Daniel L. Lau, Robert Ulichney, and Gonzalo R. Arce

Blue- and Green-Noise Halftoning Models

A review of the spatial and spectral characteristics of halftone textures.

In this article, 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 low-frequency texture but have clusters of minority pixels on blue-noise centers. Green noise is composed exclusively of midfrequency spectral components. In addition to the basic spatial and spectral characteristics of the halftoning models, this article also reviews some of the earlier work done to improve error diffusion as a noise generator. We also discuss processes to generate threshold arrays to achieve blue and green noise with the computationally efficient process of ordered dither.

©ARTVILLE

What Is Blue Noise?

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 the same-sized dots distributed as homogeneously as possible. By doing so, the spectral content of these patterns is 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 (HVS) [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 which the 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. But while isolating minority pixels makes them less visible to the human eye, it also makes the patterns far more susceptible to printer distortion and the inability of some devices, such as laser printers, to reproduce dots consistently from dot to dot [3], [4]. So as a way of providing the benefits of random dot distributions to unreliable devices while also maintaining the consistency of clustered dots, Levien [5] 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.

Green Noise

Lau et al. [6] introduced the green-noise model where the ideal patterns are composed of homogeneously distrib-

1053-5888/03/$17.00©2003IEEE

uted 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 midfrequency components, and as green is the midfrequency component to white, green noise is the midfrequency component to white noise.

Since their introduction, both blue- and green-noise models have led to innovations 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 [7], [9] 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. Given their significance, the remainder of this article details the blue- and green-noise halftoning models as well as describes the spatial and spectral metrics introduced to serve as the basis for these models. While this material offers historical perspective, today it is being realized [10], [11] 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 (CMYK) inks are superimposed. For the reader interested in these color applications, a basic fundamental understanding of the principal halftoning models will serve you well in developing and understanding future algorithms.

Spatial and Spectral Halftone Statistics

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. [6] 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 two-dimensional real space $\Re^2$ [12]. 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 events (or points) in the subset B in $\Re^2$. 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 technique for characterizing the corresponding dither patterns is to look at the relative distribution of minority pixels around one another. To do so, point processes statisticians define the pair correlation $\mathcal{R}(r)$:

$$\mathcal{R}(r) = \frac{\mathbf{E}\{\phi(R_m(r) \mid \phi[m] = 1\}}{\mathbf{E}\{\phi(R_m(r))\}}, \tag{1}$$

as the influence of a minority pixel at index m on all other pixels with index n in the annular ring $R_m(r)$ with center radius r, width Δ_r, and centered around location m. $\mathcal{R}(r)$ may be thought of as the influence at all locations a distance r away 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? For a completely uncorrelated (white) point process, the pair correlation is one for all r. 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 [13].

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 Bartlett's method of averaging periodograms [14], [15] where a periodogram is the magnitude squared of the Fourier transform of a sample output divided by the sample size. Since the spectral estimate, $\hat{P}(f)$, is a function of two dimensions and although anisotropies in the sample dither pattern can be qualitatively observed by studying three-dimensional (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

$$P_\rho(f_\rho) = \frac{1}{N_\rho(f_\rho)} \sum_{i=1}^{N_\rho(f_\rho)} \hat{P}(f). \tag{2}$$

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 $(1/2)D^{-1}$, 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 $P_\rho(f_\rho)$, these regions of cropping will be indicated along the horizontal axis, while all power spectral estimates will be divided into annular rings of radial width Δ_f such that exactly one sample along each frequency axis falls into each ring.

Blue-Noise Halftoning

A common practice for characterizing observed noise models is to assign a name based on color; 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" [16]. Blue noise is the high-frequency complement of pink noise that, due to the low-pass nature of the HVS, posses very advantageous properties for creating the illusion of continuous tone in binary halftones.

Blue-noise halftoning is 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

$$\lambda_b = \begin{cases} D/\sqrt{g}, & \text{for } 0 < g \le 1/2 \\ D/\sqrt{1-g}, & \text{for } 1/2 < g \le 1 \end{cases} \tag{3}$$

and D is the minimum distance between addressable points on the display [1], [7]. The parameter λ_b is referred to as the principal wavelength of blue noise, with its relationship to g justified by several intuitive properties:

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

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

▲ 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 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 the following:

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

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

▲ 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 Figure 1(a) where $\mathcal{R}(r)$ shows: 1) a strong inhibition of minority pixels near $r=0$, 2) a decreasing correlation of minority pixels with increasing r, and 3) a frequent occurrence of the interpoint distance λ_b, the principal wavelength, indicated by a series of peaks at integer multiples of λ_b. In Figure 1(a), the principal 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 Figure 1(b) and can be described by three unique features: 1) little or no low-frequency spectral components, 2) a flat, high-frequency (blue noise) spectral region, and 3) a spectral peak at cutoff frequency f_b, the blue-noise principal frequency, such that:

$$f_b = \begin{cases} \sqrt{g}/D, & \text{for } 0 < g \le 1/2 \\ \sqrt{1-g}/D, & \text{for } 1/2 < g \le 1, \end{cases} \tag{4}$$

as indicated in Figure 1(b) by a diamond located along the horizontal axis is the principal frequency, and 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.

