EG UK Theory and Practice of Computer Graphics (2010)
John Collomosse, Ian Grimstead (Editors)

Approximating Poisson Disk Distributions by Means of a Stochastic Dither Array

Jennifer R. Alford, Ph.D.[1] and David G. Sheppard, Ph.D.[2]

[1]Digital Teapot, Inc., Fort Worth, Texas USA; gralford@acm.org
[2]David G. Sheppard, Ph.D., L.L.C., Tucson, Arizona USA; dgsheppard@ieee.org

Abstract

Achieving blue noise point set distributions has been a common goal of two largely separate research communities: computer graphics and digital halftoning. Computer graphics research has focused largely on geometric solutions in continuous spaces. Digital halftoning research has focused on signal processing solutions in discrete image-based space. Usage of Poisson Disk point sets in computer graphics has grown beyond sampling, including object distribution and texturing, among others. The image-based field of digital halftoning can provide additional tools for graphics researchers and practitioners. It is of interest to explore the suitability of digital halftoning technology to two classic problems in computer graphics: (1) approximating Poisson Disk point distributions of constant density and (2) importance sampling of an underlying importance function. Exemplary methods from each field are implemented and, by applying well-established measures of the radially averaged power spectrum and anisotropy plots, are shown to be quite similar, although the approaches are mathematically not equivalent. Additionally, we compare the relative radius of the point sets. Further, the ability of dither array construction techniques to shape spectral characteristics of dot patterns is shown with several variations of design parameters.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image GenerationBitmap and framebuffer operations ; I.4.1 [Computer Graphics]: Image Processing and Computer VisionDigitization and Image CaptureSampling;

1. Introduction

Stochastic sampling has been of significant interest in computer graphics for over two decades. A large body of work has established that, based on characteristics of the human visual system, sample distributions with spectral characteristics that can be described and measured as "blue noise" are most desirable. The Poisson Disk distribution exhibits the desired spectral characteristics and efficient algorithms to generate such point sets have been the focus of much study. The field of digital halftoning has faced essentially the same problem. The tools that have been used to assess the quality of Poisson Disk distributions, the radially averaged power spectrum and anisotropy plots, originated in the study of digital halftone techniques [Uli87]. In digital halftoning, as in computer graphics, efficient techniques that can produce distributions of points with blue noise characteristics are desirable.

Most of the prior work in producing and analyzing Poisson Disk point sets applies to distributions of constant density. However, the related problem of adaptive sampling is of significant interest. A stochastic dither array approximates solutions to both in a very straightforward manner. This paper lays out the common interests of two separate research communities pursuing similar goals for similar reasons: achieving visual effects. Examples are included to show how a single stochastic dither array can be used to address the two problems simultaneously: 1) approximating Poisson Disk distributions, and 2) importance sampling

The computational efficiencies involved in using dither arrays warrant consideration from the graphics community. One compact, tilable array provides the means to produce a set of Poisson Disk-like points of arbitrary density and a means to implement importance sampling using only point-wise threshold operations. Further, dither array construction techniques provide an ability to shape spectral characteristics.

DOI: 10.2312/LocalChapterEvents/TPCG/TPCG10/083-090

Prior work in both computer graphics and in digital halftoning is explored. The similarities and differences of the two fields are examined by implementing exemplary techniques from each field: a Void and Cluster technique from digital halftoning and the Best Candidate dart throwing technique from computer graphics. Established analysis techniques are employed to highlight the similarity of spectral results. Results are also presented to show the flexibility of the Void and Cluster technique to shape spectral properties by filter variation, creating point distributions that are no longer "blue." The potential of using halftoning technology in computer graphics is shown by including examples of importance sampling using a stochastic dither array with pointwise thresholding.

2. Prior Work

Poisson Disk sample distributions have been of significant interest in computer graphics since Dippe and Wold [DW85] introduced their use in antialiasing. Noting the computational demands of computing a Poisson Disk distribution, Cook [Coo86] implemented jittering as an efficient approximation for use in a variety of distributed ray-tracing lighting effects. Poisson Disk distributions are useful for a growing number of problems, such as object placement, texture generation and non-photorealistic rendering. A recent comparative survey established and applied a standard framework of analysis for point sets of constant density and discusses tradeoffs among approaches for different application needs [LD08].

Early successful techniques based on applying constraints to randomly generated points [MF92, Mit91] have been improved with a variety of computational techniques [Jon06, Wei08, DH06, GM09]. Work based on polyominoes [Ost07] lends itself easily to the related problem of adaptive sampling, at the expense of slightly more visible structure. Much efficiency is gained with the use of precomputed data sets applied in conjunction with tiling schemes [CSHD03, HDK01, LD06], but tiling can introduce structured artifacts on large point sets.

