

Green-Noise Digital Halftoning

DANIEL L. LAU, GONZALO R. ARCE, SENIOR MEMBER, IEEE,
AND NEAL C. GALLAGHER, FELLOW, IEEE

In this paper, we introduce the concept of green noise—the midfrequency component of white noise—and its advantages over blue noise for digital halftoning. Unlike blue-noise dither patterns, which are composed exclusively of isolated pixels, green-noise dither patterns are composed of pixel-clusters making them less susceptible to image degradation from nonideal printing artifacts such as dot-gain. Although they are not the only techniques which generate clustered halftones, error-diffusion with output-dependent feedback and variations based on filter weight perturbation are shown to be good generators of green noise, thereby allowing for tunable coarseness. Using statistics developed for blue noise, we closely examine the spectral content of resulting dither patterns. We introduce two spatial-domain statistics for analyzing the spatial arrangement of pixels in aperiodic dither patterns, because green-noise patterns may be anisotropic, and therefore spectral statistics based on radial averages may be inappropriate for the study of these patterns.

Keywords—Blue-noise dithering, clustered point process, digital halftoning, green-noise dithering, point process, stochastic geometry.

I. INTRODUCTION

With the ultimate goal of accurately reproducing an original continuous-tone photograph without loss of tonal value or detail, printing devices such as those which rely on dye sublimation or gravure printing have been developed that are capable of rendering continuous tone. More often though, printing devices are capable of only limited-tone reproduction, and in the case of binary output, intermediate tones are represented as halftones. Examples of such devices are laser and ink-jet printers and facsimile machines.¹

In commercial printing, three printing processes that produce binary tone are letterpress or relief printing, lithography or planographic printing, and screen or porous printing. These printing processes have, in the past, relied on analog photomechanical screening methods with halftones consisting of rows of dots fixed along a grid in a regular pattern, equally spaced center-to-center. These dots vary in

size depending on the tone being rendered [2]. With the advent of digital image setters (the device which converts the original continuous-tone image to binary bitmap and then from electronic form to physical), printing has been rapidly moving to digital binary methods [3]. Originally, these digital methods attempted to imitate photomechanical screening methods [4] through clustered dot dithering. This type of digital halftoning is referred to as amplitude modulated (AM) screening.

As an alternative to AM screening, frequency modulated (FM) or stochastic screening techniques, where dots are of constant size but are variably spaced according to tone, are available to digital printers. Although relatively rare in commercial printing [5], FM screening avoids the problems associated with AM methods [4] such as moiré, the interference pattern created by superimposing two or more regular patterns. Moiré, a problem associated with color printing, is minimized in commercial printers by offsetting the orientation of the halftone screen of each color-creating a desirable circular pattern called a rosette. Great care must be taken when aligning these screens; otherwise, moiré will result [2]. A further drawback of AM screening is image contouring, a visual banding effect created by an abruptly changing halftone texture. This artifact is minimized in AM patterns by increasing the maximum number of gray levels that a cluster can represent, but in order to increase the number of achievable gray levels, the total number of pixels that compose a cluster must also be increased—reducing the spatial resolution of the resulting halftone image.²

The alternative screening technique, FM screening, eliminates the need for screen angles and screen rulings, resulting in an image without artifacts and displaying higher spatial resolution; furthermore, FM screens, by using a statistical analysis of neighboring pixels, can represent images with greater fidelity [4]. The major drawback of FM screens is dot gain—the increase in size of the printed dot relative to the intended dot size of the original halftone film.

Manuscript received February 11, 1998; revised October 5, 1998.

The authors are with the Department of Electrical and Computer Engineering, University of Delaware, Newark, DE 19717 USA.

Publisher Item Identifier S 0018-9219(98)09351-7.

¹Although current technology regarding laser and ink-jet printers does allow for multilevel output, the number of achievable gray levels is still too limited to eliminate quantization effects, and in these cases, halftones are still required in order to produce a sufficient range of intermediate tones [1].

²By dividing a large halftone cell (the sample area belonging to a single cluster of an AM pattern) into smaller cells, thereby creating a conglomerate or super cell [6], AM halftone patterns can represent the same number of gray levels but with greater spatial resolution as the size of individual clusters is decreased.

When printing halftones, dot gain creates a reduction or compression in the printed tonal range leading to a loss of definition and detail (contrast); furthermore, dot gain can lead to “plugged” or filled-in screens and cause a shift in color [7].

Due to the nature of dot gain, FM screens will typically suffer greater amounts of distortion than AM [4], but it is not the amount of dot gain which limits the use of FM halftoning. Instead, it is the reproducibility of the printed dot. In a repeatable process where the variation in dot size/shape from printed dot-to-printed dot is small, printing can rely on dot gain compensation techniques to minimize introduced distortion, but in a nonrepeatable process where the variation in dot size/shape is high, pattern robustness (clustering) becomes a desired, and in many cases a required, characteristic of the halftoning process [8]. In summary, the choice between AM and FM screens is a function of the imaging system’s ability to print individual pixels. If individual pixels can be reliably reproduced, the halftone can be based on individual pixels. Otherwise, the halftone must be composed of groups of pixels which, through clustering, form larger features that can be reliably reproduced [3].

In this paper, we look at AM–FM hybrids, stochastic halftoning techniques which create minority pixel clusters which vary, according to tone, in both their size and spacing. The advantage to using these techniques is that as stochastic processes, the superposition of two or more halftone patterns does not create moiré, thereby alleviating the need for screen angles; furthermore, by clustering minority pixels, AM–FM hybrids are also less susceptible to the effects of dot gain. Techniques which fall in this category are not new to halftoning literature and include such techniques as Velho’s and Gomes’ digital halftoning along space-filling curves (SFC’s) [9], Scheermesser’s and Bryngdahl’s digital halftoning with texture control [10], and Levien’s error diffusion with output-dependent feedback [11]. For reasons to be discussed later, this paper focuses on Levien’s technique, but the presented aspects of AM–FM halftoning will apply to many of the halftoning schemes not presented.

In order to introduce this new screen, consider first the implementations of AM and FM halftoning. The clustering of pixels in AM screens has always implied using clustered-dot dithering techniques for transforming continuous-tone images to binary form. FM screens, until recently, have always implied dispersed-dot ordered dithering. Studied in great detail by Ulichney [12], error diffusion creates dither patterns which are most pleasant in areas of uniform tone when these patterns have isotropic (radially symmetric) blue noise³ spectral characteristics. Ulichney writes the following on the subject [12].

The purpose of a dither pattern is to represent a continuous-tone level. It therefore should not have any form or structure of its own; a pattern succeeds when it is innocuous. Blue noise is visually pleasant

³Blue noise refers to the high-frequency component of white noise. Such patterns have a minimal low-frequency spectral component.

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.

Therefore compared to dispersed-dot ordered dither, error diffusion creates patterns which are visually more pleasing.

By adding an output-dependent feedback term to error diffusion, Levien [11] has shown that pixels can be clustered, forming a stochastic screen which maintains properties of both the AM and FM screens. Furthermore, by adjusting a single parameter, Levien has demonstrated the ability to increase or decrease the amount of clustering, thereby allowing for a tunable screen which can be optimized according to the reliability of an imaging system to print individual pixels. In this paper we show that the spectral characteristics of these new patterns are not blue-noise like traditional error-diffusion halftones but green, i.e., containing no low-frequency or high-frequency spectral components.

Introducing the green-noise model to digital halftoning is the primary focus of this paper. Because of the tunability of AM–FM hybrids to create varying amounts of clustering, green noise, as presented here, is a tunable model with blue noise as a limiting case. In addition to introducing a new statistical model for digital halftoning, this paper also introduces new spatial statistics for the evaluation of produced patterns. Previously, statistics for the analysis of binary halftone patterns have been based on spectral domain estimates. These spectral statistics are responsible for the name “blue noise” given to stochastic patterns created via error-diffusion and for the name “green noise” for AM–FM hybrid patterns. Like the spectral statistics, the newly introduced spatial statistics uniquely identify the spatial characteristics inherent in a particular halftoning process; furthermore, being based in the spatial domain, these new statistics have interpretations which are far more intuitive to the viewer.

II. STATISTICAL ANALYSIS OF DIGITAL HALFTONES

A. Spectral Statistics

In order to provide a mechanism for studying aperiodic patterns, Ulichney [12] developed the radially averaged power spectra along with a measure of anisotropy. Both rely on estimating the power spectrum through Bartlett’s method of averaging periodograms, i.e., the magnitude-square of the Fourier transform of the output pattern divided by the sample size, to produce the spectral estimate $\hat{P}(f)$. Although anisotropies of a dither pattern can be qualitatively observed by studying three-dimensional (3-D) plots of $\hat{P}(f)$, partitioning the spectral domain into a series of annular rings $R(f_\rho)$ of width Δ_ρ , as shown in Fig. 1, leads to two useful one-dimensional (1-D) statistics. The first statistic is the radially averaged power spectrum density (RAPSD) $P(f_\rho)$, defined for discrete $\hat{P}(f)$ as the average power in the annular ring with center radius f_ρ

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

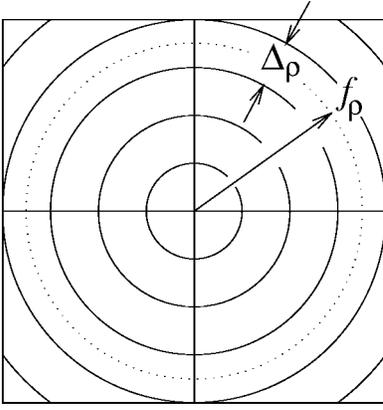


Fig. 1. By partitioning the spectral domain into a series of annular rings, the power spectrum $P(f)$ can be studied using $P(f_\rho)$ and $A(f_\rho)$.

where $N(R(f_\rho))$ is the number of frequency samples in $R(f_\rho)$. The second statistic is the anisotropy $A(f_\rho)$. It is defined as

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

the relative variance or the “noise-to-signal” ratio [12] of frequency samples of $\hat{P}(f)$ in $R(f_\rho)$.

Built on the premise that aperiodic patterns generated in areas of uniform gray are most pleasant when patterns are isotropic, $A(f_\rho)$, a measure of how isotropic a dither pattern is, implies that the pattern with lower anisotropy is the more visually pleasing of the two if both patterns have identical $P(f_\rho)$. Given that its purpose is to measure the strength of directional artifacts, $A(f_\rho)$ does not indicate the direction. It is therefore possible for a dither pattern with strong horizontal artifacts to be indistinguishable according to $P(f_\rho)$ and $A(f_\rho)$ from a pattern with strong diagonal artifacts. The human eye being less sensitive to diagonal correlations, one may find the difference in appearance far from indistinguishable.

Note that for either statistic, a rotation in the original dither pattern has no effect on either $P(f_\rho)$ or $A(f_\rho)$, and these metrics may therefore be insufficient at describing patterns uniquely when patterns are allowed and even preferred to be anisotropic. Additional information pertaining to the direction of artifacts may then be required. In order to supplement $P(f_\rho)$ and $A(f_\rho)$, we can use spatial domain statistics to characterize aperiodic dither patterns. Though largely ignored in halftone literature, the importance of spatial domain statistics cannot be overlooked. Found in the literature of stochastic geometry (the area interested in modeling complicated geometrical patterns), point processes have long been studied by statisticians to model such things as crystal formation in granite and cell locations within organic tissue. Because point process statistics are used to model the location of random points in space, they are suited perfectly for characterizing digital halftones.

B. Spatial Statistics

Though spectral domain techniques have been proposed, statisticians are more likely to rely on “nearest-neighbor” distributions, i.e., the distribution of points relative to a typical point. Several nearest-neighbor measures are extremely useful to halftoning, and in this paper we show two such spatial-statistic metrics which can be very useful in characterizing halftone patterns. We begin with a more concise definition of a point process.

The point process Φ is a stochastic model governing the location of events, or points x_i , within the space \mathfrak{R}^2 [13]. ϕ is a sample of Φ and will be written as $\phi = \{x_i \in \mathfrak{R}^2 : i = 1, \dots, N\}$. Furthermore, $\phi(B)$ is a scalar quantity defined as the number of x_i s in the subset B of \mathfrak{R}^2 . We assume that the point process Φ is simple, meaning that $i \neq j$ implies $x_i \neq x_j$, which further implies

$$\lim_{dV_x \rightarrow 0} \phi(dV_x) = \begin{cases} 1, & \text{for } x \in \phi \\ 0, & \text{else} \end{cases} \quad (3)$$

where dV_x is the infinitesimally small area around x . In terms of a discrete dither pattern, ϕ represents the set of minority pixels such that $\phi[n] = 1$, for pixel index n indicates a minority pixel at location n . Φ now represents an aperiodic halftone process with Φ_B and Φ_G representing blue-noise and green-noise halftoning processes, respectively.