Error Diffusion

Error diffusion is a neighborhood operation that quantizes the current input pixel and then transfers the quantization error onto future input pixels. Formally, Floyd and Steinberg [17] define the output pixel $y[n]$ by adjusting and thresholding the input pixel $x[n]$ such that

$$y[n] = \begin{cases} 1, & \text{if } (x[n] + x_e[n]) \ge 0 \\ 0, & \text{else} \end{cases} \tag{5}$$

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

$$x_e[n] = \sum_{i=1}^{M} b_i \cdot y_e[n-i] \tag{6}$$

with $y_e[n] = y[n] - (x[n] + x_e[n])$ and $\sum_{i=1}^{M} b_i = 1$. The original four-weight error filter specified by Floyd and

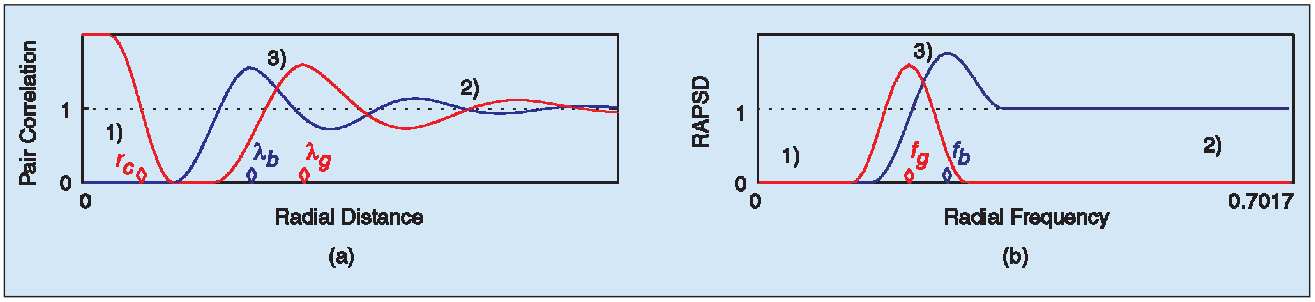

▲ *1. The (a) pair correlation and (b) radially averaged power spectrum for a (blue) blue-noise and a (red) green-noise process.*

Steinberg was selected because of the checkerboard pattern it created at gray level $g = 1/2$. There are several comprehensive reviews of error diffusion in halftoning by Evans et al. [18], Knox [19], and others that cover a great many improvements that have been introduced since the work of Ulichney but have been left out of this article.

Looking at the spatial and spectral characteristics of the resulting dither patterns shown in Figure 2(a), Ulichney noted that patterns exhibited i) correlated artifacts most noticeable at intensity levels $\mathcal{I} = 1/4$, $1/3$, and $1/2$, ii) directional hysteresis due to the order in which pixels were processed (raster order), and iii) transient behavior near edges or boundaries. 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.

Regarding the order that pixels are processed, Ulichney introduced the idea that 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 [20] and Velho and Gomes [21] employed chaotic space-filling curves. With regards to the error filter selection, Ulichney concluded that the 12-weight error filters of Jarvis et al. [22] and Stucki [23] improved some shortcomings of the original four-weight filter of Floyd and Steinberg but still showed some disturbing or noticeable artifacts at intensity levels 1/4 and 1/2 resulting in significant energy below f_b.

Regarding 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 signal-to-noise ratio, dither patterns became too uncorrelated and grainy (white). As a side note, Eschbach and Knox [24] 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

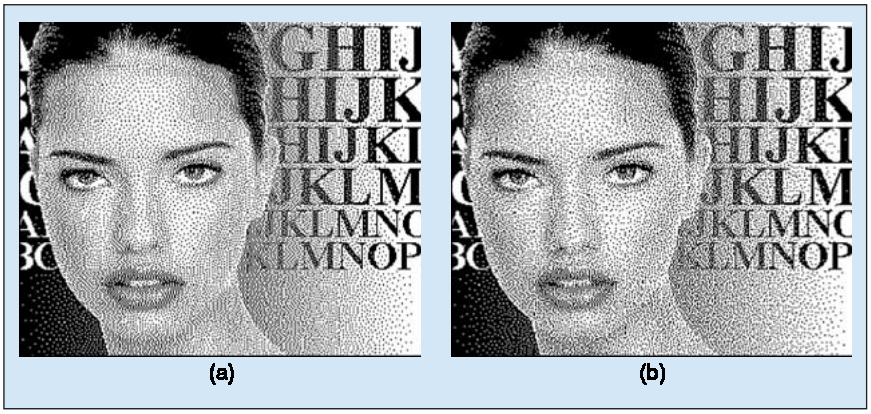

▲ *2. Grayscale images halftoned using (a) Floyd and Steinberg's original error diffusion algorithm and (b) Ulichney's perturbed filter weight scheme.*

