Fundamental Characteristics of Halftone Textures: Blue-Noise and Green-Noise

Daniel L. Lau[1], Robert Ulichney, Gonzalo R. Arce[2]
Image Systems Laboratory
HP Laboratories Cambridge
HPL-2003-62
March 31st , 2003*

E-mail: dlau@engr.uky.edu, u@hp.com, arce@ece.udel.edu

halftoning,
dithering,
blue noise,
green noise

In this paper, we review the spatial and spectral characteristics of blue and green-noise halftoning models. In the case of blue-noise, dispersed-dot dither patterns are constructed by isolating minority pixels as homogeneously as possible and by doing so, a pattern composed exclusively of high frequency spectral components is produced. Blue-noise halftoning is preferred for display devices that can accommodate isolated dots such as various video displays and some print technologies such as ink-jet. For print marking engines that cannot support isolated pixels dispersed-dot halftoning is inappropriate. For such cases, clustered-dot halftoning is used to avoid dot-gain instability. Green-noise halftones are clustered-dot blue noise patterns. Such patterns enjoy the blue-noise properties of homogeneity and lack of low frequency texture, but have clusters of minority pixels on blue-noise centers. Green noise is composed exclusively of mid-frequency spectral components. In addition to the basic spatial and spectral characteristics of the halftoning models, this paper also reviews some of the earlier work done to improve error diffusion as a noise generator. Also reviewed are processes to generate threshold arrays to achieve blue noise and green noise with the computationally efficient process of ordered dither.

* Internal Accession Date Only Approved for External Publication
[1] Department of Electrical and Computer Engineering, University of Kentucky, Lexington, KY 40506-0046
[2] Department of Electrical and Computer Engineering, University of Delaware, Newark, DE 19717

To be published in IEEE Signal Processing Magazine, 2003

Fundamental Characteristics of Halftone Textures: Blue-Noise and Green-Noise

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*

Gonzalo R. Arce
Department of Electrical and Computer Engineering
University of Delaware
Newark DE 19717 USA
email: *arce@ece.udel.edu*

March 6, 2003

Abstract

In this paper, we review the spatial and spectral characteristics of blue and green-noise halftoning models. In the case of blue-noise, dispersed-dot dither patterns are constructed by isolating minority pixels as homogeneously as possible and by doing so, a pattern composed exclusively of high frequency spectral components is produced. Blue-noise halftoning is preferred for display devices that can accommodate isolated dots such as various video displays and some print technologies such as ink-jet. For print marking engines that cannot support isolated pixels dispersed-dot halftoning is inappropriate. For such cases, clustered-dot halftoning is used to avoid dot-gain instability. Green-noise halftones are clustered-dot blue noise patterns. Such patterns enjoy the blue-noise properties of homogeneity and lack of low frequency texture, but have clusters of minority pixels on blue-noise centers. Green noise is composed exclusively of mid-frequency spectral components. In addition to the basic spatial and spectral characteristics of the halftoning models, this paper also reviews some of the earlier work done to improve error diffusion as a noise generator. Also reviewed are processes to generate threshold arrays to achieve blue noise and green noise with the computationally efficient process of ordered dither.

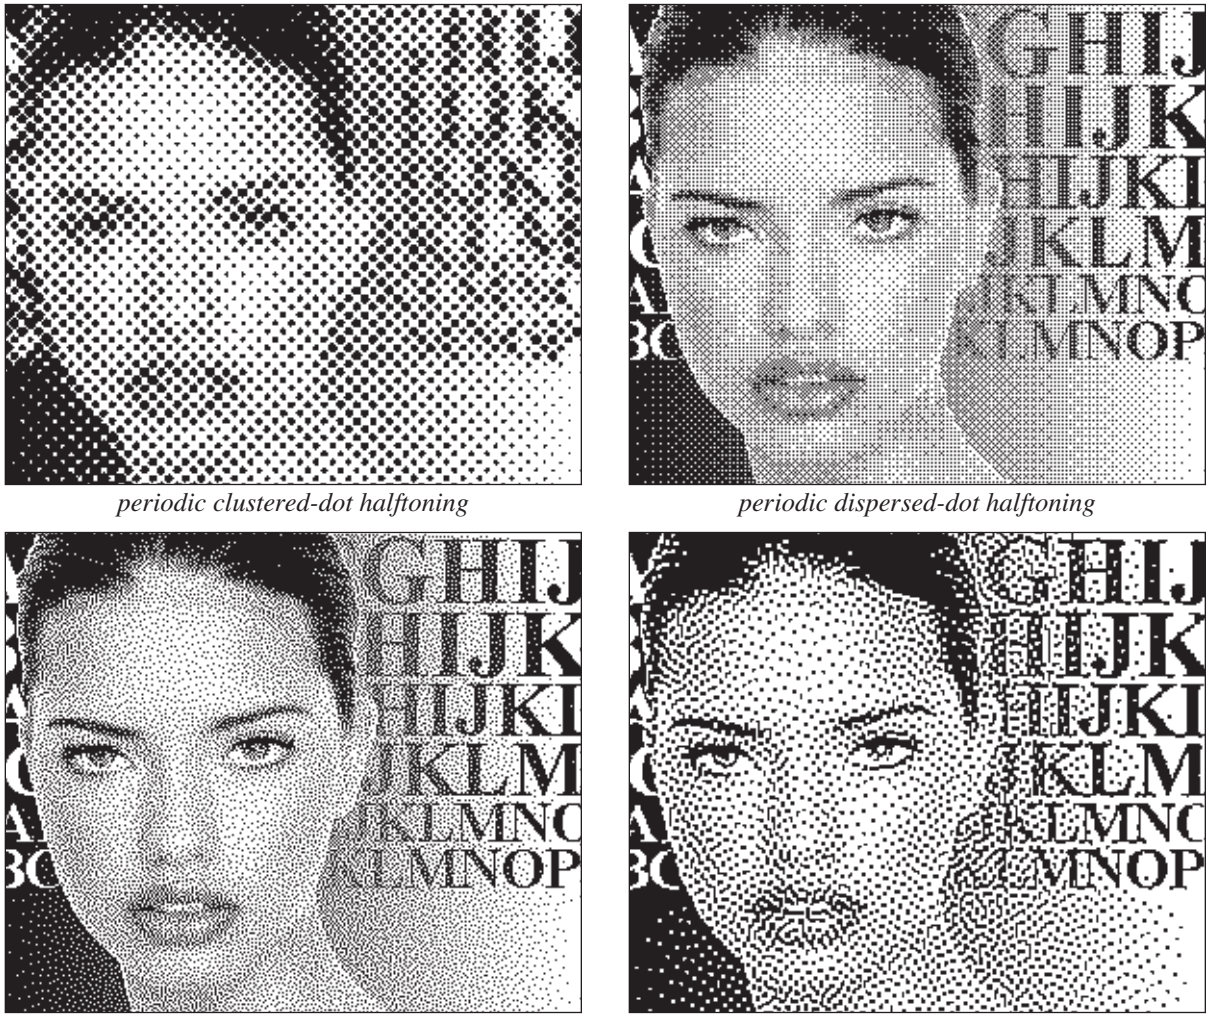

periodic clustered-dot halftoning | *periodic dispersed-dot halftoning*

aperiodic dispersed-dot halftoning | *aperiodic clustered-dot halftoning*

Figure 1: The image *Adrian* halftoned using the various combinations of aperiodic versus periodic and clustered versus dispersed-dot halftoning.

1 Introduction

In halftoning, blue-noise is the statistical model describing the ideal spatial and spectral characteristics of aperiodic dispersed-dot dither patterns [1], and in essence, the ideal blue-noise halftoning scheme produces stochastic dither patterns of same sized dots distributed as homogeneously as possible. By doing so, the spectral content of these patterns are composed entirely of high frequency spectral components. And as blue is the high frequency component to visible white light, blue-noise is the high frequency component to white-noise. Given the low-pass nature of the human visual system [2], blue-noise creates patterns visually appealing

simply because the spectral components of the pattern lie in the regions least visible to the human viewer; furthermore, the stochastic distribution of dots creates a grid-defiance illusion where the structure of the underlying grid on that pixels are aligned is no longer apparent to the viewer.

In comparison to periodic clustered-dot halftoning schemes, blue-noise maximizes the apparent resolution of printed images, creating an image that lacks the visually disturbing texture created by large clusters arranged along a regular grid. In locally periodic clustered-dot halftoning, not only are the clusters more visible than the isolated dots of aperiodic dispersed-dot halftoning, but the regular grid used to align clusters are particularly apparent to the human visual system. Now while blue-noise has become the standard approach to halftoning in ink-jet printing, it has made little headway in electrophotographic (laser) printers or commercial lithographic presses (although several exceptions exist), and the reason that blue-noise has been avoided in these devices has to do with printer distortion and the inability of these devices to reproduce dots consistently from dot to dot. As a consequence, unreliable devices have relied on locally periodic clustered-dot halftoning and the visual detriment that regular patterns pose to the human visual system.

As a way of providing the benefits of random dot distributions while maintaining the consistency of clustered dots, Levien [3] introduced error diffusion with output-dependent feedback where a weighted sum of previous output pixels is used to modulate the quantization threshold with resulting patterns composed of a random arrangement of randomly sized and shaped printed dot clusters. As was the case for error diffusion in 1985, error diffusion with output-dependent feedback lacked a defining model that characterized the spatial and spectral characteristics of the resulting *aperiodic clustered-dot* patterns. And in order to provide such a model, Lau *et al* [4] introduced the *green-noise* model where the ideal patterns are composed of homogeneously distributed pixel clusters that vary in both their size and spacing for varying shades of gray. In the Fourier domain, green-noise patterns are composed almost exclusively with mid-frequency components, and as green is the mid-frequency component to white, green-noise is the mid-frequency component to white-noise.

Since their introduction, both blue and green-noise models have lead to major breakthroughs of innovation in halftoning with improvements in both the efficiency at which

halftoning is performed and the visual fidelity of the resulting patterns. The most noteworthy of these innovations are blue and green-noise dithering arrays where a continuous-tone original is converted to a binary image using a point-process that compares the intensity of a given pixel with a threshold stored in the array. And given the significance of the roles that both blue and green-noise have played in halftoning research, we devote the remainder of this paper to detailing the blue and green-noise halftoning models as well as describing the spatial and spectral metrics introduced by Ulichney and Lau *et al* that serve as the basis for these models. While this material offers historical perspective, today it is being realized that the fundamental properties of the various dot distributions play a fundamental role in the visual pleasantness of color halftones where the monochrome halftones of cyan, magenta, yellow, and black inks are superimposed. So for the reader interested in these color applications, it is advised that a basic fundamental understanding of the principal halftoning models will serve you well in developing and understanding future algorithms.

2 Spatial and Spectral Halftone Statistics

In order to differentiate between the various schemes, a halftoning algorithm is classified according to the statistical relationship between minority pixels in the resulting dither pattern produced by halftoning images of constant intensity or gray level. The rendition of edges and other high-frequency details depend primarily on how sharp the image is or to what extent high-pass filtering is performed on the image prior to halftoning [1]. By treating the resulting dither pattern as a set of points where an event or point is said to occur at the location of a minority pixel, Lau *et al* [4] propose using the spatial statistics commonly employed in *stochastic geometry* to study point processes.

In the point process framework for continuous spaces, a point process Φ is defined as a stochastic model governing the location of events, or points x_i, within the 2-D real space $\Re^2$ [5]. We further define ϕ as a sample of Φ written as a set of randomly arranged points such that $\phi = \{x_i \in \Re^2 : i = 1, \ldots, N\}$, and we define $\phi(B)$ as a scalar quantity defined as the number of x_is in the subset B in $\Re^2$. We assume that the point process Φ is simple such

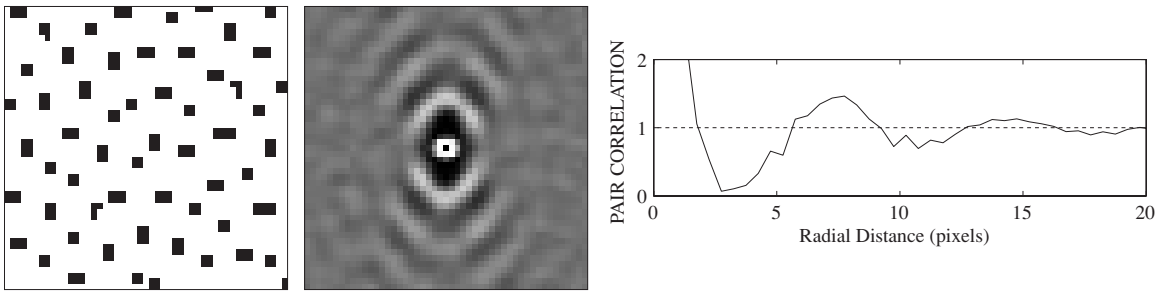

Figure 2: The (left) binary dither pattern exhibiting clusters and the corresponding (center) reduced second moment measure, $\mathcal{K}[n;m]$, and (right) pair correlation, $\mathcal{R}(r)$, derived from $\mathcal{K}[n;m]$ by dividing the spatial domain into annular rings.

that for $i \neq j$ implies $x_i \neq x_j$, which further implies that:

$$\lim_{dV_x \to 0} \phi(dV_x) = \begin{cases} 1 & \text{for } x \in \phi \\ 0 & \text{else} \end{cases}, \tag{1}$$

where dV_x is the infinitesimally small area around x. In terms of a discrete dither pattern, ϕ represents the set of minority pixels where $\phi[n] = 1$ indicates that the pixel with index n is a minority pixel in the subject dither pattern.

