Blue-Noise Halftoning for Hexagonal Grids

Daniel L. Lau[1], Robert Ulichney
Imaging Systems Laboratory
HP Laboratories Cambridge
HPL-2004-232
December 22, 2004*

E-mail: dllau@engr.uky.edu, u@hp.com

halftoning, printing, blue-noise, green-noise

In this paper, we closely scrutinize the spatial and spectral properties of aperiodic halftoning schemes on hexagonal sampling grids. Traditionally, hexagonal sampling grids have been shunned due to their inability to preserve the high frequency components of blue-noise dither patterns at gray-levels near one-half, but as will be shown, only through the introduction of diagonal correlations between dots can even rectangular sampling grids preserve these frequencies. And by allowing the sampling grid to constrain the placement of dots, a particular algorithm may introduce visual artifacts just as disturbing as excess energy below the principal frequency. If, instead, the algorithm maintains radial symmetry by introducing a minimum degree of clustering, then that algorithm can maintain its grid defiance illusion fundamental to the spirit of the blue-noise model. As such, this paper shows that hexagonal grids are preferrable because they can support gray-levels near one-half with less required clustering of minority pixels and a higher principal frequency. Furthermore, hexagonal grids allow for improved radial symmetry in the size and shape of clusters. Finally, the world's first blue-noise dither array for hexagonal grids will be demonstrated.

* Internal Accession Date Only
[1]Department of Electrical and Computer Engineering, University of Kentucky, Lexington, KY 40506-0046, USA
Approved for External Publication

Blue-Noise Halftoning for Hexagonal Grids

Daniel L. Lau
Department of Electrical and Computer Engineering
University of Kentucky
Lexington KY 40506-0046 USA
email: *dllau@engr.uky.edu*

Robert Ulichney
HP Labs
Hewlett Packard Co.
One Cambridge Center
Cambridge, MA 02142-1612 USA
email: *u@hp.com*

Abstract— **In this paper, we closely scrutinize the spatial and spectral properties of aperiodic halftoning schemes on hexagonal sampling grids. Traditionally, hexagonal sampling grids have been shunned due to their inability to preserve the high frequency components of blue-noise dither patterns at gray-levels near one-half, but as will be shown, only through the introduction of diagonal correlations between dots can even rectangular sampling grids preserve these frequencies. And by allowing the sampling grid to constrain the placement of dots, a particular algorithm may introduce visual artifacts just as disturbing as excess energy below the principal frequency. If, instead, the algorithm maintains radial symmetry by introducing a minimum degree of clustering, then that algorithm can maintain its grid defiance illusion fundamental to the spirit of the blue-noise model. As such, this paper shows that hexagonal grids are preferrable because they can support gray-levels near one-half with less required clustering of minority pixels and a higher principal frequency. Furthermore, hexagonal grids allow for improved radial symmetry in the size and shape of clusters. Finally, the world's first blue-noise dither array for hexagonal grids will be demonstrated.**

Keywords—**halftoning, printing, blue-noise, green-noise**

I. Introduction

Aperiodic, dispersed-dot halftoning is a technique for producing the illusion of continuous tone in binary display devices through a random arrangement of isolated dots. These dots are all of the same size, usually a single pixel, with their spacing defined according to tone such that dark shades of gray are produced by closely spaced dots and light shades by dots placed far apart. Relative to the human visual system, the optimal halftone patterns are composed exclusively of high frequency spectral components [1] and are commonly referred to as "blue-noise", the high frequency component of white-noise. The occurrence of low (red) frequency spectral components gives binary dither patterns a noisy appearance [2].

Although the preferred technique in inkjet printers, blue-noise is not considered a viable technique in laser printers due to the failure of the electrophotographic printing process to produce isolated minority pixels consistently. Until recently, these printers have been limited to ordered, clustered-dot halftoning algorithms, which produce a regular grid of round dots that vary in size according to tone. Within the literature, this halftoning is generally considered an undesirable halftoning technique due to its limitations in preserving spatial details and minimizing halftone visibility (the appearance of artificial textures) [3].

A new halftoning approach, green-noise halftoning produces random patterns of homogeneously distributed minority pixel clusters. Studied by Lau *et al* [4], green-noise halftone patterns are composed almost exclusively of mid-frequency spectral components. The advantages to using green-noise are well documented with respect to printer reliability [5] where clustering reduces the perimeter-to-area ratio of printed dots. In this paper, it is theorized that green-noise could unlock a host of advantages commonly associated with hexagonal sampling grids where hexagonal (a.k.a. quincuncial) grids differ from rectangular in that every other row is offset one-half pixel period. In particular, hexagonal sampling grids are well recognized for allowing a more natural radially symmetric sampling of 2-D space – preserving a circular band-limited signal with only 86% of the total number of samples used by rectangular grids.

Additional advantages to hexagonal sampling grids over rectangular are their robustness to changes in aspect ratio as illustrated in Fig. 1, which shows an asymmetric 4×4 grid with a 3.5 aspect ratio arranged rectangularly and hexagonally. The aspect ratio is the horizontal period divided by the vertical period, and here, the pixel shapes shown are defined by the perpendicular bi-sectors between neighboring pixels. From Fig. 1, Ulichney defines the covering efficiency [2] as the ratio of pixel area

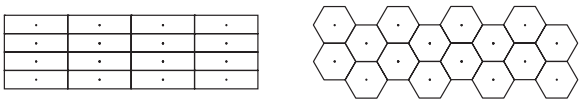

Fig. 1. Pixel shapes for a (left) 4×4 rectangular and a (right) 4×4 hexagonal grid, both with a 3.5 aspect ratio.

divided by the circumscribing circle area, which can be quantified and plotted for both types of grids as shown in Fig. 2 where the pixel shapes are indicated for selective aspect ratios. The best case for rectangular grids occurs for square grids, but what is particularly interesting, in this figure, is the wide range of aspect ratios where hexagonal grids outperform this best case: over an order of magnitude. This is important because it allows for resolution to be increased asymmetrically yet still enjoy superior radial symmetry of pixel coverage. It is very often easier to increase resolution in only one dimension, and using hexagonal grids would allow us to take advantage of that.

Given the super-high dot addressability of modern digital printers, the implementation of hexagonal grid halftoning is certainly reasonable for aperiodic screening techniques, and given the general advantages to using hexagonal sampling grids, one may wonder why hexagonal grids have not received more attention by the research community with respect to stochastic halftoning. As such, we hypothesize that much of the detraction of hexagonal grids derives from the analysis of blue-noise dithering performed by Ulichney who showed that only on a rectangular sampling grid is it possible to isolate minority pixels at all gray-levels. In contrast, minority pixels must begin to cluster as the gray-level approaches 1/2 on a hexagonal grid, which lead Ulichney to write that hexagonal sampling grids do not support blue-noise. Specifically being forced to cluster pixels, blue-noise isolates dots at some locations only to cluster at others, creating a wider range of frequencies in the spectral content of the dither pattern. This widening of the spectral content is referred to as "whitening," and as a pattern becomes more and more white, it appears more and more noisy.

In this paper, we study the application of blue and green-noise to images sampled along hexagonal grids, showing that at a critical coarseness, hexagonal sampling grids are the preferred sampling technique for stochastic dithering. In particular, this paper introduces a new model for blue-noise that emphasizes radial symmetry even in cases where radial symmetry requires the clustering of minority pixels. As will be shown, the traditional blue-noise model, by isolating minority pixels, is constrained by the sampling geometry near gray-level $g = \frac{1}{2}$ where rectangular grids force patterns into a periodic checkerboard pattern that may, in some cases, create visually disturbing artifacts. By allowing a minimum degree of clustering, it will be argued that the new blue-noise model makes an optimal trade-off to acquire a sufficient degree of flexibility in where the algorithm chooses to place dots. In order to demonstrate these concepts, this paper will rely on error-diffusion halftoning as well as an iterative technique commonly associated with blue-noise dither arrays. In concluding this paper, we will demonstrate the world's first blue-noise dither array for hexagonal sampling grids to illustrate the achievable image fidelity for halftones on these grids.