Adaptive sampling addresses the related problem of generating points sets with varying density, according to an importance function, but that exhibit blue noise characteristics locally. McCool and Fiume [MF92] identified this need and how their points, once computed, lead to a trivial solution. They note that the ordering of points resulting from incrementally reducing the radius in a dart throwing method produces a spatially distributed hierarchy of points that can be thresholded. Several other approaches have been explored, all relying on an underlying density function to produce nonuniformly distributed point sets. Those techniques have included subdividing space with polyominoes [Ost07], applying recursive Wang tiles to increase the number of points in a local region proportional to the integral of the density function in that region [KCODL06], and applying capacity constraints to weighted Voronoi regions [BSD09].

Digital halftoning seeks to simulate a continuous grayscale image using patterns of black dots. The visual quality of the result depends on the spectral characteristics of the inherent error, and here, as in Poisson Disk distributions, techniques that achieve blue noise distributions are desirable. Ulichney [Uli87] specified measures of blue noise in the context of digital halftoning as he compared the visual quality of several neighborhood error-filtering techniques [FS76, JJN76, SA85] to the computationally advantageous point-process techniques using masks designed with clustered dots and dispersed dots [Bay73]. Mitsa and Parker [MP92] first combined the desirable blue noise characteristics with the more computationally efficient mask technique to produce a blue noise mask. Ulichney presented the computationally elegant Void and Cluster technique to construct a stochastic dither array by using strictly spatial domain filtering [Uli93]. Much subsequent work has been devoted to blue noise methods in the application domain of digital halftoning [Lin94, SMS97, ABS99, JBRL08]. In both the construction of Poisson Disk distributions and in the quality assessment of halftone techniques, anisotropy plots and radially averaged power spectra [Uli87] are used to characterize spectral properties, but they are only appropriate when applied to regions of constant density and they require points be on a grid. Lagae and Dutre [LD05, LD08] added the concept of a *relative radius* for Poisson Disk distributions to quantify the concept of relaxation of an ideal packing of disks on a plane. Ishizaka [Ish09], summarizing a large body of digital halftone techniques in terms of the functions that are optimized, seeks a spatial domain measure for dot patterns. He points out the limitations of traditional geometric spatial measures for discrete methods.

3. Background

3.1. Halftoning with Masks

Let a single channel of a continuous tone image be represented by $f(i,j)$ and let b represent the number of bits used to quantize the tone in that channel. Then let the threshold array, or halftone mask, $m(i,j)$ contain 2^b distinct tone levels. Then construct a halftoned image $g(i,j)$, with each tone represented by b_g bits, by doing a pointwise threshold comparison between f and the corresponding point in the m for all (i,j). It is not necessary for m and f to be the same size because masks typically are tiled.

$$g(i,j) = \begin{cases} 0, & f(i,j) < m(i,j) \\ 1 & f(i,j) >= m(i,j) \end{cases} \quad (1)$$

This defines the most constrained case in which g is a binary image and $b_g = 1$. This is customary since halftoning grew out of printing. Similarly, with the proliferation of digital imaging, it is customary to study the situation in which

$f(i,j)$ is represented by 256 tones and, thus, $b = 8$. However, the results hold generally whenever $b_g < b$, so that halftone theory is applicable much more broadly, such as for use with high dynamic range image data and for multilevel output devices.

When $f(i,j) = d$ where d is a constant, then $g(i,j)$ is a simulation of a constant graytone, represented by a uniform density of dots, known as a dot profile. Formally, we can designate $g_d(i,j)$ to be the dot profile of density d. Commonly in practice $d \in [0, 2^b - 1]$. When working with digital images, it is often convenient for analysis to speak of the normalized density as a scaled value $d \in [0,1]$ and normalized thresholds. So if d is a constant representing a normalized graytone level and m' is a normalized threshold array, a binary array of a given density is given by

$$g_d(i,j) = \begin{cases} 0, & d < m'(i,j) \\ 1 & d >= m'(i,j) \end{cases} \quad (2)$$

In halftoning, a mask is designed to accommodate the number of graytones present in an image and its quality is assessed by its ability to simulate each possible input by means of a dot distribution. The distribution of these dots is a strong indicator of visual quality and provides the basis for quantitative quality assessments popularized by Ulichney [Uli87].

3.2. Dot Profiles and Poisson Disk point sets

The concept of "dot profile" in halftoning is analogous to a Poisson Disk distribution of constant density. For a given density level, d, our dot profile is simply $g_d(i,j)$ as defined above. To identify the spatial locations of points in a continuous (x,y) coordinate system, we identify the (i,j) locations for which $g_d(i,j) = 1$ and scale (i,j) to the appropriate coordinate system. So the set of points from a $N \times N$ binary array, g_d, mapped to a normalized (x,y) space is

$$P = \{(i/N, j/N) : where\, g_d(i,j) = 1\} \quad (3)$$

The resulting set of normalized points, P, is tilable and may be conveniently incorporated into existing code that uses existing, similarly coded Poisson Disk point sets.