From its definition in (3), $\phi(x_i)$ is a scalar random variable which can be characterized in terms of its moments. We start with the first-order moment, the intensity $\mathcal{I}(x)$

$$\mathcal{I}(x) = \lim_{dV_x \rightarrow 0} \frac{\mathbf{E}\{\phi(dV_x)\}}{dV_x}. \quad (4)$$

For a point process to be stationary, the statistical characteristics of Φ must be invariant to translation. If a process is stationary, then the intensity is constant for all $x \in \mathfrak{R}^2$ where $\mathcal{I}(x) = \mathcal{I}$ is the expected number of points per unit area. Furthermore, $\mathcal{I}[n]$ represents the unconditional probability that the sample at location n is a minority pixel.

We gain additional insight into Φ by conditioning the probability distribution of Φ given that a point lies at a given location. The result is a conditional distribution referred to as the Palm distribution [14]. A particular measure under the Palm distribution of Φ is the quantity $\mathcal{K}(x; y)$ (or for discrete space $\mathcal{K}[n; m]$)

$$\mathcal{K}(x; y) = \lim_{dV_x \rightarrow 0} \frac{\mathbf{E}\{\phi(dV_x) \mid y \in \phi\}}{\mathbf{E}\{\phi(dV_x)\}} \quad (5)$$

the ratio of the expected number of points in dV_x under the condition $y \in \phi$ to the unconditional expected number of points in dV_x . $\mathcal{K}(x; y)$, referred to as the reduced second-moment measure, may be thought of as the influence at location x of the point y . That is, is a point at x more or less likely to occur because a point exists at y ? For stationary processes, $\mathcal{K}(x; y)$ may be written as $\mathcal{K}(r, \theta)$ where r is the distance between x and y while θ is the direction to x from y . For a point process to be isotropic, the statistical

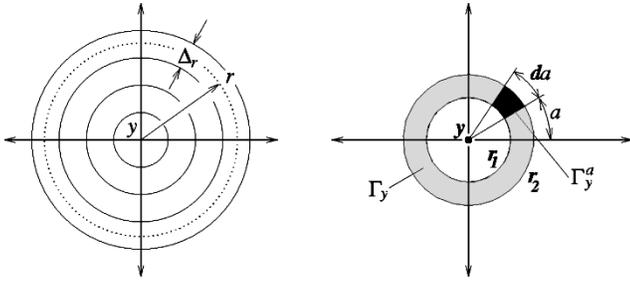


Fig. 2. By dividing the spatial domain into a series of (left) annular rings centered around location y , the spatial arrangement of points can be studied through the use of $\mathcal{R}(r)$, and by looking at the expected number of points per unit area in the segment Γ_y^a versus the ring Γ_y , the spatial arrangement can be studied using $\mathcal{D}_{r_1, r_2}(a)$.

characteristics of Φ must be invariant to rotation; therefore, if Φ is also isotropic, then $\mathcal{K}(x; y) = \mathcal{K}(r)$.

From $\mathcal{K}(x; y)$, we can derive two spatial domain statistics in a similar fashion to $P(f_\rho)$ and $A(f_\rho)$ by partitioning the spatial domain into a series of annular rings $R_y(r)$ (Fig. 2), with center radius r , width Δ_r , and centered around location y . In this paper, the annular ring $R_y(r)$ will be the set $\{x : r - \Delta_r/2 < |x - y| \leq r + \Delta_r/2\}$ where $\Delta_r = 1/2$. The first statistic for stationary and isotropic Φ is the pair correlation $\mathcal{R}(r)$, defined as

$$\mathcal{R}(r) = \frac{\mathbf{E}\{\phi(R_y(r)) \mid y \in \phi\}}{\mathbf{E}\{\phi(R_y(r))\}} \quad (6)$$

the influence that the point at y has at any x in the annular ring $R_y(r)$. Note that for a stationary point process, the unconditional expected number of points in the ring $R_y(r)$ is $\mathcal{I} \cdot N(R_y(r))$ (the intensity times the area of $R_y(r)$). $\mathcal{R}(r)$ is also the average value of $\mathcal{K}(x; y)$ in the ring $R_y(r)$, and its usefulness can be seen in the interpretation that maxima of $\mathcal{R}(r)$ indicate frequent occurrences of the interpoint distance r while minima of $\mathcal{R}(r)$ indicate an inhibition of points at r [14].

In addition to the above spatial interpretations, $\mathcal{R}(r)$ maintains an intimate link with the spectral domain. In particular, a frequent occurrence of the interpoint distance r , indicated by maxima in $\mathcal{R}(r)$, implies a peak in $P(f_\rho)$ for radial frequency $f_\rho = 1/r$ proportional in magnitude to the peak in $\mathcal{R}(r)$, meaning that a larger peak leads to a larger peak in $P(f_\rho)$. Take for instance white-noise halftoning where binary patterns are created by thresholding an input image with uniformly distributed, uncorrelated (white) noise.

Deriving its name from its spectral content, white-noise halftones have a $P(f_\rho)$ which is flat for all f_ρ . The lack of any spectral peaks implies that white-noise halftones have a pair correlation which is one for all r . This is the case since a completely random,⁴ stationary, and isotropic point process with intensity \mathcal{I} has a reduced second moment

⁴This is a completely random process if $\mathbf{E}\{\phi(dV_x) \mid y \in \phi\} = \mathbf{E}\{\phi(dV_x)\}$ for all $y \neq x, y \in \mathbb{R}$.

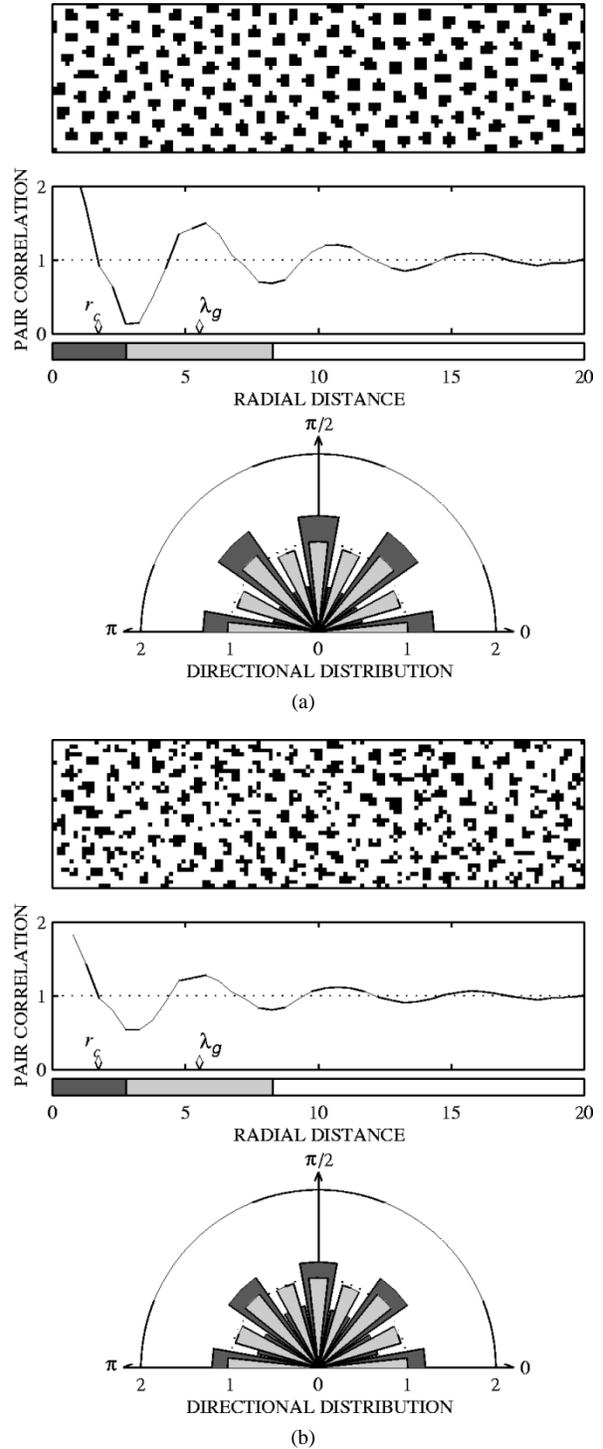


Fig. 3. Dither pattern, pair correlation and directional distribution function of (a) a dither pattern with small variation in cluster shape and (b) a dither pattern with high variation in cluster shape.

measure

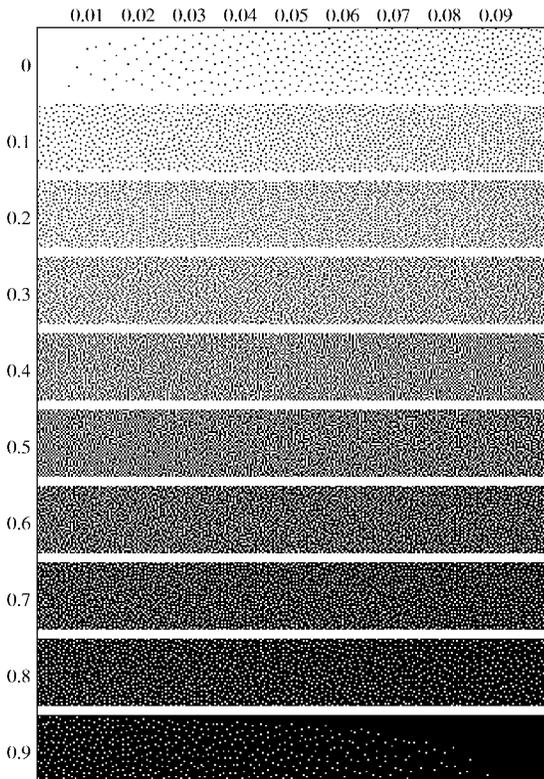
$$\mathcal{K}(x; y) = \lim_{dV_x \rightarrow 0} \frac{\mathbf{E}\{\phi(dV_x)\}}{\mathbf{E}\{\phi(dV_x)\}} = 1 \quad (7)$$

for all $x, y \in \mathbb{R}^2$, and since $\mathcal{R}(r)$ is the average $\mathcal{K}(x; y)$ in the ring $R_y(r)$ for all $r > 0$, $\mathcal{R}(r) = 1$ for all $r > 0$.

In cases where Φ is stationary but not isotropic, $\mathcal{K}(x; y)$ can be used to investigate anisotropy in Φ by defining the



(a)



(b)

Fig. 4. Blue-noise halftone of (a) a grayscale image and (b) a grayscale ramp.

directional distribution function $\mathcal{D}_{r_1, r_2}(a)$ as

$$\mathcal{D}_{r_1, r_2}(a) = \frac{\mathbf{E}\{\phi(\Gamma_y^a) \mid y \in \phi\} / N(\Gamma_y^a)}{\mathbf{E}\{\phi(\Gamma_y) \mid y \in \phi\} / N(\Gamma_y)} \quad (8)$$

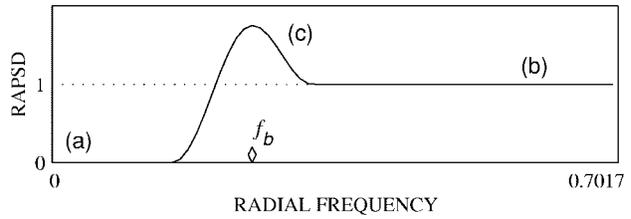


Fig. 5. The spectral characteristics of a blue-noise halftone pattern in units of $\sigma_g^2 = g(1 - g)$ with: (a) a low-frequency cutoff at principle frequency f_b ; (b) a sharp transition region; and (c) a flat high-frequency blue-noise region.

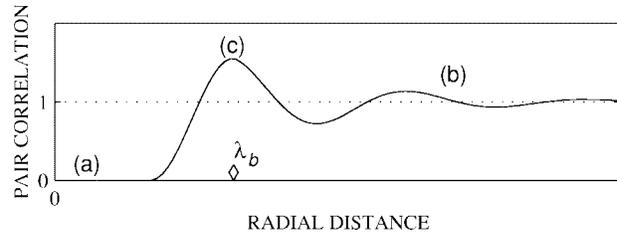


Fig. 6. The pair correlation for a blue-noise process showing: (a) a strong inhibition of minority pixels near $r = 0$; (b) a decreasing correlation of minority pixels with increasing r ; and (c) a frequent occurrence of the interpoint distance λ_b indicated by a series of peaks at integer multiples of λ_b .