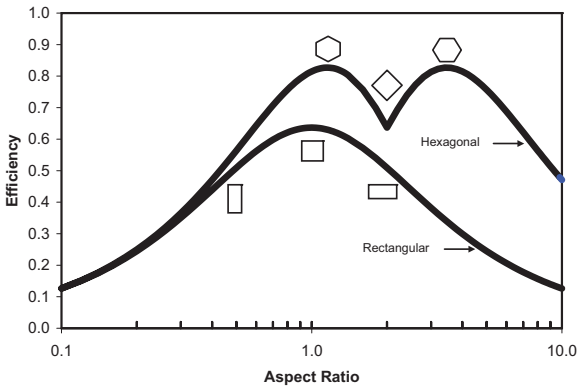

Fig. 2. The covering efficiency of printed dots for rectangular and hexagonal sampling grids versus the aspect ratio.

II. Rectangular Sampling Grids

As Ulichney [2] has shown, the optimal aperiodic, dispersed-dot halftoning schemes are the ones that distribute the minority pixels of a binary dither pattern as homogeneously as possible, trying to spread the minority pixels as far apart as they can in an isotropic manner. The resulting patterns are then composed of isolated dots separated by an average distance of λ_b such that:

$$\lambda_b = \begin{cases} 1/\sqrt{g} & , \quad \text{for } 0 < g \leq 1/2 \\ 1/\sqrt{1-g} & , \quad \text{for } 1/2 < g \leq 1 \end{cases} \tag{1}$$

where units are in terms of the pixel period of the display and g is the average gray-level of the dither pattern. Being a stochastic arrangement of dots, the actual distances between nearest neighbors are not exactly equal to λ_b but have some variation with variation being too large causing the pattern to look noisy while being too small resulting in patterns appearing periodic.

To reduce the two dimensions of a halftone pattern power spectrum to one dimension, the metric of radially averaged power spectrum (RAPS) $P_\rho(f_\rho)$ is used where the energy within thin concentric annuli is averaged for this purpose [2]. The original theory of blue noise argued that the energy of a well formed dither pattern would be of the form of Fig. 3 (left) where spectral energy is concentrated at the principal frequency, f_b, defined as a function

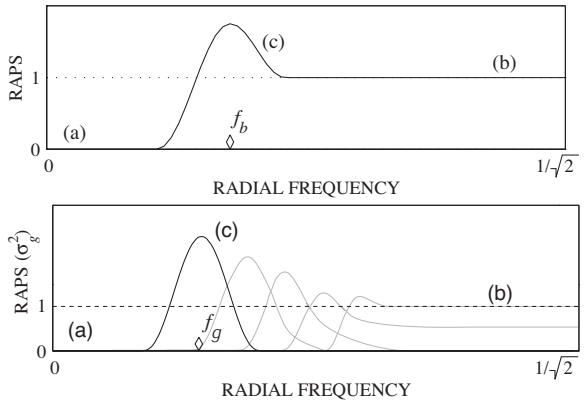

Fig. 3. The RAPS measure for (top) an ideal blue-noise dither pattern and (bottom) an ideal green-noise dither pattern.

of gray level, g, as:

$$f_b = \begin{cases} \sqrt{g} & , \quad \text{for } 0 < g \leq 1/2 \\ \sqrt{1-g} & , \quad \text{for } 1/2 < g \leq 1 \end{cases} . \qquad (2)$$

It is due to this shape that these patterns are referred to as "blue-noise," where "blue" refers to the high frequency or blue component to white light and "noise" refers to the randomness of the pattern. Looking at eqn. (2), one may note our use of the subscript b as opposed to Ulichney's use of the subscript g as an indication of the frequency's dependence on the gray-level of the dither pattern. The subscript b was later introduced by Lau *et al* [4] to differentiate f_b from the equivalent parameter for green-noise, which they indicated as f_g. We will use the later notation but caution readers to note that f_b is unique to each gray-level as just prescribed.

The parameter f_b is referred to as the principle wavelength of blue-noise and is the source of controversy with respect to hexagonal sampling. The issue at hand is aliasing and the unwanted visual artifacts that aliasing creates. An explanation begins with Fig. 4 (left) where a small area near DC of the infinite spectral plane of a rectangular sampled image is shown. Here, the spectral plane is in units of inverse pixel period and is divided into spectral annuli of radial width Δ. Taking the average power within each annulus and then plotting the average power versus the center radius f_ρ creates Ulichney's RAPS measure. Note that the maximum spectral radius within each square tile is $\frac{1}{\sqrt{2}}$. Shown in Fig. 5 is a diagram of the spectral domain for four blue-noise dither patterns with the black segments marking the principle wavelength.

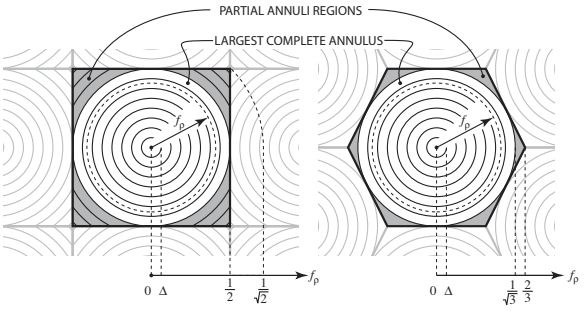

Fig. 4. The spectral planes of (left) a rectangular sampled image shown divided into annular rings of radial width Δ and center radius f_ρ, and (right) a hexagonally sampled image shown divided into annular rings with center radius f_ρ.

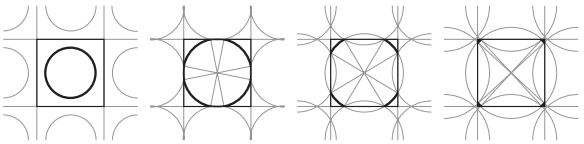

Fig. 5. The spectral rings of blue-noise dither patterns with added diagonal correlation between minority pixels for gray-levels (left) $g = 10\%$, (left-center) $g = 26\%$, (right-center), $g = 42\%$, and (right) $g = 50\%$.

As originally proposed, Ulichney envisioned the principal frequency as a circular wavefront eminating from the spectral DC origin and progressing outward as g approached $\frac{1}{2}$. At gray-level $g = \frac{1}{4}$ when the wavefront first makes contact with sides of the baseband entering the partial annuli region of Fig. 4, the wave becomes segmented into the four corners while still progressing to $f_b = 1/\sqrt{2}$, the maximum radial frequency within the baseband of a rectangular sampling grid. In the spatial domain, this packing of energy into the corners of the baseband, as depicted in Fig. 5, is achieved by adding correlation between minority pixels along the diagonal, creating a pattern where neighboring minority pixels are more likely to occur along the diagonal instead of side-by-side or above-and-below one another. If a particular halftoning scheme is especially successful at adding this diagonal correlation, then it is possible to create dither patterns at all gray-levels such that no two minority pixels occur adjacent to one another. Such a scheme would produce the familiar $g = \frac{1}{2}$ checkerboard pattern.