Having Φ for a discrete-space halftoning process, a commonly used statistic for characterizing the point process is the quantity $\mathcal{K}[n;m]$ defined as:

$$\mathcal{K}[n;m] = \frac{\mathbf{Pr}\{\phi[n] = 1 | \phi[m] = 1\}}{\mathbf{Pr}\{\phi[n] = 1\}}, \tag{2}$$

the ratio of the conditional probability that a minority pixel exists at n given that a minority pixel exists at m to the unconditional probability that a minority pixel exists at n. Referred to as the *reduced second moment measure*, $\mathcal{K}[n;m]$ may be thought of as the influence at location n of the minority pixel at m. That is, is a minority pixel at n more or less likely to occur because a minority pixel exists at m?

From $\mathcal{K}[n;m]$, we can derive a 1-D spatial domain statistic by partitioning the spatial domain into a series of annular rings $R_y(r)$with center radius r, width Δ_r, and centered around location m. This statistic for stationary and isotropic Φ is the *pair correlation* $\mathcal{R}(r)$, defined as the expected or mean value of $\mathcal{K}[n;m]$ within the ring. The usefulness of $\mathcal{R}(r)$ can be seen in the interpretation that maxima of $\mathcal{R}(r)$ indicate frequent occurrences of the inter-point distance r while minima of $\mathcal{R}(r)$ indicate an inhibition of points at r [6]. To see this behavior, Fig. 2 (right) shows the resulting pair correlation for the clustering process of

Fig. 2 (left), using annular rings $R_y(r)$ such that $\{x : r - \Delta_r/2 < |x - y| \leq r + \Delta_r/2\}$ where $\Delta_r = 1/2$, with an increased likelihood of minority pixels occurring near $r = 0$ and $r = 8$ pixels and a decreased likelihood in between 0 and 8. Because $\mathcal{K}[n; m] = 1$ for all m and n in a white-noise (uncorrelated) dither pattern, if at any time that $\mathcal{R}(r) = 1$ for a given point process, then points that are r distance apart are considered statistically uncorrelated even if they are not physically. Returning to Fig. 2 (right) as r continues to increase beyond 12 pixels, points become less and less correlated as demonstrated by the fact that $\mathcal{R}(r)$ approaches 1 with greater r.

In the Fourier domain, the power spectrum of a given dither pattern can be derived by means of spectral estimation. One technique for spectral estimation is Bartletts method of averaging periodograms [7, 8] where a periodogram is the magnitude squared of the Fourier transform of a sample output divided by the sample size. It can be shown [9] that a spectral estimate, $\hat{P}(f)$, formed by averaging K periodograms has an expectation equal to $P(f)$ smoothed by convolution with the Fourier transform of a triangle function with a span equal to the size of the sample segments and variance:

$$\mathrm{v}ar\{\hat{P}(f)\} \approx \frac{1}{K} P^2(f). \tag{3}$$

Since $\hat{P}(f)$ is a function of two dimensions and although anisotropies in the sample dither pattern can be qualitatively observed by studying 3-D plots of $\hat{P}(f)$, a more quantitative metric of spectral content is derived by partitioning the spectral domain into annular rings of width Δ_f with a central radius f_ρ, the radial frequency, and $N_\rho(f_\rho)$ frequency samples.

By taking the average value of the frequency samples within an annular ring and plotting this average versus the radial frequency, Ulichney [1] defines the *radially averaged power spectral density* (RAPSD), $P_\rho(f_\rho)$, such that:

$$P_\rho(f_\rho) = \frac{1}{N_\rho(f_\rho)} \sum_{i=1}^{N_\rho(f_\rho)} \hat{P}(f). \tag{4}$$

Because of the manner in which sampling along a rectangular grid leads to tiling of the base-band frequency on the spectral plane, rings with radial frequencies beyond $\frac{1}{2}D^{-1}$, where D is the minimum distance between samples on the display, are cropped into the corners of the spectral tile leading to fewer spectral samples in these rings. In all plots of $P_\rho(f_\rho)$,

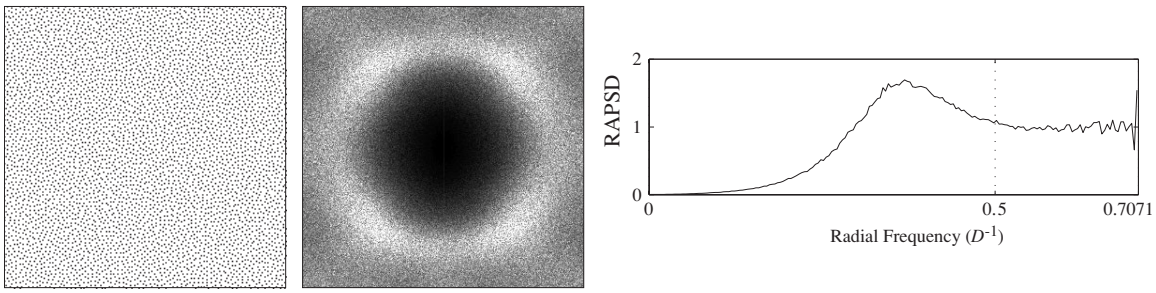

Figure 3: The (left) binary dither pattern and the corresponding (center) power spectrum, $\hat{P}(f)$, produced using Bartlett's method of averaging periodograms and (right) radially averaged power spectral density, $P_\rho(f_\rho)$, derived from $\hat{P}(f)$ by dividing the spectral domain into annular rings.

these regions of cropping will be indicated along the horizontal axis, and as a demonstration, Fig. 3 (right) shows the RAPSD for the dither pattern illustrated in Fig. 3 (left) with an increasingly chaotic behavior in the cropped rings near $f_\rho = \frac{1}{\sqrt{2}}$. Here as in the rest of this paper, the power spectral estimate is divided into annular rings of radial width Δ_f such that exactly one sample along each frequency axis falls into each ring.

3 Blue-Noise Halftoning

Common practice for characterizing observed noise models is to assign a name based on color where the most well known example is "white noise," so named because its power spectrum is flat across all frequencies, much like the visible frequencies of light. "Pink noise" is used to describe low frequency white noise, the power spectrum of which is flat out to some finite high frequency limit. The spectrum associated with Brownian motion is (perhaps whimsically) referred to as "brown noise" [10]. Blue-noise is the high frequency compliment of pink noise that, due to the low-pass nature of the human visual system, posses very advantageous properties for creating the illusion of continuous-tone in binary halftones.

Blue-noise halftoningis characterized by a distribution of binary pixels where the minority pixels are spread as homogeneously as possible [1] such that when applied to an image of constant gray level g, minority pixels are separated by an average distance λ_b where:

$$\lambda_b = \begin{cases} D/\sqrt{g} & , \text{ for } 0 < g \leq 1/2 \\ D/\sqrt{1-g} & , \text{ for } 1/2 < g \leq 1 \end{cases} \tag{5}$$

and D is the minimum distance between addressable points on the display [1, 11]. The

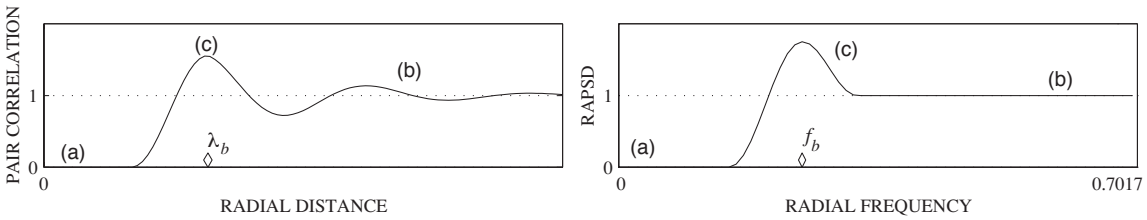

Figure 4: The (left) pair correlation and (right) radially averaged power spectrum for a blue-noise process.

parameter λ_b is referred to as the *principle wavelength* of blue-noise, with its relationship to g justified by several intuitive properties:

1. As the gray value approaches perfect white ($g = 0$) or perfect black ($g = 1$), the principle wavelength approaches infinity.

2. Wavelength decreases symmetrically with equal deviations from black and white toward middle gray ($g = 1/2$).

3. The square of the wavelength is inversely proportional to the number of minority pixels per unit area.

In terms of point processes, Φ_B is an inhibitive or soft-core point process that minimizes the occurrence of any two points falling within some distance λ_b of each other. These types of point processes are most commonly thought of as Poisson point processes where all points are approximately equally distant apart, and as a Poisson point process, we can characterize blue-noise halftones in terms of the pair correlation, $\mathcal{R}(r)$, by noting that:

1. Few or no neighboring pixels lie within a radius of $r < \lambda_b$.

2. For $r > \lambda_b$, the expected number of minority pixels per unit area approaches g for $0 \leq g \leq \frac{1}{2}$ or $1 - g$ for $\frac{1}{2} < g \leq 1$ with increasing r.

3. The average number of minority pixels within the radius r increases sharply near $r = \lambda_b$.

The resulting pair correlation for blue-noise is therefore of the form in Fig. 4 (left) where $\mathcal{R}(r)$ shows: (a) a strong inhibition of minority pixels near $r = 0$, (b) a decreasing correlation

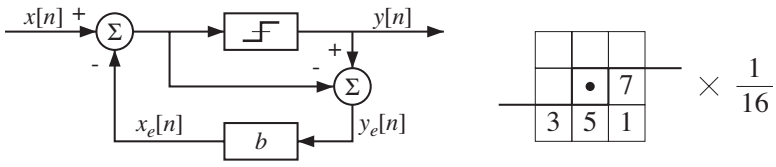

Figure 5: The (left) error diffusion algorithm and (right) four-weight error filter first introduced by Floyd and Steinberg.

of minority pixels with increasing r, and (c) a frequent occurrence of the inter-point distance λ_b, the principle wavelength, indicated by a series of peaks at integer multiples of λ_b. In Fig. 4 (left), the principle wavelength is indicated by a small diamond located along the horizontal axis.

Turning to the spectral domain, the spectral characteristics of blue-noise in terms of $P_\rho(f_\rho)$ are shown in Fig. 4 (right) and can be described by three unique features: (a) little or no low frequency spectral components, (b) a flat, high frequency (blue-noise) spectral region and (c) a spectral peak at cutoff frequency f_b, the blue-noise *principle frequency*, such that:

$$f_b = \begin{cases} \sqrt{g}/D & , \text{ for } 0 < g \leq 1/2 \\ \sqrt{1-g}/D & , \text{ for } 1/2 < g \leq 1 \end{cases} . \tag{6}$$

Indicated in Fig. 4 (right) by a diamond located along the horizontal axis is the principle frequency, and please note that $P(f_\rho)$ is plotted in units of $\sigma_g^2 = g(1-g)$, the variance of an individual pixel in the subject dither pattern.

3.1 Error Diffusion

Error-diffusion is depicted in Fig. 5 (left) and is a neighborhood operation that quantizes the current input pixel and then transfers the quantization error onto future input pixels. Formally, Floyd and Steinberg [12] define the output pixel $y[n]$ by adjusting and thresholding the input pixel $x[n]$ such that:

$$y[n] \;=\; \begin{cases} 1 & , \text{ if } (x[n] + x_e[n]) \geq 0 \\ 0 & , \text{ else} \end{cases} \tag{7}$$

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

$$x_e[n] \;=\; \sum_{i=1}^{M} b_i \cdot y_e[n-i] \tag{8}$$

Figure 6: Gray-scale images halftoned using (left) Floyd and Steinberg's original error diffusion algorithm and (right) Ulichney's perturbed filter weight scheme.

with $y_e[n] = y[n] - (x[n] + x_e[n])$ and $\sum_{i=1}^{M} b_i = 1$. Shown in Fig. 5 (right) is the original four weight error filter specified by Floyd and Steinberg, which was selected because of the checker board pattern it created at gray level $g = \frac{1}{2}$.