• 7/16

1/16 5/16 3/16

Fig. 7. Floyd–Steinberg filter weights.

the expected number of points 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 around the point $y \in \phi$ such that $a \leq \angle(x - y) < a + \Delta_a$ (see Fig. 2), to the expected number of points per unit area in Γ_y itself. Note that for isotropic point patterns, $\mathcal{D}_{r_1, r_2}(a) = 1.0$ for all a , and that $\mathcal{D}_{r_1, r_2}(a) > 1.0$ indicates a favoring of points at angles near a while $\mathcal{D}_{r_1, r_2}(a) < 1.0$ indicates an inhibition of points. Furthermore, the parameters r_1 and r_2 allow us to look at point distributions for various ranges from a point y . Most notably, a near distribution for the range $r \in [0, r_1)$, a far distribution for $r \in [r_1, \infty)$, and an overall distribution for $r \in [0, \infty)$.

Compared to the spectral statistics $P(f_\rho)$ and $A(f_\rho)$, the spatial statistics $\mathcal{R}(r)$ and $\mathcal{D}_{r_1, r_2}(a)$ together offer a more intuitive understanding of the underlying point process Φ . This is perhaps best witnessed in Fig. 3 where minority pixels form clusters. In this case, clustering results in a nonzero pair correlation for r near zero with $\mathcal{R}(r) = 1$ occurring at $r = r_c = 1.73$. This parameter r_c is proportional to the size of clusters as $\pi r_c^2 = \bar{M}$. That is, the area covered by the average sized cluster of \bar{M} pixels is equal to the area covered a circle of radius r_c . Furthermore, the peaks of $\mathcal{R}(r)$, which occur at integer multiples of $\lambda = 5.53$, indicate that pixel clusters are separated center-to-center by an average distance of λ with the sharpness of peaks inversely proportional to the amount of variation in cluster size. Fig. 3(a) and (b) shows that when the variation in cluster size and shape is small, the

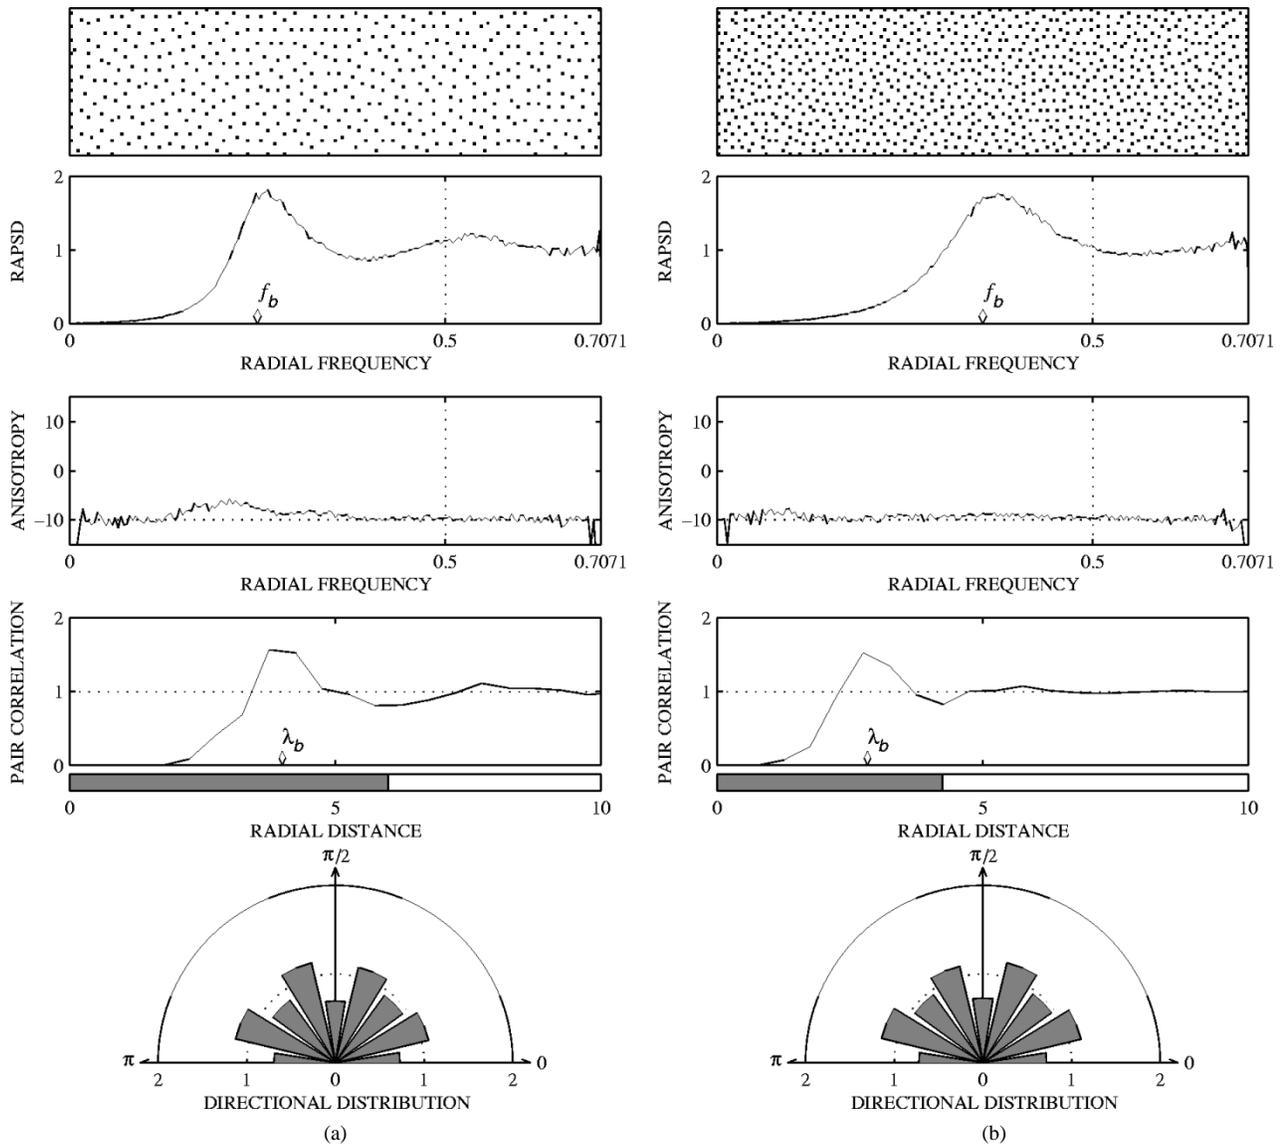


Fig. 8. Analysis of blue-noise halftone patterns for (a) $g = 1/16$ and (b) $g = 1/8$.

pair correlation has very sharp peaks, while large variations lead to a pair correlation which is blurred (lower contrast between maxima and minima).

III. BLUE-NOISE HALFTONING

Blue-noise halftoning, labeled as Φ_B and demonstrated in Fig. 4, is characterized by a distribution of binary pixels where the minority pixels are spread as homogeneously as possible [12]. Distributing pixels in this manner creates a pattern which is aperiodic, isotropic, and does not contain any low-frequency spectral components. The result of halftoning a continuous-tone discrete-space monochrome image with blue noise is a pattern which, as Ulichney [12] describes, “is ‘pleasant’—it (blue noise) does not clash with the structure of an image by adding one of its own, or degrade it by being too ‘noisy’ or uncorrelated.”

A. Spectral Statistics

Blue noise, when applied to an image of constant gray level g , spreads the minority pixels of the resulting binary

image such that pixels are separated by an average distance λ_b

$$\lambda_b = \begin{cases} D/\sqrt{g}, & \text{for } 0 < g \leq 1/2 \\ D/\sqrt{1-g}, & \text{for } 1/2 < g \leq 1 \end{cases} \quad (9)$$

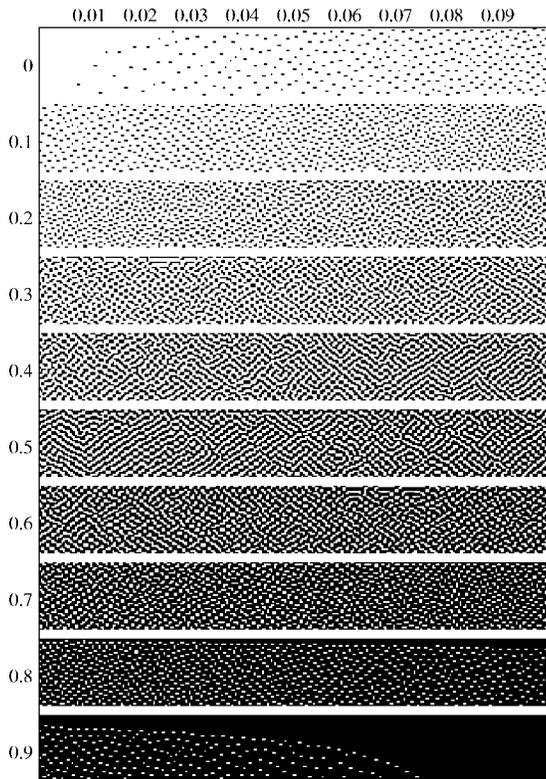
where D is the minimum distance between addressable points on the display [12], [15]. The parameter λ_b is referred to as the principle wavelength of blue noise, with its relationship to g justified by several intuitive properties.

- i) As the gray value approaches perfect white ($g = 0$) or perfect black ($g = 1$), the principle wavelength approaches infinity.
- ii) Wavelength decreases symmetrically with equal deviations from black and white toward middle gray ($g = 1/2$).
- iii) The square of the wavelength is inversely proportional to the number of minority pixels per unit area.

Again we note that the distribution of minority pixels is assumed to be stationary and isotropic.



(a)



(b)

Fig. 9. Green-noise halftone of (a) a grayscale image and (b) a grayscale ramp.

Turning to the spectral domain, the spectral characteristics of blue noise in terms of $P_\rho(f_\rho)$ are shown in Fig. 5 and can be described by three unique features: 1) little or no low

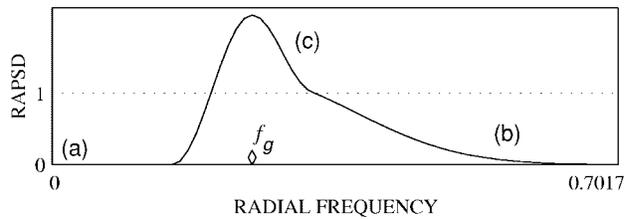


Fig. 10. The spectral characteristics of a green-noise halftone pattern in units of $\sigma_g^2 = g(1-g)$ with: (a) little or no low frequency spectral components; (b) high-frequency spectral components which diminish with increased clustering; and (c) a spectral peak at $f_\rho = f_g$.

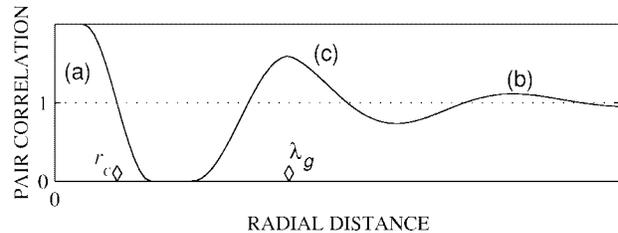


Fig. 11. The pair correlation for green noise showing: (a) pixel clustering with cluster radius r_c ; (b) decreasing influence as r increases; and (c) the green-noise principle wavelength λ_g .

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 principle frequency, such that

$$f_b = \begin{cases} \sqrt{g}/D, & \text{for } 0 < g \leq 1/2 \\ \sqrt{1-g}/D, & \text{for } 1/2 < g \leq 1. \end{cases} \quad (10)$$

As will be the convention for this paper, the principle frequency is indicated in Fig. 5 by a diamond located along the horizontal axis. Also note that $P(f_\rho)$ is plotted in units of $\sigma_g^2 = g(1-g)$, the variance of an individual pixel in I_g .