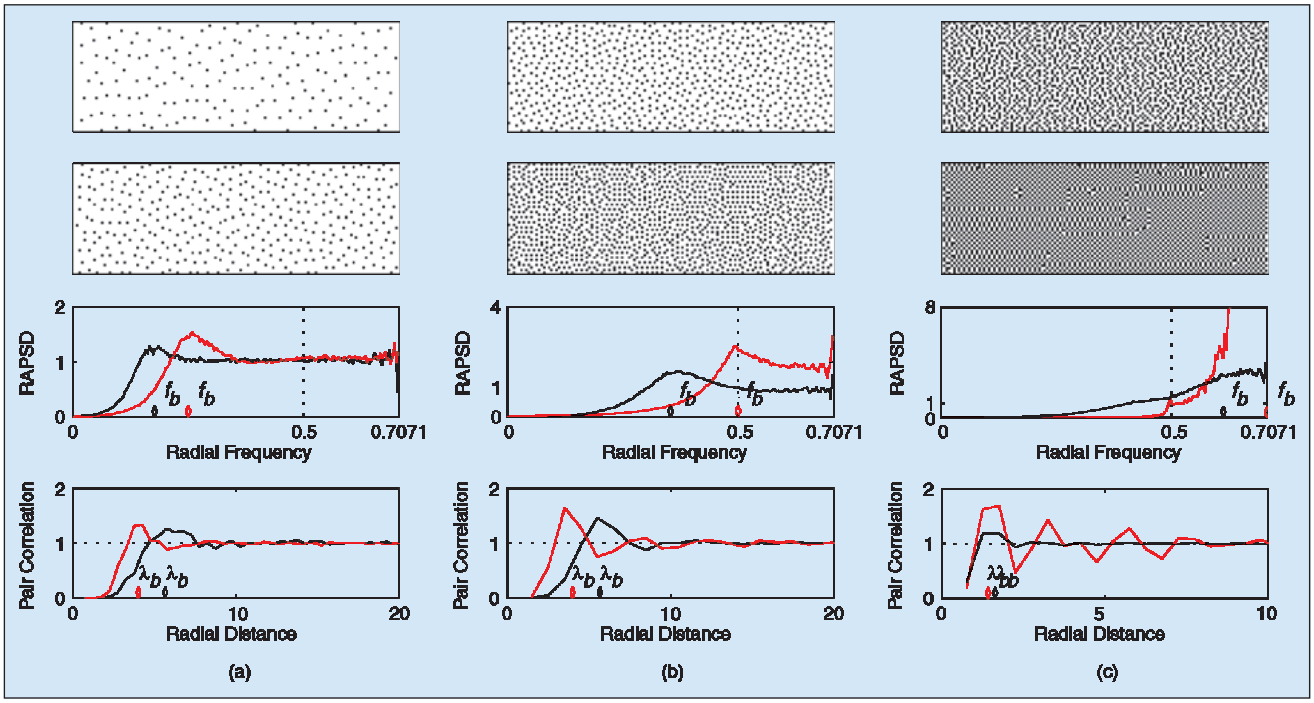

▲ *3. The spatial and spectral statistics for Ulichney's perturbed filter weight scheme at gray levels (a) $g = 1/32$ and $1/16$ (red), (b) $1/8$ and $1/4$ (red), and (c) $1/3$ and $1/2$ (red).*

the other within the pair. This guaranteed that the sum of error weights would always be equal to one, and it was this scheme of perturbing filter weights that led to a much improved error diffusion scheme where 50% noise added to Floyd and Steinberg's original four-weight error filter and a serpentine raster led to the results demonstrated in Figure 2(b). Figure 3 shows the corresponding spatial and spectral characteristics of resulting dither patterns with all metrics exhibiting improved blue-noise characteristics relative to those produced by Floyd and Steinberg's error diffusion.

Blue-Noise Dither Arrays

Blue-noise dithering can also be achieved with the point process of ordered dither. The trick 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 is presented in [25]. 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 [7]. Mitsa and Parker [7] 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 [8] that 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 the intermediate pattern such that minority pixels along the left edge of the pattern will influence pixels along the right side such that no seams are visible in the resulting halftones.

