

Digital halftoning by means of green-noise masks

Daniel L. Lau, Gonzalo R. Arce, and Neal C. Gallagher

Department of Electrical and Computer Engineering, University of Delaware, Newark, Delaware 19716

Received August 11, 1998; revised manuscript received February 3, 1999; accepted February 16, 1999

We introduce a novel technique for generating green-noise halftones—stochastic dither patterns composed of homogeneously distributed pixel clusters. Although techniques employing error diffusion have been proposed previously, the technique here employs a dither array referred to as a green-noise mask, which greatly reduces the computational complexity formerly associated with green noise. Compared with those generated with blue-noise masks, halftones generated with green-noise masks are less susceptible to printer distortions. Because green noise constitutes patterns with widely varying cluster sizes and shapes, the technique introduced here for constructing these green-noise masks is tunable; that is, it allows for specific printer traits, with small clusters reserved for printers with low distortion and large clusters reserved for printers with high distortion. Given that blue noise is a limiting case of green noise, this new technique can even create blue-noise masks.

© 1999 Optical Society of America [S0740-3232(99)01807-4]

OCIS codes: 100.0100, 100.2810.

1. INTRODUCTION

Digital halftoning is the technique employed to convert images from continuous tone to binary. Desktop printers such as laser and ink-jet, as well as commercial planographic and screen presses capable of only binary output, rely on digital halftoning to produce the illusion of continuous tone. Before the advent of digital image setters, halftoning was done by using analog photomechanical screening methods, which were later imitated digitally through clustered-dot dithering.¹ Commonly referred to as amplitude-modulated (AM) halftoning, clustered-dot dithering produces the illusion of continuous tone by varying the size of round dots, which are printed along a regular lattice.

Widely-used today, AM halftoning suffers from several major drawbacks. The first is moiré, the interference pattern created by superimposing two or more regular patterns. Moiré occurs in the printing of color images, which requires superimposing cyan, magenta, yellow, and black screens. Minimizing the effects of moiré requires special care to properly offset the orientation of each screen—creating a desirable circular pattern called a rosette. The second drawback of AM halftoning is spatial resolution. By clustering the smallest indivisible printed dots (pixels) into larger macrodots, AM halftoning, in many cases, achieves a far lower spatial resolution in printed images than the resolution of the printer itself.

In 1976 Floyd and Steinberg² introduced error diffusion, which is today the foundation of frequency-modulated (FM) halftoning. In FM halftoning, the illusion of continuous tone is achieved by varying the distance between printed dots, which are themselves held constant in size. Because the arrangement is a random pattern of dots, FM halftoning is commonly referred to as stochastic screening.

FM halftoning offers several advantages over AM halftoning. The first is that when printed dots are arranged in a random fashion, superimposing two or more patterns

does not create moiré—eliminating the need for screen angles.³ The second advantage is spatial resolution. By not forming macrodots, FM halftoning maximizes the spatial resolution of printed images relative to the resolution of the printer. Figure 1 makes the comparison between AM (top left) and FM (top right) halftoning for a continuous-tone image, illustrating the superiority of FM halftoning.

The technique itself, error diffusion, also poses several advantages over alternative FM techniques, such as Bayer's dither,⁴ which creates patterns with "strong periodic structure that imparts an unnatural appearance to resulting images" (Ref. 5, p. 322). Instead, error diffusion, by diffusing the quantization error to neighboring pixels, is able to reproduce images with greater fidelity.³ In his landmark paper on error diffusion, Ulichney⁶ shows that the patterns produced by error diffusion closely approximate blue noise (the high-frequency component of white noise), with stochastic patterns having a power spectrum that closely approximates blue noise visually more pleasing than patterns that do not.

The drawback of error diffusion, though, is its high computational complexity, which requires both processing and storage of neighboring pixels. Ordered dither techniques such as clustered-dot and Bayer's dither are simple and require minimal processing,⁵ but again, these techniques suffer from low spatial resolution and strong periodic structure. To overcome these limitations, Mitsa and Parker⁷ introduced the blue-noise mask, a dither array designed to generate halftone patterns that closely approximate blue noise. As a demonstration, Fig. 1 (middle left) also shows the resulting image halftoned by using a blue-noise dither array. Since its introduction, halftoning with the use of blue-noise dither arrays has become the preferred stochastic halftoning technique in printers that reliably reproduce individual pixels.

In printers that do not reproduce pixels reliably, FM halftoning suffers its major drawback: FM halftones are more susceptible than AM halftones to printer distortions




Fig. 1. Binary tomato image generated by using (top left) clustered-dot dithering, (top right) error diffusion, (middle left) a blue-noise dither array, (middle right) error diffusion with printer distortion, (bottom left) error diffusion with output-dependent feedback with printer distortion, and (bottom right) a green-noise mask with printer distortion.

such as dot gain, the increase in size of a printed dot from its intended size. Dot gain, in the printing of black on white, can cause isolated white pixels to fill, or plug,⁸ with ink; furthermore, by creating patterns that are darker than the original ratio of white pixels to black pixels, dot gain can cause a shift in color.³ When the printing process is repeatable, these distortions can be anticipated, but in instances such as screen printing, where the pro-

cess is not repeatable,⁹ compensating for printer distortions is a far more difficult task. In these instances methods that resist printer distortions and are more robust to variations in the printing process are desirable.¹⁰

The major relationship between halftone patterns and the amount of dot gain seems to be the perimeter-to-area ratio of printed dots. That is, the halftone screen having the greatest perimeter-to-area ratio of printed dots will be

far more susceptible to the distortions caused by dot gain.⁸ FM halftoning, having a much higher ratio than AM halftoning, is, therefore, more susceptible. For this reason many printers and commercial presses still rely on AM techniques.

Recently, several variations of error diffusion have been introduced, such as those using space-filling curves¹¹ or output-dependent feedback.¹² These techniques do not produce a homogeneous arrangement of isolated pixels as suggested by blue noise but, instead, form pixel clusters that are themselves spaced homogeneously throughout a dither pattern. In their paper¹⁰ Lau *et al.* describe these patterns as green noise (the mid-frequency component of white noise). Still a stochastic patterning of dots, green noise enjoys the same advantages as those of blue noise with respect to moiré, but by allowing adjustable coarseness, the resulting screens offer a tunable halftoning technique that can be optimized for specific printer characteristics such that patterns are composed of large clusters in printers with high dot gain and small clusters in printers with low dot gain. The objective of using green noise is to capture the maximum-dispersion attributes of blue noise along with the attributes of clustering associated with AM halftoning. Figure 1 shows the resulting images for printing in the presence of distortion, where a black dot is modeled as a round dot with a diameter 10:6 times greater than the minimum distance between sample points, with the use of (middle right) error diffusion and (bottom left) error diffusion with output-dependent feedback, as described by Ref. 12.

The drawback of these variations is again the high computational complexity associated with error diffusion, which requires both processing and storage of neighboring pixels. In this paper we introduce the concept of halftoning by using green-noise masks, dither arrays designed to generate halftone screens that closely approximate green noise. While earlier attempts at green-noise dither arrays have been made,^{13–16} this is the first to directly address the desirable characteristics of the green-noise model as presented by Lau *et al.*¹⁰ What the new dither arrays being introduced here offer is a technique for generating halftone patterns that can be tuned to specific printer characteristics with minimal computational complexity; furthermore, patterns generated with the green-noise mask need not be radially symmetric. As suggested by the model of the human visual system,¹⁷ diagonal correlations are less objectionable to the human viewer, and although blue-noise masks have been designed to have increased diagonal correlation,¹⁸ green-noise patterns typically have significantly higher anisotropic attributes.