Looking at the spatial and spectral characteristics of the resulting dither patterns shown in Figs. 6 (left), Ulichney noted that patterns exhibited (i) correlated artifacts most noticeable at intensity levels $\mathcal{I} = \frac{1}{4}$, $\frac{1}{3}$, and $\frac{1}{2}$, (ii) directional hysteresis due to the order in which pixels were processes (raster order), and (iii) transient behavior near edges or boundaries. In order to improve the visual quality of dither patterns produced by error diffusion, Ulichney experimented with modifications to the algorithm that included changing the raster order, using different error filters, adding a white-noise component to the quantization threshold, and perturbing the error filter coefficients at each pixel.

With regards to the order that pixels are processed, Ulichney introduced the idea that

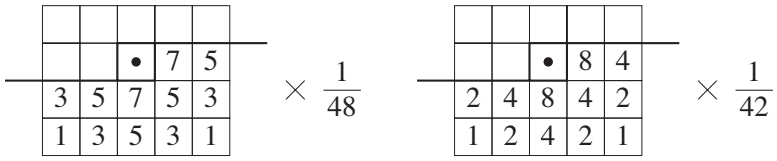

Figure 7: The (left) Jarvis *et al* and (right) Stucki 12-weight error filters.

a serpentine left-to-right and then right-to-left raster scan was far superior to the originally prescribed left-to-right raster. In other works, Witten and Neal [13] and Velho and Gomes [14] even employed chaotic space-filling curves. With regards to the error filter selection, Ulichney concluded that the 12 weight error filters of Jarvis *et al* [15] and of Stucki [16], Fig. 7 (left) and (right), improved some shortcomings of the original 4 weight filter of Floyd and Steinberg but still showed some disturbing or noticeable artifacts at intensity levels $\frac{1}{4}$ and $\frac{1}{2}$ resulting in significant energy below f_b. With regards to threshold modulation, Ulichney experimented with varying the quantization threshold at each pixel by low-variance white-noise and found that while this tended to relax disturbing patterns it did so at the expense of added low frequency textures. Beyond that SNR, dither patterns became too uncorrelated and grainy (white). As a side note, Eschbach and Knox [17] later determined that perturbing the quantization threshold as equivalent to adding low-variance white-noise to the input image prior to halftoning.

For perturbing error filter weights, Ulichney proposed pairing filter weights of similar magnitude with a noise component proportional to a percentage of the smaller weight added to one weight and subtracted from the other within the pair. This guaranteed that the sum of error weights would always be equal to 1, and it was this scheme of perturbing filter weights that lead to arguably the best error diffusion scheme where Ulichney added 50% noise to Floyd and Steinberg's original 4 weight error filter and a serpentine raster as demonstrated in Fig. 6 (right). Figure 8 shows the corresponding spatial and spectral characteristics of resulting dither patterns with all metrics exhibited improved blue-noise characteristics relative to those produced by Floyd and Steinberg's error diffusion.

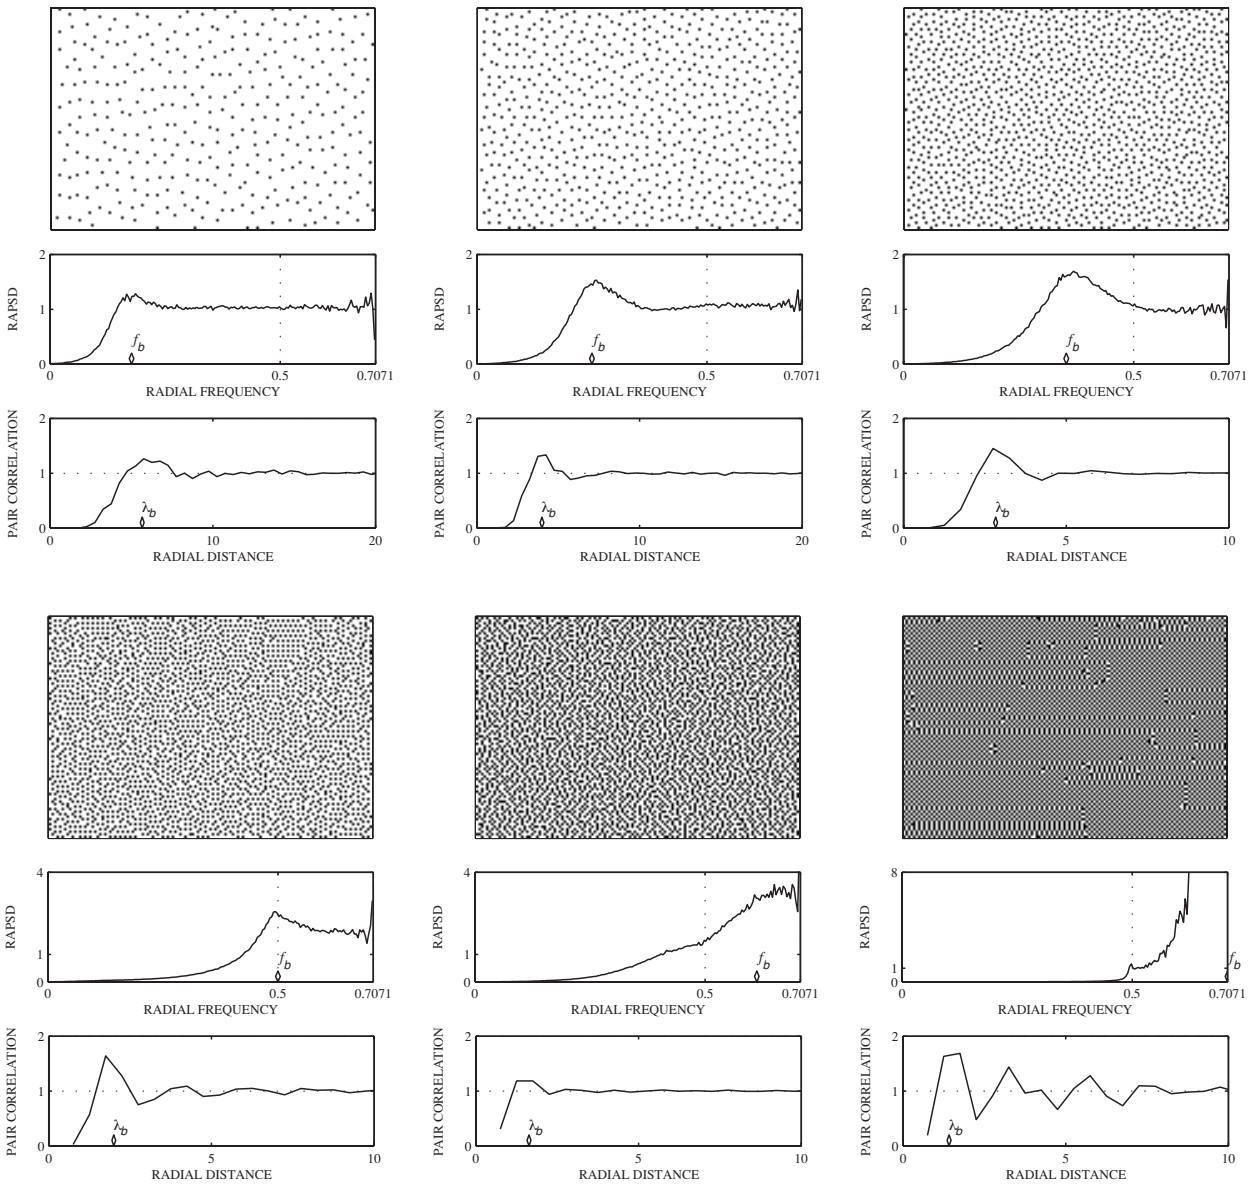

Figure 8: The spatial and spectral statistics for Ulichney's perturbed filter weight scheme at intensity levels $\mathcal{I} = \frac{1}{32}$, $\frac{1}{16}$, $\frac{1}{8}$, $\frac{1}{4}$, $\frac{1}{3}$, and $\frac{1}{2}$.

3.2 Blue-Noise Dither Arrays

Blue-noise dithering can also be achieved with the point process of ordered dither. The trick of course is using an appropriate dither array. Because of the implementation advantages of ordered dither over neighborhood processes, this has become an active area of research. In the printing industry, ordered dither arrays used for this purpose are often referred to as "stochastic screens." An overview of approaches to generating blue-noise dither templates

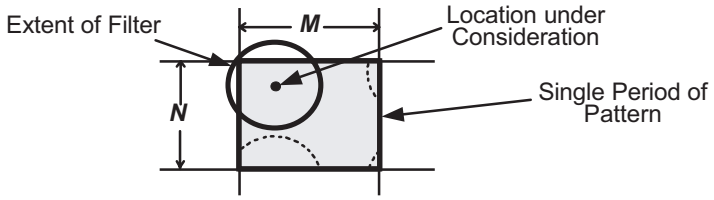

Figure 9: Two-dimensional periodicity of ordered dither pattern, and wrap-around property of void-and-cluster-finding filters.

is presented in [18]. One approach would be to build a template by directly shaping the spectrum of binary patterns by an iterative process so as to force blue-noise characteristics [11]. Mitsa and Parker [11] introduced the BInary Pattern Spectral density Manipulation Algorithm (BIPPSMA) as a process that builds an ordered dither threshold array by directly manipulating the frequency domain shape of the candidate binary pattern. They referred to the resulting ordered dither array as a "Blue Noise Mask."

A very straightforward and effective approach to generating relatively small dither templates of this type is the Void-and-Cluster algorithm [19], and will be outlined here. As with all ordered dither, the array and resulting binary patterns are periodic. Figure 9 illustrates this periodicity. This algorithm looks for voids and clusters in prototype binary patterns by applying a void- or cluster-finding filter at the area under consideration. Because of this implied periodicity, a filter extent will effectively wrap around as shown.

A void-finding filter considers the neighborhood around every majority pixel in a prototype binary pattern, and a cluster-finding filter considers the neighborhood around every minority pixel. The algorithm uses these filters to identify the biggest void or tightest cluster in the pattern. We start by relaxing an arbitrary initial pattern to form one that is homogeneously distributed. In Fig. 10 (a) a 16×16 binary pattern is shown with 26 minority pixels randomly positioned. The purpose of the algorithm is to move minority pixels form tight clusters into large voids. With each iteration the voids should be smaller and the clusters looser. This is done one pixel move at a time until both the voids and stop getting smaller, and the clusters stop getting looser. It turns out that the condition of convergence is quite simple; processing is complete when removing the pixel from the tightest cluster creates the largest void.

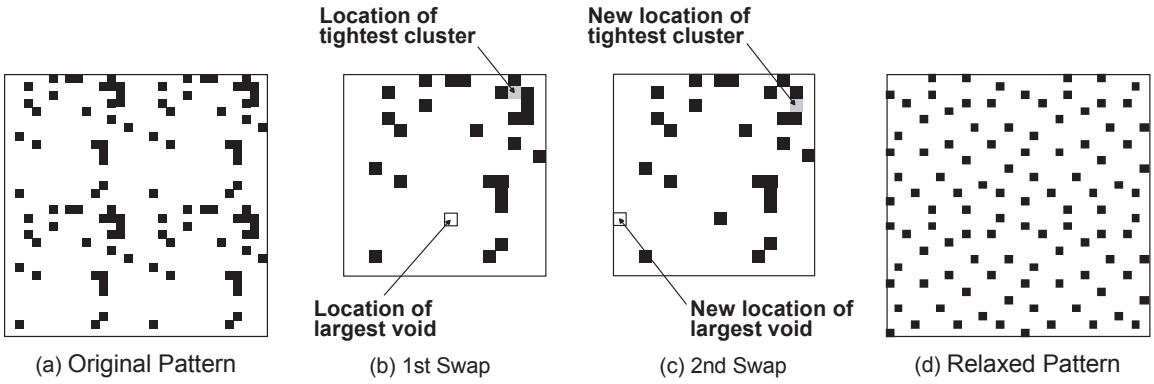

Figure 10: The Void-and-Cluster algorithm showing (a) four periods of an initial 16×16 pattern with 26 minority pixels, (b and c) the first two iterations of the initial binary pattern generator, and (d) four periods of the resulting rearranged, or relaxed, pattern after 12 iterations.