Figure 4 shows the result of dithering an image with a 64×64 void-and-cluster generated dither array.

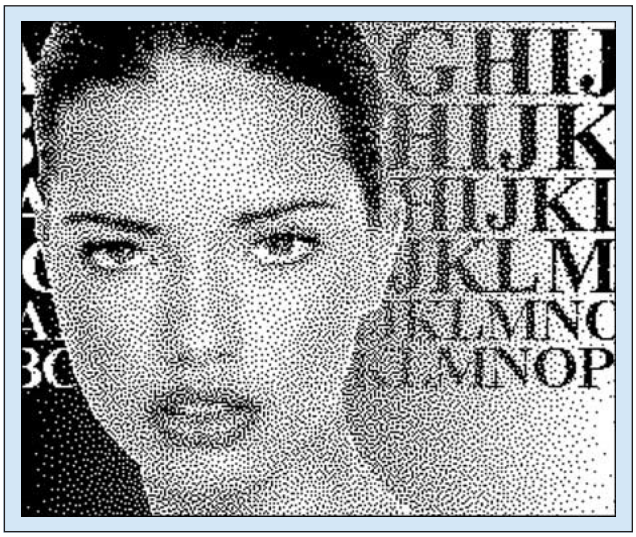

▲ *4. Dithering with a* 64×64 *void-and-cluster array.*

Green-Noise Halftoning

Just as blue noise is the high-frequency component of white noise, green noise is the midfrequency component that, like blue, benefits from aperiodic, uncorrelated structure without low-frequency graininess. But unlike blue noise in halftoning, green-noise dither 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 patterns. Hence the term "green."

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 aperiodic dispersed-dot 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

$$T_x(B) = \{y_i - x : i = 1,2,\ldots\} \tag{7}$$

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

$$\phi_G = \sum_{x_i \in \phi_p} T_{x_i}(\phi_{d_i}) = \sum_{x_i \in \phi_p} \{y_{ji} - x_i : j = 1,\ldots,M_i\}, \tag{8}$$

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.

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

$$\mathcal{I} = \frac{\phi_G(I_g)}{N(I_g)} = \begin{cases} g, & \text{for } 0 < g \le 1/2 \\ 1-g, & \text{for } 1/2 < g \le 1 \end{cases} \tag{9}$$

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 (9), $\overline{M}$, the average number of minority pixels per cluster in I_g, is

$$\overline{M} = \frac{\phi_G(I_g)}{\phi_p(I_g)} = \frac{\mathcal{I} \cdot N(I_g)}{\phi_p(I_g)}, \tag{10}$$

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

Although obvious, (10) 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 $\overline{M}$ is held constant for varying $\mathcal{I}$. In addition, (10) says that the total number of clusters per unit area is proportional to $\mathcal{I} / \overline{M}$. 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 $\mathcal{I} = 1/2$, 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, $\mathcal{I} / \overline{M}$. 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

$$\lambda_g = \begin{cases} D / \sqrt{(g) / \overline{M}}, & \text{for } 0 < g \le 1/2 \\ D / \sqrt{(1-g) / \overline{M}}, & \text{for } 1/2 < g \le 1 \end{cases} \tag{11}$$

the green-noise principal wavelength. This implies that the parent process, ϕ_p, is itself a blue-noise point process with intensity $\mathcal{I} / \overline{M}$. If we assume a stationary and isotropic green-noise pattern, the pair correlation will have the form of Figure 1(a) given that:

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

▲ neighboring clusters are located at an average distance of λ_g apart

▲ as r increases, the influence that clusters have on neighboring clusters decreases.

The result is a pair correlation that has a nonzero component for $0 \le r < r_c$ due to clustering, a decreasing influence as r increases, and peaks at integer multiples of 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 Figure 1(a).

Assuming that the variation in cluster size is small for a given I_g, the placement of clusters λ_g apart leads to a strong spectral peak in $P(f_\rho)$ at $f_\rho = f_g$, the green-noise principal frequency:

$$f_g = \begin{cases} \sqrt{(g) / \overline{M}} / D, & \text{for } 0 < g \le 1/2 \\ \sqrt{(1-g) / \overline{M}} / D, & \text{for } 1/2 < g \le 1. \end{cases} \tag{12}$$

From (12) we make two intuitive observations: as the average size of clusters increases, f_g approaches DC; and as the size of clusters decreases, f_g approaches f_b. Figure 1(b) illustrates the desired characteristics of $P(f_\rho)$ for ϕ_G showing three distinct features: 1) little or no low-frequency spectral components, 2) high-frequency spectral components that diminish with increased clustering, and 3) a spectral peak at $f_\rho = f_g$.

The sharpness of the spectral peak in $P(f_\rho)$ at the green-noise principal 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 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.