B. Spatial Statistics

In view of Fig. 5, we can begin to characterize blue-noise halftones in terms of the pair correlation $\mathcal{R}(r)$ by noting that:

- 1) few or no neighboring pixels lie within a radius of $r < \lambda_b$;
- 2) for $r > \lambda_b$, the expected number of minority pixels per unit area approaches \mathcal{I} with increasing r ;
- 3) the average number of minority pixels within the radius r increases sharply near $r = \lambda_b$.

The resulting pair correlation for blue noise is therefore of the form in Fig. 6 where $\mathcal{R}(r)$ shows: 1) a strong inhibition of minority pixels⁵ 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 principle wavelength, indicated by a series of peaks at integer multiples of λ_b . Similar to $P(f_\rho)$, the principle wavelength is indicated in Fig. 6 with a diamond located along the horizontal axis.

⁵Such processes are commonly referred to as hard-core, where no two points of ϕ are within some distance r of each other.

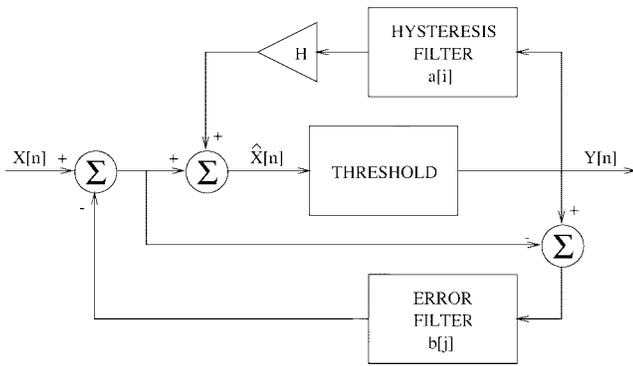


Fig. 12. Error diffusion with hysteresis algorithm.

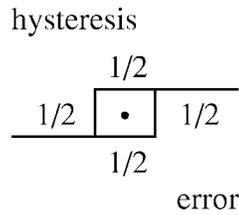


Fig. 13. An arrangement of two hystereses and two error-diffusion coefficients.

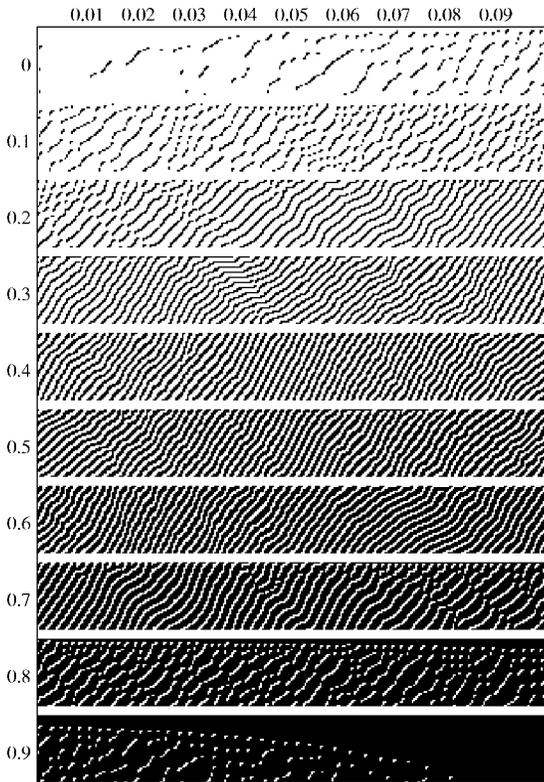
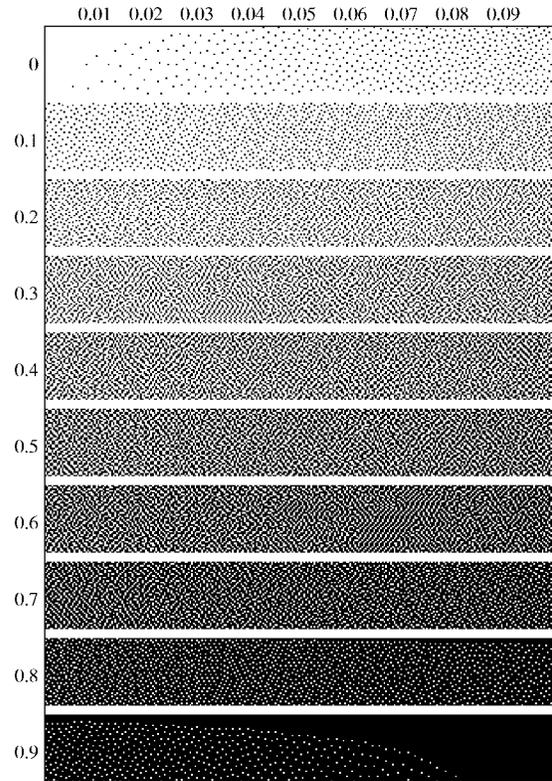


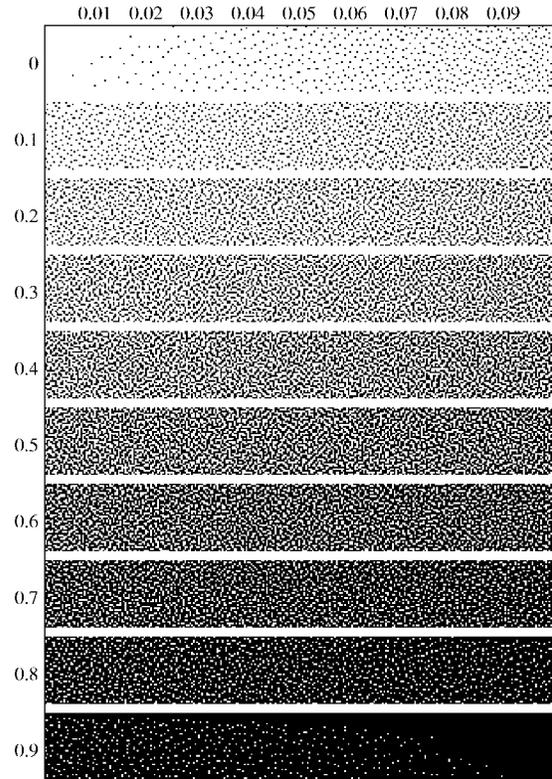
Fig. 14. Green-noise halftones of a grayscale ramp using two hystereses and two error coefficients with $H = 1$ and a left-to-right raster scan.

C. Error Diffusion

Although the original error-diffusion algorithm of Floyd and Steinberg [16] does create blue-noise patterns, Ulichney [12] shows that with some variations, error diffusion is an improved generator of blue noise. These variations in-



(a)



(b)

Fig. 15. Green-noise halftones of a grayscale ramp using two hystereses and two error coefficients with: (a) $H = 1/2$ and (b) $H = 1/2$ with 50% random error filter weights.

clude using a serpentine scan, randomized weight positions, and perturbed threshold. Another variation of particular importance involves perturbing the filter weights instead

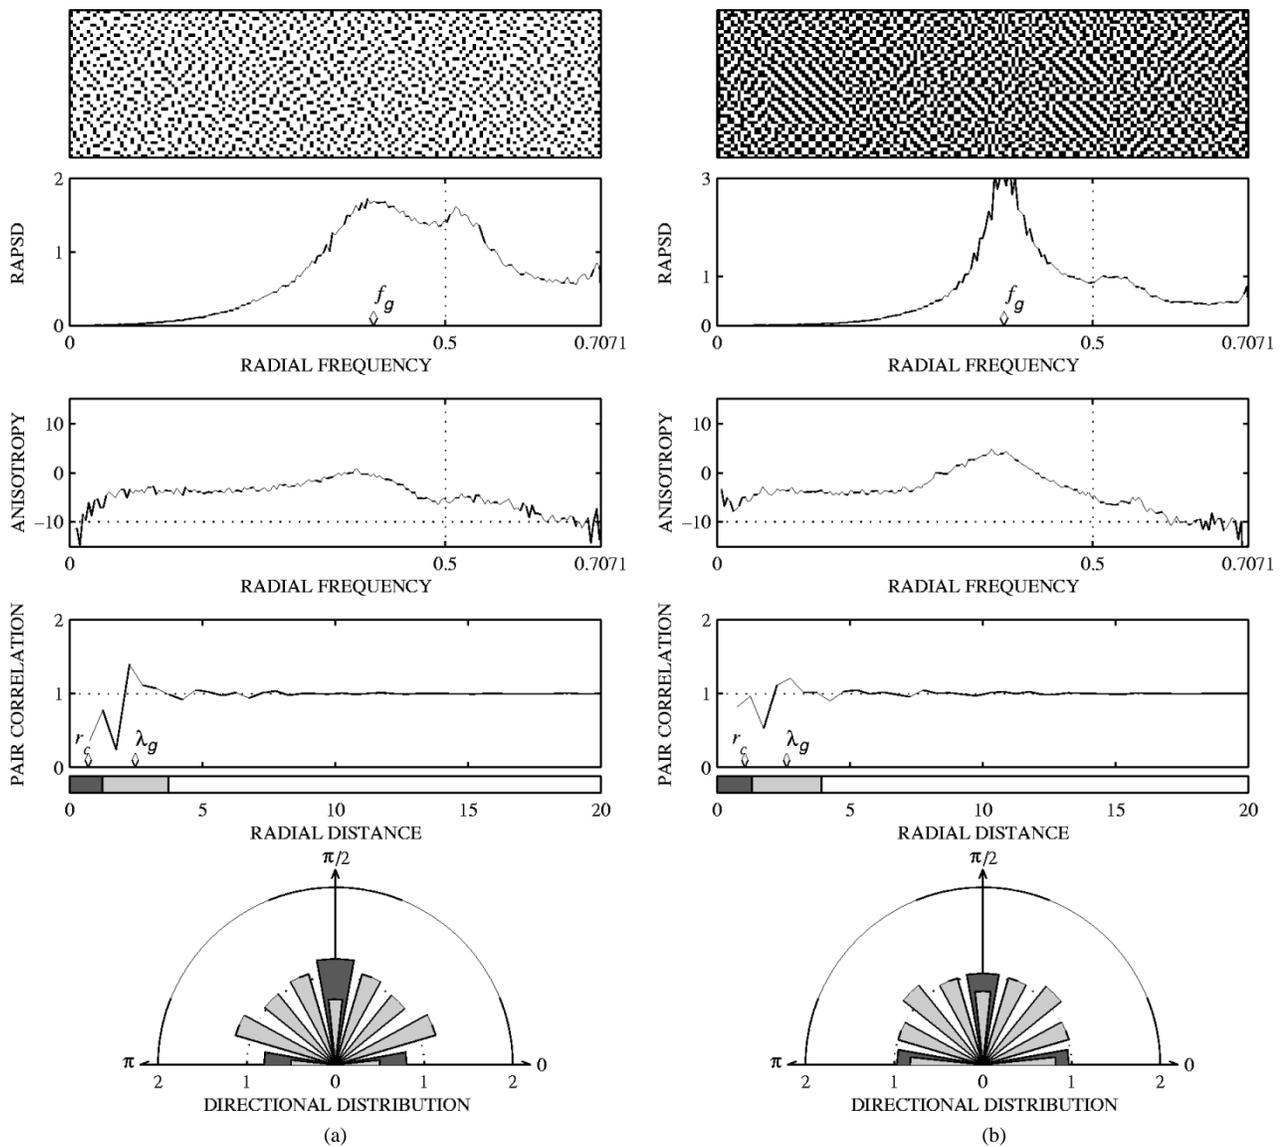


Fig. 16. Green-noise halftone statistics for: (a) $g = 1/4$ and (b) $g = 1/2$ using two hystereses and two error coefficients with $H = 1/2$.

of the threshold. This perturbation of filter weights is accomplished by first pairing weights of comparable value. Then for each pair of weights, a scaled random value is added to one and subtracted from the other. In order to prevent negative values, the maximum noise amplitude (100%) is the value of the smaller weight in each pair.

Using the Floyd–Steinberg filter weights (Fig. 7), Ulichney shows that perturbing each of the two pairs (7/16, 5/16) and (3/16, 1/16) creates a good blue-noise process. In particular, adding 50% noise to each pair appears to optimize the tradeoff between graininess and stable texture [12]. Fig. 8 shows the resulting spatial and spectral characteristics for $g = 1/16$ and $1/8$ using this scheme.