The minority pixel in the tightest cluster, and the majority pixel in the largest void are identified in Fig. 10 (b). After the first iteration, the minority pixel in the tightest cluster is moved to the largest void, resulting in the pattern shown in Fig. 10 (c). Once again, the locations of the new tightest cluster and new largest void are identified. It should be noted that it is entirely possible for minority pixels to be moved more than once; the search for voids and clusters at each iteration is independent of past moves.

The results of this example are summarized in Fig. 10 (d) where 12 iterations were needed before convergence. Four periods are shown of both the (a) input pattern, and (d) the relaxed or rearranged pattern to illustrate the wrap-around or edge-abutting consequences of tiling twospace with such patterns. Note how homogeneously distributed the resulting pattern is. Next, starting with this relaxed pattern as a starting point, a dither template is ordered in parallel. Elements of increasing value in the dither template are entered as minority pixels are inserted into the voids. Then returning to this starting pattern, elements of decreasing value are entered as minority pixels are removed from the tightest clusters.

Figure 11 shows the result of dithering an image with a 64×64 void-and-cluster generated dither array. It should be noted that the image does not appear as sharp as those produced by error diffusion. As mentioned earlier, the added run-time complexity of error diffusion does afford the side benefit of serving as a sharpening filter, even if uncontrollable. As will be shown in the next section, a prefilter as part of a rendering system can make up for this.

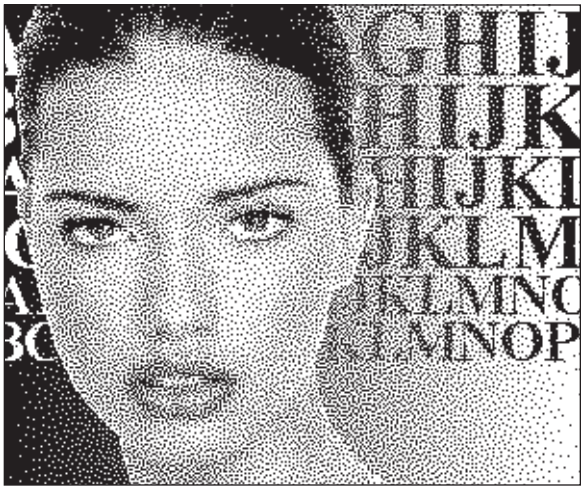

Figure 11: Dithering with a 64×64 void-and-cluster array.

It is interesting to note that if a completely empty (all white) pattern is used as a starting point, this algorithm will generate recursive tessellation dither templates. This will in fact also result if the starting point is any of the recursive tessellation patterns.

4 Green-Noise Halftoning

Just as blue-noise is the high frequency component of white-noise, green-noiseis the mid-frequency component that, like blue, benefits from aperiodic, uncorrelated structure without low frequency graininess. But unlike blue, green-noise patterns exhibit clustering (a collection of consecutive four-neighborhood pixels all of the same value). The result is a frequency content that lacks the high frequency component characteristic of blue-noise. Hence the term "green". Furthermore, green-noise forms aperiodic patterns that are not necessarily radially symmetric. Since the contrast sensitivity function of the human visual system is not radially symmetric, we allow green-noise to have asymmetric characteristics. The objective is to combine the maximum dispersion attributes of blue-noise with that of clustering of periodic clustered-dot halftone patterns.

Point process statisticians have long described cluster processes such as those seen in green-noise by examining the cluster process in terms of two separate processes: (i) the *parent process* that describes the location (centroid) of clusters, and (ii) the *daughter process* that describes the shape of clusters. In periodic clustered-dot halftoning, clusters are placed

along a regular lattice, and therefore, variations in periodic clustered-dot patterns occur in the cluster shape. In FM halftoning, cluster shape is deterministic, a single pixel. It is the location of clusters that is of interest in characterizing FM patterns. Green-noise patterns, having variation in both cluster shape and cluster location, require an analysis that looks at both the parent and daughter processes.

Looking first at the parent process Φ_p, ϕ_p represents a single sample such that $\phi_p = \{x_i : i = 1, \ldots, N_c\}$ where N_c is the total number of clusters. For the daughter process Φ_d, ϕ_d represents a single sample cluster of Φ_d such that $\phi_d = \{y_j : j = 1, \ldots, M\}$ where M is the number of minority pixels in cluster ϕ_d. By first defining the translation or shift in space $T_x(B)$ of a set $B = \{y_i : i = 1, 2, \ldots\}$ by x, relative to the origin, as:

$$T_x(B) = \{y_i - x : i = 1, 2, \ldots\} \tag{9}$$

and then defining ϕ_{d_i} as the ith sample cluster for $i = 1, \ldots, N_c$, a sample ϕ_G of the green-noise halftone process Φ_G is defined as:

$$\phi_G = \sum_{x_i \in \phi_p} T_{x_i}(\phi_{d_i}) = \sum_{x_i \in \phi_p} \{y_{ji} - x_i : j = 1, \ldots, M_i\}, \tag{10}$$

the sum of N_c translated clusters. The overall operation is to replace each point of the parent sample ϕ_p, of process Φ_p, with its own cluster ϕ_{d_i}, of process Φ_d.

In order to derive a relationship between the total number of clusters, the size of clusters, and the gray level of a binary dither pattern, I_g is defined as the binary dither pattern resulting from halftoning a continuous-tone discrete-space monochrome image of constant gray level g, and $I_g[n]$ is defined as the binary pixel of I_g with pixel index n. From the definition of $\phi(B)$ as the total number of points of ϕ in B, $\phi_G(I_g)$ is the scalar quantity representing the total number of minority pixels in I_g, and $\phi_p(I_g)$ is the total number of clusters in I_g with $\phi_p(I_g) = N_c$. The intensity, $\mathcal{I}$, being the expected number of minority pixels per unit area can, now, be written as:

$$\mathcal{I} = \frac{\phi_G(I_g)}{N(I_g)} = \begin{cases} g & , \text{ for } 0 < g \leq 1/2 \\ 1 - g & , \text{ for } 1/2 < g \leq 1 \end{cases}, \tag{11}$$

the ratio of the total number of minority pixels in I_g to $N(I_g)$, the total number of pixels composing I_g. Given (11), $\bar{M}$, the average number of minority pixels per cluster in I_g, is:

$$\bar{M} = \frac{\phi_G(I_g)}{\phi_p(I_g)} = \frac{\mathcal{I} \cdot N(I_g)}{\phi_p(I_g)}, \tag{12}$$

the total number of minority pixels in I_g divided by the total number of clusters in I_g.

Although obvious, (12) shows the very important relationship between the total number of clusters, the average size of clusters and the intensity for I_g. Periodic clustered-dot halftoning is the limiting case where $\phi_p(I_g)$ is held constant for varying $\mathcal{I}$, while FM halftoning is the limiting case where $\bar{M}$ is held constant for varying $\mathcal{I}$. In addition, (12) says that the total number of clusters per unit area is proportional to $\mathcal{I}/\bar{M}$. For isolated minority pixels (blue-noise), the square of the average separation between minority pixels (λ_b) is inversely proportional to $\mathcal{I}$, the average number of minority pixels per unit area [1]. By determining the proportionality constant using $\lambda_b = \sqrt{2}$ for $\mathcal{I} = \frac{1}{2}$, the relationship between λ_b and $\mathcal{I}$ is determined as $\lambda_b = D/\sqrt{\mathcal{I}}$.

In green-noise, it is the minority pixel clusters that are distributed as homogeneously as possible, leading to an average separation (center-to-center) between clusters (λ_g) whose square is inversely proportional to the average number of minority pixel clusters per unit area, $\mathcal{I}/\bar{M}$. Using the fact that $\lim_{M\to 1} \lambda_g = \lambda_b$, the proportionality constant can be determined such that λ_g is defined as:

$$\lambda_g = \begin{cases} D/\sqrt{(g)/\bar{M}} & , \text{ for } 0 < g \leq 1/2 \\ D/\sqrt{(1-g)/\bar{M}} & , \text{ for } 1/2 < g \leq 1 \end{cases} , \tag{13}$$

the *green-noise principle wavelength*. This implies that the parent process, ϕ_p, is itself a blue-noise point process with intensity $\mathcal{I}/\bar{M}$.

If we assume a stationary and isotropic green-noise pattern, the pair correlation will have the form of Fig. 12 (left) given that:

1. Daughter pixels, on average, will fall within a circle of radius r_c centered around a parent point such that $\pi r_c^2 = \bar{M}$ (the area of the circle with radius r_c is equal to the average number of pixels forming a cluster).

2. Neighboring clusters are located at an average distance of λ_g apart.

3. As r increases, the influence that clusters have on neighboring clusters decreases.

The result is a pair correlation that has: (a) a non-zero component for $0 \leq r < r_c$ due to clustering, (b) a decreasing influence as r increases, and (c) peaks at integer multiples of

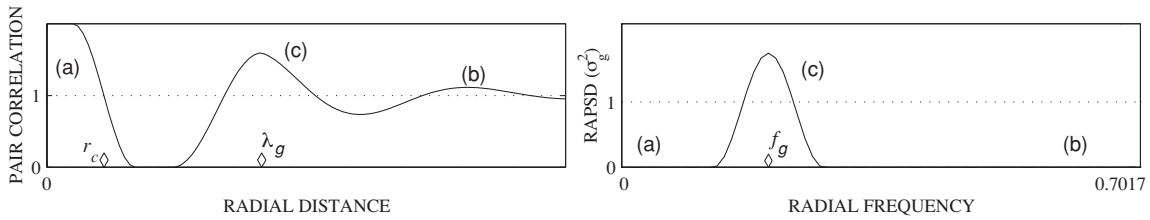

Figure 12: The (left) pair correlation and (right) radially averaged power spectrum for a green-noise process.

λ_g indicating the average separation of pixel clusters. Note that the parameter r_c is also indicated by a diamond placed along the horizontal axis in Fig. 12 (left).

In the case of stationary and anisotropic green-noise patterns, the pair correlation will also be of the form of Fig. 12 (left), but because clusters are not radially symmetric, blurring occurs in $\mathcal{R}(r)$ near the cluster radius r_c. In a similar fashion because the separation between clusters will also vary with direction, blurring will occur at each peak in $\mathcal{R}(r)$ located at integer multiples of λ_g. Assuming that the variation in cluster size is small for a given I_g, this placement of clusters λ_g apart leads to a strong spectral peak in $P(f_\rho)$ at $f_\rho = f_g$, the *green-noise principle frequency*:

$$f_g = \begin{cases} \sqrt{(g)/\bar{M}}/D & , \text{ for } 0 < g \leq 1/2 \\ \sqrt{(1-g)/\bar{M}}/D & , \text{ for } 1/2 < g \leq 1 \end{cases} . \tag{14}$$

From (14) we make several intuitive observations: (i) as the average size of clusters increases, f_g approaches DC, and (ii) as the size of clusters decreases, f_g approaches f_b. Figure 12 (right) illustrates the desired characteristics of $P(f_\rho)$ for ϕ_G showing three distinct features: (a) little or no low frequency spectral components, (b) high-frequency spectral components that diminish with increased clustering, and (c) a spectral peak at $f_\rho = f_g$.

The sharpness of the spectral peak in $P(f_\rho)$ at the green-noise principle frequency is affected by several factors. Consider first blue-noise where the separation between minority pixels should have some variation. The wavelengths of this variation, in blue-noise, should not be significantly longer than λ_b as this adds low-frequency spectral components to the corresponding dither pattern I_g [1], causing I_g to appear more white than blue. The same holds true for green-noise with large variations in cluster separation leading to a spectral peak at $f_\rho = f_g$ that is not sharp but blurred as the variation in separation adds new

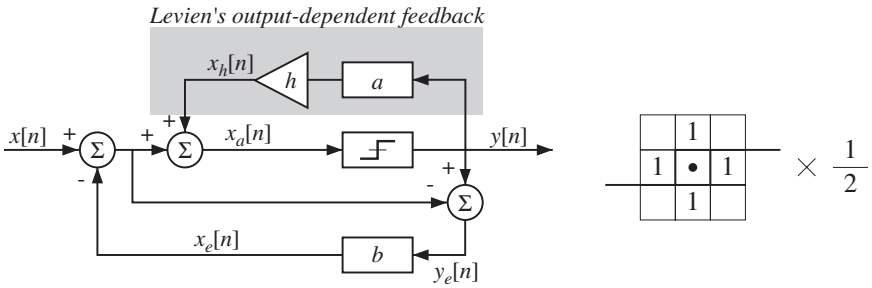

Figure 13: The (left) error diffusion with output-dependent feedback algorithm and (right) two-error and two-feedback coefficient filter first introduced by Levien.

spectral components to I_g. This whitening effect on I_g is also created by increased variation in the size of clusters with excessively large clusters leading to low-frequency components and excessively small clusters leading to high. In summary, the sharpest spectral peak at f_g will be created when I_g is composed of round (isotropic) clusters whose variation in size is small and whose separation between nearest clusters is also isotropic with small variation.