Error Diffusion with Output-Dependent Feedback

Like Floyd and Steinberg's error diffusion, Levien's error-diffusion with output-dependent feedback (EDODF) predates 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

$$y[n] = \begin{cases} 1, & (x_a[n] = x[n] + x_e[n] + x_b[n]) \ge 0 \\ 0, & \text{else} \end{cases} \tag{13}$$

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

$$x_h[n] = h\sum_{i=1}^{N} a_i \cdot y[n-i] \tag{14}$$

with $\sum_{i=0}^{N} a_i = 1$ and h is an arbitrary constant. Called the hysteresis constant, h acts as a tuning parameter with larger h leading to coarser output textures [5] 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 ones and zeros with larger hysteresis constants leading to coarser halftones as in Figure 6 where the error/feedback coefficients are as shown in Figure 5 for $h = 0.5$, 1.0, 2.0, and 3.0.

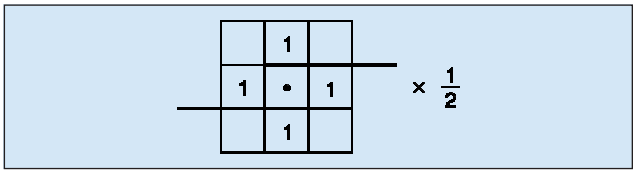

▲ *5. The two-error and two-feedback coefficient filter first introduced by Levien.*

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 becoming 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 four-weight filter for

▲ *6. The image Adrian halftoned using Levien's error-diffusion with output-dependent feedback for hysteresis values (a) h = 1.0 and (b) h = 2.0.*

▲ *7. The image Adrian halftoned using Levien's error-diffusion with output-dependent feedback for hysteresis values (a) h = 1.0 and (b) h = 2.0 using balanced weights.*

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 [26] 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 of h, and in so doing, they specified the optimal values of a_1 and a_2 versus h. Shown in Figure 7 are the halftone images produced using these balanced weights for $h = 0.5$, 1.0, 2.0, and 3.0, and shown in Figure 8 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).

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. [9] 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 iteratively adds points to a dither pattern such that the resulting pattern has a pair correlation matching that of the ideal green-noise pattern.

From BIPPCCA, Lau et al. [9], [27] 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 [27], multiple dither arrays were constructed in a correlated fashion to regulate the amount of dot overlap in color halftones. Figure 9 is an example of a monochrome binary halftone pattern created using a green-noise dither array.

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.