IV. GREEN-NOISE HALFTONING

Just as blue noise is the high-frequency component of white noise, green noise, labeled as Φ_G and demonstrated in Fig. 9, is the mid-frequency component which, like blue,

benefits from aperiodic, uncorrelated structure without low-frequency graininess, but unlike blue, green-noise patterns exhibit clustering.⁶ The result is a frequency content which lacks the high-frequency component characteristic of blue noise. Hence the term “green.” Furthermore, green noise forms aperiodic patterns that are not necessarily radially symmetric. Since the contrast sensitivity function of the human visual system is not radially symmetric, we allow green noise to have asymmetric characteristics. The objective is to combine the maximum dispersion attributes of blue noise with that of clustering of AM halftone patterns.

Point process statisticians have long described cluster processes such as those seen in green noise by examining the cluster process in terms of two separate processes: 1)

⁶The use of the word cluster refers to a collection of consecutive four-neighborhood pixels all of the same value. The definition is the same as that used in [17] for a clump of marked vertices in a square lattice.

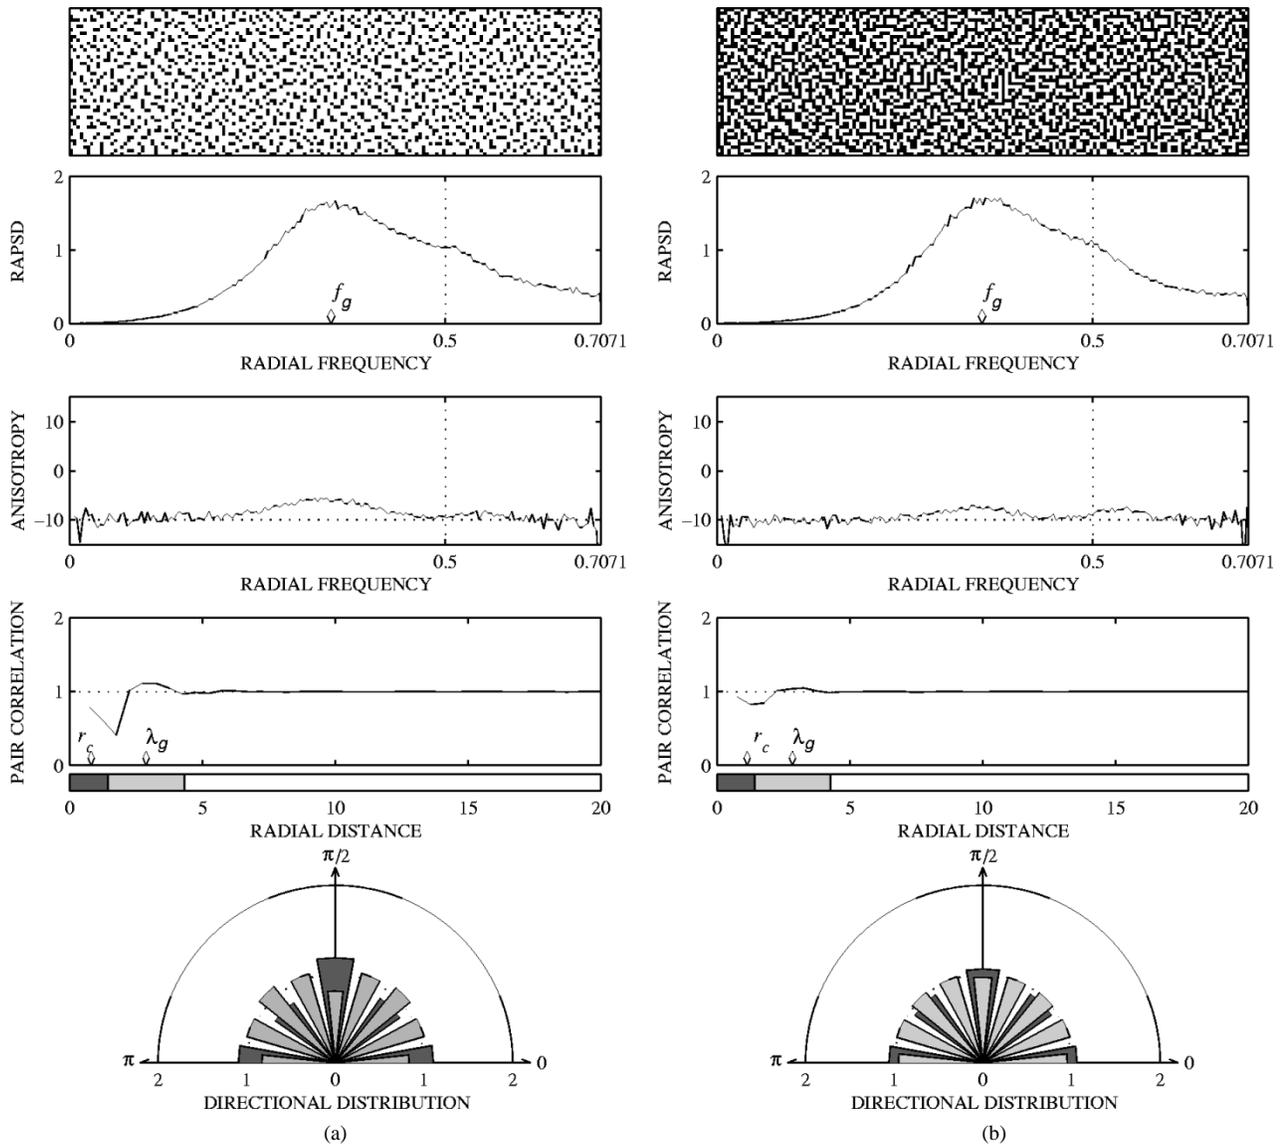


Fig. 17. Green-noise halftone statistics for: (a) $g = 1/4$ and (b) $g = 1/2$ using two hystereses and two error coefficients with $H = 1/2$ and 50% random error filter weights.

the parent process which describes the location of clusters⁷ and 2) the daughter process, which describes the shape of clusters. In AM halftoning, clusters are placed along a regular lattice, and therefore variations in AM patterns occur in the cluster shape. In FM halftoning, cluster shape is deterministic, i.e., 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 which looks at both the parent and daughter processes.

Looking first at the parent process Φ_p , ϕ_p represents a single sample of the parent process such that $\phi_p = \{x_i : i = 1, \dots, N_c\}$ 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 $\phi_d = \{y_j : j = 1, \dots, M\}$ 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

⁷The location being defined as the centroid of all points within the cluster.

$B = \{y_i : i = 1, 2, \dots\}$ by x , relative to the origin, as

$$T_x(B) = \{y_i - x : i = 1, 2, \dots\} \quad (11)$$

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

$$\phi_G = \sum_{x_i \in \phi_p} T_{x_i}(\phi_{d_i}) = \sum_{x_i \in \phi_p} \{y_{ji} - x_i : j = 1, \dots, M_i\} \quad (12)$$

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

In order to derive a relationship between the total number of clusters, the size of clusters, and the gray level of a binary dither pattern, I_g is defined as the binary dither pattern resulting from halftoning a continuous-tone discrete-space monochrome image of constant gray level g , and $I_g[n]$ is defined as the binary pixel of I_g with pixel index n . From the definition of $\phi(B)$ as the total number of points of ϕ

in B , $\phi_G(I_g)$ is the scalar quantity representing the total number of minority pixels in I_g , and $\phi_p(I_g)$ is the total number of clusters in I_g with $\phi_p(I_g) = N_c$. The intensity \mathcal{I} being the expected number of minority pixels per unit area now can be written as

$$\mathcal{I} = \frac{\phi_G(I_g)}{N(I_g)} = \begin{cases} g, & \text{for } 0 < g \leq 1/2 \\ 1 - g, & \text{for } 1/2 < g \leq 1 \end{cases} \quad (13)$$

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 (13), \bar{M} , the average number of minority pixels per cluster in I_g is

$$\bar{M} = \frac{\phi_G(I_g)}{\phi_p(I_g)} = \frac{\mathcal{I} \cdot N(I_g)}{\phi_p(I_g)} \quad (14)$$

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

A. Spectral Statistics

Although obvious, (14) shows the very important relationship between the total number of clusters, the average size of clusters, and the intensity for I_g . AM halftoning is the limiting case where $\phi_p(I_g)$ is held constant for varying \mathcal{I} , while FM halftoning is the limiting case where \bar{M} is held constant for varying \mathcal{I} . In addition, (14) says that the total number of clusters per unit area is proportional to \mathcal{I}/\bar{M} . For isolated minority pixels (blue noise), the square of the average separation between minority pixels (λ_b) is inversely proportional to \mathcal{I} , the average number of minority pixels per unit area [12]. 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 which are distributed as homogeneously as possible, leading to an average separation (center-to-center) between clusters (λ_g) whose squares are inversely proportional to the average number of minority pixel clusters per unit area \mathcal{I}/\bar{M} . Using the fact that $\lim_{M \rightarrow 1} \lambda_g = \lambda_b$, the proportionality constant can be determined such that λ_g is defined as

$$\lambda_g = \begin{cases} D/\sqrt{(g)/\bar{M}}, & \text{for } 0 < g \leq 1/2 \\ D/\sqrt{(1-g)/\bar{M}}, & \text{for } 1/2 < g \leq 1 \end{cases} \quad (15)$$

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

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

$$f_g = \begin{cases} \sqrt{(g)/\bar{M}}/D, & \text{for } 0 < g \leq 1/2 \\ \sqrt{(1-g)/\bar{M}}/D, & \text{for } 1/2 < g \leq 1. \end{cases} \quad (16)$$

From (16) we make several intuitive observations: 1) as the average size of clusters increases, f_g approaches DC and 2) as the size of clusters decreases, f_g approaches f_b . Fig. 10 illustrates the desired characteristics of $P(f_\rho)$ for ϕ_G showing three distinct features: 1) little or no low-

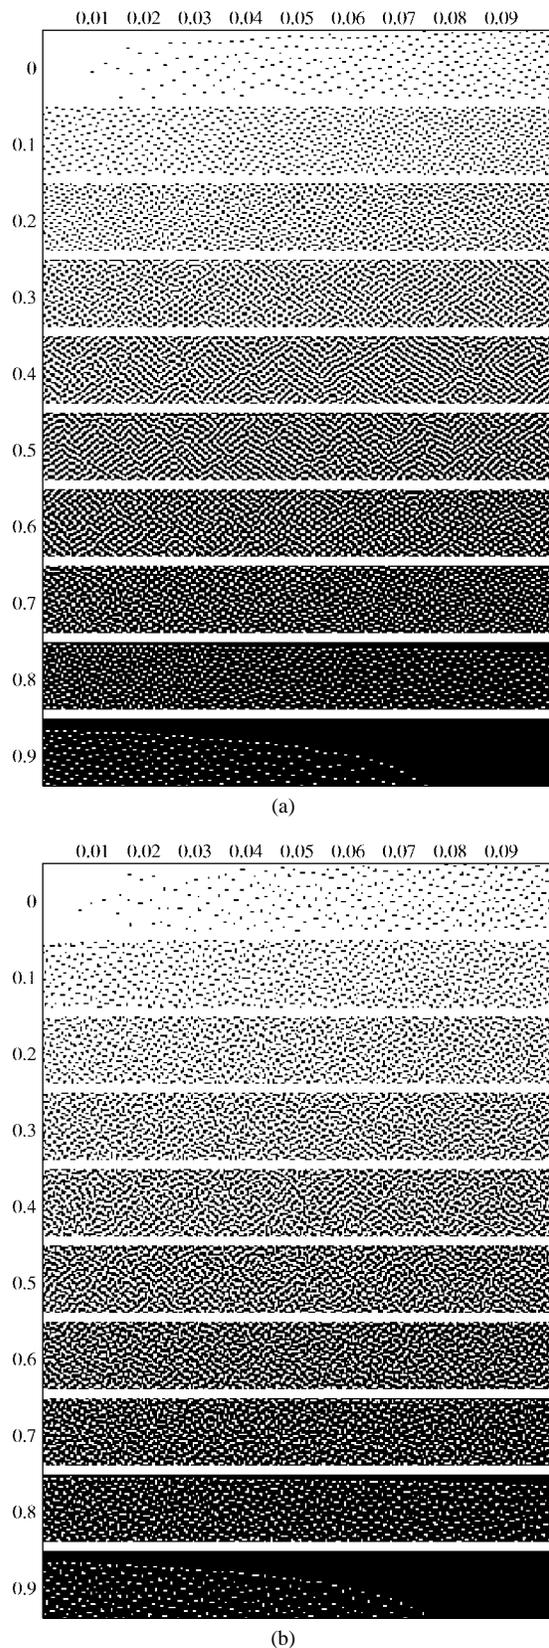


Fig. 18. Green-noise halftones of a grayscale ramp using two hystereses and two error coefficients with: (a) $H = 1$ and (b) $H = 1$ with 75% random hysteresis filter weights.

frequency spectral components; 2) high-frequency spectral components which diminish with increased clustering; and 3) a spectral peak at $f_\rho = f_g$.