4.1 Error Diffusion with Output-Dependent Feedback

Like Floyd and Steinberg's error-diffusion, Levien's error-diffusion with output-dependent feedback (EDODF), shown in Fig. 13 (left), pre-dates the stochastic model describing the halftones that it creates. Here, the weighted sum of previous output pixels is added back to the accumulated pixel value, $x_a[n]$, such that:

$$y[n] \;=\; \begin{cases} 1 \;, & (x_a[n] = x[n] + x_e[n] + x_h[n]) \geq 0 \\ 0 \;, & \text{else} \end{cases} \tag{15}$$

where $x_h[n]$ is the hysteresis or feedback term defined as:

$$x_h[n] \;=\; h \sum_{i=1}^{N} a_i \cdot y[n-i] \tag{16}$$

with $\sum_{i=0}^{N} a_i = 1$ and h is an arbitrary constant. Referred to as the *hysteresis constant*, h acts as a tuning parameter with larger h leading to coarser output textures [3] as h increases ($h > 0$) or decreases ($h < 0$) the likelihood of a minority pixel if the previous outputs were also minority pixels. The result is a pattern with clustered 1s and 0s with larger hysteresis constants leading to coarser halftones as in Fig. 14 where the error/feedback coefficients are as shown in Fig. 13 (right) for $h = 0.5$, 1.0, 2.0, and 3.0.

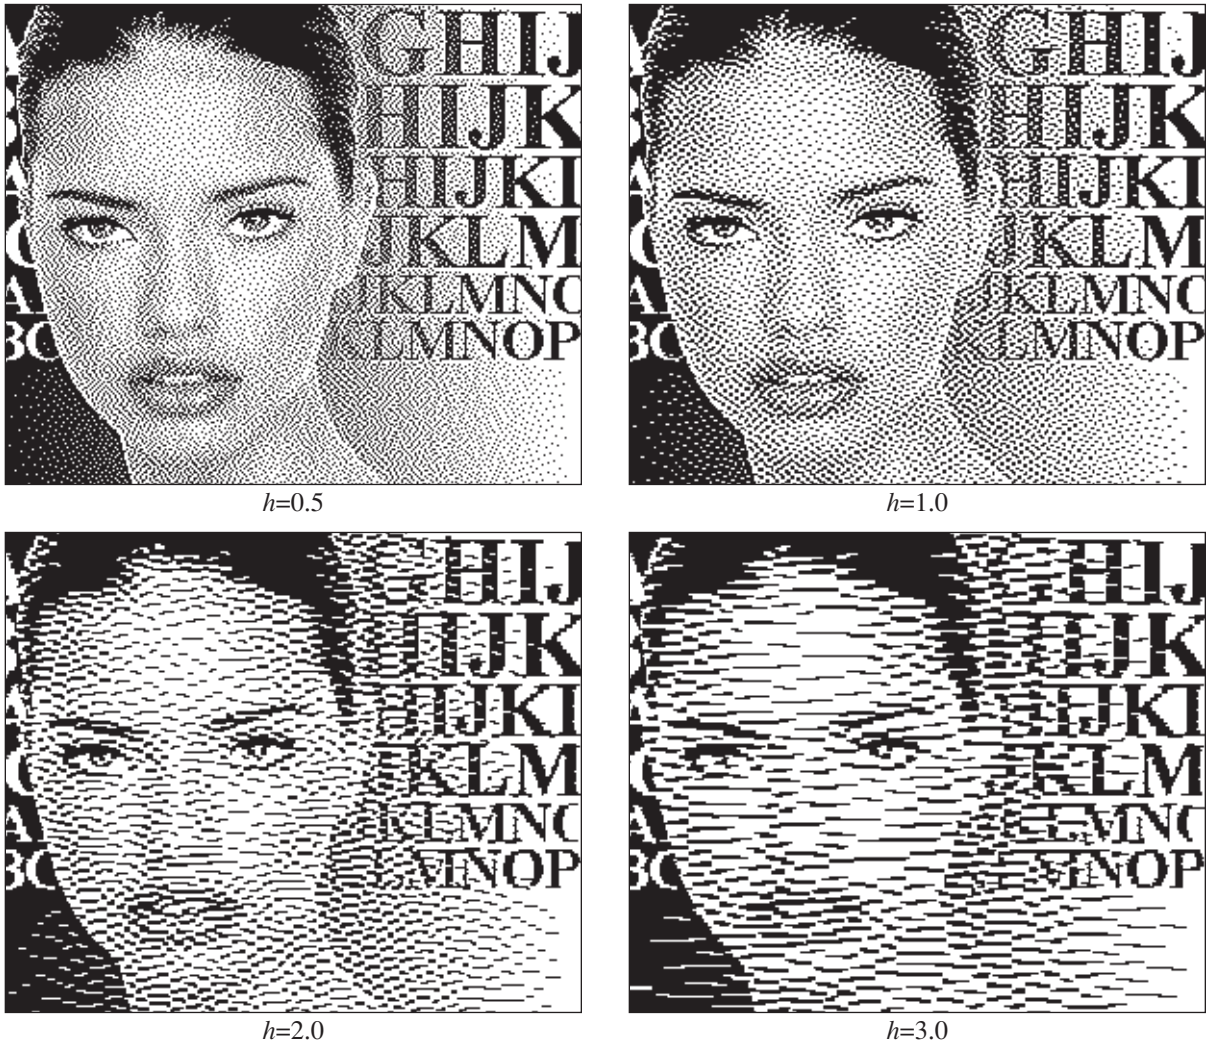

Figure 14: The image *Adrian* halftoned using Levien's error-diffusion with output-dependent feedback for hysteresis values $h = 0.5$, 1.0, 2.0, and 3.0.

Looking at the dither patterns created at extreme values of h near 0 and 3.0, Lau *et al* noted that patterns exhibited strong anisotropic features with clusters become too elongated along either the vertical ($h = 0$) or horizontal axis ($h = 3.0$). In responses to these artifacts, Lau *et al* looked at applying the various modifications first proposed by Ulichney and found that combining Floyd and Steinberg's 4 weight filter for feedback and Stucki's 12 weight filter for error with 30% perturbation on each filter gave acceptable results that broke up the long clusters formed at extreme values of h above 2.0. But in a later paper, Lau and Arce [20] looked at changing the proportion of feedback from the horizontal and vertical feedback weights such that they could improve the radial symmetry of patterns at the various values

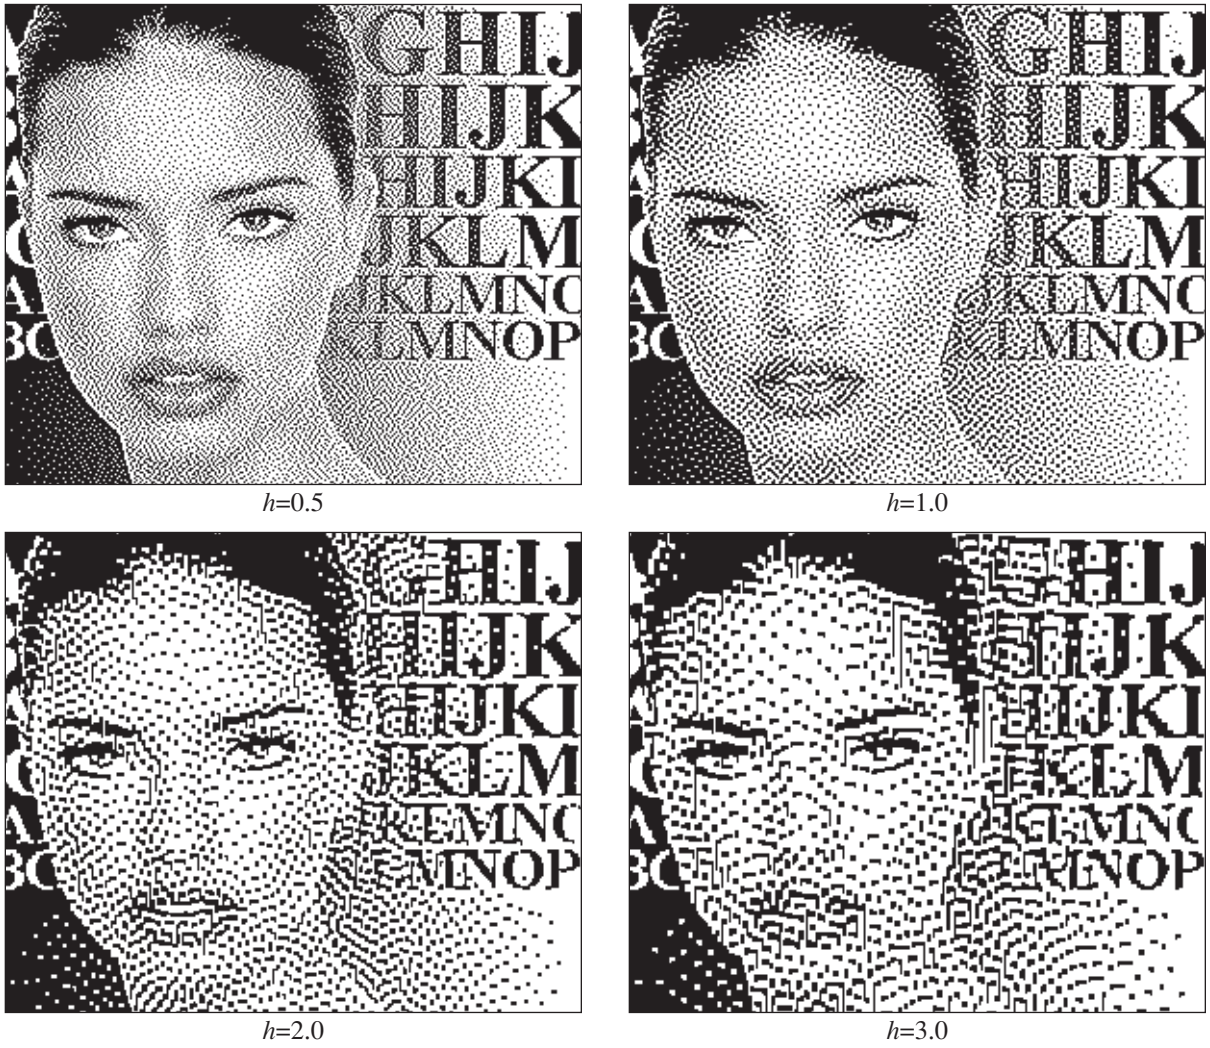

Figure 15: The image *Adrian* halftoned using Levien's error-diffusion with output-dependent feedback for hysteresis values $h = 0.5$, 1.0, 2.0, and 3.0 using balanced weights.

of h, and in so doing, they specified the optimal values of a_1 and a_2 versus h. Shown in Fig. 15 are the halftone images produced using these balanced weights for $h = 0.5$, 1.0, 2.0, and 3.0, and shown in Fig. 16 are the spatial and spectral statistics for $h = 2.5$. As illustrated, optimizing a_1 and a_2 to h greatly improves the radial symmetry of resulting pattern. As such, Lau and Arce even suggested employing an adaptive hysteresis/feedback parameter that changes either according to the gray level of the current input pixel (tone dependent hysteresis) or according to the local frequency content (frequency dependent hysteresis).