A further significant characteristic of halftone masks is that they exhibit a "stacking" property, so that dots (points) in one density level form subsets of dots (points) at higher density levels. Formally stated, if $g_d(i,j) = 1$ then $g_{d'}(i,j) = 1$ when $d < d'$ and if $g_{d'}(i,j) = 0$ then $g_d(i,j) = 0$. This is a corollary to the hierarchical property found to be desirable in Poisson Disk distributions.

Because the dither array is toroidal, a smaller array can be computed, stored and tiled to extend infinitely. However, it is well known that tiling introduces periodic artifacts in the period of the size of the tile, which may or may not be tolerable, depending on the application. In halftoning, this effect is studied with respect to its visual impact, converting the period to units of cycles per degree of visual angle. Acceptable quality is determined by perceptual impact, which depends on resolution of display device, viewing distance and other application-dependent factors. Here, we will follow the convention of the halftoning literature and assess characteristics of a single mask without tiling.

3.3. Stochastic Dither Array Construction

Since Mitsa and Parker [MP92] first showed how to construct a blue noise mask that was designed to exhibit blue noise properties at every dot profile, single mask construction in digital halftoning has been implemented with a variety of optimization approaches [SMS97, Ish09]. A number of them would have been a reasonable starting point for our work. We chose to implement the Void and Cluster algorithm [Uli93] to construct a stochastic array because its three phase algorithmic structure is simple and powerful, because it allows dot-by-dot processing, because it implements a concept of rank that leads directly to a hierarchical point set and because control over spectral characteristics can be effected by variations to the Gaussian filter parameters.

The algorithm begins by starting with an initial binary pattern in an $M \times N$ array, with a fixed distribution of 1's and 0's at the desired starting normalized density of less than 50%. The initial binary pattern is repeatedly altered by filtering to locate the positions of largest open spaces, or "void" locations and the positions of densest energy, or "clusters" of 1's. The algorithm then swaps the 1's and 0's at those locations and repeats until convergence. Then, from that seed pattern, the algorithm first increases the density of 1's by using the filter to identify the locations of the largest open spaces, or "void" locations, then decreases density by removing 1's by filtering to identify the location of the largest "clusters." Thus the algorithm proceeds by choosing specific candidate locations, one by one, that maximize or minimize the filter, thus eliminating a reliance on random point generation and questions of convergence.

Note that in the process of adding and removing 1's, representing points, a unique integer is retained, corresponding to the order in which the point was chosen. In this way, the resulting stochastic dither array represents an ordering of $M \times N$ integers stored in a two dimensional array, m, with the location of each successive integer in the array corresponding directly to the spatial location.

4. Methods

The Void and Cluster algorithm was implemented using Matlab 7.90529 on a Dell Precision M2400, with Intel Core Duo 2.54 GHz Processor and 8GB RAM, running 64-bit Windows 7. A "distinct" mask is defined by both the size