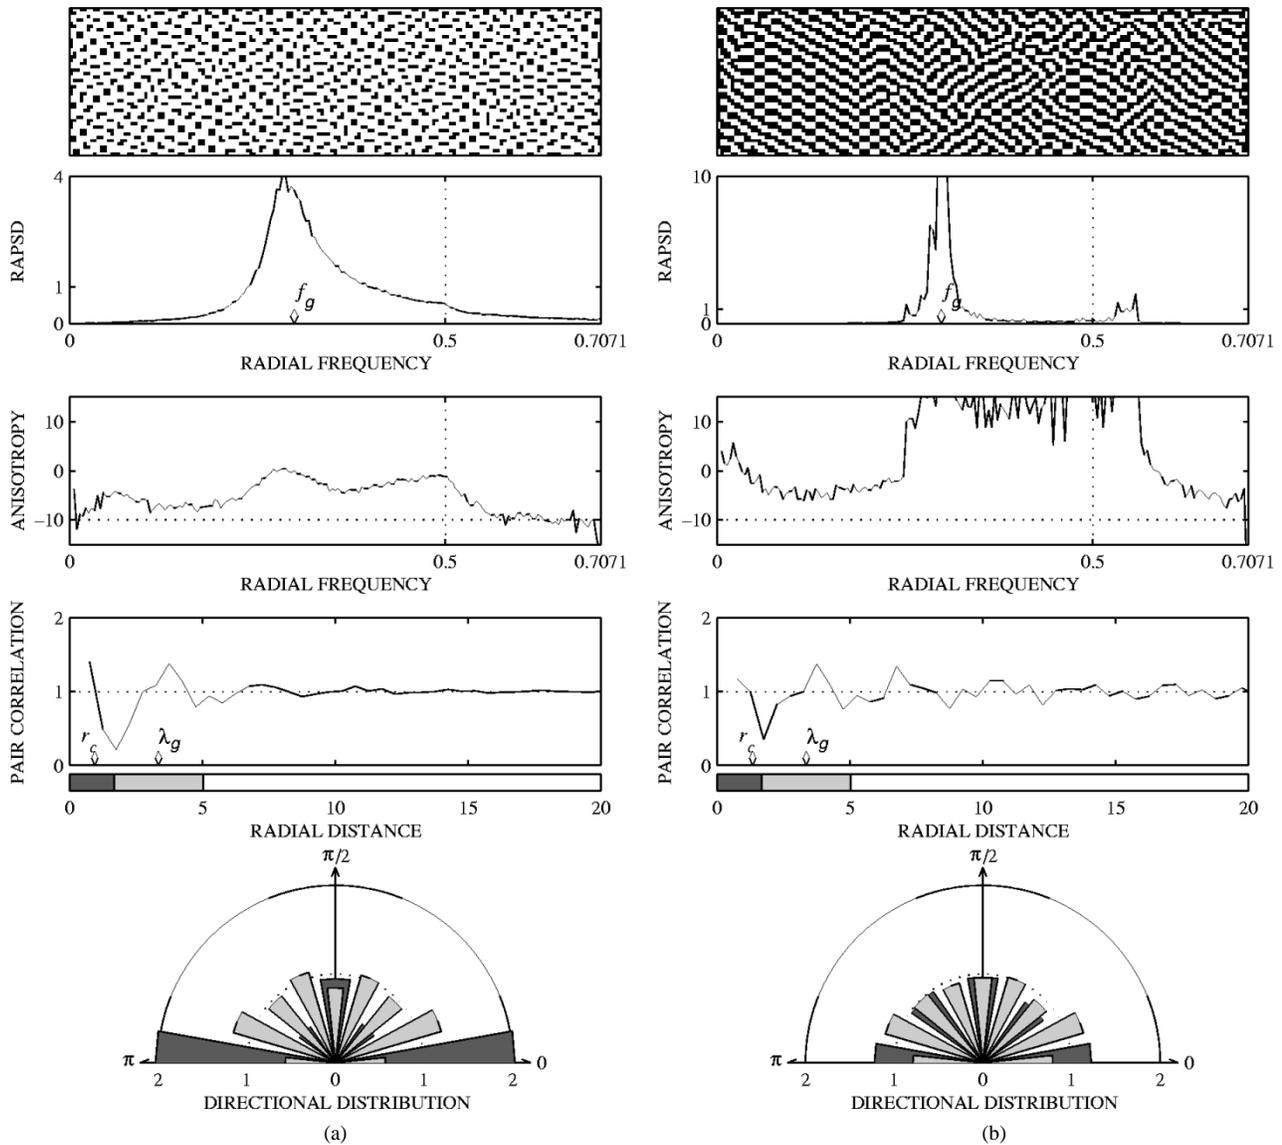


Fig. 19. Green-noise halftone statistics for: (a) $g = 1/4$ and (b) $g = 1/2$ using two hystereses and two error coefficients with $H = 1$.

The sharpness of the spectral peak in $P(f_p)$ at the green-noise principle frequency is affected by several factors. Consider first blue noise, where the separation between minority pixels should have some variation. The wavelengths of this variation, in blue noise, should not be significantly longer than λ_b as this adds low-frequency spectral components to the corresponding dither pattern I_g [12], 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 $f_p = f_g$ which 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-frequency components. 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.

B. Spatial Statistics

If we assume a stationary and isotropic green-noise pattern, the pair correlation will have the form of Fig. 11 given the following:

- 1) daughter pixels, on average, will fall within a circle of radius r_c centered around a parent point such that $\pi r_c^2 = \bar{M}$ (the area of the circle with radius r_c is equal to the average number of pixels forming a cluster);
- 2) neighboring clusters are located at an average distance of λ_g apart;
- 3) as r increases, the influence that clusters have on neighboring clusters decreases.

The result is a pair correlation which has: 1) a nonzero component for $0 \leq r < r_c$ due to clustering; 2) a decreasing

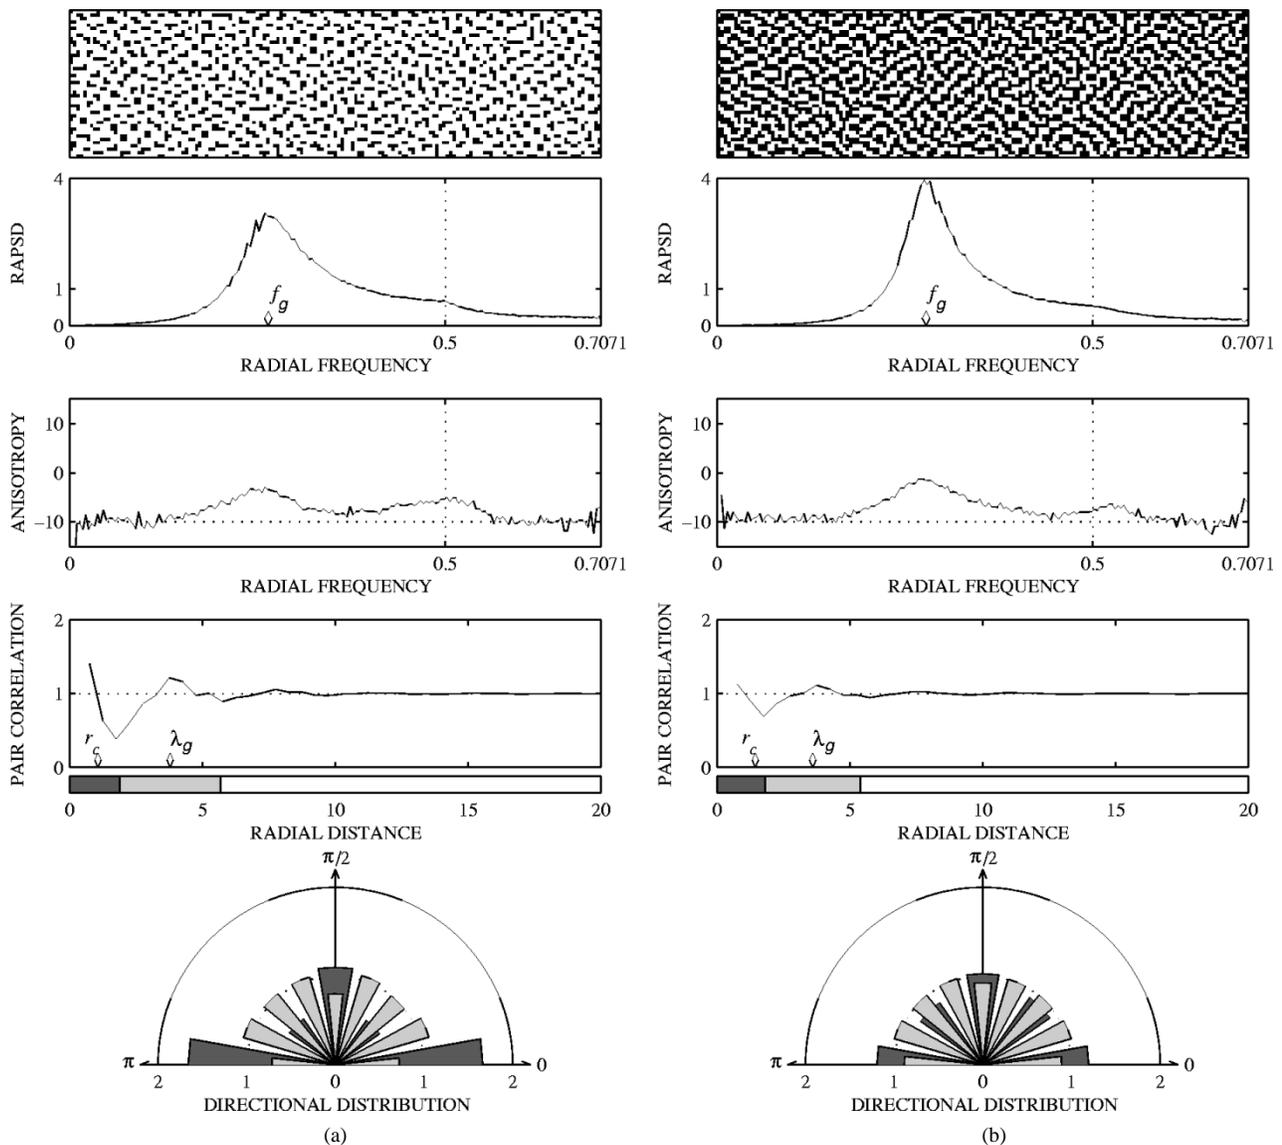


Fig. 20. Green-noise halftone statistics for: (a) $g = 1/4$ and (b) $g = 1/2$ using two hystereses and two error coefficients with $H = 1$ and 75% random hysteresis filter weights.

influence as r increases; and 3) 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 Fig. 11.

In the case of stationary and anisotropic green-noise patterns, the pair correlation will also be of the form of Fig. 11, but because clusters are not radially symmetric, blurring occurs in $\mathcal{R}(r)$ near the cluster radius r_c . In a similar fashion, because the separation between clusters will also vary with direction, blurring will occur at each peak in $\mathcal{R}(r)$ located at integer multiples of λ_g .

C. Error Diffusion with Hysteresis

Although error diffusion is a good generator of blue noise, the nature of green noise to cluster pixels makes error diffusion inappropriate. As an alternative, Levien [11] has proposed error diffusion with output-dependent feedback (Fig. 12) where a weighted sum of the previous output

pixels is used to vary the threshold—making minority pixels more likely to occur in clusters. Furthermore, the amount of clustering is controlled through the scalar constant H , the hysteresis constant, with large values of H leading to large clusters and small values of H leading to small clusters.

As mentioned previously, other techniques for creating binary dither patterns with adjustable coarseness include Velho's and Gomes' digital halftoning along SFC's [9] and Scheermesser's and Bryngdahl's digital halftoning with texture control [10]. SFC is a technique where a two-dimensional (2-D) image is halftoned using a 1-D clustered-dot dithering approach which traverses the image along a space filling curve such as the Peano, Hilbert, or Sierpinsky curve. By manipulating the maximum number of pixels that can form a cluster, the SFC technique can control the amount of coarseness in resulting images, and unlike AM halftoning where the maximum number of pixels that can

form a cluster limits the number of gray levels that the pattern can represent, SFC diffuses quantization error from one cluster to the next. The result is a technique which combines the benefits of aperiodic structure with those of clustered dots.