In essence, green-noise masks offer far more control of the spatial arrangement of dots than the static blue-noise model—allowing for tunable dot size and dot shape. This paper, along with introducing the concept of green-noise masks, also introduces a technique for their construction. Unlike the blue-noise mask, which is constructed based on spectral domain statistics, the green-noise mask construction algorithm relies on spatial statistics.^{19–21} Figure 1 shows the resulting images for printing in the presence of distortion with the use of (bottom right) a green-noise mask.

2. SPECTRAL ANALYSIS OF HALFTONES

Discussed in detail by Ulichney,⁶ the term “blue” denotes the high-frequency component of the visible-light spectrum that makes up white light. To measure the spectral content of these blue-noise dither patterns, Ulichney first defines I_g as a binary dither pattern generated by halftoning a continuous-tone monochrome image of constant gray level g . An estimate $\hat{P}(f)$ of the spectral content of I_g is then generated through Bartlett’s method of averaging periodograms, the squared magnitude of the Fourier transform of the output pattern divided by the sample size. In this paper $\hat{P}(f)$ is generated by averaging ten periodograms of size 256×256 .

Given $\hat{P}(f)$, Ulichney then defines two one-dimensional statistics as a means of quantitatively measuring the spectral content of I_g . The first metric is the radially averaged power spectrum density (RAPSD), defined as

$$P(f_\rho) = \frac{1}{N(R(f_\rho))} \sum_{f \in R(f_\rho)} \hat{P}(f), \quad (1)$$

the average power in a series of annular rings $R(f_\rho)$, as shown in Fig. 2, which partition the spectral domain such that each ring has center radius f_ρ and radial width Δ_ρ and is composed of $N(R(f_\rho))$ frequency samples. The second statistic is the anisotropy $A(f_\rho)$, defined as

$$A(f_\rho) = \frac{1}{N(R(f_\rho)) - 1} \sum_{f \in N(R(f_\rho))} \frac{[\hat{P}(f) - P(f_\rho)]^2}{[P(f_\rho)]^2}, \quad (2)$$

the relative variance or the noise-to-signal ratio of samples within the annular ring of center radius f_ρ . $A(f_\rho)$ is a measure of how isotropic a dither pattern is, with low $A(f_\rho)$ indicating an isotropic dither pattern and high $A(f_\rho)$ being the result of directional artifacts in I_g .

A. Spectral Statistics of Blue Noise

In blue-noise halftoning, the spatial arrangement of pixels is such that minority pixels are spread as homogeneously as possible—leading to an isotropic arrangement in which minority pixels are separated by an average distance of λ_b . The parameter λ_b is referred to as the principal wavelength of blue noise, and given that the square of this wavelength is inversely proportional to the gray level ($1 = \text{white}$, $0 = \text{black}$), λ_b follows the relationship




Fig. 2. With the spectral domain partitioned into a series of annular rings, the power spectrum $P(f)$ can be studied by using $P(f_\rho)$ and $A(f_\rho)$.




Fig. 3. Spectral characteristics of (top) a blue-noise halftone pattern and (bottom) a green-noise halftone pattern.

$$\lambda_b = \begin{cases} D/\sqrt{g} & \text{for } 0 < g \leq 1/2 \\ D/\sqrt{1-g} & \text{for } 1/2 < g \leq 1 \end{cases} \quad (3)$$

where D is the minimum distance between addressable points on the display. Ulichney⁶ notes that this separation should vary in an uncorrelated manner and that the wavelengths of this variation should not be significantly larger than λ_b . If variations are allowed to be significantly larger, the dither pattern may appear too white relative to the optimal blue-noise pattern.

With respect to the spectral content of blue noise, the optimal blue-noise dither pattern will have a $P(f_\rho)$ as shown in Fig. 3 (top) with (a) a minimal low-frequency content, (b) a flat blue-noise region, and (c) a spectral peak at cutoff frequency $f_\rho = f_b$, the principal frequency of blue noise for gray level g such that

$$f_b = \begin{cases} \sqrt{g}/D & \text{for } 0 < g \leq 1/2 \\ \sqrt{1-g}/D & \text{for } 1/2 < g \leq 1 \end{cases} \quad (4)$$

(In this and other figures below, a small diamond placed along the horizontal axis represents principal frequencies and wavelengths.) Being isotropic, the optimal blue-noise pattern will also have minimal $A(f_\rho)$, since, as Ulichney writes (paraphrased), an important property of visually pleasing dither patterns is isotropy: directional artifacts are perceptually disturbing (Ref. 6, p. 58).

B. Spectral Statistics of Green Noise

In green-noise halftoning, the spatial arrangement of pixels is such that minority pixels form clusters that are themselves spread as homogeneously as possible—leading to an arrangement in which clusters are separated (center to center) by an average distance of λ_g .¹⁰ The parameter λ_g is the principal wavelength of green noise, and given that the square of this wavelength is inversely proportional to the gray level divided by \bar{M} , the average number of pixels per cluster, λ_g is such that

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

Lau *et al.*¹⁰ note that in green noise the cluster size/shape greatly influences the separation between clusters, with large variation in cluster size/shape leading to large variations in separation between clusters. Just as with

blue noise, variations significantly longer than λ_g lead to a dither pattern that may appear too white relative to the optimal green-noise pattern.

The spectral content of green-noise patterns is such that the optimal green-noise pattern has a $P(f_\rho)$ of the form of Fig. 3 (bottom), showing (a) a minimal low-frequency content, (b) a minimal high-frequency content, and (c) a spectral peak at radial frequency $f_\rho = f_g$, the principal frequency of green noise for gray level g such that

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

Thus, for a given gray level g , the green-noise and blue-noise principal frequencies are different. Just as with λ_g , cluster size/shape greatly influences isotropy, with nonisotropic cluster shape leading to nonisotropic λ_g , but unlike for blue noise, isotropy is not a prerequisite of visually pleasing green-noise patterns. Typically, green-noise patterns have a $A(f_\rho)$ significantly higher than that of a blue-noise pattern with equivalent g .

3. SPATIAL ANALYSIS OF HALFTONES

Based on radial averages, $P(f_\rho)$ and $A(f_\rho)$ are invariant to rotation. That is, it is possible for two patterns, one with diagonal correlations and one with horizontal correlations, to have identical $P(f_\rho)$ and $A(f_\rho)$. According to the model of the human visual system proposed by Sullivan *et al.*,¹⁷ the pattern with diagonal correlations will appear less objectionable to a human viewer; furthermore, the model suggests that isotropy may not be a requirement of visually pleasing patterns. Green-noise patterns are such an instance where a visually pleasing dither pattern need not be isotropic. So, as a supplement to the spectral analysis proposed by Ulichney, Lau *et al.*¹⁰ propose using the spatial analysis of point processes commonly associated with stochastic geometry, the area interested in modeling complicated geometrical patterns.