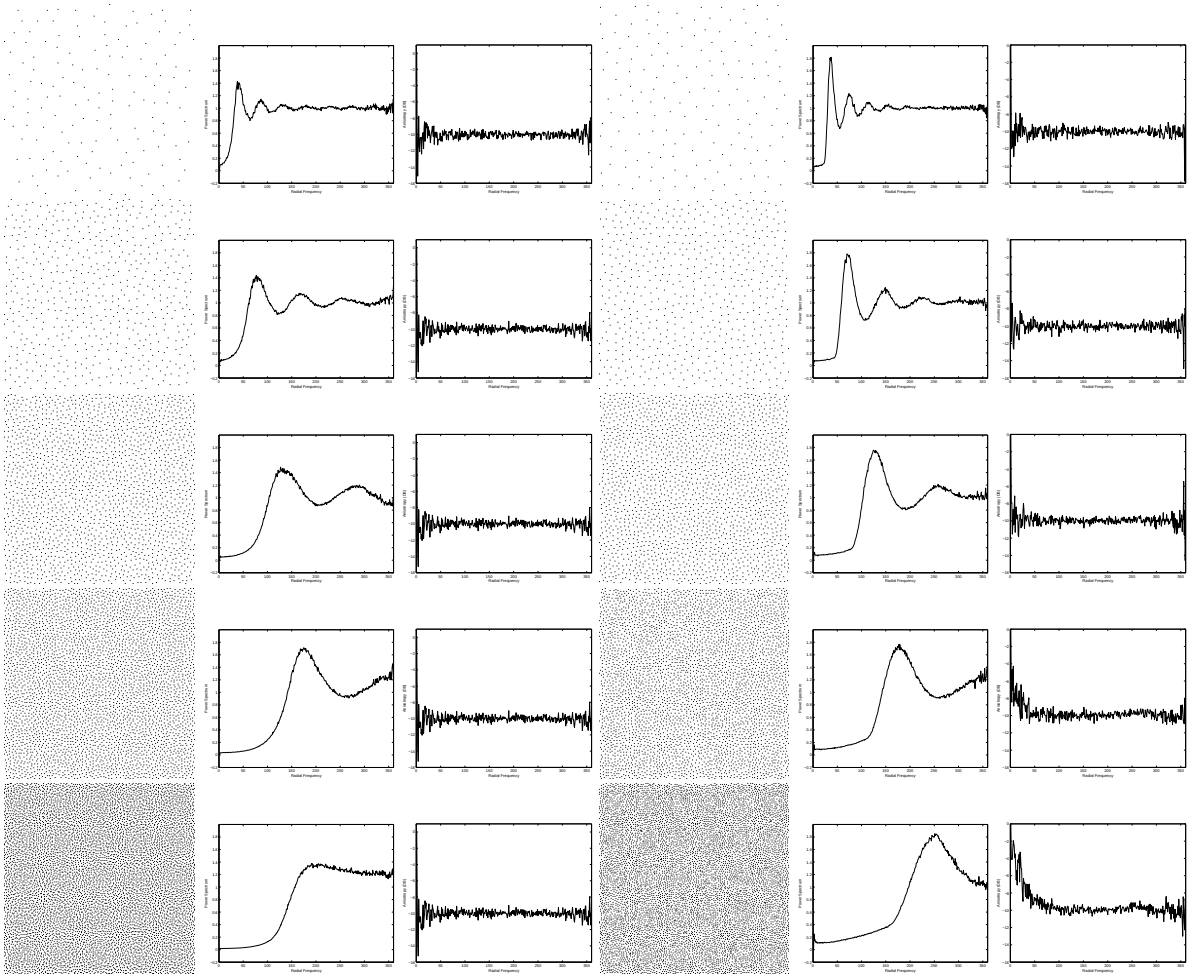

Figure 1: *Comparison of points generated using a stochastic dither array (columns 1-3) and a Best Candidate dart-throwing technique (columns 4-6). Each row corresponds to point sets of cardinalities 1K, 4K, 13K, 26K, and 52K, respectively*

of the mask and the filter used. Ten instances of each distinct mask were constructed to provide adequate frequency domain analysis data, starting with a different initial pattern each time. In all the experiments, we used $N \times N$ images, where $N = 512$. Production of a single stochastic dither array required approximately 6 hours.

Begin with a random, white noise initial binary pattern with 20% density. That is, we place $n = floor(512 * 512 * 0.2) = 52,428$ points randomly among the $512 * 512 = 262,144$ locations. A radially symmetric Gaussian filter was defined in the spatial domain as

$$h(i,j) = e^{-r^2/2\ ^2} \text{where } r = \sqrt{(i-\frac{N}{2})^2 + (j-\frac{N}{2})^2} \quad (4)$$

The value of was varied throughout our experiments. In this implementation, filtering was performed in the frequency domain by applying the Discrete Fourier Transform (DFT) to both the filter and the current binary pattern and taking the product. Ulichney performed filtering by modified convolution in the spatial domain.

Two sets of experiments were conducted. The first explored issues of equivalence are between classic dart-throwing and dot profiles generated from a stochastic dither array. Per Ulichney, was set to 1.5, a value known to produce good overall results in the Void and Cluster algorithm. For comparison, the Best Candidate dart throwing algorithm was used to produce several distinct point sets. Here, a "distinct" point set is defined by the approximate density of points. Five distinct sets of points, P_{bc}, were generated with cardinalities of approximately 1K, 4K, 13K, 26K, and 52K. For each distinct set, ten instances were generated to provide for adequate frequency domain analysis.

For comparison to a dither array, point sets of the same cardinality were used. Recall that the sequential output of the ranks effectively reorders the $N \times N$ locations in the array and the first T reordered locations define a hierarchical point set. So, it is sufficient to find all (i, j) such that $m(i, j) <= T$ and then scale (i, j) to a (x,y) location in a given domain. Several point sets were produced from a single dither array using a single threshold operation for each set. That is, a point set, P_T, from a dither array, m, is constructed for comparison with a Best Candidate point set, P_{bc}, as follows:

$$P_T = \{(i/N, j/N) : \text{where } m(i, j) \leq T \text{and } T = |P_{bc}|\} \quad (5)$$

The second experiment explored the impact of filter design during mask construction on spectral characteristics. For this, six distinct masks were constructed by using $= 0.75, 1.5, 3.0, 6.0, 12.0,$ and 24.0. From each, five point sets, $P_{(\ ,T\ )}$, were computed using Eq. 5 with thresholds of 1K, 4K, 13K, 26K, and 52K.