Scheermesser's and Bryngdahl's technique attempts to minimize the cost associated with a particular arrangement of dots by iteratively turning pixels "on" and "off." The cost associated with a particular arrangement of dots is determined by two factors. The first is an image metric which measures the difference between the perceived images of the binary halftone pattern and the continuous tone original. The second cost is a numerical texture metric which measures the relative orientation of minority pixels. Scheermesser's and Bryngdahl's technique is able to produce adjustable coarseness by adjusting the weight of the cost of the texture metric versus the perceived image metric.

There are, of course, many other algorithms which produce stochastic arrangements of minority pixel clusters (green noise). The reason that this paper chooses to look at Levien's technique is because it is ideally suited, based so much on Floyd's and Steinberg's original error-diffusion algorithm [16], to present an evolution of the blue-noise model originally described by Ulichney [12]—having, like the green-noise model itself, blue noise as a limiting case ($H = 0$). Furthermore, as will be shown, error diffusion with output-dependent feedback creates patterns which may be described as "good" green-noise patterns, i.e., exhibiting sharp peaks in both the spectral $P(f_\rho)$, and spatial $\mathcal{R}(r)$ domains.

Mathematically, Levien's algorithm is defined as follows:

$$\hat{X}[n] = X[n] + \sum_j b[j] \cdot e[n - j] \quad (17)$$

$$Y[n] = \hat{X}[n] + H \cdot \sum_i a[i] \cdot Y[n - i] \geq T \quad (18)$$

$$e[n] = Y[n] - \hat{X}[n] \quad (19)$$

where $a[i]$ and $b[j]$ are the hysteresis and error-diffusion coefficients, respectively, such that $\sum_i a[i] = \sum_j b[j] = 1$. In the following sections, we look at some variations to error diffusion with output-dependent feedback and compare their results in terms of the spatial and spectral domain metrics discussed previously. These variations include the following.

- 1) *Hysteresis constant.* Of paramount importance, the hysteresis constant controls the amount of clustering, and therefore it influences greatly the green-noise characteristics of resulting dither patterns. It is this parameter H that allows for tunable coarseness.
- 2) *Choice of error/hysteresis filter.* Within the context of Levien's definition of error diffusion with output-dependent feedback, many choices exist for the error filter, and coupled with the hysteresis filter, even more possibilities can occur. For computational efficiency, as small a filter as possible is always preferred.

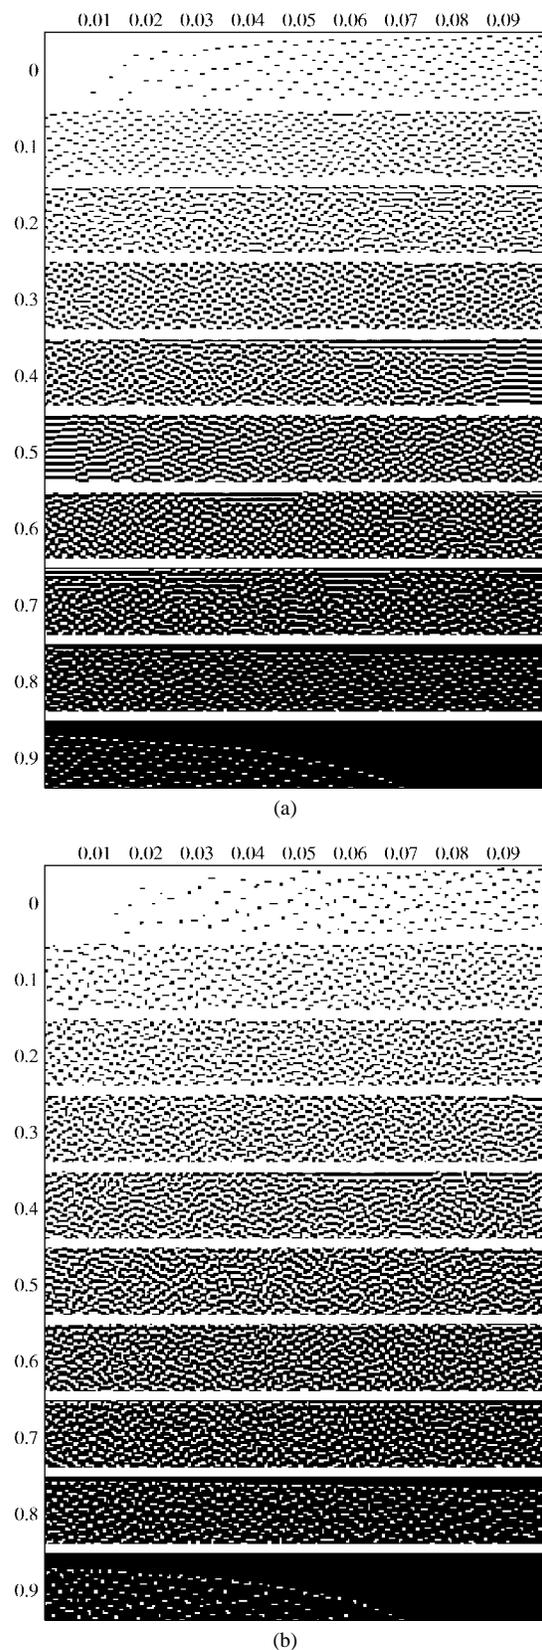


Fig. 21. Green-noise halftones of a grayscale ramp using two hystereses and two error coefficients with: (a) $H = 3/2$ and (b) $H = 3/2$ with 50% random hysteresis filter weights.

- 3) *Stochastic error/hysteresis filter perturbation.* Introduced by Ulichney [12] to improve the blue-noise characteristics of error diffusion, random noise can be

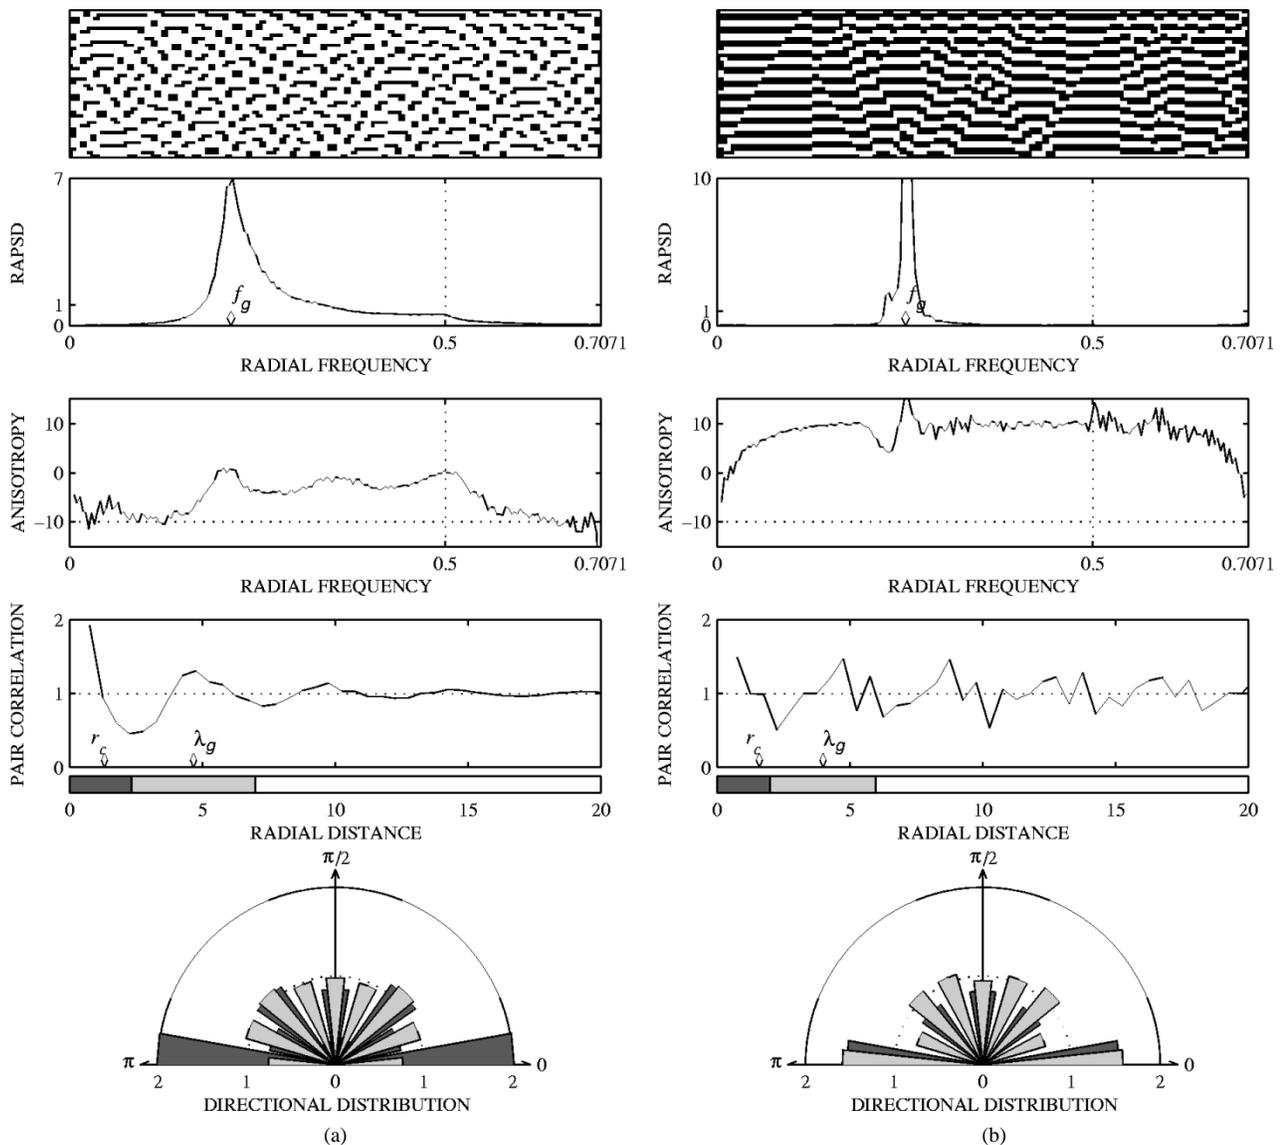


Fig. 22. Green-noise halftone statistics for: (a) $g = 1/4$ and (b) $g = 1/2$ using two hystereses and two error coefficients with $H = 3/2$.

added to the weights of both the error and hysteresis filters. This same process is described in Section III-C for error diffusion with four weights, and unlike the previous two variations, the application of filter perturbation to either the error diffusion or the feedback coefficients is new to Levien's technique.

1) *Two Hystereses and Two Error Weights:* The scheme originally proposed by Levien [11], which uses only two hystereses and two error weights (Fig. 13), has low computational complexity but coupled with a left-to-right raster scan (Fig. 14) yields poor results due to strong diagonal texture patterns, thereby making mandatory alternate scanning paths such as the serpentine (left-to-right then right-to-left) raster scan, which is used for this and all other arrangements of filter weights. Note that in these and the following figures, the binary dither patterns are shown under ideal printing conditions unless otherwise noted.