In this framework an aperiodic halftoning process is defined as a stochastic model Φ governing the location of events, or points x_i , within the space \mathfrak{R}^2 . As a sample of Φ , ϕ is written as $\phi = \{x_i, i = 1, \dots, N\}$, with $\phi(B)$, a scalar quantity, defined as the number of x_i 's in the subset B of \mathfrak{R}^2 . In digital halftoning, ϕ represents a binary dither pattern, with $\phi(x) = 1$ indicating that the pixel at sample x is a minority pixel (pixel x is 1 for $0 \leq g < 1/2$ and 0 for $1/2 \leq g \leq 1$) and $\phi(x) = 0$ indicating that the pixel at sample x is not.

As a scalar quantity, the first-order moment or the expected value of $\phi(x)$ is the intensity $\mathcal{I}(x)$, which is the unconditional probability that sample x is a minority pixel. For a binary pattern representing gray level g , $\mathcal{I}(x) = g$ for $0 \leq g < 1/2$ and $1 - g$ for $1/2 \leq g \leq 1$. A second statistic for characterizing Φ is the quantity $\mathcal{K}(x; y)$, defined as

$$\mathcal{K}(x; y) = \frac{\mathbf{E}\{\phi(x)|y \in \phi\}}{\mathbf{E}\{\phi(x)\}}, \quad (7)$$

the ratio of the conditional expectation that sample x is a minority pixel given that sample y is a minority pixel to




Fig. 4. (a) With the spatial domain divided into a series of annular rings centered on location y , the spatial arrangement of points can be studied through the use of $\mathcal{R}(r)$; (b) from the expected number of points per unit area in the segment Γ_y^a versus that in the ring Γ_y , the spatial arrangement can be studied by using $\mathcal{D}_{r_1,r_2}(a)$.

To address the need for a spatial measure of isotropy, Lau et al.¹⁰ propose $\mathcal{D}_{r_1,r_2}(a)$, the directional distribution function, which is defined as

$$\mathcal{D}_{r_1,r_2}(a) = \frac{\mathbf{E}\{\phi(\Gamma_y^a)|y \in \phi\}/N(\Gamma_y^a)}{\mathbf{E}\{\phi(\Gamma_y)|y \in \phi\}/N(\Gamma_y)}, \quad (9)$$

the expected number of minority pixels *per unit area* in a segment Γ_y^a of the ring $\Gamma_y = \{x: r_1 \leq |x - y| < r_2, y \in \phi\}$, centered on the minority pixel located at point y such that $a \leq \angle(x - y) < a + \Delta_a$, to the expected number of minority pixels *per unit area* in Γ_y itself [Fig. 4(b)]. Similar to the case of $\mathcal{R}(r)$, values of $\mathcal{D}_{r_1,r_2}(a)$ less than 1 indicate an inhibition of minority pixels in the direction of a for pixels separated by a distance greater than r_1 and less than r_2 . Likewise, $\mathcal{D}_{r_1,r_2}(a)$ greater than 1 suggests an increased likelihood of the direction a between minority pixels. For isotropic dither patterns, $\mathcal{D}_{r_1,r_2}(a) = 1$ for all a .

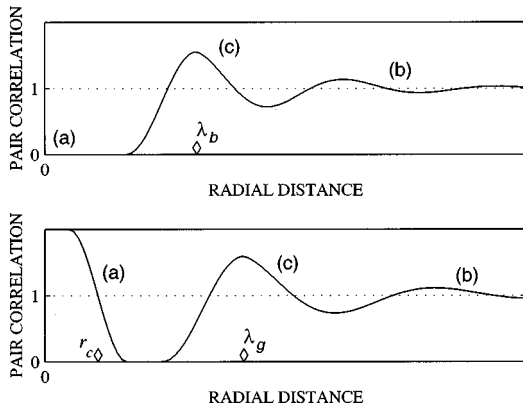


Fig. 5. Pair correlation for (top) a blue-noise process and (bottom) a green-noise process.

the unconditional expectation that sample x is a minority pixel. Referred to as the reduced second-moment measure, $\mathcal{K}(x; y)$ is a measure of the influence that a minority pixel at sample y has on pixel x . $\mathcal{K}(x; y) > 1$ indicates that sample x is more likely to be a minority pixel given y , and $\mathcal{K}(x; y) < 1$ indicates that sample x is less likely given y .

For a stationary point process Φ , $\mathcal{K}(x; y) = \mathcal{K}(r, \theta)$, where r is the distance between samples x and y and θ is the direction from y to x . For stationary point processes that have statistical properties invariant to rotation (isotropic), $\mathcal{K}(r, \theta)$ is independent of θ and is commonly referred to as the pair correlation $\mathcal{R}(r)$, which is defined explicitly as

$$\mathcal{R}(r) = \frac{\mathbf{E}\{\phi(R_y(r))|y \in \phi\}}{\mathbf{E}\{\phi(R_y(r))\}}, \quad (8)$$

the ratio of the expected number of minority pixels located in the ring $R_y(r) = \{x: r < |x - y| \leq r + dr\}$ [Fig. 4(a)] under the condition that a minority pixel exists at sample y to the unconditional expected number of minority pixels located in the set $R_y(r)$. Like $P(f_p)$, $\mathcal{R}(r)$ offers a one-dimensional measure of a two-dimensional quantity, but being invariant to rotation, $\mathcal{R}(r)$ offers no information on the isotropy of Φ .

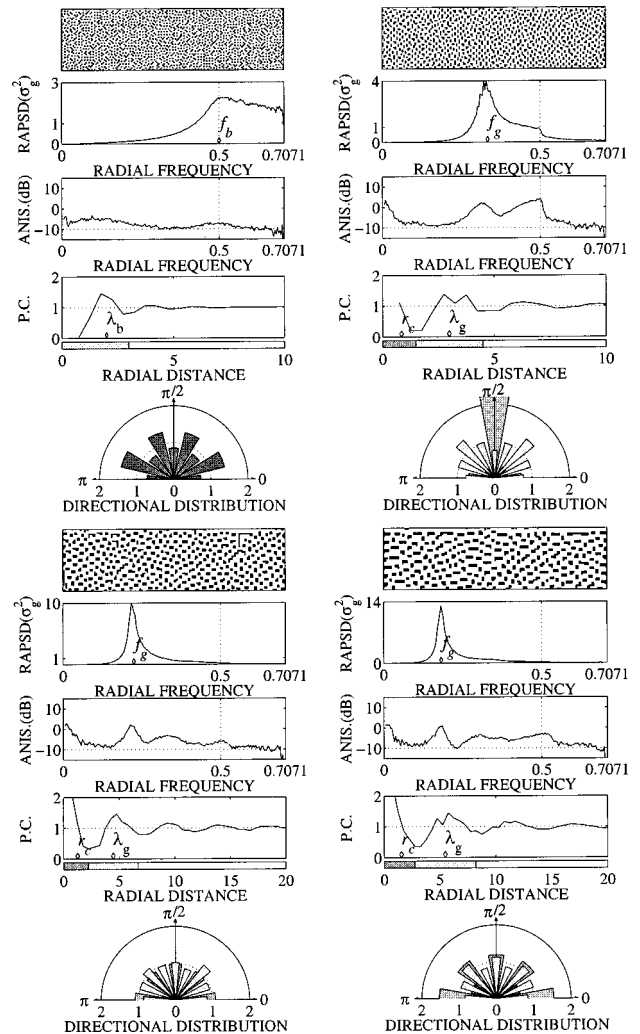


Fig. 6. Resulting metrics for dither patterns generated with the use of error diffusion representing gray level $g = 3/4$ for (top left) blue noise, (top right) green noise with small clusters, (bottom left) green noise with medium clusters, and (bottom right) green noise with large clusters.