Floyd's and Steinberg's [6] error-diffusion is a classic example of a blue-noise generating halftoning algorithm that adds such correlation. In error-diffusion, the output pixel $y[n]$ is determined by adjusting and thresholding the input pixel $x[n]$ such that

$$y[n] \quad = \quad \begin{cases} 1 & , \quad \text{if } (x[n] + x_e[n]) \geq 0 \\ 0 & , \quad \text{else} \end{cases} \qquad (3)$$

where $x_e[n]$ is the diffused quantization error accumulated during previous iterations as

$$x_e[n] \quad = \quad \sum_{i=1}^{M} b_i \cdot y_e[n-i] \qquad (4)$$

with $y_e[n] = y[n] - (x[n] + x_e[n])$. The diffusion coefficients b_i, which regulate the proportions to which the

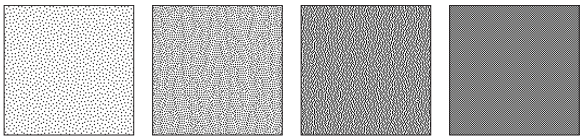

Fig. 6. The blue-noise dither patterns created by error-diffusion using the Floyd-Steinberg error filter for gray-levels (left) $g = 10\%$, (left-center) $g = 26\%$, (right-center), $g = 42\%$, and (right) $g = 50\%$.

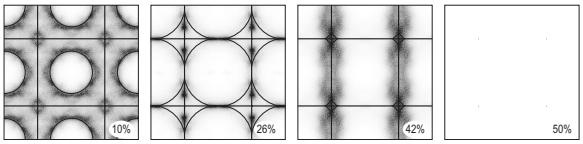

Fig. 7. The power spectra for dither patterns of Fig. 6 for gray-levels (left) $g = 10\%$, (left-center) $g = 26\%$, (right-center), $g = 42\%$, and (right) $g = 50\%$. The superimposed black circles mark the location of the principle frequency (lines have been ommitted for $g = 50\%$).

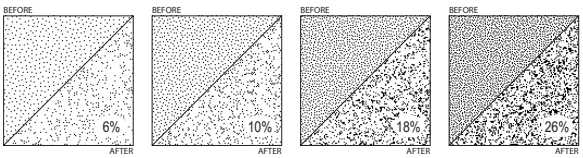

Fig. 9. The spatial blue-noise dither patterns, before and after down-sampling by 2, representing gray-levels (left) $g = 6\%$, (center) $g = 26\%$, and (right) $g = 50\%$.

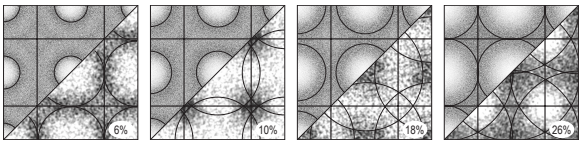

Fig. 10. The power spectra corresponding to the blue-noise dither patterns of Fig. 9 before (top-left) and after (bottom-right) down-sampling with the principle frequency marked in black. The images corresponding to after down-sampling were enhanced to better illustrate the distribution of energy.

quantization error at pixel n transfers or diffuses into neighboring pixels, are such that $\sum_{i=1}^{M} b_i = 1$. Floyd and Steinberg specifically chose their 4-filter weights because of their behavior near gray-level $g = \frac{1}{2}$. Shown in Figs. 6 and 7 are the spatial dither patterns and their corresponding power spectra as g progresses from 0 to $\frac{1}{2}$.

Now while Ulichney originally believed packing energy into the corners of the power spectrum to be the ideal behavior for blue-noise, we ultimately see that adding diagonal correlation, especially to the degree of Floyd and Steinberg's error-diffusion, violates the two basic characteristics of blue-noise: radial symmetry combined with aperiodicity. Maintaining an average distance between minority pixels of λ_b, near $g = \frac{1}{2}$, forces the minority pixels to lock into a fixed and periodic pattern whose only saving grace, in terms of visual pleasantness, is that its high spatial frequency makes it less visible than similar patterns at lesser gray-levels. As will be discussed in Sec. III, maintaining a cut-off frequency of f_b, by adding directional correlation, also creates a significant dilemma for hexagonal grids where the maximum spatial frequency occurs at $g = \frac{1}{3}$.

Looking back at Ulichney's original definition of blue-noise dither patterns as being radially symmetric while also having some variation in the distance between minority pixels, we begin to wonder if there are alternative behaviors for blue-noise, in this gray-level range $\frac{1}{4} < g \leq \frac{3}{4}$ where the baseband constrains the placement of dots, such that dither patterns can maintain their grid-defiance illusion and not adopt a periodic or textured appearance. As a first attempt at such a blue-noise model, we can try to enforce both the principal frequency f_b, defined according to eqn. (2), as well as radial symmetry – creating the spectral behavior depicted in Fig. 8 where, as g exceeds $\frac{1}{4}$, the principal frequency ring extends beyond the sides of the baseband. Spectral energy from neighboring rings will then extend into the baseband and, hence, introduce alias artifacts into the dither pattern. To see the effects of this aliasing, we can down-sample blue-noise dither patterns by a factor of two to double the radius of the spectral ring.

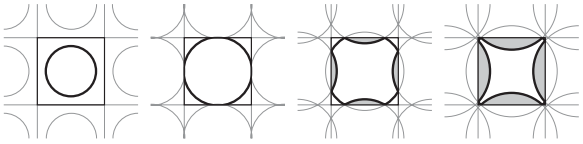

Fig. 8. The spectral rings of blue-noise dither patterns for gray-levels (left) $g < 25\%$, (left-center) $g = 25\%$, (right-center), $g = 40\%$, and (right) $g = 50\%$.

The intensity or gray-level of the spatial dither patterns should not be affected by the down-sampling operation, and what we see in Figs. 9-10 is that at gray-levels beyond 6.25% ink coverage, where the corresponding principal frequencies overlap neighboring rings (after down-sampling), the resulting dither patterns will exhibit light ($g = 10\%$) to moderate ($g = 26\%$) and then severe ($g = 50\%$, not shown) clustering of minority pixels, causing the pattern to take on an unpleasant appearance. For $g < 6.25\%$, aliasing will occur due to the high frequency spectral content ($f_\rho > f_b$) characteristic of blue-noise, resulting in a noisy appearance of its own, but this particular aliasing leads only to the high variation in the distance between minority pixels that is not so high as to cause minority pixels to touch. Clustering of minority pixels only seems to occur when the spectral rings intersect.

So under the premise that aliasing of the principal frequency leads to unwanted clustering of minority pixels in the spatial domain, we can now look at specific error-diffusion techniques known to exhibit clustering at gray-

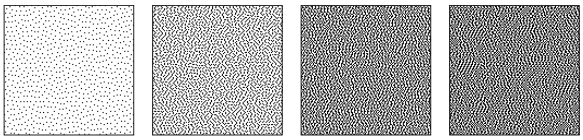

Fig. 11. The spatial dither patterns created by error-diffusion using the Jarvis *et al* error filter for 10%, 42%, and 50% ink coverage.

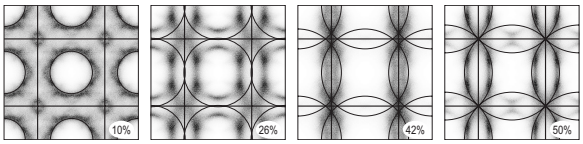

Fig. 12. The power spectra for dither patterns of Fig. 11 for gray-levels (left) $g = 10\%$, (left-center) $g = 26\%$, (right-center), $g = 42\%$, and (right) $g = 50\%$. The superimposed black circles mark the location of the principle frequency.

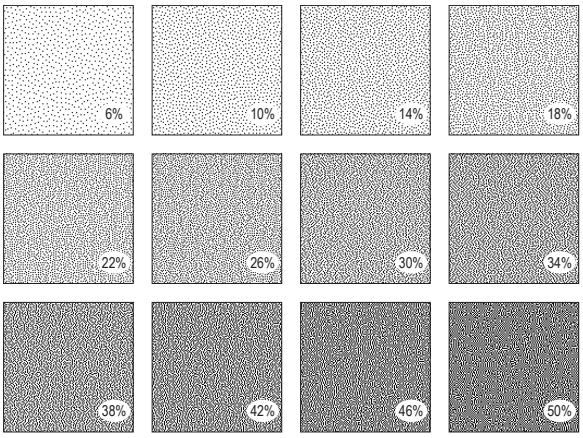

Fig. 13. The binary dither patterns for Ulichney's perturbed filter weight scheme on a rectangular sampling grid as g transitions from 6% to 50% coverage.