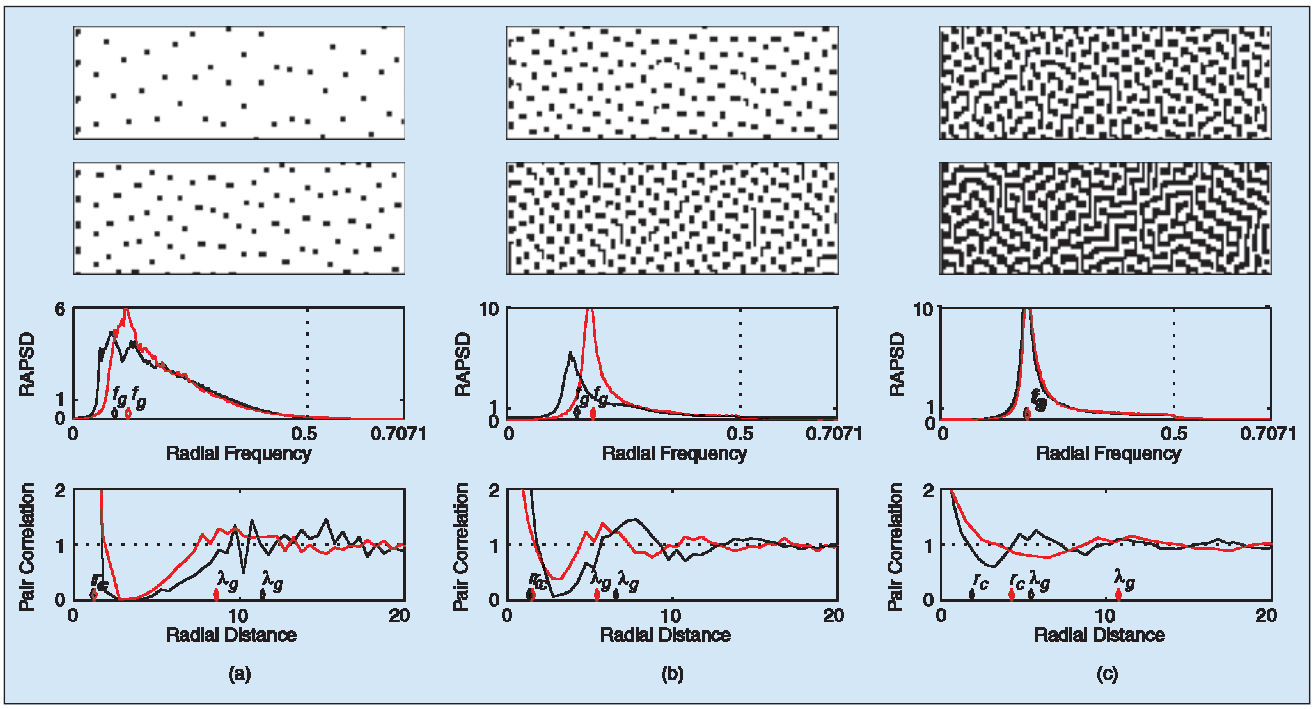

▲ *8. The spatial and spectral statistics for Levien's error-diffusion with output-dependent feedback using balanced weights with h = 2.5 at gray levels (a) g = 1/32 and 1/16 (red), (b) 1/8 and 1/4 (red), and (c) 1/3 and 1/2 (red).*

Visual Pleasantness

Figure 10 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 amplitude modulated (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].

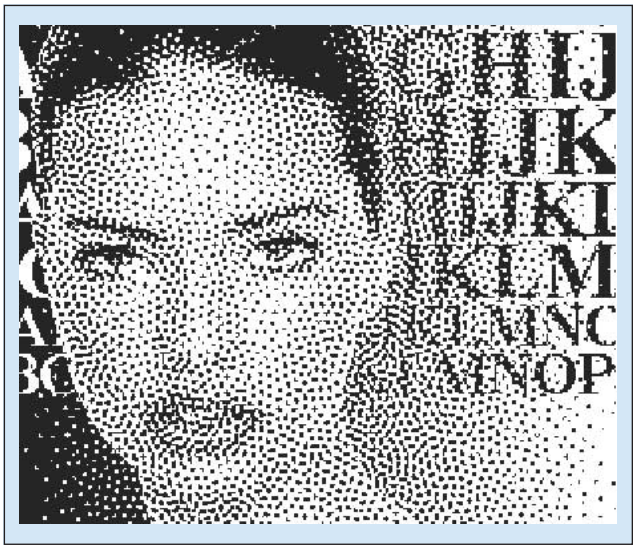

▲ *9. A binary halftone image created using a green-noise dither array.*

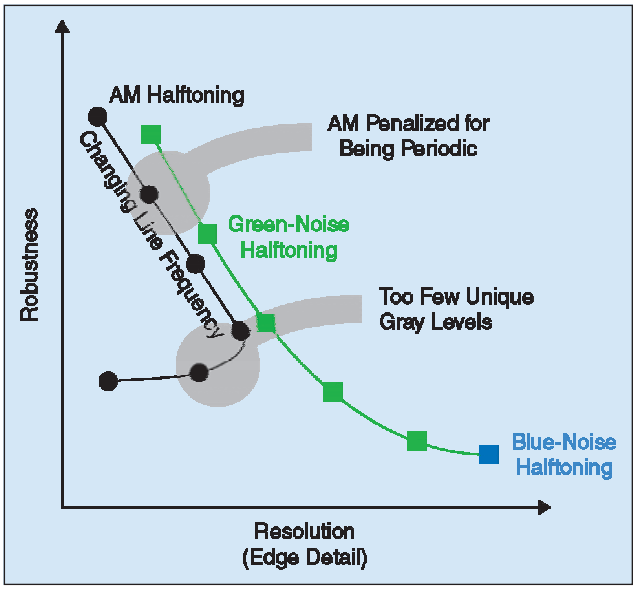

▲ *10. The relationships between spatial resolution and halftone robustness for the various halftoning paradigms for a fixed print resolution.*

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, 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 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 Figure 10, the line representing green noise is shown as always having a visual fidelity higher than AM at the same measure of robustness.

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' [21] 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 [28] or is estimated using a printed dot model [29].

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 [30], [31]. 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 [32]. The patterns are, therefore, inappropriate in devices that do not produce isolated pixels reliably such as in laser printers [33], where compensating for distortion is much more complicated as isolated dots are more sensitive to process variation [3]. 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 [6]. How robustness is achieved is through clustering [35] 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 [3], [4], 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 Figure 10 with AM halftoning providing the most robust patterns with green noise bridging the gap with blue noise.