In the case of tone dependent hysteresis, Lau *et al* note that optimizing EDODF for a given printing process is achieved by specifying the parameter h according to the desired

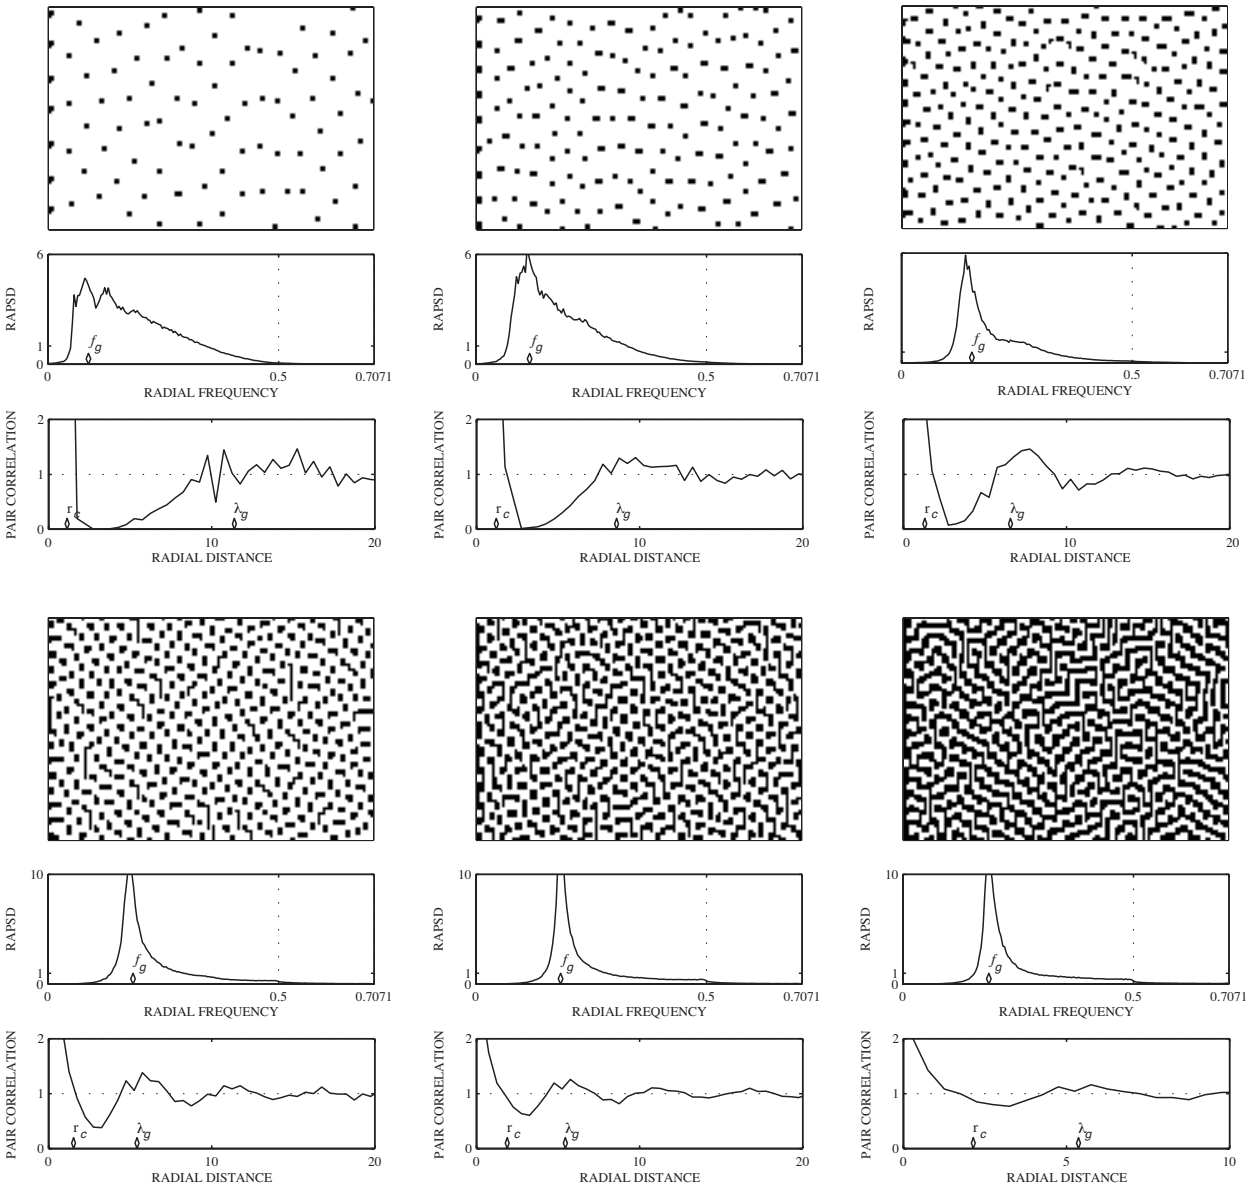

Figure 16: The spatial and spectral statistics for Levien's error-diffusion with output-dependent feedback using balanced weights with $h = 2.5$ at intensity levels $\mathcal{I} = \frac{1}{32}$, $\frac{1}{16}$, $\frac{1}{8}$, $\frac{1}{4}$, $\frac{1}{3}$, and $\frac{1}{2}$.

robustness, but as a constant, EDODF may, like periodic clustered-dot halftoning, sacrifice spatial resolution at certain gray-levels for pattern robustness at other levels. And it may, therefore, by advantageous to employ an adaptive hysteresis parameter that varies according to the input gray-level where, for each gray-level, h is the minimum value such that the output tone is within a pre-specified tolerance of the input. Further improvements in spatial resolution may be gained by varying h according to the frequency content of the input

image. In this scheme, the resulting halftoned image will be composed of large clusters in DC regions where distortions are most noticeable to the human eye and small clusters near edges where distortions are least noticeable and spatial details require small clusters in order to be preserved.

4.2 Green-Noise Dither Arrays

Having a statistical model describing the spatial and spectral properties of visually pleasing dither patterns that cluster minority pixels, Lau *et al* [21] introduced the BInary Pattern Pair Correlation Construction Algorithm (BIPPCCA) as an iterative process for constructing dither patterns that mimic green-noise. By specifying a pair correlation shaping filter, BIPPCCA can construct dither patterns with arbitrary pair correlations and, in the case of green-noise, with varying coarseness. Like BIPPSMA for blue-noise, BIPPCCA also represented a major milestone for the green-noise model because it was the first direct application of the ideal spatial characteristics of green-noise as BIPPCCA iterative adds points to a dither pattern such that the resulting pattern has a pair correlation matching that of the ideal green-noise pattern.

Techniques that indirectly constructed green-noise patterns include Velho and Gomes' digital halftoning along space filling curves (SFC) [14] and Scheermesser and Bryngdahl's digital halftoning with texture control [22]. SFC is a technique where a two-dimensional image is halftoned using a one-dimensional clustered-dot dithering approach that traverses the image along a space filling curve such as the Peano, Hilbert or Sierpinsky curves. By manipulating the maximum number of pixels that can form a cluster, the SFC technique can control the amount of coarseness in resulting images, and unlike periodic clustered-dot halftoning where the maximum number of pixels that can form a cluster limits the number of gray levels that the pattern can represent, SFC diffuses quantization error from one cluster to the next. The result is a technique that combines the benefits of aperiodic structure with those of clustered dots.

Scheermesser and Bryngdahl's technique attempts to minimize the cost associated with a particular arrangement of dots by iteratively turning pixels "on" and "off". The cost associated with a particular arrangement of dots is determined by two factors. The first is

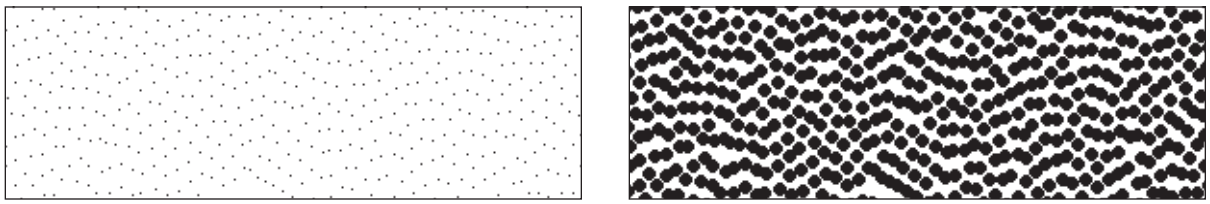

Figure 17: The binary halftones produced using (left) blue-noise for $\mathcal{I} = \frac{1}{64}$ and the (right) corresponding green-noise pattern created by replacing each point of the blue-noise pattern with a round disk cluster.

an image metric that measures the difference between the perceived images of the binary halftone pattern and the continuous tone original. The second cost is a numerical texture metric that measures the relative orientation of minority pixels. Scheermesser's and Bryngdahl's technique is able to produce adjustable coarseness by adjusting the weight of the cost of the texture metric versus the perceived image metric. But neither this or SFC try to address the specific characteristics defined by Lau *et al* for visually pleasing green-noise.

In the case of Velho and Gomes, patterns exhibit an unreasonable amount of white-noise content that imparts a distracting texture onto the image, while patterns constructed by Scheermesser and Bryngdahl show too little variation in cluster size/shape as patterns adopt an artificial texture created by the texture metric. In the green-noise model, there exists an inherent relationship between the cluster and the arrangement of clusters such that the shape of clusters varies according to the arrangement of nearest neighbors to fill in those regions where clusters are far apart. Clusters should appear more like pieces of a puzzle. If, instead, minority pixels of a blue-noise pattern were simply replaced with a single pre-determined cluster, visually pleasant textures would only result from well formed blue-noise patterns. Any imperfections in the blue-noise pattern such as a smaller average spacing between points along the horizontal axis would become obvious if the cluster shape was not similarly shaped to be narrower along the horizontal. Figure 17 demonstrates the phenomenon where an otherwise well formed blue-noise pattern shows its disparities when convolved with a round disk cluster. While Fig. 17 (left) has it's problems with hysteresis, these defects are only amplified in appearance in Fig. 17 (right) after replacing each minority pixel with its own cluster. A good green-noise process would change the shape of clusters to fit between neighboring clusters and would, hence, compensate for any defects in the parent

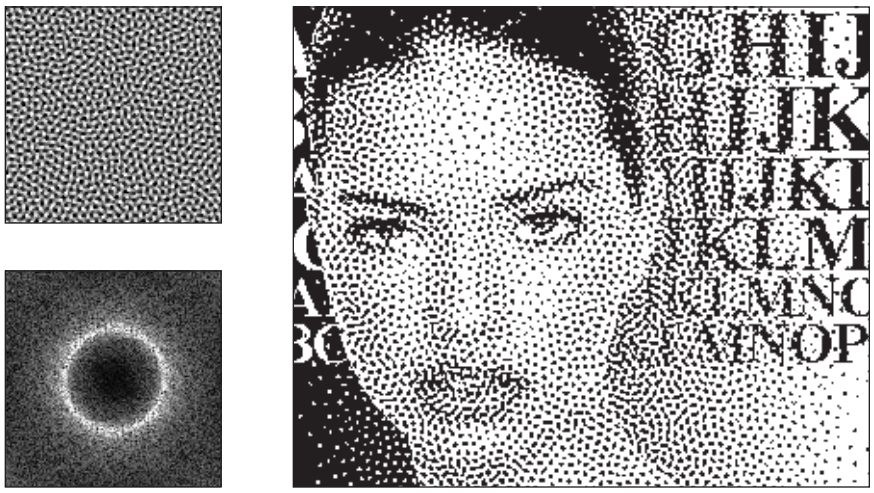

Figure 18: A (top-left) 128×128 green-noise dither array along with (bottom-left) the corresponding power spectrum showing the predominantly mid-frequency content (DC point at the center of the figure) and (right) the corresponding halftone image.

process.

From BIPPCCA, Lau *et al* [21, 23] constructed the first green-noise dither arrays that, like blue-noise dither arrays, convert a continuous-tone image to binary using a pixelwise thresholding operation between a pixel in the original image and the corresponding pixel within the array. In [23], multiple dither arrays were constructed in a correlated fashion to regulate the amount of dot overlap in color halftones. Shown in Fig. 18 are several example monochrome green-noise dither arrays along with the power spectrums produced using each that show the predominantly mid-frequency (green) spectral content. In the case of Fig. 18 (right most), this dither array was constructed to be asymmetric such that the average size of clusters for black clusters on a white background is smaller than for white clusters on a black background. Such a dither array addresses the varying behavior of printed dots on either side of $g = \frac{1}{2}$ and, in the power spectrum, results in two separate principle frequencies as illustrated.