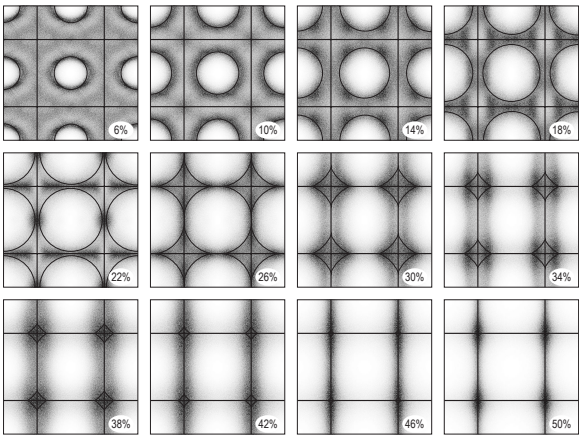

Fig. 14. The power spectra for Ulichney's perturbed filter weight scheme on a rectangular sampling grid as g transitions from 6% to 50% coverage.

levels between $g = \frac{1}{4}$ and $\frac{3}{4}$ to see if the clustering that these algorithms introduce are, in fact, the product of aliasing. We saw a little bit of this behavior in Fig. 7 at gray-level $g = 42\%$ using Floyd and Steinberg's error filter where there is a clear correlation between the distribution of energy in the power spectrum and the spectral ring at f_b. Looking at the spatial dither patterns and power spectra produced by error-diffusion using Jarvis, Judice, and Ninke's [7] 12-weight filter in Figs. 11 and 12, we clearly see this clustering/aliasing behavior as evidenced by the strong spectral components for $g = \frac{1}{2}$ shown in perfect alignment with the spetral rings at radial frequency f_b from neighboring replications of the baseband frequency.

A. Perturbed Filter Weights

If we now look at the clustering found in Ulichney's perturbed filter weight scheme [2] in Figs. 13 and 14, where the spatial dither patterns and corresponding spectra for g transitioning from 6% to 50% coverage is shown, we see some differences with both Floyd and Steinberg's error filter and with Jarvis *et al*'s. From visual inspection, one can see that, by perturbing filter weights, the resulting dither patterns better maintain radial symmetry by moving some of the spectral energy inside the principal frequency ring through a small, controlled degree of clustering. That is, by allowing a small degree of clustering, Ulichney's perturbed filter scheme is able to reduce the principal frequency of the pattern, breaking up some of the periodic textures that would otherwise form due to the added diagonal correlation. But given that the observed clustering is only slight, we would describe the the perturbed filter weight scheme as generally behaving in the manner first prescribed by Ulichney in that it adds diagonal correlation and packs spectral energy into the corners of the baseband as g approaches $\frac{1}{2}$. If there is a disturbing artifact to be found in the patterns of Fig. 13, it is the discontuities in texture created by clusters within an otherwise periodic texture – leading us to wonder if it is the clustering or the periodic textures that are most to blame for the noisy appearance.

B. Void-And-Cluster

Given the disturbing artifacts created by discontuities in texture, we can look at alternatives to error-diffusion where we note that while, in theory, error-diffusion should diffuse error in a homogeneous fashion and hence minimize low-frequency graininess at all gray-levels, Figs. 7-12 show that not all filters are created equal. Furthermore, we note that it was a trail and error technique used by Ulichney to discover a perturbed error filter scheme that maintained radial symmetry without aliasing artifacts near $g = \frac{1}{2}$. So for a scheme that generates halftones in a more intuitive fashion, we can use Ulichney's iterative Void-and-Cluster initial pattern technique (VACip) where,

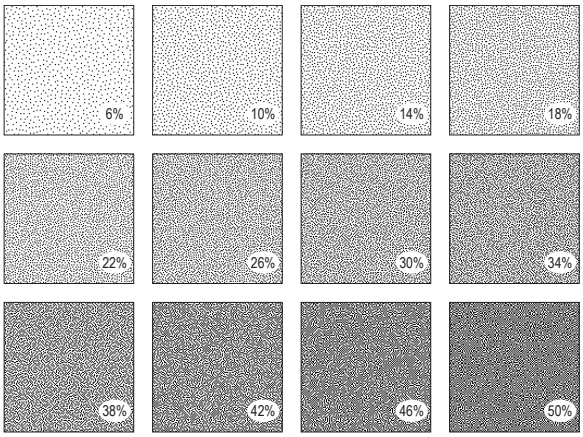

Fig. 15. The binary dither patterns for Void-and-Cluster initial patterns on a rectangular sampling grid as g transitions from 6% to 50% coverage.

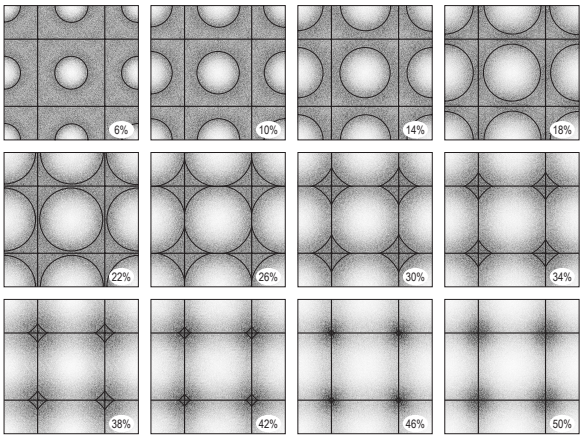

Fig. 16. The power spectra for Void-and-Cluster initial patterns on a rectangular sampling grid as g transitions from 6% to 50% coverage.

in this iterative algorithm, a white-noise dither pattern of appropriate gray-level is filtered using a low-pass FIR filter to obtain a measure of minority pixel density. The minority pixel with the highest corresponding density is replaced with a majority pixel, and the dither pattern is then filtered again by the same low-pass filter to obtain an updated measure of minority pixel density. The majority pixel with the lowest corresponding density is then replaced with a minority pixel, returning the dither pattern to the proper ratio of minority to majority pixels as defined by the gray-level g. The process is then repeated until, during a particular iteration, the majority pixel with the lowest density is the same pixel as the previous minority pixel with the highest density. If this is the case, the algorithm has converged, and the process is complete.

Because VACip iteratively swaps pixels according to an analysis of the entire local neighborhood around a subject pixel, and not just from half of the local neigborhood as in error-diffusion, VACip can more readily guarantee spatial homogeniety. And using appropriate low-pass filters, we expect VACip to maintain radial symmetry while minimizing low-frequency graininess for any gray-level. In this regard using a Gaussian low-pass filter with variance $\sigma^2 = 0.32$, Figs. 15 and 16 show the spatial dither patterns and corresponding spectra as g transitions from 6% to 50% coverage where, from visual inspection, it is confirmed that VACip behaves very similar to the perturbed filter weight scheme of Figs. 13 and 14 in that it allows some spectral energy to exist inside the principal frequency ring for gray-levels beyond $g = \frac{1}{4}$. What VACip does, beyond perturbed filter weights, is achieve much better radial symmetry given the lack of a deterministic raster path.