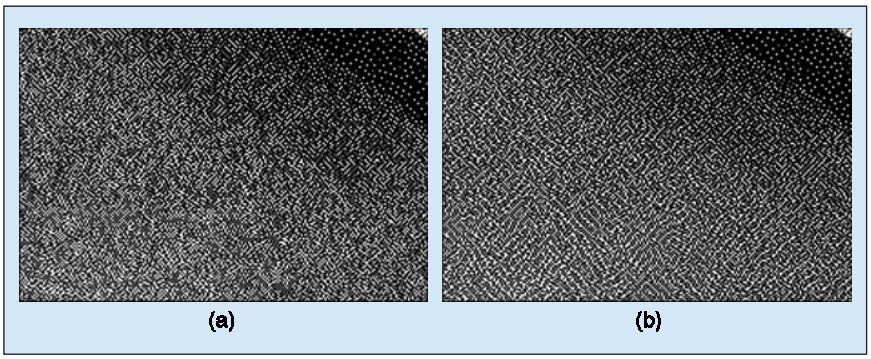

▲ *11. Color halftones produced using (a) blue and (b) green noise with no correlation between colors showing the varying effects of stochastic moiré where the patchiness of moiré is reduced using green noise by varying the coarseness of patterns between colors.*

Color Halftoning: Stochastic Moiré

In traditional color halftoning devices, the binary halftones of CMYK are superimposed [36], 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°, 75°, 0°, and 45°, respectively [37], [38]. Because it is commonly believed that moiré is a result of superimposing regular patterns, stochastic halftones such as those produced by blue noise are believed to avoid the moiré phenomenon allowing for the introduction of low-cost color ink-jet printers [34]. But superimposing uncorrelated dispersed-dot patterns does lead to low-frequency graininess commonly referred to in halftoning literature as color-noise [39].

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 [10]. 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 interlocking screens [38], [40]. 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 misregistration 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. [10], [11] 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, as in Figure 11. 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.

Daniel L. Lau received his Ph.D. degree from the University of Delaware in 1999 and is an assistant professor in the Department of Electrical and Computer Engineering at the University of Kentucky. Prior to joining the University of Kentucky, he performed research at Lawrence Livermore National Laboratory and served as a DSP engineer at Aware Inc, Bedford, Massachusetts, developing ADSL modems. His research interests include 3-D video processing and visual perception of 3-D images, and he is currently involved in research projects involving telecollaboration for minimally invasive surgery and telemedicine for home health care using 3-D video and automated surveillance.

Robert Ulichney received a Ph.D. from the Massachusetts Institute of Technology in electrical engineering and computer science. He was with Digital Equipment Corp. from 1978-1998 where he focused on image and video implementations for a number of different organizations. From 1998-2002 he was with Compaq Computer Corporation's Cambridge Research Lab where he led a number of research efforts in video and image processing. He is now a researcher at the Printing and Imaging Center at HP Labs, Massachusetts. He has several issued patents for contributions to a range of products in the areas of hardware and software-only motion video, graphic controllers, and hardcopy.

Gonzalo R. Arce received the Ph.D. degree from Purdue University, Indiana, in 1982 and has since been with the faculty of the Department of Electrical and Computer Engineering at the University of Delaware where he is the Charles Black Evans Professor and chairman of Electrical and Computer Engineering. His interests include statistical/nonlinear signal processing, multimedia security, electronic imaging and display, and digital communications. He received the NSF Research Initiation Award and is a Fellow of the IEEE. He was the cochair of the 2001 EUSIPCO/IEEE Workshop on Nonlinear Signal and Image Processing (NSIP'01), the 1991 SPIE's Symposium on Nonlinear Electronic Imaging, and the 2002 and 2003 SPIE ITCOM conferences. He was an associate editor for *IEEE Transactions for Signal Processing*, senior editor of the *Applied Signal Processing Journal*, and guest editor for *IEEE Transactions on Image Processing* and *Journal Optics Express*.

References

[1] R.A. Ulichney, "Dithering with blue noise," *Proc. IEEE*, vol. 76, pp. 56-79, Jan. 1988.

[2] F.W. Campbell, "The human eve as an optical filter," *Proc. IEEE*, vol. 56, pp. 1009-1014, June 1968.

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

[4] M. Rodriguez, "Promises and pitfalls of stochastic screening in the graphic arts industry," in *Proc. IS&T's 8th Int. Congr. Advances in Non-Impact Printing Technologies*, 1992, pp. 34-37.

[5] R. Levien, "Output dependant feedback in error diffusion halftoning," in *Proc. IS&T's 8th Int. Congr. Advances in Non-Impact Printing Technologies*, Williamsburg, VA, Oct. 25-30, 1992, pp. 280-282.

[6] D.L. Lau, G.R. Arce, and N.C. Gallagher, "Green-noise digital halftoning," *Proc. IEEE*, vol. 86, pp. 2424-2444, Dec. 1998.

[7] T. Mitsa and K.J. Parker, "Digital halftoning technique using a blue noise mask," *J. Opt. Soc. Amer.*, vol. 9, pp. 1920-1929, 1992.

