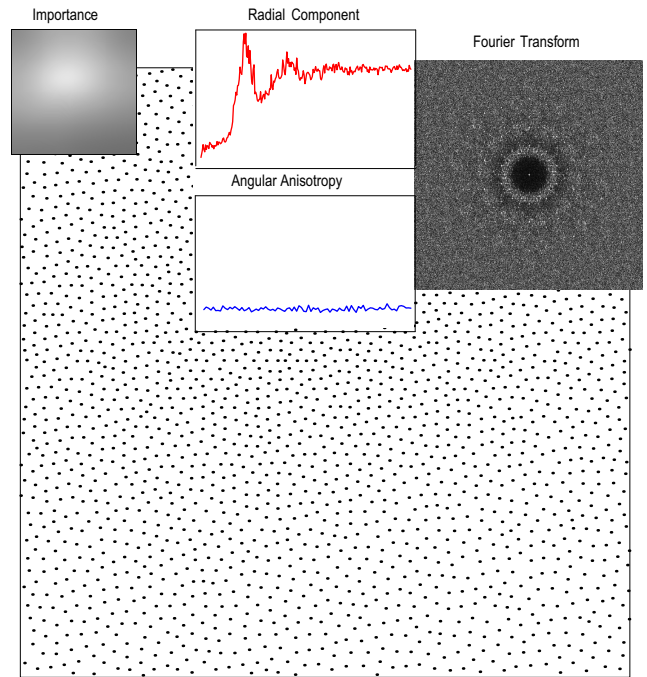


Figure 10: Fourier amplitude spectrum (DFT) of a smooth importance density image sampled with 2600 points, together with its radial component and angular anisotropy of the spectrum. Please notice the typical blue noise profile of the radial component.

Several examples of lookup tables of various sizes can be found on¹. Please notice that the lookup table of size 8×21 described here is compact, yet still gives satisfactory results. Larger size lookup tables are applicable as well.

5 Case Study: Environment Map Sampling

One of many applications of our technique in computer graphics is the sampling of HDR environment maps. The idea is to reduce the environment map to a relatively small number of point light sources, thus speeding up the integration of the incoming illumination. We compare our sampling system with those used in recent incarnations of the above idea, notably *Structured Importance Sampling* [Agarwal et al. 2003] and *LightGen* [Cohen and Debevec 2001]. See the companion video¹.

This problem essentially reduces to the k-centers problem, to which there is no known polynomial-time solution. To achieve a fast solution, an approximation must be used. [Agarwal et al. 2003] use the Hochbaum-Shmoys algorithm, and [Cohen and Debevec 2001] use k-means clustering; both are iterative searches. In [Kollig and Keller 2003], a modified Lloyd's relaxation scheme is used to distribute the sampling points, which is also an iterative process. In our system, the sampling points are deterministic and the lookup table is pre-calculated; only a thresholding operation must be performed during the sampling. Thus, we can obtain an empirical linear time approximation (time on a 2.6 GHz P4 processor):

No. of points	236	343	455	690	930	1847	3006
Time in ms	6	9	12	17	22	42	64

The resulting sample distribution compares well with the other techniques. The running times, though, are several orders of magnitude lower. In order to obtain running times similar to our system,

¹Siggraph 2004 Full Conference DVD-ROM; also available on the web site of the first author: www.iro.umontreal.ca/~ostrom/ImportanceSampling




Figure 11: Galileo’s Tomb environment map, sampled with **left:** *LightGen*, **center:** *Structured Importance Sampling*, and **right:** our system. The map size is 1024×512 , the number of samples is 300. On a 2.6 GHz P4, sampling times were: **left:** 45 minutes, **center:** 25 seconds, **right:** 8 milliseconds.

one could use the cumulative importance density function sampled with a stratified Monte-Carlo sampling pattern or low-discrepancy sequences, as shown in Figure 13. As mentioned before, the resulting sample points do not exhibit a blue-noise distribution. Whether a blue-noise distribution of the lights is better in this context can be debated, but that is beyond the scope of this paper.

With the relative speed of our system, the bottleneck quickly becomes the rendering process. With a rendering system that could handle a few hundred lights in real time, it would be possible to simulate distant illumination with dynamic maps at interactive frame-rates, using our sampling system.

See Figures 1 and 11 and the companion video.

6 Discussion and Future Work

Several important notions are given in this paper without formal proof. For example, the property of uniformity of point distribution introduced in Section 3 must be thoroughly studied and adequately presented. This is a considerable work that goes far beyond the scope of this paper. Adequate mathematical tools for this study should be developed.