Now even though the amount of clustering is only slight, the resulting patterns from VACip and from Ulichney's perturbed filter scheme offer some evidence that, perhaps, clustering of minority pixels will have desirable properties for halftoning if not done to too much of an extreme. In particular, these algorithms move spectral energy inside the radial frequency f_b creating what Lau *et al* [4, 8] referred to as green-noise where the optimal halftoning schemes distribute minority pixel *clusters* as homogeneously as possible. Doing so creates a pattern where clusters of average size $\bar{M}$ pixels are separated (centroid-to-centroid) by an average distance of λ_g where:

$$\lambda_g = \begin{cases} 1/\sqrt{(g/\bar{M})}, & \text{for } 0 < g \leq 1/2 \\ 1/\sqrt{((1-g)/\bar{M})}, & \text{for } 1/2 < g \leq 1 \end{cases} . \quad (5)$$

The name "green-noise" derives from the predominantly mid-frequency content of the corresponding RAPS metric as illustrated in Fig. 3 (right) where increased clustering leads to a transition from the limiting case of blue-noise ($\bar{M} = 1$) to the mid-frequency only band of coarse patterns (black line). Here the primary spectral component is centered around the green-noise principle frequency, f_g, where:

$$f_g = \begin{cases} \sqrt{(g/\bar{M})}, & \text{for } 0 < g \leq 1/2 \\ \sqrt{((1-g)/\bar{M})}, & \text{for } 1/2 < g \leq 1 \end{cases} . \quad (6)$$

Unlike blue-noise where the randomness in the pattern is derived from variations in the separation between minority pixels, green-noise also exhibits variation in the size/shape of clusters. Too much variation in either parameter leads to spectral whitening with excessively large clusters leading to low-frequency artifacts and excessively small clusters leading to high.

Now while perturbed filter weights and VACip only introduced a small degree of clustering, the green-noise model tells us that it is possible to eliminate diagonal correlation without introducing unwanted aliasing artifacts and, hence, maintain radial symmetry at all gray-levels. Specifically, aliasing can be eliminated if the amount of clustering, at gray-levels $\frac{1}{4} < g \leq \frac{3}{4}$, is sufficiently high

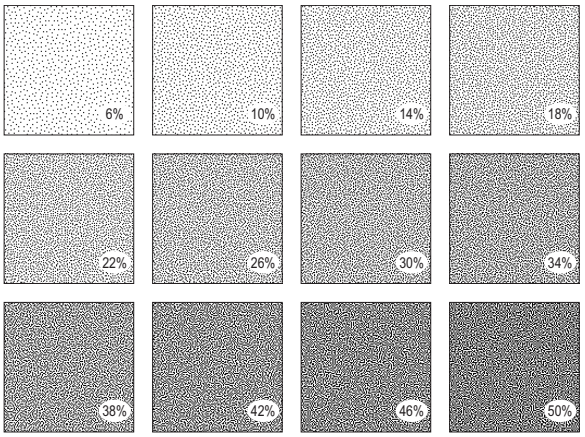

Fig. 17. The binary dither patterns for Void-and-Cluster initial patterns with σ^2 optimized to maintain radial symmetry on a rectangular sampling grid as g transitions from 6% to 50% coverage.

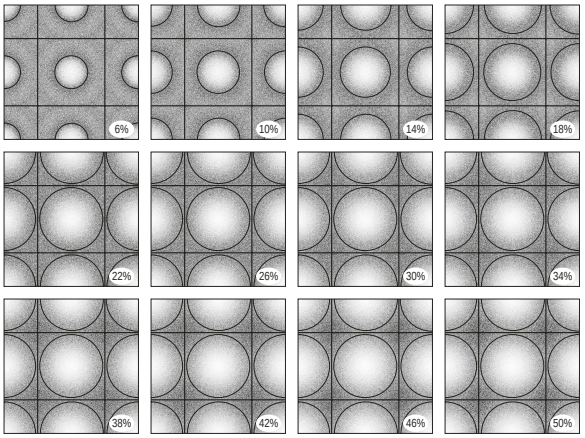

Fig. 18. The power spectra for Void-and-Cluster initial patterns with σ^2 optimized to maintain radial symmetry on a rectangular sampling grid as g transitions from 6% to 50% coverage. The black lines indicate the principal frequencies progress according to eqn. (2) up to 22% and then remain constant up to 50%.

as to reduce the principal frequency to that for gray-level $g = \frac{1}{4}$, where the principal frequency ring is the largest complete ring that can fit inside the baseband. So in an attempt to see the effects of adding this clustering, we can repeat the experiment of Figs. 15 and 16 using VACip but where the variance of the low-pass Gaussian filter is defined as $\sigma^2 = 0.6$. In this manner, we expect to see a cut-off frequency that increases with g, as prescribed by eqn. (2) for $0 \leq g \leq \frac{1}{4}$, but that levels off to a constant for $\frac{1}{4} < g \leq \frac{3}{4}$. Shown in Figs. 17 and 18 are the corresponding spatial and spectral dither patterns that clearly show this behaviour with power spectra almost identical for $\frac{1}{4} < g \leq \frac{1}{2}$. Visual inspection will show that while patterns are coarser than before, the lack of periodic texture components near $g = \frac{1}{2}$ eliminates the disturbing artifacts created by discontinuities in texture found in previous figures.

In light of the results demonstrated in Figs. 13-18, we propose a new model for blue-noise that places an increased emphasis on the need for maintaining radial symmetry and avoiding periodic textures by modifying the notion of the blue-noise principal frequency from being a wavefront progressing into the corners of the baseband to, instead, a wave progressing outward until gray-level $g = \frac{1}{4}$. Beyond $g = \frac{1}{4}$, the wavefront stops its progression as a complete, unbroken ring. This new model characterizes the ideal blue-noise dither patterns as having a principal frequency defined as:

$$f_b = \begin{cases} \sqrt{g}, & \text{for } 0 < g \leq \frac{1}{4} \\ 1/2, & \text{for } \frac{1}{4} < g \leq \frac{3}{4} \\ \sqrt{1-g}, & \text{for } \frac{3}{4} < g \leq 1 \end{cases} . \qquad (7)$$

Given the above property, we note that the patterns of Fig. 17 succeed at modeling ideal blue-noise, and while one may note that these patterns are not visually optimal, we respond by saying that the blue-noise model was never meant to define the *visually optimal* distribution of dots. Instead, blue-noise was meant only to characterize the statistical properties of error-diffusion patterns that, in 1988, made these patterns visually more appealing than those produced by periodic techniques. To characterize true visual supremacy, one needs to take into account properties of the human visual system such as its reduced sensitivity to diagonal correlation. Results of such studies have since lead us to model-based halftoning techniques such as direct binary search [9].

III. Hexagonal Sampling Grids

As depicted in Fig. 1, a regular hexagonal sampling lattice is characterized by samples placed a horizontal distance apart equal to some sample period D_x and a vertical distance apart of $\frac{\sqrt{3}}{2}D_x$. By shifting every second row of the lattice by half a pixel ($\frac{D_x}{2}$), a sample point is separated from its six neighboring samples by an equal distance D_x. In order to sample an image using hexagonal grids with the same number of samples per unit area as a rectangular grid, the spacing D_x should be equal to $(\frac{2}{\sqrt{3}})^{\frac{1}{2}}D_r$ where D_r is the sample period for the rectangular grid. For display purposes, an ideal printing device would print each pixel as a regular hexagon, but because this paper will be printed on a traditional rectangular grid device, we will up-sample our hexagonal grid halftones by a factor of two using nearest-neighbor interpolation and print these dither patterns on a rectangular grid, shifting each other pair of rows by a single pixel. The images will then be further scaled along the vertical axis by a factor of $\frac{\sqrt{3}}{2}$ to create a symmetric sampling grid.