A. Spatial Statistics of Blue Noise

Given the goal of blue noise to space minority pixels as homogeneously as possible, the pair correlation of the ideal blue-noise dither pattern is of the form of Fig. 5 (top), showing (a) a strong inhibition of points near $r = 0$, (b) a decreasing correlation of minority pixels with increasing r , and (c) a frequent occurrence of the interpoint distance λ_b , the principal wavelength of blue noise, as illustrated by a series of peaks at integer multiples of λ_b .

Figure 6 (top left) shows the resulting metrics of a blue-noise halftone pattern for $g = 3/4$ generated by error diffusion. Note that $P(f_\rho)$ is shown in units of $\sigma^2 = g(1 - g)$ and that $A(f_\rho)$ is shown in units of decibels, with -10 dB considered background noise.⁶ The pair correlation is taken with $\Delta_r = D/2$, and the directional distribution function is calculated for $da = 2\pi/16$, with the range $(r_1, r_2]$ indicated by color according to the bar located across the horizontal axis of the pair correlation (P.C.) plot.

B. Spatial Statistics of Green Noise

In contrast to blue noise, green noise has a pair correlation of the form of Fig. 5 (bottom), showing (a) clustering of points for $r \leq r_c$, the cluster radius, (b) a decreasing correlation of minority pixels with increasing r , and (c) a frequent occurrence of the intercluster distance λ_g , the principal wavelength of green noise. The parameter r_c , the cluster radius, is related to the average number of minority pixels per cluster as

$$\pi r_c^2 = \bar{M}. \quad (10)$$

Equation (10) states that the area enclosed by a circle of radius r_c is equivalent to the area covered by a cluster of size \bar{M} pixels. Figure 6 shows the resulting metrics for several green-noise halftone patterns with $g = 3/4$ generated by error diffusion with output-dependent feedback, with Fig. 6 (top right) having the smallest clusters, with an average size of 2.2 pixels, and Fig. 6 (bottom right) having the largest, with an average size of 7.5 pixels.

4. BINARY PATTERN PAIR CORRELATION CONSTRUCTION ALGORITHM

In their paper Mitsa and Parker⁷ introduced halftoning by means of the blue-noise mask, where a binary image is generated by comparing on a pixel-by-pixel basis a continuous-tone monochrome image with a dither array or a mask. A pixel of the resulting binary image is set to 1 if the continuous-tone image is greater than the dither array; otherwise, the pixel is set to 0. The blue-noise mask derives its name from the fact that given a continuous-tone monochrome image of constant gray level, the resulting dither pattern has spatial and spectral statistics that approximate blue noise.

To generate this blue-noise mask, Mitsa and Parker introduced the algorithm referred to as the binary pattern power spectrum manipulation algorithm (BIPPSMA), which takes a white-noise dither pattern and iteratively swaps pixels according to the pattern's spectral content. Minimizing the low-frequency content, BIPPSMA at-




Fig. 7. Isotropic shaping function $\tilde{\mathcal{K}}(r)$ used to construct green-noise dither patterns in BIPPCCA.

tempts to generate the optimal blue-noise pattern; furthermore, because BIPPSMA limits which pixels can be swapped, subsequent dither patterns can be generated from old patterns until finally all gray levels from 0 to 1 are represented within the mask.

In this paper we introduce an iterative procedure called the binary pattern pair correlation construction algorithm (BIPPCCA). Its basic premise is to construct binary dither patterns with arbitrary reduced second-moment measures and a given intensity by randomly converting pixels of an arbitrarily sized array from a majority (0) to a minority (1) value. Progressively building on the previous iteration, BIPPCCA begins with an $M \times N$ all-zero array ϕ , with one pixel selected at random and converted to a minority pixel. Given the dither pattern ϕ with minority pixels $\{x_i, i = 1, 2, \dots\}$, BIPPCCA assigns a probability of becoming a minority pixel to each majority pixel in ϕ . BIPPCCA then replaces the maximum likely majority pixel (the majority pixel with the highest corresponding probability) with a minority pixel. The process is then repeated until the dither pattern ϕ of size $M \times N$ has $\mathcal{I} \times M \times N$ minority pixels, where $\mathcal{I} = g$ for $0 \leq g \leq 1/2$ or $\mathcal{I} = 1 - g$ for $1/2 < g \leq 1$.

BIPPCCA is able to construct ϕ such that the resulting dither pattern has a desired reduced second-moment measure by adjusting the probabilities of majority pixels being converted to minority pixels at each iteration according to the current set of minority pixels in ϕ and $\tilde{\mathcal{K}}(r, \theta)$, the reduced second-moment measure shaping function. $\tilde{\mathcal{K}}(r, \theta)$ is a user-specified function derived from $\mathcal{K}(r, \theta)$, with values of $\tilde{\mathcal{K}}(r, \theta) > 1$ increasing the likelihood of minority pixels being placed a distance r and a direction θ apart and values of $\tilde{\mathcal{K}}(r, \theta) < 1$ decreasing this likelihood. Recall from Section 3 that $\mathcal{K}(r, \theta) > 1$ indicates that given a minority pixel at location y , all samples x for which $|y - x| = r$ and $\angle(x - y) = \theta$ are more likely to be a minority pixel than any point z for which $\mathcal{K}(z - y) < \mathcal{K}(r, \theta)$. So given that a minority pixel is placed at y , BIPPCCA increases the likelihood of a pixel becoming 1 for all pixels x for which $\tilde{\mathcal{K}}(r, \theta) > 1$. It also decreases this likelihood for all z for which $\tilde{\mathcal{K}}(z - y) < 1$.

The function $\tilde{\mathcal{K}}(r, \theta)$ is referred to as the spatial shaping function because of its influence in shaping the reduced second-moment measure of the resulting output. To be used in this paper for $\tilde{\mathcal{K}}(r, \theta)$, in Fig. 7 a simple approximation is shown of the ideal pair correlation for isotropic green noise with peaks at all integer multiples of λ_g , the principal wavelength of green noise as defined in Eq. (5), and valleys midway between peaks. We note

that although more elaborate $\tilde{\mathcal{K}}(r, \theta)$ could be proposed, this model was selected because of its simple structure. In simulations 1.01 proved to be a good value as the maximum amplitude (labeled G in Fig. 7) for $\tilde{\mathcal{K}}(r, \theta)$.

In addition to $\tilde{\mathcal{K}}(r, \theta)$, the probabilities of majority pixels being converted are also influenced at each iteration by the current concentration of minority pixels in ϕ through C , the concentration array. C ensures homogeneity in ϕ by decreasing the probabilities of becoming a minority pixel for majority pixels in areas of dense minority-pixel concentration and increasing the probabilities for majority pixels in areas of sparse minority-pixel concentration.

In BIPPCCA the concentration of minority pixels is measured as the output after the low-pass filter H_{LP} is applied to ϕ . In the selection of a low-pass filter, an obvious choice for H_{LP} as suggested in Ref. 5 is the Gaussian filter such that

$$H_{LP}(r) = \exp\left(\frac{-r^2}{2\sigma^2}\right) \quad (11)$$

for some constant σ . Note that for a minority pixel to have an influence on neighboring clusters, the filter H_{LP} should have a higher σ for small \mathcal{I} , where clusters are far apart, than for large \mathcal{I} , where clusters are close together. In this paper such a relationship between H_{LP} and \mathcal{I} is ensured by setting $2\sigma^2 = \lambda_g^2$. The concentration matrix is then constructed from $\{H_{LP} \otimes \phi\}$, the output after ϕ is filtered with the low-pass filter H_{LP} by using circular convolution, according to the mapping of Fig. 8.