As we mentioned in Section 2, three-dimensional extension of Penrose tiling was proposed by Mackay more than twenty years ago; it played an important role in the discovery of quasicrystals. Consequently, one may expect to build a 3D construction, similar to our construction in Section 3, which would result in isotropic 3D point distribution, modulated by 3D importance density functions.

Such hierarchical constructions would be useful in various computer graphics applications. Moreover, a multi-dimensional variant of our construction may exist as well. It would be helpful in various light transport calculations where multi-dimensional importance density functions must be sampled (see also [Veach 1997], [Kollig and Keller 2001], [Kollig and Keller 2002]).

Sampling point sets produced with our method are not perfectly isotropic. One possible way to improve our results would be studying different available aperiodic tilings with the methodology introduced in this paper. Among the known aperiodic tilings that share various properties with Penrose tiling, we are considering Ammann’s octagonal tiling [Grünbaum and Shephard 1986], Socolar’s dodecagonal tiling [Socolar 1989], and Wang tiling [Grünbaum and Shephard 1986], [Hiller et al. 2001], [Cohen et al. 2003].

Another way to improve isotropy would be to use our method as a starting point for some other methods, such as weighted Lloyd’s relaxation, which will then converge in much fewer iterations.

7 Conclusions

Let us summarize the contributions of this paper.

First, we have proposed an original method for sequentially numbering all vertices of Penrose tiling, based on the Fibonacci number system. The Penrose tiling is used as an underlying structure for a recursive subdivision. The numbers associated with the tiles are used as thresholds in the sampling process.

Second, we improve the above system with corrective vectors to en-

sure blue noise properties of the sampling point distribution. This is achieved through an off-line Lloyd relaxation scheme. The corrective vectors obtained in the optimization are stored in a two-dimensional lookup table.

Our technique is very fast because the required processing per sampling point is simple. The processing time grows linearly with the number of sampling points. Typical processing time for sampling a scene with thousands of sampling points can be measured in milliseconds. Because each tile is processed independently, the proposed system is parallelizable and therefore can be efficiently implemented with hardware.

Because of its speed, simplicity, and multi-resolution properties combined with good quality of point distribution, our importance sampling technique may be applied in a large variety of graphical applications.

8 Acknowledgements

For the first author, this paper represents a culmination of a very long-term research project that goes back to his Ph.D. thesis some ten years ago. We would like to express our gratitude to many people who were involved at different degrees in discussions related to the present paper. Thanks to Roger Hersch, Isaac Amidror, David Salesin, Branko Grünbaum, Przemyslaw Prusinkiewicz, Craig Kaplan, Douglas Zongker, Eric Stollnitz, Daniel Wood, Julie Dorsey, Pierre Alliez, Mathieu Desbrun, Alexander Keller, Jiri Patera, Mark Grundland, Pierre McKenzie, Jean Vaucher, and Neil Stewart. We would like to thank the anonymous SIGGRAPH reviewers for their constructive and pertinent comments. Very special thanks to Frédo Durand and Pierre Poulin for their unconditional support and help. Finally, we are indebted to our families for their support during the hard days of the paper submission.