5 Conclusions: Blue versus Green-Noise

In devices that can robustly accommodate isolated pixels, such as various video displays and many ink-jet printers, blue noise is the preferred halftoning technique. However, for

the many print marking engines that cannot robustly accommodate isolated pixels various factors must be taken into account including the visual pleasantness and edge detail of the halftone, the variation in the printed dot and the distortion introduced by the printing device, and in the case of color the interaction of the overlapping patterns/ink. It is also important to remember that we need not consider the decision as choosing between blue or green-noise, but we should, instead, focus on the tunability of green-noise to range from a fine blue-noise to a coarse pattern closer to locally periodic clustered-dot and to optimize the coarseness of green-noise to a specific printing device.

5.1 Visual Pleasantness

Shown in Fig. 19 is a diagram illustrating the relationships between perceived resolution, a measure of visual pleasantness, and halftone robustness (resistance to printer distortion) for the three halftoning models of AM (a term sometimes used to for locally periodic classical screen halftoning), blue-noise, and green-noise halftoning. Given that the purpose of a binary dither pattern is to represent a continuous-tone level, a dither pattern should not have any form or structure of its own, and a pattern succeeds when it is innocuous. Blue-noise is visually pleasant because it does not clash with the structure of an image by adding one of its own or degrade it by being too "noisy" or uncorrelated. Blue-noise even defies the structure of the underlying grid such that even though the dots are perfect squares with each precisely aligned to a given position on a rectangular grid, the collective ensemble tends to destroy this rigorous alignment creating what can be called a grid defiance illusion [1].

In instances where because of printer distortion minority pixels must be clustered, the green-noise model has many benefits for printer distortion such as minimizing the perimeter-to-area ratio (see Sec.5.2), but primarily, it describes the spatial and spectral properties of the optimal halftones in terms of visual pleasantness. In its essence, the green-noise model describes the halftone pattern most like blue-noise under the constraint that minority pixels must be clustered with an average cluster size greater than one pixel. At one pixel, the green-noise model is equivalent to blue, and it is, therefore, said that green-noise has blue-noise as a limiting case. On the opposite end of the coarseness spectrum, green-noise benefits from its stochastic arrangement of clusters, and because the eye is less sensitive to the artifacts

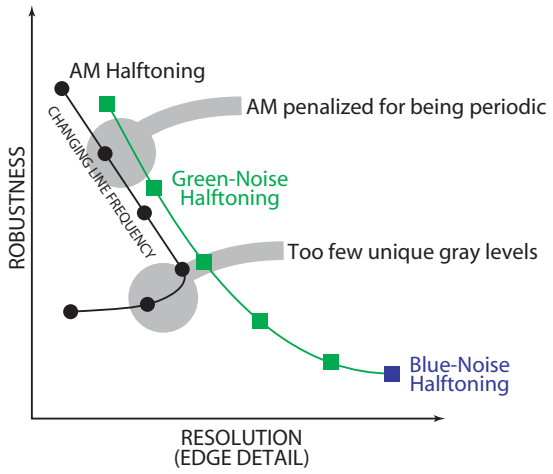

Figure 19: The relationships between spatial resolution and halftone robustness for the various halftoning paradigms for a fixed print resolution.

created by stochastic halftones, green-noise has better visual fidelity than periodic clustered-dot halftones with the same average cluster size. Green-noise maintains the grid defiance illusion, and so in Fig. 19, the line representing green-noise is shown as always having a visual fidelity higher than AM at the same measure of robustness.

5.2 Printer Distortion

Due to various distortions to the printed dot, the gray level produced by a printed halftone does not equal the ratio of black to white pixels in the dither pattern. As a means of correcting this disparity in gray level, a means of tone correction is applied such that each pixel of the input image with gray level g is replaced by a pixel with some gray level g' [14] where the mapping of gray levels from g to g' is determined by direct measurement of the input versus output reflectance curve for a given printer [24] or is estimated using a printed dot model [25]. The underlying assumption is that when trying to reproduce gray level g', the printer will consistently produce the apparent gray level g. The problem for tone correction occurs when the printer does not produce dot distortions consistently.

Instead of applying tone correction to correct for printer distortions, model based halftoning algorithms have been proposed that take into account a model of the printed dot to decide when and when not to print a dot in the halftone. Roetling and Holiday [26] offered one of the first examples of model based halftoning when they adjusted the thresholds of ordered dither

arrays based on the dot overlap created by the hard circular dot model. Pappas [27] also used the hard circular dot when he first included a dot model into error diffusion by introduced modified error diffusion. But the problem with many model based techniques is that they assume that dots are printed consistently by relying on an average dot size/shape, and by assuming reliability in the printing device, resulting patterns are typically computationally expensive tone-corrected blue-noise patterns [28]. The patterns are, therefore, inappropriate in devices that do not produce isolated pixels reliably such as in laser printers [29].

When looking at the problem of producing consistent and accurate tone, printer distortion is not, in general, considered a "bad" thing, and its occurrence, either high or low, does not limit the choice in halftoning techniques for a given printing process [30]. So blue-noise may be just as applicable as periodic clustered-dot halftoning for a given printer. Dot-loss can be bad if it limits the extent to that perfect black can be produced, but in many instances though, dot-loss only occurs for small, isolated black dots, and therefore, does not limit the choice in halftoning techniques either. What does limit the choice, in halftoning, is the repeatability of dot-gain/loss where if a printer consistently reproduces dots with little variation, accurate tone reproduction can be achieved through tone correction [31]. But in a printing process that is not repeatable, compensating for distortion is much more complicated as isolated dots are more sensitive to process variation [32]. This is clearly evident in laser printers, and hence the reason that these devices have relied upon periodic clustered-dot halftoning for so long.

Noting that periodic clustered-dot halftoning produces patterns with far less noticeable variation in tone in laser printers, it is far more advantageous, in unreliable printing devices, to use a halftoning scheme that resists distortion, making the output more robust to variations in the printing process [4]. How robustness is achieved is through clustering [33] as it is the perimeter-to-area ratio of printed dot clusters that is the most tell-tale characteristic to the impact of printer distortion on the halftone [31,32], and a technique like green-noise will always be more robust than blue with greater robustness achieved through increased clustering. Increased robustness is, therefore, indicated in Fig. 19 with AM halftoning providing the most robust patterns with green-noise bridge the gap with blue-noise.

5.3 Rectangular Versus Hexagonal Display Grids

This paper has, so far, assumed a symmetric rectangular (square) display grid, but examples do exist where the display device aligns pixels on a hexagonal sampling grid. In direct comparison, hexagonal sampling grids have the distinct advantage of being able to preserve a circular band-limited signal with 13.4% fewer samples per unit area than square grids [34, 35], and hexagonal sampling grids have been shown to produce samples with greater intersample dependency that allows "lost" samples to be more accurately recovered from its neighbors [36]. And in a binary display device such as an ink-jet printer that prints round dots that overlapping neighboring dots, Ulichney [24] showed that hexagonal girds have better covering efficiency (ratio of pixel area to printed dot area) that results in: (i) a more linear tone scale rendition, (ii) more similarly sized isolated black and white pixels, and (iii) less spectral overlap (aliasing) for circularly bandlimited images.

Hexagonal grids are also preferred over rectangular when halftoning is done using an order dither array since hexagonal grids can be arranged such that the maximum frequencies along the horizontal and vertical axes are greater than that of rectangular grids. Given the human visual system's increased sensitivity along the horizontal and vertical axes, this increase leads to less sensitivity to the regular patterns that ordered dithering produces. But when halftoning is done using blue-noise, hexagonal display grids have the distinct disadvantage of introducing aliasing artifacts for gray levels $\sqrt{1/3} < g < \sqrt{2/3}$ when the principal frequency exceeds $1/\sqrt{3}$, Fig. 20. While some level of relief from aliasing can be achieved by adding correlation along the symmetry axes in the spectral plane, aliasing is unavoidable above $\frac{1}{3}$. In a binary halftone, aliasing results in clustering and on a hexagonal sampling grid, some level of clustering must occur beyond $g = \frac{1}{3}$ as illustrating in Fig. 20 (right).

On a rectangular display grid, aliasing first occurs at intensity level $\mathcal{I} = \frac{1}{4}$, but by adding diagonal correlation between minority pixels, it is possible to move the spectral energy of a dither pattern into the corners of the spectral tile avoiding any overlap with neighboring tiles and thereby avoiding any clustering in the dither pattern, Fig. 21. So a technique like Floyd and Steinberg's original error diffusion scheme avoids clustering at gray level $g = \frac{1}{2}$ simply because the inventors specifically selected their error weights to create the regular

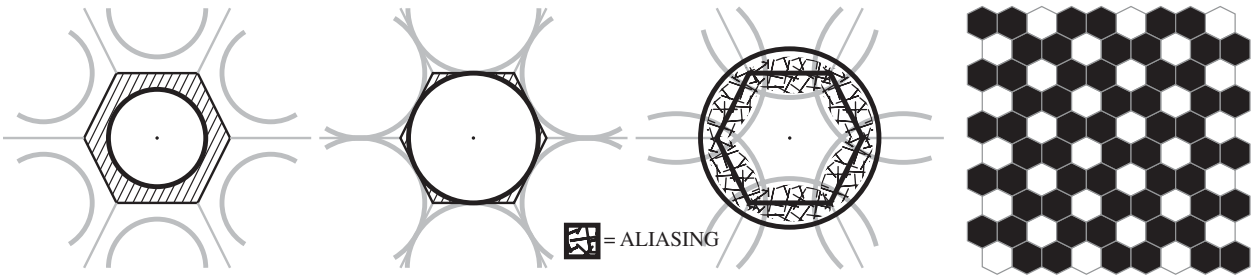

Figure 20: Diagram showing the introduction of aliasing artifacts in blue-noise halftones on a hexagonal display where (left) shows the power spectrum prior to aliasing, (center-left) shows the power spectrum at intensity level $\sqrt{1/3}$ where aliasing first occurs, (center-right) shows full aliasing at $\mathcal{I} \geq \frac{1}{3}$, and (right) shows the dither pattern at $g = \frac{1}{3}$ where no space exists such that white pixels can be added without introducing clusters.

checkerboard pattern. In contrast, Jarvis *et al* used a filter that created patterns that were much more radially symmetric, and this symmetry created the overlapping with neighboring spectral tiles that lead to the aliasing/clustering of minority pixels at gray levels near $g = \frac{1}{2}$.

In summary, all blue-noise halftoning leads to aliasing unless rectangular sampling grids are employed and the halftoning algorithm introduces enough diagonal correlation as to push power into the diagonals of the spectral plane. But by clustering minority pixels with sufficient coarseness, green-noise eliminates spectral overlap while also containing the spectral content of dither patterns to a narrow band of radial frequencies. Corresponding green-noise patterns will then appear less noisy than blue-noise patterns containing aliasing artifacts would appear. On hexagonal grids, the denser packing of the base band frequency in the spectral domain increases the amount of overlap for the rings formed by blue-noise halftoning algorithms. Blue-noise patterns will, therefore, suffer from greater degrees of aliasing on hexagonal displays. Aliasing that cannot be avoided by adding directional correlation. By employing green-noise, though, aliasing concerns are abated with green-noise patterns profiting from the advantages already associated with hexagonal grids over rectangular.

As already mentioned, 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 that rectangular sampling grids are preferable to hexagonal because only they support blue-noise

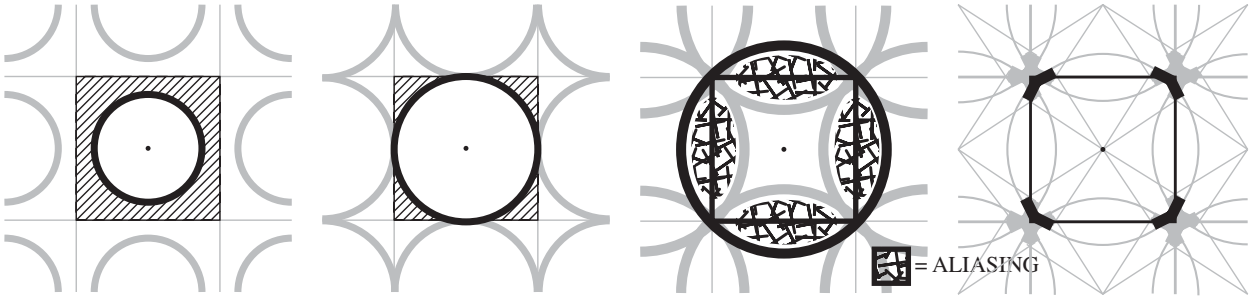

Figure 21: Diagram showing the introduction of aliasing artifacts in blue-noise halftones on a rectangular display where (left) shows the power spectrum prior to aliasing, (center-left) shows the power spectrum at intensity level $\frac{1}{4}$ where aliasing first occurs, (center-right) shows full aliasing at $\mathcal{I} = \frac{1}{2}$, and (right) shows the impact of adding diagonal correlation between minority pixels such that aliasing is avoided.

is invalid for many printing devices anyway. And the overall conclusion is that since green-noise halftoning 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.