The steps of BIPPCCA for generating an $M \times N$ binary dither pattern representing intensity level \mathcal{I} are described as follows:

1. Initialize all pixels of an $M \times N$ array ϕ to 0.
2. Randomly select one pixel of ϕ , and convert that pixel, $\phi[m, n]$, to 1.
3. Create an $M \times N$ array U of uniformly distributed random numbers such that $U[i, j] \in (0, 1]$ is the probability that $\phi[i, j]$ will become a minority pixel.
4. Given the most recently converted pixel $\phi[m, n]$, scale the value $U[i, j]$ for all pixels $\phi[i, j] = 0$ by $\tilde{\mathcal{K}}(r, \theta)$ such that $(U[i, j])_{\text{new}} = (U[i, j])_{\text{old}} \times \tilde{\mathcal{K}}(r, \theta)$, where $r^2 = [\min(|m - i|, M - |m - i|)]^2 + [\min(|n - j|, N - |n - j|)]^2$ (the minimum wraparound distance between the two pixels $\phi[m, n]$ and $\phi[i, j]$) and θ is the direction from $\phi[m, n]$ to $\phi[i, j]$.




Fig. 8. Mapping function used to construct the concentration matrix C from the output after ϕ is filtered with the low-pass filter H_{LP} by using circular convolution.

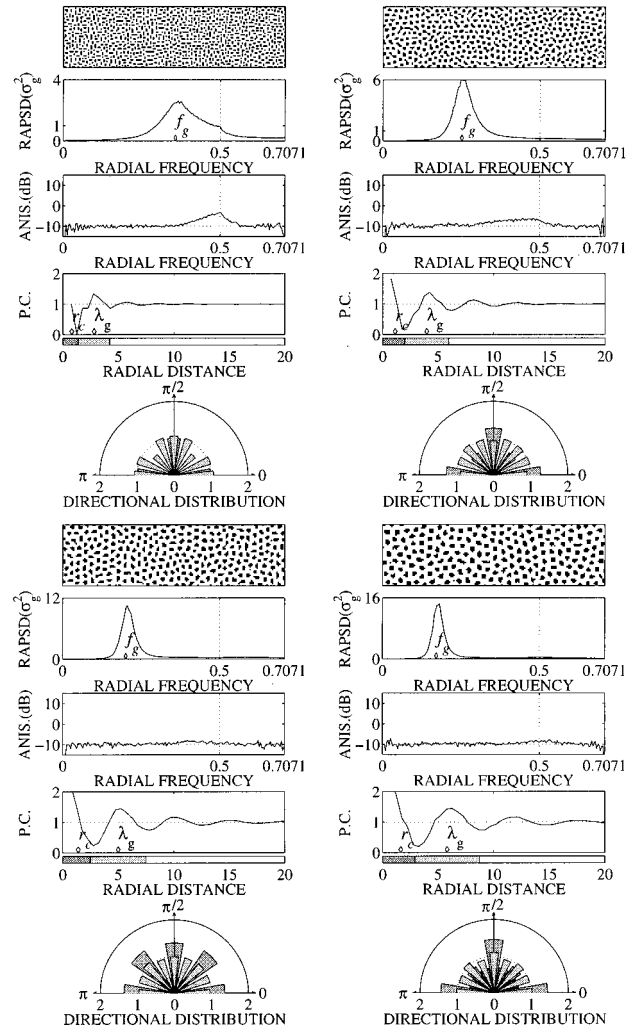


Fig. 9. Resulting metrics for dither patterns generated with the use of BIPPCCA representing gray level $g = 3/4$ with average cluster sizes of (top left) 2.0 pixels, (top right) 3.9 pixels, (bottom left) 6.0 pixels, and (bottom right) 7.9 pixels.

5. Construct the concentration matrix C with use of the mapping of Fig. 8 from $\{H_{LP} \otimes \phi\}$, the output after ϕ is filtered with the low-pass filter H_{LP} by using circular convolution.

6. Locate the majority pixel in ϕ with the highest probability (the pixel $\phi[m, n] = 0$ such that $U[m, n] \times C[m, n] > U[i, j] \times C[i, j]$ for all $1 \leq i \leq M$ and $1 \leq j \leq N$ and that $\phi[i, j] = 0$), and set that pixel, $\phi[m, n]$, to 1.

7. If the number of minority pixels in ϕ is equal to $\mathcal{I} \times M \times N$, then the algorithm quits with the output pattern given by ϕ ; otherwise, continue at step 4.

As demonstration, Fig. 9 shows the resulting output patterns and corresponding metrics after the use of BIPPCCA with the isotropic shaping function of Fig. 7, with Fig. 9 (top left) having the smallest clusters, with an av-