References

- AGARWAL, S., RAMAMOORTHI, R., BELONGIE, S., AND JENSEN, H. 2003. Structured importance sampling of environment maps. *ACM Trans. on Graphics* 22, 3 (July), 605–612.
- ALLIEZ, P., MEYER, M., AND DESBRUN, M. 2002. Interactive geometry remeshing. *ACM Trans. on Graphics* 21, 3, 347–354.
- BAYER, B. 1973. An optimum method for two-level rendition of continuous-tone pictures. In *IEEE Int. Conf. on Communications*, 11–15.
- COHEN, J., AND DEBEVEC, P. 2001. *LightGen, HDRShop plugin*. <http://www.ict.usc.edu/~jcohen/lightgen/lightgen.html>.
- COHEN, M., SHADE, J., HILLER, S., AND DEUSSEN, O. 2003. Wang tiles for image and texture generation. *ACM Trans. on Graphics* 22, 3 (July), 287–294.
- COOK, R. 1986. Stochastic sampling in computer graphics. *ACM Trans. on Graphics* 5, 1 (Jan.), 51–72.
- DEBEVEC, P. 1998. Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proc. SIGGRAPH '98*, 189–198.
- DU, Q., FABER, V., AND GUNZBURGER, M. 1999. Centroidal Voronoi tessellations: Applications and algorithms. *SIAM Review* 41, 4 (Dec.), 637–676.
- FOLEY, J., VAN DAM, A., FEINER, S., AND HUGHES, J. 1990. *Computer Graphics, Principles and Practice*, 2nd ed. Addison-Wesley.
- GARDNER, M. 1977. Extraordinary nonperiodic tiling that enriches the theory of tiles. *Scientific American* 236, 110–121.
- GLASSNER, A. 1995. *Principles of Digital Image Synthesis*. Morgan Kaufmann.
- GLASSNER, A. 1998. Andrew Glassner's notebook: Penrose tiling. *IEEE Computer Graphics & Applications* 18, 4, 78–86.
- GRAHAM, R., KNUTH, D., AND PATASHNIK, O. 1994. *Concrete Mathematics: a Foundation for Computer Science*, 2nd ed. Chapter 6.6. Addison-Wesley.
- GRÜNBAUM, B., AND SHEPHARD, G. 1986. *Tilings and Patterns*. W.H. Freeman.
- HILLER, S., DEUSSEN, O., AND KELLER, A. 2001. Tiled blue noise samples. In *Proc. Vision Modeling and Visualization*, 265–272.
- HOFF, K., CULVER, T., KEYSER, J., LIN, M., AND MANOCHA, D. 1999. Fast computation of generalized voronoi diagrams using graphics hardware. In *Proc. SIGGRAPH '99*, 277–286.
- KNUTH, D. 1997. *The Art of Computer Programming, Volume 1, Fundamental Algorithms*, 3rd ed. page 86. Addison-Wesley.
- KOLLIG, T., AND KELLER, A. 2001. Efficient bidirectional path tracing by randomized quasi-monte carlo integration. *Niederreiter, K. Fang, and F. Hickernell, Eds., Monte Carlo and Quasi-Monte Carlo Methods 2000*, 290–305.
- KOLLIG, T., AND KELLER, A. 2002. Efficient multidimensional sampling. *Computer Graphics Forum* 21, 3, 557–564.
- KOLLIG, T., AND KELLER, A. 2003. Efficient illumination by high dynamic range images. In *Eurographics Symposium on Rendering: 14th Eurographics Workshop on Rendering*, 45–51.
- LLOYD, S. 1983. An optimization approach to relaxation labeling algorithms. *Image and Vision Computing* 1, 2, 85–91.
- MACKAY, A. 1982. Crystallography and the Penrose pattern. *Physica* 114A, 609–613.
- MCCOOL, M., AND FIUME, E. 1992. Hierarchical poisson disk sampling distributions. In *Proc. Graphics Interface '92*, 94–105.
- MITCHELL, D. 1991. Spectrally optimal sampling for distributed ray tracing. In *Proc. SIGGRAPH '91*, vol. 25, 157–164.
- NIEDERREITER, H. 1992. *Random Number Generation and Quasi-Monte-Carlo Methods*. Soc. for Industrial and Applied Mathematics.
- OSTROMOUKHOV, V., HERSCH, R., AND AMIDROR, I. 1994. Rotated dispersion dither: a new technique for digital halftoning. In *Proc. SIGGRAPH '94*, 123–130.
- OSTROMOUKHOV, V. 2001. A simple and efficient error-diffusion algorithm. In *Proc. SIGGRAPH 2001*, 567–572.
- PENROSE, R. 1974. The role of aesthetics in pure and applied mathematical research. *Bull. Inst. Math. & its Applns.* 10, 266–271.
- PENROSE, R. 1979. Pentaplexity, a class of non-periodic tilings of the plane. *The Mathematical Intelligencer* 2, 32–37.
- SECORD, A., HEIDRICH, W., AND STREIT, L. 2002. Fast primitive distribution for illustration. In *13th Eurographics Workshop on Rendering*, 215–226.
- SHIRLEY, P. 1991. Discrepancy as a quality measure for sample distributions. In *Proc. Eurographics '91*, 183–194.
- SOCOLAR, J. 1989. Simple octagonal and dodecagonal quasicrystals. *Physical Revue B* 39, 10519–10551.
- STEINHARDT, P., AND OSTLUND, S. 1987. *The Physics of Quasicrystals*. World Scientific.
- SURAZHISKY, V., ALLIEZ, P., AND GOTSMAN, C. 2003. Isotropic remeshing of surfaces: a local parameterization approach. In *Proc. of 12th Int. Meshing Roundtable*.
- ULICHNEY, R. 1987. *Digital Halftoning*. MIT Press.
- ULICHNEY, R. A. 1988. Dithering with blue noise. *Proc. of the IEEE* 76, 56–79.
- VEACH, E. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis. Stanford University.
- ZHOU, B., AND FANG, X. 2003. Improving mid-tone quality of variable-coefficient error diffusion using threshold modulation. *ACM Trans. on Graphics* 22, 3 (July), 437–444.