Now as described by Ulichney [2], a well formed blue-noise dither pattern will be such that minority pixels will be separated by an average distance λ_b as defined in eqn. (1). For reasons relating to the derivation of the

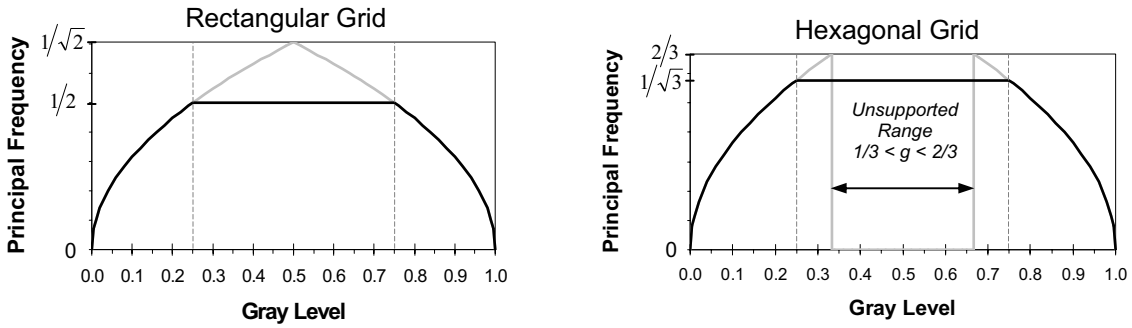

Fig. 19. The principal frequency, f_b, versus gray-level, g, for (top) rectangular and (bottom) hexagonal for both the (gray) old and (black) new models.

Fourier transform of a regular hexagonal sampling lattice, the principal frequency, f_b, of a blue-noise pattern will be defined according to:

$$f_b = \begin{cases} \frac{2}{\sqrt{3}}\sqrt{g}, & \text{for } 0 < g \leq 1/2 \\ \frac{2}{\sqrt{3}}\sqrt{1-g}, & \text{for } 1/2 < g \leq 1 \end{cases} . \qquad (8)$$

Noting from Fig. 4 (right) that the maximum radial frequency that fits within the bounds of the baseband occurs at $f_b = \frac{2}{3}$ when the gray-level reaches $g = \frac{1}{3}$, it is not possible to have a dither pattern adhering to the blue-noise RAPS of Fig. 3 (left) on a hexagonal sampling gird in the range of $\frac{1}{3} < g \leq \frac{2}{3}$. This is, perhaps, better illustrated in Fig. 19 where the principal frequency is shown plotted versus gray-level for both rectangular and hexagonal sampling grids.

For both grid geometries, Fig. 19 shows the relationship between f_b and g where the gray lines indicate the original relationship proposed by Ulichney in 1988. Looking specifically at hexagonal grids shows that, for gray levels between $\frac{1}{3}$ and $\frac{2}{3}$, f_b exceeds the maximum radial frequency that fits inside the baseband and was, hence, undefined by Ulichney is this so-called, "unsupported region." Based upon trial-and-error experiments with error-diffusion, Ulichney further theorized that because stochastic dither patterns would always have significant energy below f_b in this region, patterns would always have excessively large variations in the spacing between dots, and under no circumstances, would patterns not look noisy and uncorrelated.

We now see that Ulichney's theory was wrong, and that if we apply the notion of an expanding spectral ring that stalls its growth at the point where the ring first contacts the limits of the baseband, our new model says that beyond the point of first contact, the sampling grid restricts the placement of dots and can, therefore, introduce artifacts just as disturbing as any excess energy below f_b. In response, we are defining f_b according to eqn. (7) as the optimal trade-off between pattern coarseness and sample grid dot placement where we allow the pattern to exhibit a minimum degree of clustering in order to maintain radial symmetry at all gray-levels. This new definition is illustrated by the black line relationship between f_b and g in Fig. 19, which levels off at the maximum complete annuli that fits inside the baseband.

A particularly elegant property of our new blue-noise model is that if we set f_b equal to $\frac{1}{\sqrt{3}}$ and solve for g, we see that the spectral ring first contacts the limits of the hexagonal baseband at gray-level $g = \frac{1}{4}$, exactly where it occured for rectangular sampling grids!!! The implications of this fact are enormous. First, it implies that our new blue-noise model of eqn. (7) applies to hexagonal grids without modification except for the previously mentioned $\frac{\sqrt{3}}{2}$ factor such that:

$$f_b = \begin{cases} \frac{2}{\sqrt{3}}\sqrt{g}, & \text{for } 0 < g \leq \frac{1}{4} \\ \frac{2}{\sqrt{3}}(1/2), & \text{for } \frac{1}{4} < g \leq \frac{3}{4} \\ \frac{2}{\sqrt{3}}\sqrt{1-g}, & \text{for } \frac{3}{4} < g \leq 1 \end{cases} . \qquad (9)$$

Second and more importantly, it implies that, in the range from $g = 0$ to $\frac{1}{4}$, blue-noise dither patterns on a hexagonal grid have a 15.47% higher cut-off frequency than those corresponding to rectangular grids with the same number of samples per unit area. As such, blue-noise dither patterns, on hexagonal grids, are less visible than those on rectangular grids for these gray-levels and, if we use the new blue-noise model, at all gray-levels. Finally in situations where clustering does occur, either purposely for minimizing the effects of printer distortions or forcibly near gray-level $g = \frac{1}{2}$, hexagonal sampling grids can form pixel pairs in three directions as opposed to two, allowing for improved radial symmetry in the size and distribution of minority pixel clusters. Coarse halftone patterns should, therefore, form smoother visual textures on hexagonal grids than rectangular.

So assuming that hexagonal is the preferred sampling geometry, we are in the familiar position of trying to find a means by which to generate optimal dot distributions and to do so in a computationally efficient manner. Looking at the first published study of error-diffusion on hexagonal sampling grids, Figs. 20 and 21 show the spatial dither

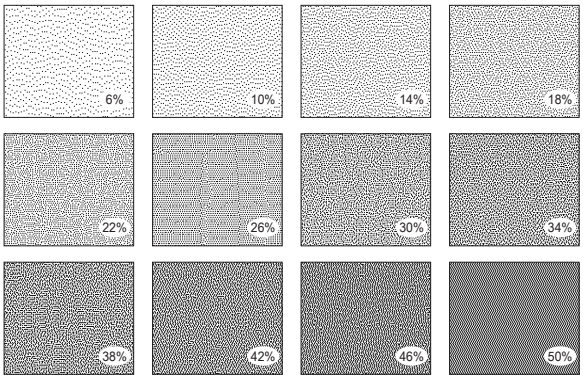

Fig. 20. The binary dither patterns for error-diffusion using the Stevenson and Arce filter on a hexagonal sampling grid as g transitions from 6% to 50% coverage.

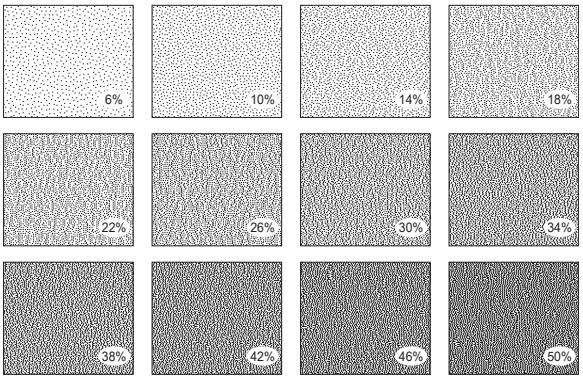

Fig. 22. The binary dither patterns for Ulichney's perturbed filter weight scheme on a hexagonal sampling grid as g transitions from 6% to 50% coverage.

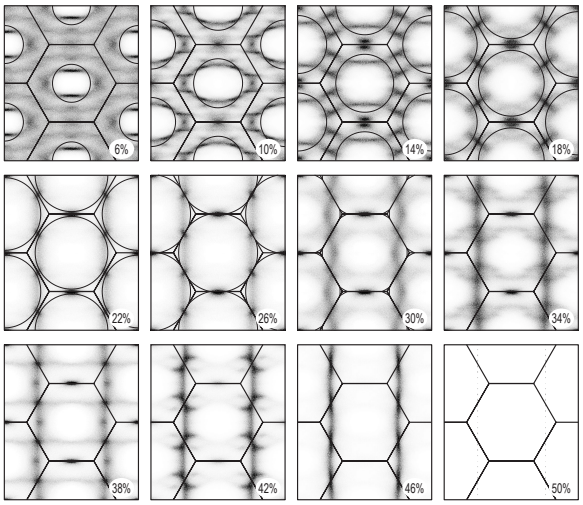

Fig. 21. The power spectra for error-diffusion using the Stevenson and Arce filter on a hexagonal sampling grid as g transitions from 6% to 50% coverage.