5. Results: Comparing Void and Cluster Dither Arrays and Best Candidate Dart Throwing

Results from the first experiment appear in Figure 1 arranged with columns 1-3 showing results for Void and Cluster and columns 4-6 showing results for Best Candidate. Rows are arranged to reflect a fixed point set size, with rows 1-6 providing data for point set sizes 1K, 4K, 13K, 26K, and 52K, respectively. Presented for each technique, for each point set size, are a 196×196 magnified portion of a 512×512 dot pattern, the radially averaged power spectrum and an anisotropy plot. Additionally, Table 1 presents the relative radius, , for each.

Spatial Distribution: Referencing columns 1 and 4 of Figure 1, each row shows point sets with varying cardinalities: (row 1) 1.035, (row 2) 4,096, (row 3) 13,108, (row 4) 26,215, (row 5) 52,429, or, as a density percentage on a grid of 512×512: 0.5%, 2%, 5%, 10%, and 20% respectively. Both techniques produce point distributions that are quite similar and represent "blue noise" distributions for a variety of point set sizes. A keen observer may notice more regularity in the points generated with the Best Candidate dart throwing technique and note the increased uniformity of spacing.

Lagae and Dutre [LD08] devised a meaningful dot pattern metric based on the concept of relative radius, denoted by , and noted that the best patterns are those in which $0.65 \leq \ \leq 0.85$. We compute an relative radius for each dot pattern in each distinct set. The average for each distinct set is presented in Table 1. Data in rows 1 and 2 show the relative radius of the two exemplary methods. The Best Candidate method has a larger relative radius, suggesting it approaches an ideal packing pattern more closely than points generated with a dither array. Points generated with Void and Cluster dither arrays still fall near the range for good dot patterns as asserted by Lagae and Dutre for lower densities, then move steadily farther away from an ideal packing as grid constraints limit point placement.

Table 1: *Relative Radius Comparison*

Method	1K	4K	13K	26K	52K
BC	0.72	0.73	0.74	0.74	075
VC	0.64	0.65	0.58	0.42	0.42
binned BC	0.68	0.63	0.58	0.41	0.41

An obvious limitation of using image based methods is the binning error introduced when restricting points to a grid. To gain a feel for the impact of the loss of precision from binning errors, the Best Candidate point sets were mapped to a 512×512 grid, followed by computation of the relative radius on the resulting degraded point set. The values are shown in the row 3 of Table 1.

The impact of the binning error must be considered within the context of the particular application. At some density level, points will begin to exhibit structural artifacts as points necessarily begin to be placed in adjacent locations, with a minimum distance of $1/N$. Note that when normalized density $d > 0.5$, the relationship between the 1's and 0's mirror each other in terms of density and frequency content. In the context of halftoning, dot profiles for which $d > 0.5$ correspond to reproducing darker tones. In the context of generating point sets, choosing all $N \times N$ points is most likely inappropriate. However, the full set of $N \times N$ numbers, arranged in the mask, is necessary for implementing importance sampling.

Frequency domain analysis: Frequency domain analysis helps to quantify the visual characteristics and provide insight into subtle differences. The widely used metrics of the radially averaged power spectrum and anisotropy plots defined by Ulichney in the context of digital halftoning are also used in computer graphics [Uli87]. The results of computing the metrics on each of the sets are shown in Figures 1 columns 2-3, for the Void and Cluster technique, and 5-6, for the Best Candidate technique.

The graphs of frequency content exhibit many similarities, as do the dot profiles themselves; but an interesting difference is seen in the amplitude of the RAPS plots. For both the 4K points and 1K point sets, shown in rows 1 and 2 of Figure 1, the oscillating effect of the peaks at and beyond the transition region is more exaggerated. This recalls the $sinc(x) = sin(\ x\ )/\ x\ $ function, which corresponds to an ideal low pass filter in the spatial domain. We believe that this effect is achieved by the strict radial distance constraint used in Poisson Disk generation techniques.

The radially averaged power spectrum is shown for each of the dot profiles in Figure 1, which are each produced by thresholding the same stochastic dither array. For example,