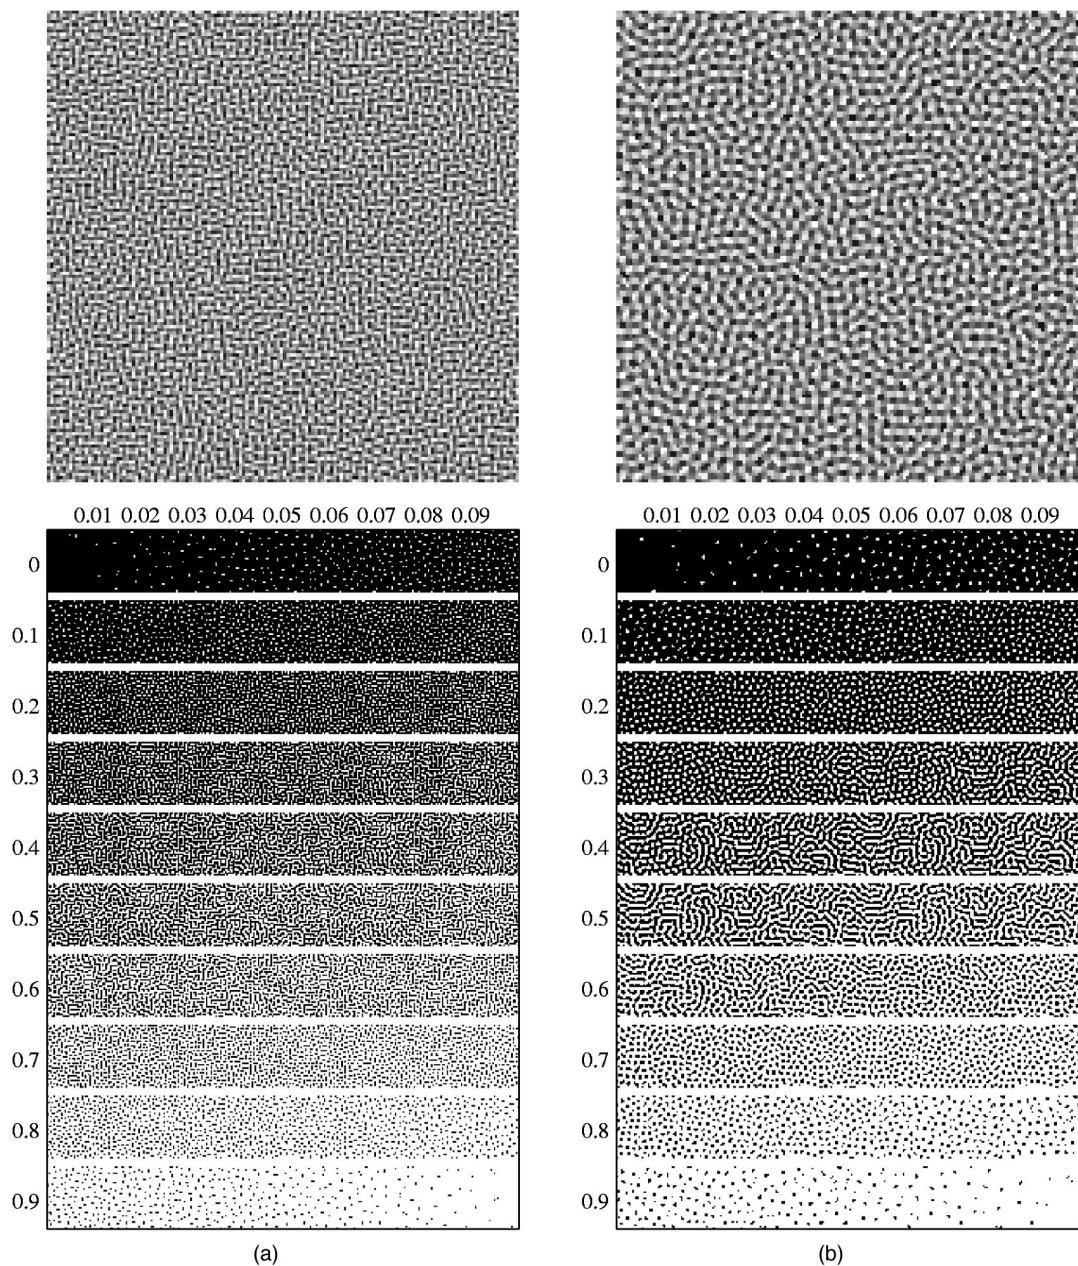


Fig. 10. Continues on facing page.

erage size of 2.0 pixels, and Fig. 9 (bottom right) having the largest, with an average size of 7.9 pixels.

The above algorithm, BIPPCCA, shares some similarities with the previously introduced BIPPSMA. One similarity is the two-dimensional wraparound property described by Ulichney⁵; that is, minority pixels placed near an edge of the dither array influence the placement of pixels near the opposing edge. This property allows the output dither pattern to be tiled edge to edge to form larger images that appear aperiodic; furthermore, the algorithms are not limited to specific array sizes, thereby allowing for arbitrary dimensions.

5. CONSTRUCTING THE GREEN-NOISE MASK

In this section we describe an algorithm for constructing green-noise masks, dither arrays that, when thresholded

at gray level g , produce binary green-noise dither patterns appropriate to g . Constructed with the use of BIPPCCA, these masks can be of any size, where large images are halftoned by using new masks formed by tiling edge to edge the original dither array. The basic premise of their construction is to generate a set $\{\phi_g, 0 \leq g \leq 1\}$ of dither patterns by using BIPPCCA from a set $\{\tilde{\mathcal{K}}(r, \theta; g), 0 \leq g \leq 1\}$ of shaping functions, with one pattern and one shaping function for each possible discrete gray level g (256 levels for 8-bit gray-scale images). The dither array is then constructed by assigning to each pixel a threshold according to the spatial arrangement of binary pixels within $\{\phi_g, 0 \leq g \leq 1\}$. Note that because BIPPCCA generates patterns such that minority pixels are represented by pixels equal to 1, the dither pattern ϕ_g for $1/2 < g \leq 1$ is generated by inverting the dither pattern created by BIPPCCA for gray level $1 - g$.