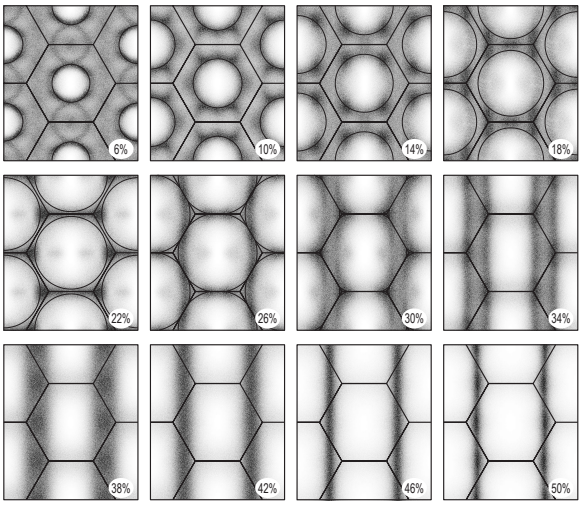

Fig. 23. The power spectra for Ulichney's perturbed filter weight scheme on a hexagonal sampling grid as g transitions from 6% to 50% coverage.

patterns and power spectral densities corresponding to the Stevenson and Arce [10] error filter. From visual inspection, one sees a consistent blue-noise appearance for gray-levels below $g = \frac{1}{3}$, but strong vertical artifacts, deriving from the raster scan, seem to dominate near $g = \frac{1}{2}$. Looking specifically at the power spectra, one sees the spectral lines running vertically that intersect the limits of the baseband at the same points where the spectral rings make first contact at $g = \frac{1}{4}$. Furthermore, these lines of energy do not seem to become prominant components until after gray-level $g = \frac{1}{4}$, adding credence to our claim that it is these gray-levels where the sampling lattice begins to constrain the distribution of dots.

A. Perturbed Filter Weights

Noting the relatively poor performance of the Stevenson and Arce filter, Ulichney proposed using the same perturbed filter weight scheme demonstrated in Figs. 13 and 14. Shown in Fig. 22 and 23 are the binary dither patterns and corresponding power spectra for this technique on hexagonal sampling grids. Like the Stevenson and Arce filter, the perturbed filter scheme produces visually pleasing patterns below gray-level $g = \frac{1}{4}$ but is forced to cluster pixels as g approaches $\frac{1}{2}$. While it clearly does a better job in this range, the deterministic raster leads to strong vertical artifacts very similar to those produced by the Stevenson and Arce filter.

B. Void-And-Cluster

Given the poor performance of error-diffusion, we can attempt to create dither patterns using VACip where we expect to achieve the preferred behavoir of only introducing clustering when necessary as to avoid aliasing. Like the discrete Fourier transform, convolution of a binary dither pattern with a linear, FIR filter can be achieved using techniques for matrices from rectangular sampling

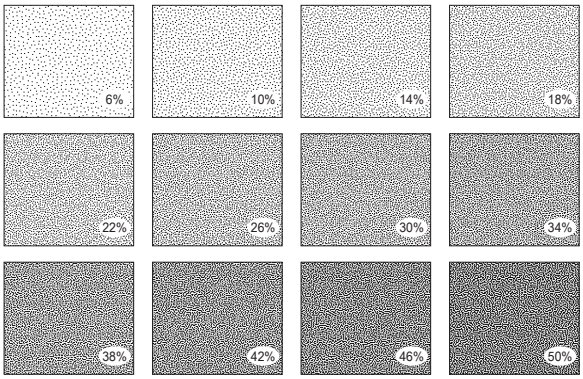

Fig. 24. The binary dither patterns for Void-and-Cluster initial patterns on a hexagonal sampling grid as g transitions from 6% to 50% coverage.

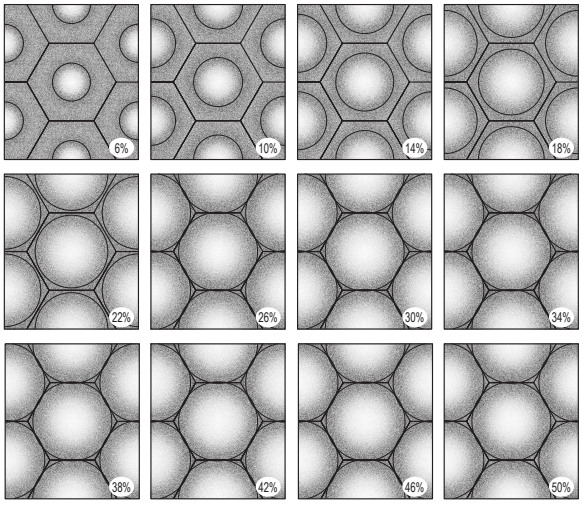

Fig. 25. The power spectra for Void-and-Cluster initial patterns on a hexagonal sampling grid as g transitions from 6% to 50% coverage.

grids if the subject matrices store the skewed versions of the hexagonal sampled data [2]. Shown in Figs. 24 and 25 are the binary dither patterns and corresponding power spectra generated by VACip using Gaussian low-pass filters where $\sigma^2 = 0.32$ as g transitions from 6% to 50%. In the case of Figs. 15 and 16, this same filter variance resulted in dither patterns with significant energy packed into the corners of the baseband on a rectangular sampling grid. But here, results show identical spectral distributions beyond $g = \frac{1}{4}$ where there is no room in the corners of the baseband as there was for rectangular grids. From visual inspectrion, it is clear that these spectral distributions are achieved through the clustering of minority pixels in the spatial domain that, near gray-level $g = \frac{1}{2}$, create worm patterns.

While worm textures/patterns are traditionally thought of as disturbing artifacts, we note that the radial symmetry of these particular patterns creates a twisting and turning path from pixel to pixel. This constant spiraling creates a smooth, almost invisible texture. We would further note that the worm patterns found here are far less objectionable than the strong directional patterns created by error-diffusion in either Fig. 20 or 22. In fact, seeing Fig. 24 finally offers some insight into what the *visually* optimal stochastic dither pattern may look like for gray-level $g = \frac{1}{2}$ on a hexagonal sampling grid – something that has yet to be determined. In seeing Fig. 24, it is hoped that deriving optimal halftoning schemes for these grids will be easier to do.

C. Dither Arrays

Noting the succesful conversion of the VACip algorithm to hexagonal grids in this paper, we have taken the added time to generate the hexagonal dither array shown in Fig. 26 and demonstrated versus a similar array on a rectangular grid in Fig. 27. Dither arrays refer to a halftoning technique where a continuous-tone image is converted to binary by a pixelwise comparison with thresholds stored in a dither array matrix or screen. Input pixels with intensity values greater than the corresponding threshold value are set to one while pixels below are set to zero. For large images, dither arrays are tiled end-to-end until all input pixels have a corresponding pair within the screen. In order to avoid discontinuities in the halftone texture near boundaries of the screen, dither arrays are designed to satisfy a wrap-around property. For a thorough description of the construction algorithms used for building these screens, we refer the reader to the original works by Mitsa and Parker [11] and Ulichney [3].