APPENDIX A: Pseudo-code of the Adaptive Subdivision and Sampling

ADAPTIVESAMP(*t* of type *tile*)

```

1   ▷ Structure tile contains the fields:
2   ▷ type: ['a'..'f']
3   ▷ LOS: Level of Subdivision
4   ▷ vertices: center if 'a' or 'b', triangle otherwise
5   ▷ Fcode
6   ▷ mag and importance are global variables
7   local_LOS ← GETMAXLOSWITHINTILE(t)
8   if t.LOS ≥ local_LOS
9     then ▷ Terminal: don't need more subdivisions
10    local_importance ← mag · GETLOCALIMPORTANCE(t)
11    if t.type = ('c' or 'd' or 'e' or 'f')
12      then return ▷ Not a "sampling tile": do nothing
13    if local_importance ≥ FIBOTODECIMAL(t.Fcode)
14      then OUTPUTSAMPLE(t.center)
15    return
16  else ▷ Need more subdivisions
17    {t1, ..., tn} ← SUBDIVIDEUSINGPRODUCTIONRULES(t)
18    return {ADAPTIVESAMP(t1), ..., ADAPTIVESAMP(tn)}
```

It is worth mentioning that the routines GETMAXLOSWITHINTILE(*t*) and GETLOCALIMPORTANCE(*t*) play a very important, even crucial role in the algorithm. If, for any reason (*e.g.*, because of a singularity in the importance density), they fail to evaluate the local importance/max importance, it may result in locally erroneous sampling density.

APPENDIX B: ϕ - and F- Number Systems

Details about ϕ - and F- (Fibonacci) number systems can be found in [Knuth 1997] and [Graham et al. 1994].

The ϕ -system is a positional number system in base ϕ , where $\phi = \frac{1+\sqrt{5}}{2}$ is the *Golden Ratio*. Any rational number x can be expressed in this system exactly as in our conventional binary or decimal systems, except that instead of using powers of two or ten, this system employs powers of ϕ . For example, the number $(101.001)_\phi$ in base ϕ is

$$(101.001)_\phi = \phi^2 + \phi^0 + \phi^{-3} \approx 3.8541_{10}$$

The ϕ -system is closely related to the F-system (the abbreviation for Fibonacci system). The F-system is also a positional system. Any integer n can be presented in the F-system as a sum of Fibonacci numbers F_j multiplied by their positional coefficients, which may be 0's or 1's. Thus, a number n can be expressed by its F-code $(b_m b_{m-1} \dots b_3 b_2)_F$:

$$n = (b_m b_{m-1} \dots b_3 b_2)_F \iff n = \sum_{j=2}^m b_j F_j. \quad (4)$$

The first index in the summation is $j = 2$ because of the convention used for Fibonacci numbers F_j :

$$F_0 = 0, F_1 = 1, F_2 = 1, F_3 = 2, F_4 = 3, F_5 = 5, F_6 = 8, F_7 = 13, \dots$$

The representation of numbers is not unique in the F-system, but it becomes unique if the rule of *normal form* is imposed: two adjacent 1's are not permitted. The procedure of conversion from an arbitrary sequence of 0's and 1's to the normal form, along with many other technical details, can be found in [Graham et al. 1994]. Here are the first twelve integers expressed in the F-system in normal form:




Figure 12: Function $\Psi(x)$. Note that $\Psi(x)$ “jumps” every F_{2j}




Figure 13: A ramp importance density image sampled with (top) our system and the cumulative importance density function sampled with (middle) a stratified Monte-Carlo pattern, or (bottom) Sobol low-discrepancy sequence.

$$\begin{aligned}
1 &= (00001)_F, & 2 &= (00010)_F, & 3 &= (00100)_F, \\
4 &= (00101)_F, & 5 &= (01000)_F, & 6 &= (01001)_F, \\
7 &= (01010)_F, & 8 &= (10000)_F, & 9 &= (10001)_F, \\
10 &= (10010)_F, & 11 &= (10100)_F, & 12 &= (10101)_F.
\end{aligned}$$

More examples of interpretation of such sequences of 0's and 1's, which we shall call F-codes, as integer numbers, are shown in Figure 6.

The routine FIBOTODECIMAL converts F-codes to the conventional decimal representation.

FIBOTODECIMAL(*Fcode*)

```

1  accumulator ← 0
2  for i ← 0 to LENGTH(Fcode) - 1
3    do
4      accumulator ← accumulator + Fcode[i] · Fi+2
5  return accumulator
```

Function $\Psi(x)$ that maps a real positive x onto interval $[0..1]$, as shown in Figure 12, is defined as follows:

$$\Psi(x) = (\log_{\phi^2} \sqrt{5} \cdot x) \bmod 1.$$

It can be easily derived from the well-known Binet's formula

$$F_n = [\phi^n / \sqrt{5}],$$

where $[]$ is the usual notation for the *nint* (Nearest Integer) function.