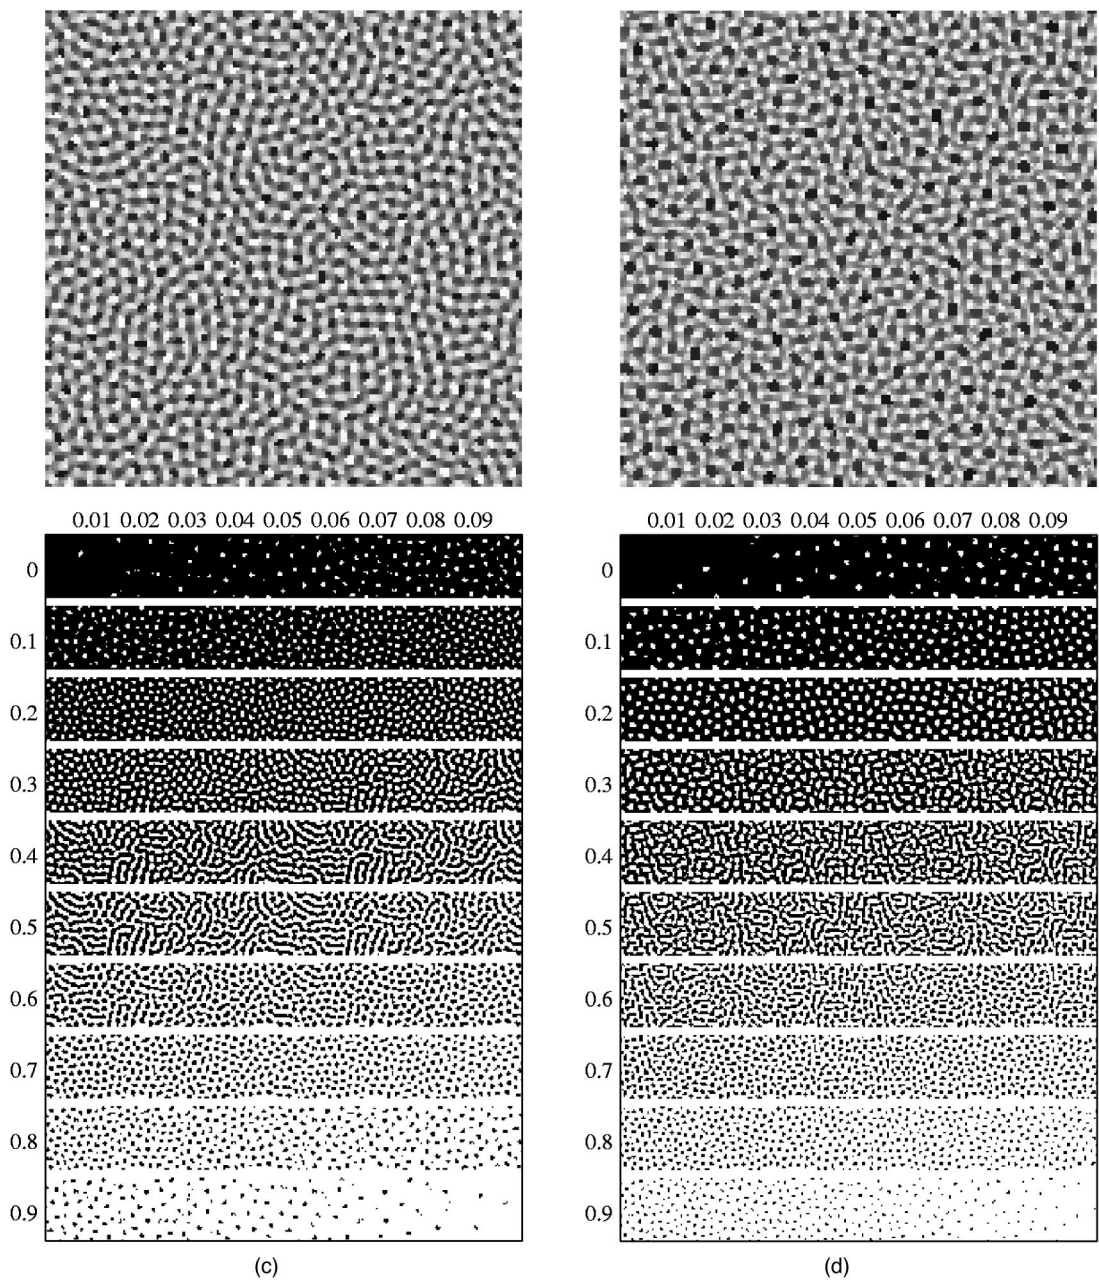


Fig. 10. (a)–(c) Masks generated by BIPPCCA having cluster sizes that increase symmetrically with equal deviations from black and white toward middle gray, (d) mask generated by BIPPCCA with nonsymmetric deviations with larger clusters for $g < 1/2$.

To avoid ambiguities in the assignment of thresholds to the dither array, the dither patterns ϕ_g are constructed under the stacking constraint that $\phi_k \subset \phi_g$ for all $k < g$ or that if given $\phi_k[m, n] = 1$, then $\phi_g[m, n] = 1$ for all $g > k$. As a consequence, the threshold assigned to each dither array pixel $DA[m, n]$ is equal to the minimum g for which $\phi_g[m, n] = 1$. In BIPPCCA a dither pattern ϕ_k can be constructed, given ϕ_g such that $\phi_k \subset \phi_g$, by constraining step 6 of BIPPCCA to consider only the most likely pixel (in ϕ_k) as the majority pixel $\phi_k[m, n]$, such that $\phi_g[m, n]$ is a minority pixel (=1); furthermore, a dither pattern ϕ_g can be constructed, given ϕ_k such that $\phi_k \subset \phi_g$, if, in step 1, ϕ_g is initialized to ϕ_k and each value $U[m, n]$ is scaled according to the location of minority pixels in ϕ_g as defined by $\tilde{\mathcal{K}}(r, \theta)$. In addition to

the above modifications to BIPPCCA, the shaping function $\tilde{\mathcal{K}}(r, \theta)$ must also observe the stacking constraint, as minority-pixel clusters in pattern ϕ_g cannot decrease in size from those in ϕ_k .

Note that the patterns forming the set $\{\phi_g, 0 \leq g \leq 1\}$ need not be generated in any particular order and that generating patterns in a random order may offer better results than those obtained by constructing ϕ_g by consecutive gray levels. Under the stacking constraint, a pattern ϕ_g is constrained to be a subset of all patterns ϕ_h for which $h > g$ and is constrained to have as a subset the pattern ϕ_k , where $k < g$. So the ordering in which patterns are generated defines which patterns ϕ_g has constraining its construction. As an example of this, Mitsa and Parker⁷ construct the blue-noise mask beginning

with gray level 1/2, with all subsequent patterns ϕ_g for $g < 1/2$ independent of all ϕ_g for $g > 1/2$. In this paper masks are generated by first constructing the patterns for gray levels 0, 1, 1/4, 3/4, and 1/2 in succession. The next patterns to be constructed have gray levels that fall midway between these constructed patterns. This routine is continued, with the current set of constructed patterns falling midway between the existing set until 256 unique patterns have been constructed.

Without loss of generality, the steps for creating a green-noise mask described here assume that images are composed of a discrete set of N gray levels defined by the monotonically increasing sequence $\{g_i, i = 1, 2, \dots, N\}$, with $g_1 = 0$ (black) and $g_n = 1$ (white). The following list describes the steps for generating the set $\{\phi_g, 0 \leq g \leq 1\}$:

1. Given $g_1 = 0$ and $g_N = 1$, define ϕ_{g_1} as an all-zero matrix and ϕ_{g_N} as an all-unity matrix.
2. Define the sequence $\{g_{k_i}, i = 1, 2, \dots, N\}$ as a rearrangement or a reordering of the sequence $\{g_i, i = 1, 2, \dots, N\}$ such that every gray level g_i appears once and only once in the sequence $\{g_{k_i}, i = 1, 2, \dots, N\}$, with $g_{k_1} = 0$ and $g_{k_2} = 1$ ($k_1 = 1$ and $k_2 = N$).




Fig. 11. Principal wavelengths for masks (a)–(d), with the blue-noise principal wavelength indicated by a dotted curve.

3. Set $j = 3$.
4. Construct the dither pattern $\phi_{g_{k_j}}$ under the stacking constraint for all patterns $\phi_{g_{k_i}}$ for $i = 1, 2, \dots, j - 1$ (for all patterns that have been generated up to the current iteration).
5. Set $j = j + 1$. If $j = N + 1$, the process is complete; otherwise, continue at step 4.

Figure 10 shows four masks of size 150×150 created by using the above algorithm for 8-bit gray-scale images ($N = 256$). These masks, labeled (a)–(d), were each designed according to their own isotropic shaping function sets $\{\tilde{\mathcal{K}}(r; g), 0 \leq g \leq 1\}$, with each shaping function $\tilde{\mathcal{K}}(r; g)$ defined by Fig. 7 and by the plots of λ_g versus gray level in Fig. 11. Table 1 lists the average size of clusters, the principal wavelength, and the principle frequency for each mask at several values of g .

For masks (a)–(c), the average size of clusters [related to λ_g by Eq. (5)] were chosen such that $\bar{M}(g)$, the average size of clusters for gray level g , equals $\bar{M}(1 - g)$. Specifically, mask (a) was designed for printers with low dot-gain characteristics as clusters range in size from 1.3 to 4.1 pixels. Masks (b) and (c) were designed for printers with medium to high dot-gain characteristics as clusters range in size from 2.5 to 10.0 pixels for mask (b) and from 3.7 to 14.0 pixels for mask (c). Mask (d) is the resulting mask when clusters are smaller for black minority pixels ($1/2 < g \leq 1$) than for white minority pixels. As defined by the dot-overlap model proposed by Pappas and Neuhoff,²² isolated black dots are not at risk of being lost because of overlap, and therefore black dots may form smaller clusters at g near 1 than clusters formed by white dots at g near 0. Under this criterion mask (d) has, at $g = 0.25$, clusters with an average size of 10.8 pixels and has, at $g = 0.75$, clusters with an average size of 5.3 pixels.