In Fig. 26, we show just a 128×128 cropped section of a 256×256 mask along with the magnitude of the its corresponding Fourier transform, showing the uniquely high frequency components of the dither array. In Fig. 27, both screens were generated using VAC with identical low-pass filter variances of $\sigma^2 = 0.32$, and the gray-scale images used were the same as that used by Ulichney [2] to demonstrate the effects of high-pass filtering an original gray-scale image prior to dithering. To present a fair comparison, both images in this figure have the same number of pixels per unit area. As we are using a regular (square) rectangular grid and a regular hexagonal grid, the number of rows and columns will not be the same. The rectangular grid has an aspect ratio of 1, and the hexagonal grid has an aspect ratio of $2/\sqrt{3}$. To maintain our constant pixel density constraint, the horizontal period used for the hexagonal case is $(2/\sqrt{3})^{\frac{1}{2}}$ that of the rectangular case.

From visual inspection, we would argue that the hexagonal grid dither array is far superior to any hexagonal grid halftoning scheme that we are aware in terms of maintaining radial symmetry while simultaneously spreading minority pixels as homogenuously as possible. Seeing both the rectangular and hexagonal grid dither arrays side-by-side, it should also be clear from visual inspection that

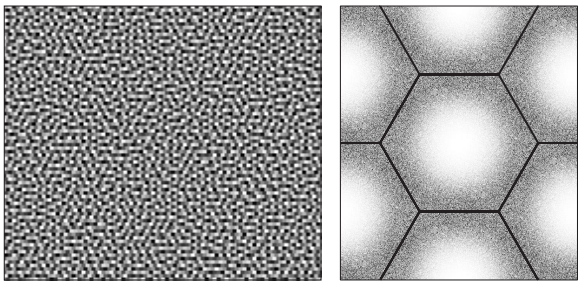

Fig. 26. The (left) blue-noise dither array for hexagonal sampling grids along with (right) the corresponding magnitude of its Fourier transform.

the hexagonal mask clearly creates the all important grid-defiance illusion as one cannot, without very close inspection, determine which of the images is printed on a hexagonal sampling grid.

IV. Conclusions

By introducing a means by which dither patterns could be quantitatively evaluated, Ulichney's blue-noise model has played a fundamental role in halftoning research, and it is, perhaps, one of the most often cited works in halftoning as the term, "blue-noise," has become synonomous with visually pleasing. But the model is not without its short-comings. In particular, Ulichney argued that hexagonal sampling grids were inferior to rectangular for displaying visually pleasing stochastic dither patterns inspite of the fact that there are numerous advantages to using hexagonal grids for other image processing purposes. Specifically, Ulichney determined that it was not possible to generate dither patterns at all gray-levels as a pattern of isolated dots on hexagonal grids and that, at $\frac{1}{3} < g \leq \frac{2}{3}$, minority pixels would be forced to occur in clusters. This clustering, Ulichney thought, would always lead to a large variation in the spacing between dots, resulting in a noisy, uncorrelated appearance.

Under the green-noise model, it is possible to create visually pleasing dither patterns when the halftoning algorithm intentionally clusters pixels, and in fact, a small degree of clustering can be beneficial to producing visually pleasing halftones by maintaining radial symmetry whenever the sampling grid would otherwise constrain the placement of dots. As such, this paper has introduced a new blue-noise model that incorporates clustering in the gray-level range $\frac{1}{4} < g \leq \frac{3}{4}$ regardless of the sampling grid geometry. By doing so, an optimal halftoning scheme will maintain its grid-defiance illusion at all gray-levels producing a radially symmetric distribution of stochastically arranged dots. Such a dither pattern will be void of the patches of periodic textures found in previously ideal techniques, of the traditional blue-noise model, that pack spectral energy into the corners of the rectangular power spectrum through the introduction of directional correlation.

Based on observations made in previous papers [5, 12], many printing devices are unable to produce blue-noise dither patterns without introducing severe tonal distortion and without introducing strong spatial variations in DC regions. These devices are, therefore, incapable of attaining the high spatial frequencies found only in rectangular sampling grids, not hexagonal. So the argument, under the previous model, that rectangular sampling grids are preferable to hexagonal because only they support blue-noise, is invalid for many printing devices. The overall conclusion of this paper is, therefore, that since clustering algorithms may, in many cases, be the only way to produce reliably printed, stochastic dither patterns, there is no conclusive reason for using rectangular sampling grids instead of hexagonal.

Going into the future, one would hope that the results presented here will reinvigorate members of the research community into looking again at the problem of halftoning on hexagonal sampling grids as there are significant challenges to overcome. In particular, we are now back in a state faced for rectangular sampling grids in the early 1990s where the challenge was to derive the optimal distribution of dots in a computationally efficient manner. One such result of that work was the creation of blue-noise dither arrays. As is the case for dither arrays [3, 11], we expect that the solutions to the various problems will be the same as those proposed for rectangular sampling grids and that the bulk of the work will focus on tuning various halftoning parameters. In particular, Pappas and Neuhoff's model-based error-diffusion [13] will require modified printer models that take into account the new geometry, and Allebach's direct binary search [9] will, of course, require a new low-pass filter to model the human visual system. The list can, of course, go on, but it is our desire, in presenting this work, to see these many problems solved sooner rather than later as it has been 19 years since Stevenson and Arce published their paper on the subject and 16 years since Ulichney presented his analysis. Yet little has transpired otherwise.

References

[1] J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 21, no. 1, pp. 33–38, 1991.

[2] R. A. Ulichney, *Digital Halftoning*. MIT Press, Cambridge, MA, 1987.

[3] R. A. Ulichney, "The void-and-cluster method for dither array generation," in *Proceedings SPIE, Human Vision, Visual Processing, Digital Displays IV* (B. E. Rogowitz and J. P. Allebach, eds.), vol. 1913, pp. 332–343, 1993.

[4] D. L. Lau, G. R. Arce, and N. C. Gallagher, "Green-noise digital halftoning," *Proceedings of the IEEE*, vol. 86, no. 12, pp. 2424–2444, 1998.

[5] D. L. Lau and G. R. Arce, *Modern Digital Halftoning*. Signal Processing and Communications, New York, New York, USA: Marcel Dekker, Inc., 2001.

[6] R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial gray-scale," *Proceedings Society Information Display*, vol. 17, no. 2, pp. 75–78, 1976.

Fig. 27. A comparison of blue-noise dither arrays on a (left) rectangular and a (right) hexagonal sampling grid where each image has the same total number of pixels but not the same dimensions.

[7] J. F. Jarvis, C. N. Judice, and W. H. Ninke, "A survey of techniques for the display of continuous-tone pictures on bilevel displays," *Computer Graphics and Image Processing*, vol. 5, pp. 13–40, 1976.

[8] D. L. Lau, G. R. Arce, and N. C. Gallagher, "Digital halftoning by means of green-noise masks," *Journal of the Optical Society of America*, vol. 16, pp. 1575–1586, July 1999.

[9] M. Analoui and J. P. Allebach, "Model based halftoning using direct binary search," in *Proceedings SPIE, Human Vision, Visual Processing, and Digital Display III* (B. E. Rogowitz, ed.), vol. 1666, pp. 96–108, August 1992.

[10] R. L. Stevenson and G. R. Arce, "Binary display of hexagonally sampled continuous-tone images," *Journal of the Optical Society of America*, vol. 2, no. 7, 1985.

[11] T. Mitsa and K. J. Parker, "Digital halftoning technique using a blue noise mask," *Journal of the Optical Society of America*, vol. 9, pp. 1920–1929, 1992.

[12] B. E. Cooper and D. L. Lau, "An evaluation of green-noise masks for electrophotographic halftoning," in *Proceedings of the IS&T's and SPIE's Electronic Imaging*, (San Jose, California, USA), January 22-28 2000.

[13] T. N. Pappas and D. L. Neuhoff, "Model-based halftoning," in *Proceedings of SPIE, Human Vision, Vision Processing and Digital Display II* (B. E. Rogowitz, M. H. Brill, and J. P. Allebach, eds.), vol. 1453, pp. 244–255, June 1991.