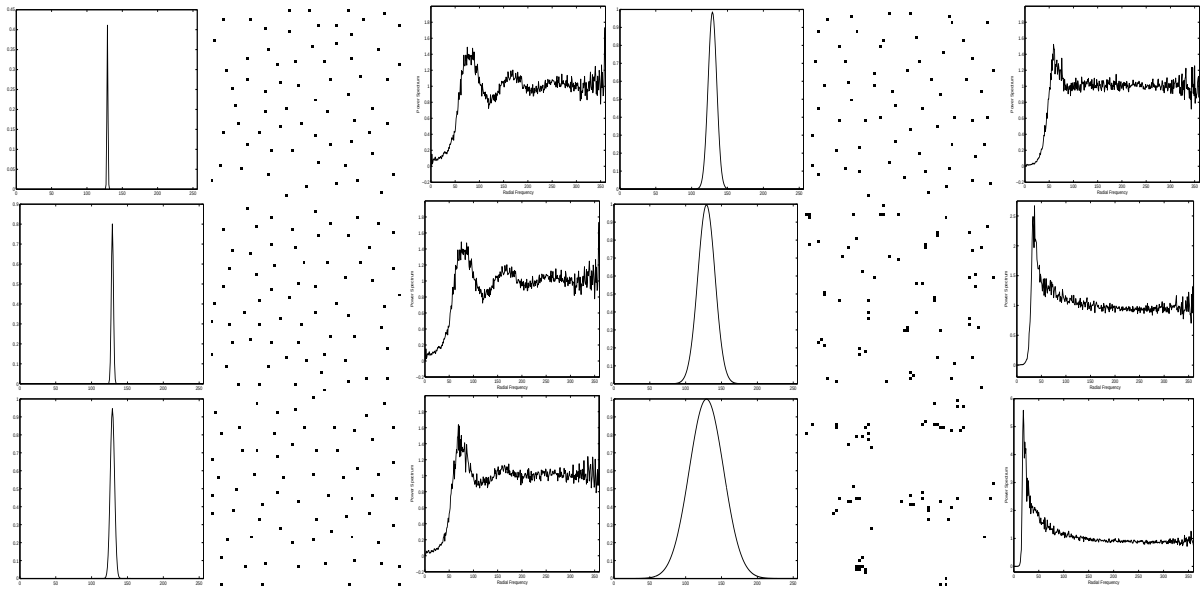

Figure 2: *Results from varying sigma uniformly in a radially symmetric Gaussian filter. Shown for each are a spatial domain representation of the filter, a dot pattern of 4096 points, magnified to* 196×196 *and the corresponding RAPS.*

structures that are apparent in the dot profiles for Figures 1 in rows 5-6, column 1 might be objectionable for some applications, but the frequency characteristics for these patterns are still blue noise. Although the RAPS characteristics of rows 5-6, column 4, are similar here, it is important to understand that the frequency domain measures of Best Candidate points require binning before computation. Conclusions drawn about the spectral characteristics of binned Best Candidate points are limited by the binning resolution. Evaluation of Best Candidate points in the frequency domain depend on resolution of grid used in the metric. It would be inappropriate to conclude from these graphs that the Best Candidate characteristics vary with density in the same way that the Void and Cluster patterns do. This graph suggests that, within inherent constraints present in a discrete system, the Void and Cluster technique is comparable enough to existing Poisson Disk algorithms to merit consideration.

5.1. Results: Varying Spectral Shape with Filter Parameters

Results of the second experiment appear in Figure 2. For each distinct mask created with a variation in filter parameter, we show a radial cross section of each filter in the spatial domain, a $196 \times$ 196 magnified dot pattern for produced using a thresholded of T = 4096, and the radially averaged power spectrum. These items are shown in columns 1-3 for $= 0.75, 1.5,$ and 3.0, in rows 1-3, respectively. Similarly, these items are shown in columns 4-6 for $= 6.0, 12.0,$ and 24.0.

Reading down columns 1 and 4, note that the filter is a low-pass filter and that, as increases, the filter widens in the spatial domain. By properties of the Fourier transform, a decrease of filter width in the spatial domain results in an increase of filter with in the frequency domain. Columns 2 and 5 show the resulting patterns in which dots begin to cluster near one another. Columns 4 and 6 show interesting trends regarding the RAPS. The oscillating characteristics of the high frequency region is first dampened for $= 0.75, 1.5, 3.0, 6.0$; but energy remains uniformly distributed throughout the high frequency region. However, when $= 12.0,$ and 24.0., the high frequency is greatly diminished, which combined with the increased peak at lower frequencies, recalls a bandpass filter effect. The clumps in the dot patterns are very noticeable when $= 12.0,$ and 24.0 (column 2, rows 2 and 3) and create interesting patterns. This clustering behavior may be of use in some applications, e.g. in object location or biological system simulation.

6. Adaptive Sampling with a Stochastic Dither Array

As hierarchical set of Poisson Disk-distributed points can be used to adaptively sample a space with a known probability density function [MF92]. We define an importance image as discrete samples of a two-dimensional density function. In Figure 3 (top) the importance function is defined as a gray ramp image, with white representing the maximum value. Below that are shown three corresponding sets of sample points of varying density. Sample locations are represented as black dots. In each case, the sample points were achieved