Table 1. Average Number of Pixels per Cluster \bar{M} , Principal Wavelength λ_g , and Principal Frequency f_g for Masks of Fig. 1^a

g	Mask (a)			Mask (b)		
	\bar{M}	λ_g	f_g	\bar{M}	λ_g	f_g
1/255	1.3(2.0)	18.2(22.6)	0.05(0.06)	2.5(4.0)	25.7(31.9)	0.04(0.03)
64/255	2.0(2.0)	2.8(2.8)	0.35(0.35)	4.0(4.0)	4.0(4.0)	0.23(0.25)
128/255	4.1(4.0)	2.9(2.8)	0.35(0.35)	10.0(8.0)	4.5(4.0)	0.23(0.25)
192/255	2.0(2.0)	2.8(2.8)	0.35(0.35)	3.9(4.0)	4.0(4.0)	0.22(0.25)
254/255	1.3(2.0)	18.2(22.6)	0.05(0.06)	2.5(4.0)	25.7(31.9)	0.04(0.03)
g	Mask (c)			Mask (d)		
	\bar{M}	λ_g	f_g	\bar{M}	λ_g	f_g
1/255	3.7(6.0)	30.6(39.1)	0.03(0.03)	8.8(10.0)	47.4(50.5)	0.02(0.02)
64/255	6.3(6.0)	5.0(4.9)	0.20(0.20)	10.8(10.0)	6.6(6.3)	0.15(0.16)
128/255	14.0(12.0)	5.3(4.9)	0.19(0.20)	16.0(10.0)	5.7(4.5)	0.18(0.22)
192/255	6.3(6.0)	5.0(4.9)	0.20(0.20)	5.3(5.5)	4.6(4.7)	0.22(0.21)
254/255	3.5(6.0)	30.0(39.1)	0.03(0.03)	1.3(1.0)	18.0(16.0)	0.06(0.06)

^aShown in parentheses are the ideal parameters used to construct $\tilde{\mathcal{K}}(r; g)$.




Fig. 12. Tomato image halftoned by use of the green-noise masks of Fig. 10 [labeled (a)–(d)].

6. CONCLUSIONS

Unlike blue-noise halftoning, which distributes the minority pixels of a binary dither pattern as homogeneously as possible, green-noise halftoning forms clusters of minority pixels that are themselves distributed as homogeneously as possible. Typically, green-noise patterns are generated with error-diffusion-based techniques. One such technique, Levien's¹² error diffusion with output-dependent feedback, offers tunable coarseness, with small clusters reserved for printers with low dot gain and large clusters reserved for printers with high dot gain. Dot gain is the increase in size of a printed dot relative to the intended dot size. The distortion caused by the overlap of printed dots as a result of dot gain is minimized by clustering pixels. Hence green-noise patterns are more robust than blue-noise patterns to these distortions.

In this paper we have constructed green-noise dither patterns through the algorithm BIPPCCA, which constructs binary patterns according to a desired reduced second-moment measure. Although demonstrated here for green noise, BIPPCCA can be used to construct blue-noise patterns as well. Under a stacking constraint, we have also demonstrated the use of BIPPCCA to build green-noise masks, dither arrays designed to produce green-noise halftones by thresholding pixel by pixel a continuous-tone original. Far less computationally com-

plex than error-diffusion-based algorithms, these green-noise masks can also be tuned to specific printer characteristics by adjusting pattern coarseness. Figure 12 makes the comparison between green-noise masks for varying printer characteristics, with images (a)–(c) having cluster sizes that increase symmetrically with equal deviations from black and white toward middle gray. Image (d) illustrates an instance where white minority pixels form larger clusters than black minority pixels.

Regardless of cluster size, the masks of images (a)–(d) in Fig. 12 were all designed according to an isotropic shaping function. As Ulichney⁶ writes that directional artifacts are perceptually disturbing, green-noise patterns constructed with isotropic shaping functions achieve a far greater degree of isotropy than those typically associated with error diffusion. Future work with green noise will address the vector green-noise mask for efficient halftoning of color images with green noise.

ACKNOWLEDGMENT

This research was supported in part by the National Science Foundation under grant CDA-9703088.

Address correspondence to Gonzalo R. Arce at the location on the title page or by phone, 302-831-8030, or

e-mail, arce@ece.udel.edu. Daniel L. Lau and Neal C. Gallagher may also be reached by e-mail at, respectively, lau@ece.udel.edu and gallaghe@ece.udel.edu.

REFERENCES

1. J. E. Adamcewicz, "A study on the effects of dot gain, print contrast and tone reproduction as it relates to increased solid ink density on stochastically screened images with conventionally screened images," M.S. thesis (Rochester Institute of Technology, Rochester, N.Y., 1994).
2. R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial gray-scale," *Proc. Soc. Inf. Disp.* **17**, 75–78 (1976).
3. K. Laughlin, "An investigation of amplitude and frequency modulated screening on dot gain and variability," M.S. thesis (Rochester Institute of Technology, Rochester, N.Y., 1994).
4. B. Bayer, "An optimum method for two level rendition of continuous-tone pictures," in *IEEE International Conference on Communications, Conference Record* (Institute of Electrical and Electronics Engineers, New York, 1973), pp. 11–15.
5. R. A. Ulichney, "The void-and-cluster method for dither array generation," in *Human Vision, Visual Processing, and Digital Display IV*, J. P. Allebach and B. E. Rogowitz, eds., *Proc. SPIE* **1913**, 332–343 (1993).
6. R. A. Ulichney, "Dithering with blue noise," *Proc. IEEE* **76**, 56–79 (1988).
7. T. Mitsa and K. J. Parker, "Digital halftoning technique using a blue-noise mask," *J. Opt. Soc. Am. A* **9**, 1920–1929 (1992).
8. M. Rodriguez, "Graphic arts perspective on digital halftoning," in *Human Vision, Visual Processing, and Digital Display V*, B. E. Rogowitz and J. P. Allebach, eds., *Proc. SPIE* **2179**, 144–149 (1994).
9. M. A. Coudray, "Causes and corrections of dot gain on press," *Screen Print. J. Technol. Manage.* **86**, 18–26 (1996).
10. D. L. Lau, G. R. Arce, and N. C. Gallagher, "Green-noise digital halftoning," *Proc. IEEE* **86**, 2424–2444 (1998).
11. L. Velho and J. M. Gomes, "Digital halftoning with space filling curves," *Comput. Graph.* **25**, 81–90 (1991).
12. R. Levien, "Output dependent feedback in error diffusion halftoning," in *IS&T's Eighth International Congress on Advances in Non-Impact Printing Technologies* (Society for Imaging Science and Technology, Springfield, Va., 1992), pp. 280–282.
13. J. P. Allebach, "Random nucleated halftone screening," *Photograph. Sci. Eng.* **22**, 89–91 (1978).
14. V. Ostromoukhov, "Pseudo-random halftone screening for color and black & white printing," in *Recent Progress in Digital Halftoning*, R. Eschbach, ed. (Society for Imaging Science and Technology, Society for Imaging Science and Technology, Springfield, Va., 1995), pp. 130–134.
15. S. Wang, "Stochastic (stochastic clustered) halftone screen design," in *IS&T's NIP 13: International Conference on Digital Printing Technologies* (Society for Imaging Science and Technology, Springfield, Va., 1997), pp. 516–521.
16. S. Aoki, "New halftoning method using adaptive cell," in *IS&T's NIP 14: International Conference on Digital Printing Technologies* (Society for Imaging Science and Technology, Springfield, Va., 1998), pp. 277–280.
17. J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Trans. Syst. Man Cybern.* **21**, 33–38 (1991).
18. M. Yao and K. J. Parker, "Modified approach to the construction of a blue noise mask," *J. Electron. Imaging* **3**, 92–97 (1994).
19. N. A. C. Cressie, *Statistics for Spatial Data* (Wiley, New York, 1983).
20. P. J. Diggle, *Statistical Analysis of Spatial Point Patterns* (Academic, London, 1983).
21. D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic Geometry and Its Applications* (Wiley, New York, 1987).
22. T. N. Pappas and D. L. Neuhoff, "Printer models and error diffusion," *IEEE Trans. Image Process.* **4**, 66–79 (1995).