5.4 Color Halftoning: Stochastic Moiré

In traditional color halftoning devices, the binary halftones of cyan, magenta, yellow, and black are superimposed [37], and in the case of periodic clustered-dot halftoning, the regular grids of the halftone screens interact to form periodic moiré patterns whose visibility are minimized by aligning the CMYK halftones to screen angles 15^o, 75^o, 0^o, and 45^o respectively [38, 39]. Because it is commonly believed that moiré is a result of superimposing regular patterns, stochastic halftones such as those produced by blue-noise are also believed to avoid the moiré phenomenon allowing for the introduction of low-cost color ink-jet printers [30]. But superimposing uncorrelated dispersed-dot patterns does lead to low-frequency graininess commonly referred to in halftoning literature as color-noise [40].

The low-frequency graininess created by superimposing dispersed-dot halftones is, in fact, a product of the same moiré phenomenon found in periodic clustered-dot halftones, but because the component colors are stochastic, the resulting moiré is an aperiodic texture referred to as stochastic moiré [41]. And the visibility of this moiré is its worst when overlapping halftones are uncorrelated and have the same principal wavelength. As a response

to stochastic moiré, considerable effort is being made to correlate the CMYK halftones such that patterns form either perfectly overlapping or perfectly inter-locking screens [39, 42]. A major problem that these overlapping and interlocking screens create is that they require a very high degree of control over the alignment of screens in the final print as even slight mis-registration can lead to dramatic shifts in color/texture.

When looking at the superposition of green-noise halftones, the visibility of stochastic moiré is increased by the lower frequency of the resulting moiré textures. While interlocking the CMYK screens minimizes the visibility of the extraneous textures, Lau *et al* [41,43] write that the visibility of stochastic moiré is minimized in uncorrelated screens by varying the coarseness between colors such that the probability of overlapping screens having the same principal wavelength is minimized, Fig. 22. Lau *et al* argue that the optimal coarseness of screens is ordered according to the luminance of each colorant such that, for a given printer, black has the smallest average cluster size followed by magenta, cyan, and then yellow with the largest. Because these overlapping green-noise screens minimize stochastic moiré visibility without correlation, the constraints regarding registration are greatly alleviated, and Lau *et al* further argue that, in cases where perfect registration cannot be guaranteed, green-noise halftoning is the preferred technique even in an ideal printing device where traditionally blue-noise would otherwise be considered optimal.

References

[1] R. A. Ulichney, "Dithering with blue noise," *Proceedings of the IEEE*, vol. 76, no. 1, pp. 56–79, 1988.

[2] F. W. Campbell, "The human eye as an optical filter," *Proceedings of the IEEE*, vol. 56, no. 6, pp. 1009–1014, June 1968.

[3] R. Levien, "Output dependant feedback in error diffusion halftoning," in *IS&T's Eighth International Congress on Advances in Non-Impact Printing Technologies*, Williamsburg, Virginia, USA, October 25-30 1992, pp. 280–282.

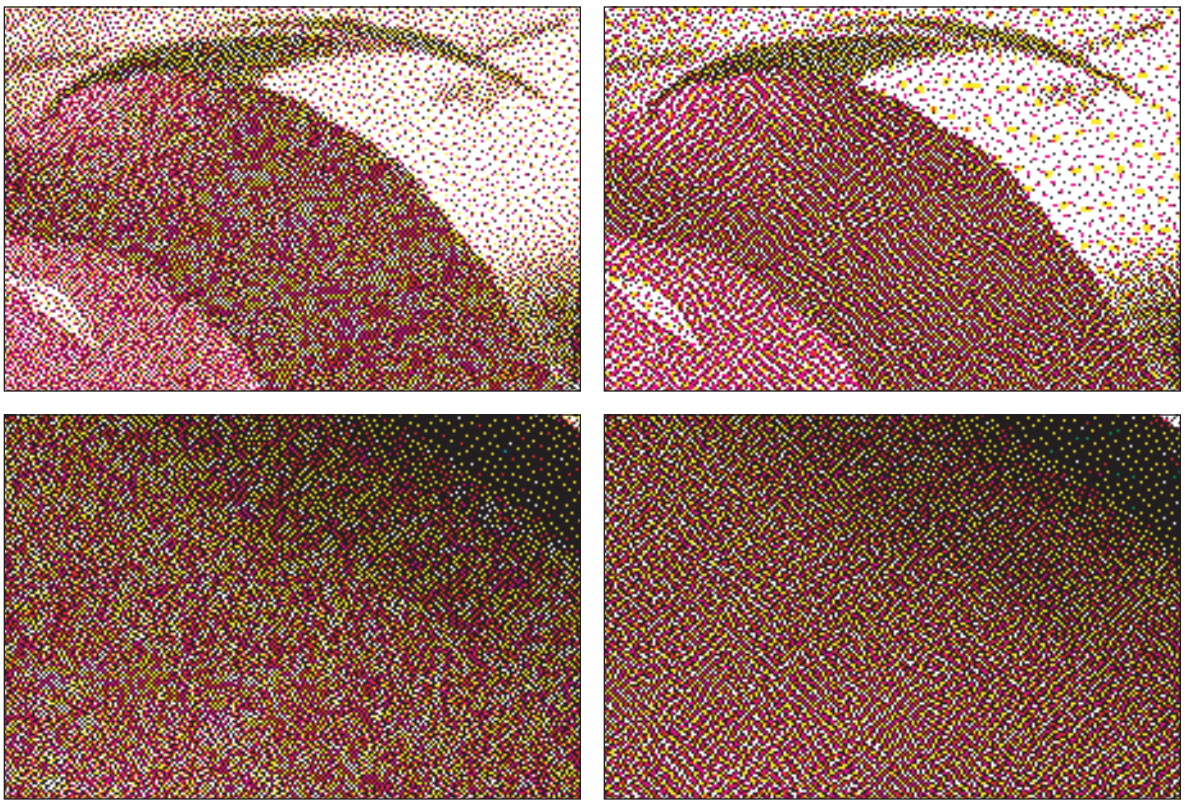

Figure 22: Color halftones corresponding to different images (top versus bottom) produced using (left) blue and (right) green-noise with no correlation between colors showing the varying effects of stochastic moiré where the patchiness of moiré is reduced using green-noise by varying the coarseness of patterns between colors.

[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] N. A. C. Cressie, *Statistics for Spatial Data*, John Wiley and Sons, New York, 1983.

[6] D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic Geometry and Its Applications*, John Wiley and Sons, New York, 1987.

[7] M. S. Bartlett, "The spectral analysis of a point process," *Journal of the Royal Statistical Society, Series B*, vol. 25, no. 2, pp. 264–280, 1964.

[8] M. S. Bartlett, "The spectral analysis of two-dimensional point processes," *Biometrika*, vol. 51, pp. 299–311, 1964.

[9] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing*, Prentice Hall, New York, 1975.

[10] M. Gardner, "White and brown music, fractal curves and one-over-f fluctuations," *Scientific America*, pp. 16–32, April 1978.

[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] 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.

[13] I. H. Whitten and M. Neal, "Using peano curves for bilevel display of continuous-tone images," *IEEE Computer Graphics and Applications*, pp. 47–52, May 1982.

[14] L. Velho and J. M. Gomes, "Digital halftoning with space filling curves," *Computer Graphics*, vol. 25, no. 4, pp. 81–90, 1991.

[15] 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.

[16] P. Stucki, "Mecca-a multiple-error correcting computation algorithm for bilevel image hardcopy reproduction," Tech. Rep. RZ1060, IBM Research Laboratory, Zurich, Switzerland, 1981.

[17] R. Eschbach and K. T. Knox, "Error-diffusion algorithm with edge enhancement," *Journal of the Optical Society of America*, vol. 8, no. 12, pp. 1844–1850, December 1991.

[18] K. E. Spaulding, R. L. Miller, and J. Schildkraut, "Methods for generating blue-noise dither matrices for digital halftoning," *Journal of Electronic Imaging*, vol. 6(2), 1997.

[19] 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., 1993, vol. 1913, pp. 332–343.

[20] D. L. Lau and G. R. Arce, "Robust halftoning with green-noise," in *IS&T's Image Processing Image Quality Image Capture Systems Conference (PICS'99)*, Savanah, Georgia, USA, April 26-28 1999.

[21] 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, no. 7, pp. 1575–1586, July 1999.

[22] T. Scheermesser and O. Bryngdahl, "Control of texture in image halftoning," *Journal of the Optical Society of America A*, vol. 13, no. 8, 1996.

[23] D. L. Lau, G. R. Arce, and N. C. Gallagher, "Digital color halftoning via generalized error-diffusion and vector green-noise masks," *IEEE Transactions on Image Processing*, Submitted February 1999.

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

[25] C. J. Rosenberg, "Measurement-based evaluation of a printer dot model for halftone algorithm tone correction," *Journal of Electronic Imaging*, vol. 2, no. 3, pp. 205–212, 1993.

[26] P. G. Roetling and T. M. Holladay, "Tone reproduction and screen design for pictorial electrophotographic printing," *Journal of Applied Photographic Engineering*, vol. 15, no. 4, pp. 179–182, 1979.

[27] 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 Jan P. Allebach, Eds., June 1991, vol. 1453, pp. 244–255.

[28] M. Yao and K. J. Parker, "Dot gain compensation in the blue noise mask," in *Proceedings of SPIE, Human Vision, Visual Processing, and Digital Display VI*, B. E. Rogowitz and J. P. Allebach, Eds., 1995, vol. 2411, pp. 221–227.

[29] S. Aoki, "New halftoning method using adaptive cell," in *IS&T's NIP 14: International Conference on Digital Printing Technologies*, Toronto, Ontario, Canada, October 18-23 1998, pp. 277–280.

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

[31] M. Rodriguez, "Promises and pitfalls of stochastic screening in the graphic arts industry," *IS&T's Eighth International Congress on Advances in Non-Impact Printing Technologies*, 1992.

[32] M. Rodriguez, "Graphic arts perspective on digital halftoning," in *Proceedings of SPIE, Human Vision, Visual Processing, and Digital Display V*, B. E. Rogowitz and J. P. Allebach, Eds., 1994, vol. 2179, pp. 144–149.

[33] K. T. Knox, "Introduction to digital halftones," in *Recent Progress in Digital Halftoning*, R. Eschbach, Ed., pp. 30–33. IS&T, 1994.

[34] R. M. Mersereau, "Two-dimensional signal processing from hexagonal rasters," in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, 1978, pp. 739–742.

[35] R. M. Mersereau, "The processing of hexagonally sampled two-dimensional signals," *Proceedings of the IEEE*, vol. 67, no. 6, pp. 930–949, 1979.

[36] R. J. Marks, "Multidimensional-signal sample dependency at nyquist densities," *Journal of the Optical Society of America A*, vol. 3, no. 2, pp. 268–273, 1986.

[37] R. W. G. Hunt, *The Reproduction of Color in Photography, Printing and Television*, Fountain Press, Tolworth, England, 1987.

[38] I. Amidror, R. D. Hersch, and V. Ostromoukhov, "Spectral analysis and minimization of moiré patterns in color separation," *Journal of Electronic Imaging*, vol. 3, no. 3, pp. 295–317, 1994.

[39] H. R. Kang, *Digital Color Halftoning*, SPIE-International Society for Optical Engineering, Bellingham, WA USA, 1999.

[40] J. S. Liu and F. H. Cheng, "Color halftoning - a non-separable model," *Proceedings of Internation Conference on Image Processing*, pp. 561–564, 1996.

[41] D. L. Lau, A. M. Khan, and G. R. Arce, "Minimizing stochastic moiré by means of green-noise masks," *Journal of the Optical Society of America A*, vol. 19, no. 11, pp. 2203–2217, November 2002.

[42] M. Yao and K. J. Parker, "Application of the blue-noise mask in color halftoning," in *Proceedings of SPIE, Visual Communications and Image Processing*, R. Ansari and M. J. Smith, Eds., 1996, vol. 2727, pp. 876–880.

[43] D. L. Lau, A. M. Khan, and G. R. Arce, "Stochastic moiré in color halftoning," in *IS&T's Image Processing Image Quality Image Capture Systems Conference (PICS'01)*, Montreal Quebec Canada, April 2001.