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

[9] D.L. Lau, G.R. Arce, and N.C. Gallagher, "Digital halftoning by means of green-noise masks," *J. Opt. Soc. Amer.*, vol. 16, pp. 1575-1586, July 1999.

[10] D.L. Lau, A.M. Khan, and G.R. Arce, "Minimizing stochastic moire by means of green-noise masks," *J. Opt. Soc. Amer. A*, vol. 19, pp. 2203-2217, Nov. 2002.

[11] D.L. Lau, A.M. Khan, and C.R. Arce, "Stochastic moir6 in color halftoning," in *Proc. IS&T's Image Processing Image Quality Image Capture Systems Conf. (PICS'O1)*, Montreal, Quebec, Canada, Apr. 2001, pp. 96-100.

[12] N.A.C. Cressie, *Statistics for Spatial Data*. New York: Wiley, 1983.

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

[14] M.S. Bartlett, "The spectral analysis of a point process," *J. R. Stat. Soc. B*, vol. 25, no. 2, pp. 264-280, 1964.

[15] M.S. Bartlett, "The spectral analysis of two-dimensional point processes," *Biometrika*, vol. 51, no. 3/4, pp. 299-311, Dec. 1964.

[16] M. Gardner, "White and brown music, fractal curves and one-over-f fluctuations," *Sci. Amer.*, vol. 328, pp. 16-32, Apr. 1978.

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

[18] B.L. Evans, V. Monga, and N. Damera-Venkata, "Variations on error diffusion: Retrospectives and future trends," in *Proc. SPIE/IS&T Conf. Color Imaging: Processing, Hardcopy, and Applications*, Jan. 2003, pp. 371-389.

[19] K. Knox, "Evolution of error liffusion," *J. Electron. Imaging*, vol. 8, no. 4, pp. 422 429, 1999.

[20] I.H. Whitten and M. Neal, "Using peano curves for bilevel display of continuous-tone images," *IEEE Comput. Graph. Appl.*, vol. 2, pp. 47-52, May 1982

[21] L. Velho and J.M. Gomes, "Digital halftoning with space filling curves," *Comput. Graph.*, vol. 25, no. 4, pp. 81-90, 1991.

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

[23] P. Stucki, "Mecca-a multiple-error correcting computation algorithm for bilevel image hardcopy reproduction," IBM Res. Lab., Zurich, Switzerland, Tech. Rep. RZ1O6O, 1981.

[24] R. Eschbach and K.T. Knox, "Error-diffusion algorithm with edge enhancement," *J. Opt. Soc. Amer.*, vol. 8, pp. 1844-1850, Dec. 1991.

[25] K.E. Spaulding, R.L. Miller, and J. Schildkraut, "Methods for generating blue-noise dither matrices for digital halftoning," *J. Electron. Imaging*, vol. 6, no. 2, pp. 208-230, 1997.

[26] D. L. Lau, G. R. Arce, and N. C. Gallagher, "Robust halftoning with green-noise," in *Proc. IS&T's Image Processing Image Quality Image Capture Systems Conf. (PICS '99)*, Savannah, GA, Apr. 26-28 1999, pp. 315-320.

[27] D.L. Lau, G.R. Arce, and N.C. Gallagher, "Digital color halftoning via generalized error-diffusion and vector green-noise masks," *IEEE Trans. Image Processing*, vol. 9, no. 5, pp. 923-935, May 2000.

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

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

[30] P.G. Roetling and T.M. Holladay, "Tone reproduction and screen design for pictorial electrophotographic printing," *J. Appl. Photoqraph. Eng.*, vol. 15, no. 4, pp. 179-182, 1979.

[31] T.N. Pappas and D.L. Neuhoff, "Model-based halftoning," in *Proc. SPIE, Human Vision, Vision Processing and Digital Display II*, vol. 1453, B.E. Rogowitz, M.H. Brill, and Jan P. Allebach, Eds., June 1991, pp. 244-255.

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

[33] S. Aoki, "New halftoning method using adaptive cell," in *IS&T's NIP 14: Int. Conf. Digital Printing Technologies*, Toronto, Ontario, Canada, Oct. 18-23 1998, pp. 277-280.

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

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

[36] R.W.G. Hunt, *The Reproduction of Color in Photography, Printing and Television*. Tolworth, U.K.: Fountain Press, 1987.

[37] I. Amidror, R.D. Hersch, and V. Ostromoukhov, "Spectral analysis and minimization of moire patterns in color separation," *J. Electron. Imaging*, vol. 3, no. 3, pp. 295-317, 1994.

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

[39] J.S. Liu and F.H. Cheng, "Color halftoning—A non-separable model," in *Proc. Int. Conf. Image Processing*, 1996, pp. 561-564.

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