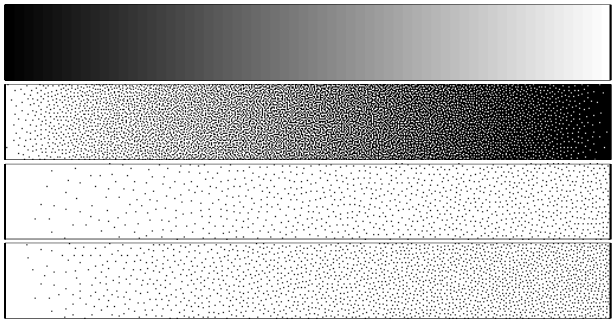

Figure 3: *Grayscale ramp as importance image (top) and three sample point sets generated by thresholding a stochastic data array. Reductions of sampling density by 99% and 90% achieved by applying a constant scaling factor.*

with a pointwise threshold operation. Variation of density was achieved by a constant operation at each point. The first set shows the sample points generated by thresholding with the full dynamic range of the stochastic array. As the importance image approaches full value, shown on the right as the white region, the number of sample points approaches full saturation, so that every point is selected to be sampled and the right region appears solid black. Because choosing every sample point may be impractical in computer graphics application, the number of sample points can be easily controlled by the importance image. The bottom two pictures of Figures 3 show how a hierarchical reduction in sampling density can be achieved by applying the dither array. These were produced by scaling each point of the input image by dividing by 32 and 64, respectively.

Without loss of generality, we start with the same two-dimensional density function used by Balzer [BSD09] defined as $= e^{-20x^2-20y^2} + 0.2\sin^2(x)\sin^2(y)$ over the range [-1,1]. We use this to define discrete values over a 512×512 array, f, which we show as an intensity image in Figure 4(left). This image is thresholded with a normalized stochastic array, m', defining a sample point set P in (x,y) space as

$$P = \{(x,y) : \text{where } x = (i-256)/512, \; y = (j-256)/512 \text{ and } f(i,j) > m'(i,j)\} \quad (6)$$

The sample points that result are shown in Figure 4(right). Figure 5(top) shows a portion of a magnified cross section of the continuous density function along a diagonal passing through the center. Figure 5(bottom) show the points along the corresponding magnified portion of the axis which represent sample locations. The density of sample points adapts to the amplitude of the density function as seen in the closely spaced points along the left correspond to the high curve, and the sparsely spaced points in the middle correspond to the dip in the curve. More sample points are selected at the (x,y) locations that correspond to areas of greater intensity.

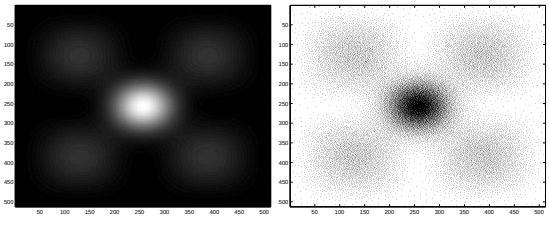

Figure 4: *(left) 2D density function* $= e^{-20x^2-20y^2} + 0.2\sin^2(x)\sin^2(y)$ *as a* 512×512 *importance image and (right) the sampled image after thresholding with a stochastic data array. White regions in (left) represent maximum importance. Black dots represent sample points in (right).*

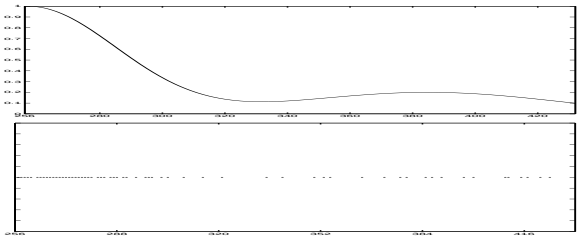

Figure 5: *(top) Diagonal cross section of 2D density function* $= e^{-20x^2-20y^2} + 0.2\sin^2(x)\sin^2(y)$ *and (bottom) corresponding samples selected after thresholding.*

7. Summary

We have explored an image-based approach to the classical computer graphics problem of generating Poisson Disk sample point distributions. Generation of the stochastic dither array is computationally intensive but, once computed, allows for uniform and adaptive sampling in a straightforward and very efficient manner. Comparisons between exemplary techniques in each field show quantitatively and qualitatively similar results. However, the image-based approach of digital halftoning techniques imposes restrictions on point placement that are not present in continuous space techniques for computing Poisson Disk point sets.

Dither arrays are precomputed data sets and, thus, significantly faster to implement than "on the fly" techniques. They are inherently parallelizable, scalable, and their toroidal nature makes them well suited for adapting to a variety of architectures, including embedded systems and other constrained platforms. The construction of the data set is driven by intentional filter design. This allows for control of spectral characteristics. Lastly, the precomputed data set is hierarchical and allows for adaptive sampling of importance images and density functions in a very direct manner.

Acknowledgements

We thank Dunbar and Humphreys for providing source code on their website that was helpful in implementing a Best Candidate algorithm. We thank Thomas Schlömer for providing point set analysis source code in the public domain.

References

[ABS99] ANCIN H., BHATTACHARJYA A. K., SHU J.: New void-and-cluster method for improved halftone uniformity. *Journal of Electronic Imaging 8*, 104 (1999). 2

[Bay73] BAYER B. E.: An optimum method for two-level rendition of continuous-tone pictures. vol. 1, IEEE International Conference on Communications, IEEE, pp. 2611–2615. 2

[BSD09] BALZER M., SCHLÖMER T., DEUSSEN O.: Capacity-constrained point distributions: A variant of lloyd's methods. *ACM Transactions on Graphics 28*, 3 (August 2009). 2, 7

[Coo86] COOK R. L.: Stochastic sampling in computer graphics. *ACM Transactions on Graphics 5*, 1 (January 1986), 51–72. 2

[CSHD03] COHEN M. F., SHADE J., HILLER S., DEUSSEN O.: Wang tiles for image texture generation. *ACM Transactions on Graphics* (2003). 2

[DH06] DUNBAR D., HUMPHREYS G.: A spatial data structure for fast Poisson-disk sample generation. *ACM Transactions on Graphics 25* (2006), 503–508. 2

[DW85] DÍPPE M. A., WOLD E. H.: Antialiasing through stochastic sampling. vol. 19, ACM Siggraph, ACM, pp. 69–78. 2

[FS76] FLOYD R. W., STEINBERG L.: An adaptive algorithm for spatial grey scale. vol. 17, Society for Information Display, pp. 75–77. 2

[GM09] GAMITO M. N., MADDOCK S.: Accurate multidimensional poisson-disk sampling. *ACM Transactions on Graphics 29*, 1 (December 2009), 19. 2

[HDK01] HILLER S., DEUSSEN O., KELLER A.: Tiled blue noise samples. In *Proceedings of Vision Modeling and Visualization* (2001), pp. 265–272. 2

[Ish09] ISHIZAKA K.: New spatial measure for dispersed-dot halftoning assuring good point distribution in any density. *IEEE Transactions on Image Processing 18*, 9 (September 2009), 2030–2047. 2, 3

[JBRL08] J. BACCA RODIGUEZ G. R. A., LAU D.: Blue-noise multitone dithering. *IEEE Transactions on Image Processing 17*, 8 (August 2008), 1368–1382. 2

[JJN76] JARVIS J., JUDICE C., NINKE W.: A survey of techniques for the display of continuous-tone images on a bilevel display. *Computer Graphics and Image Processing 5* (1976), 13–40. 2

[Jon06] JONES T. R.: Efficient generation of Poisson-disk sampling patterns. *journal of graphics, gpu, and game tools 11*, 2 (2006), 27–36. 2

[KCODL06] KOPF J., COHEN-OR D., DEUSSEN O., LISCHINSKI D.: Recursive Wang tiles for real-time blue noise. *ACM Transactions on Graphics* (2006). 2

[LD05] LAGAE A., DUTRÉ P.: A procedural object distribution function. *ACM Transactions on Graphics 24*, 4 (October 2005), 1442–161. 2

[LD06] LAGAE A., DUTRÉ P.: An alternative for wang tiles. *ACM Transactions on Graphics 25*, 4 (2006), 1442–1459. 2

[LD08] LAGAE A., DUTRÉ P.: A comparison of methods for generating poisson disk distributions. *Computer Graphics Forum 27*, 1 (2008), 114–129. 2, 5

[Lin94] LIN Q.: Improving halftone uniformity and tonal response. IST Tenth Congress on Advances in Non-Impact Printing TechnologiesT Tenth Congress on Advances in Non-Impact Printing Technologies, pp. 377–380. 2

[MF92] MCCOOL M., FIUME E.: Heirarchical Poisson disk sampling distributions. *Graphics Interface '92* (1992), 94–105. 2, 6

[Mit91] MITCHELL D. P.: Spectrally optimal sampling for distribution ray tracing. *Computer Graphics Forum 25*, 4 (July 1991), 157–164. 2

[MP92] MITSA T., PARKER K.: Digital halftoning technique using a blue-noise mask. *Journal of the Optical Society of America A 9*, 11 (1992), 1920–1929. 2, 3

[Ost07] OSTROMOUKHOV V.: Sampling with polyominoes. *ACM Transactions on Graphics 26*, 3 (July 2007). 2

[SA85] STEVENSON R. L., ARCE G. R.: Binary display of hexagonally sampled continuous-tone images. *Journal of the Optical Society of America A 2*, 7 (1985), 1009–1013. 2

[SMS97] SPAULDING K. E., MILLER R. L., SCHILDKRAUT J.: Methods for generating blue-noise dither matrices for digital halftoning. *Journal of Electronic Imaging 6*, 2 (1997), 208–230. 2, 3

[Uli87] ULICHNEY R.: *Digital Halftoning*. The MIT Press, 1987. 1, 2, 3, 5

[Uli93] ULICHNEY R.: The void-and-cluster method for dither array generation. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* (September 1993), Allebach J., Rogowitz B. E., (Eds.), vol. 1913 of *Human Vision, Visual Processing, and Digital Display IV*, Proceedings of the SPIE, SPIE, pp. 332–343. 2, 3

[Wei08] WEI L.-Y.: Parallel poisson disk sampling. *ACM Transactions on Graphics 27*, 3 (August 2008), 1–9. 2