In addition to the serpentine raster scan, a small hysteresis constant ($H = 1/2$) produces patterns for small g which are very similar to traditional blue noise [Fig. 15(a)], but as g increases and approaches $1/2$, the amount of clustering also increases, thereby shifting the spectral content from blue to green (Fig. 16). This behavior is indicated by $P(f_\rho)$ for $g = 1/4$ and $1/2$ where, for blue noise, $f_b = 0.50$ and 0.71 , and for this arrangement of weights, $f_g = 0.40$ and 0.38 (strictly less than f_b). Furthermore, as seen in Fig. 15(b), adding a perturbation of 50% randomness to the error filter weights breaks up many of the directional artifacts, which creates patterns with reduced diagonal correlation and increased variation in cluster size. Resulting patterns (Fig. 17), therefore begin to appear noisy as the spectral peaks in $P(f_\rho)$ at $f_\rho = f_g$ are blurred. Also note that as a result of increased variation in cluster size, the peaks and valleys of $\mathcal{R}(r)$ are also blurred. This further indicates a dither pattern with spatial

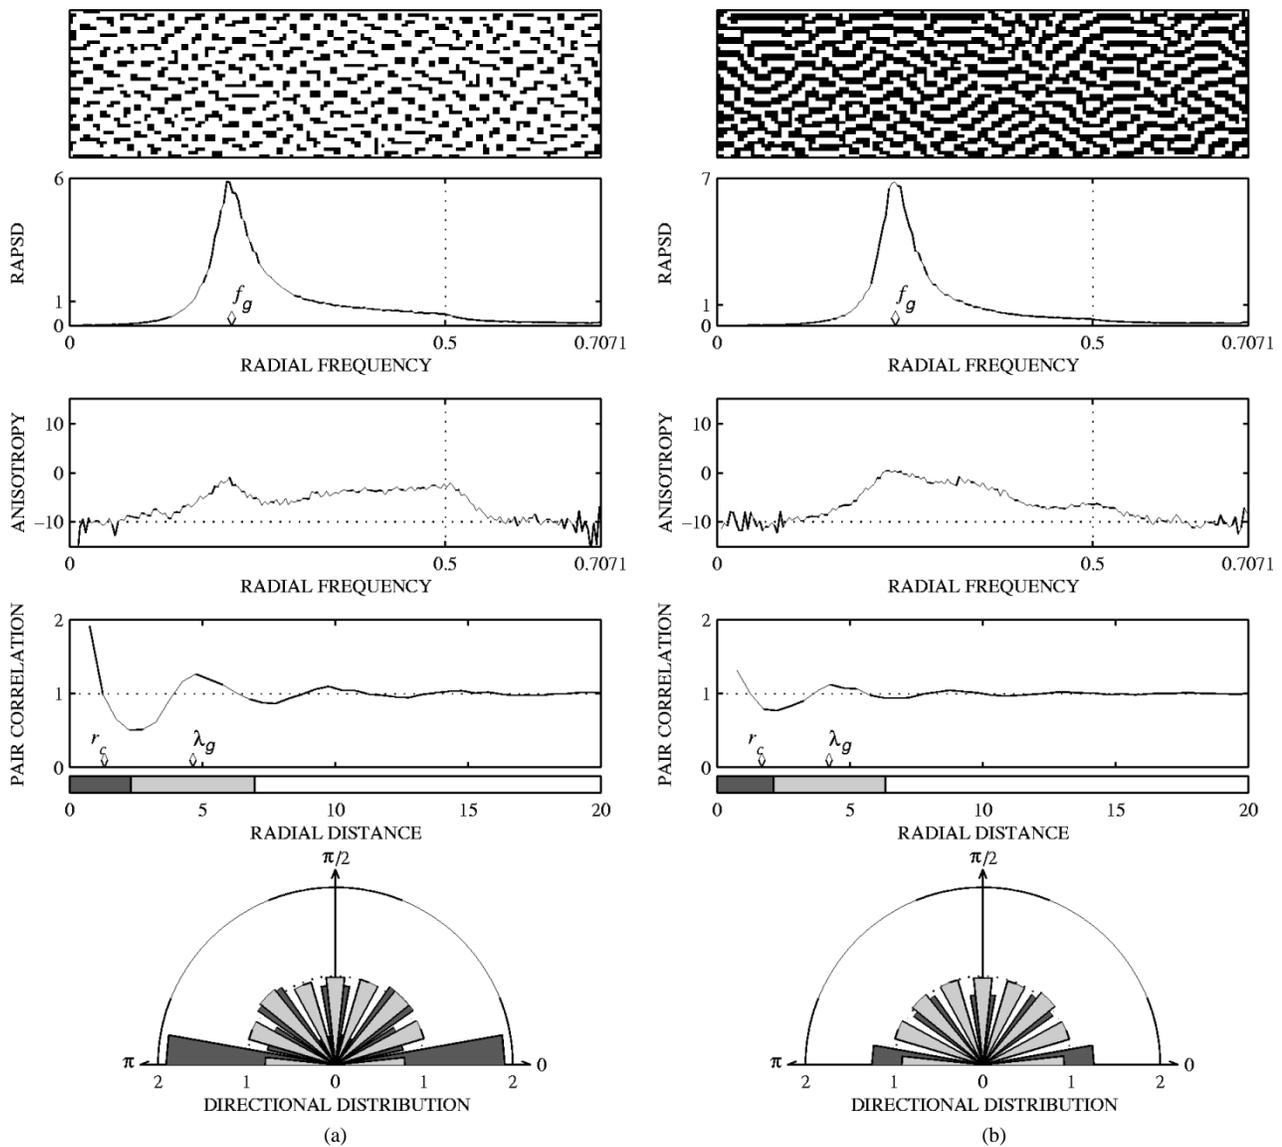


Fig. 23. Green-noise halftone statistics for: (a) $g = 1/4$ and (b) $g = 1/2$ using two hystereses and two error coefficients with $H = 3/2$ and 50% random hysteresis filter weights.

and spectral characteristics more similar to white-noise halftones.

In contrast to a low hysteresis constant, Fig. 18 shows the resulting patterns generated using a moderate hysteresis constant ($H = 1$). Specifically, in Fig. 18(a), patterns exhibit clustering at all gray levels with a well-defined boundary formed between clusters. As is exhibited in Fig. 19 for $g = 1/4$ where $f_g = 0.30$ ($\bar{M} = 2.80$), this clustering behavior is indicated by the sharp drop in $\mathcal{R}(r)$ for $r_c < r < \lambda_g$ and by the ripple formed in $\mathcal{R}(r)$ with peaks occurring at integer multiples of λ_g .

Figs. 18(b) and 20 illustrate the whitening effect of adding a perturbation of 75% randomness to the hysteresis weights. For each instance of g , an increased variation in cluster size is clearly visible, most noticeably for $g = 1/16$ where clusters of 1, 2, and 3 pixels can be seen throughout the pattern. Without perturbation at $g = 1/16$, clusters are commonly composed of two pixels with only a few composed of one. Rarely do clusters of three pixels

occur without perturbation at this low gray level, but also without this perturbation, patterns exhibit strong direction artifacts—creating diagonal textures. The perturbation of hysteresis weights breaks up many of these textures, especially near $g = 1/2$.

As a final demonstration of using two hystereses and two error weights, Fig. 21 shows the resulting dither patterns using a high hysteresis constant ($H = 3/2$). Illustrated in Fig. 22 for $g = 1/4$ and $1/2$, the variation in cluster size has increased relative to $H = 1/2$ and $H = 1$ as clusters are beginning to form long horizontal bars—creating textures which are not diagonal but horizontal. Adding a perturbation of 50% randomness to the hysteresis weights [Fig. 21(b)] in this scheme breaks up many of these horizontal artifacts—reducing $A(f_\rho)$ for all gray levels but also whitening the spectral content slightly (Fig. 23).

2) *Four Hysteresis and 12 Error Weights:* As an investigation of the effects of different filter weights, Fig. 24 shows the combination of the Floyd–Steinberg weights

hysteresis

	3/16	5/16	1/16	
	7/16	•	8/42	4/42
	2/42	4/42	8/42	4/42
	1/42	2/42	4/42	2/42
				1/42
				error

Fig. 24. An arrangement of four hysteresis and 12 error-diffusion coefficients.

[16] as the hysteresis filter and the Stucki weights [18] as the error filter. Using a small hysteresis constant ($H = 1/2$), Fig. 25(a) shows the resulting gray scale ramp which exhibits a larger degree of clustering relative to using two hystereses and two error weights with equivalent H . Due to very small variation in cluster size, this method leads to patterns with very sharp spectral peaks in $P(f_\rho)$, as seen in Fig. 26, but at $g = 1/2$, the clusters form a pattern with very distinct horizontal artifacts.

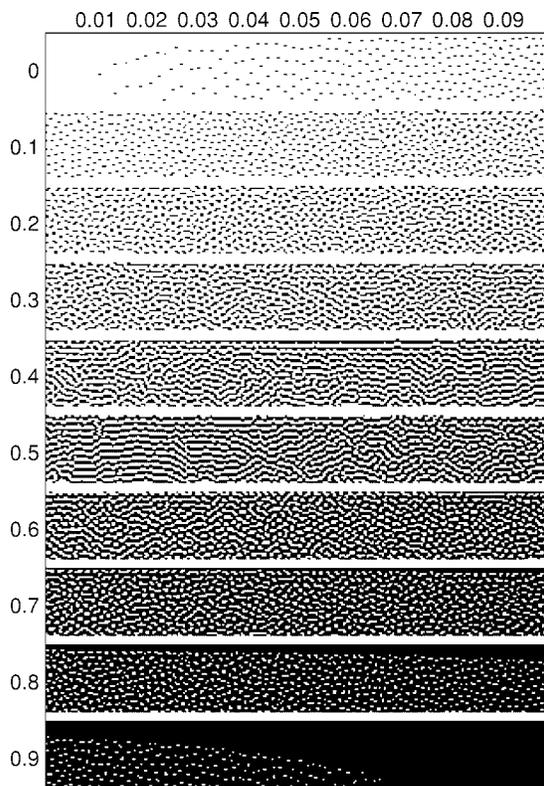
Figure 25(b) shows the same combination of weights but with an added perturbation of 30% randomness to the hysteresis weights and 30% randomness to the error weights. In this case, as seen in Fig. 27, many of the horizontal artifacts are broken up, and at the same time only a small amount of spectral whitening has occurred in patterns for $g = 1/4$ and $1/2$.

For a comparison of the principle frequencies and average cluster sizes for patterns generated using the described variations, see the table of Fig. 28.

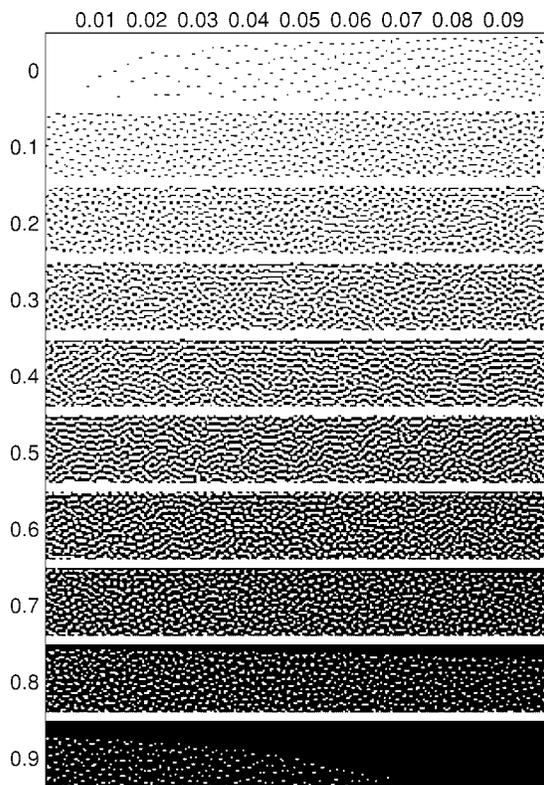
V. DOT GAIN

As stated in the introduction, dot gain is the increase in size of the printed dot relative to the intended size, and whether a function of the mechanical printing process or the optical properties of paper, dot gain is present in all printing processes. Therefore, it cannot be eliminated. Assuming that the process is repeatable, dot gain can be anticipated and controlled [7], but in printing processes which are not repeatable [7], compensating for dot gain is a far more difficult task. In these instances, it may be more desirable to use a scheme which resists dot gain, thereby making the printed output more robust to variations in the printing process. An example where dot gain is anticipated and therefore compensated for is Pappas' and Neuhoff's [19] modified error-diffusion algorithm employing a circular dot-overlap model which models printed dots as round dots which overlap neighboring pixels, thereby causing printed halftones to appear darker than the original fraction of ones.

Although for many printers the circular dot-overlap model is invalid [20], it illustrates an important aspect of dot gain with regards to digital halftoning. That being the relationship between dot gain and the perimeter-to-area ratio of printed dots with higher ratios leading to higher average dot gain. In dispersed dot halftones (FM), this ratio is much higher than that for clustered dot halftones, which creates higher gain for dispersed dots [4]. It is for this



(a)



(b)

Fig. 25. Green-noise halftones of a grayscale ramp using four hysteresis and 12 error coefficients with: (a) $H = 1/2$ and (b) $H = 1/2$ with 30% random error filter weights and 30% random hysteresis filter weights.

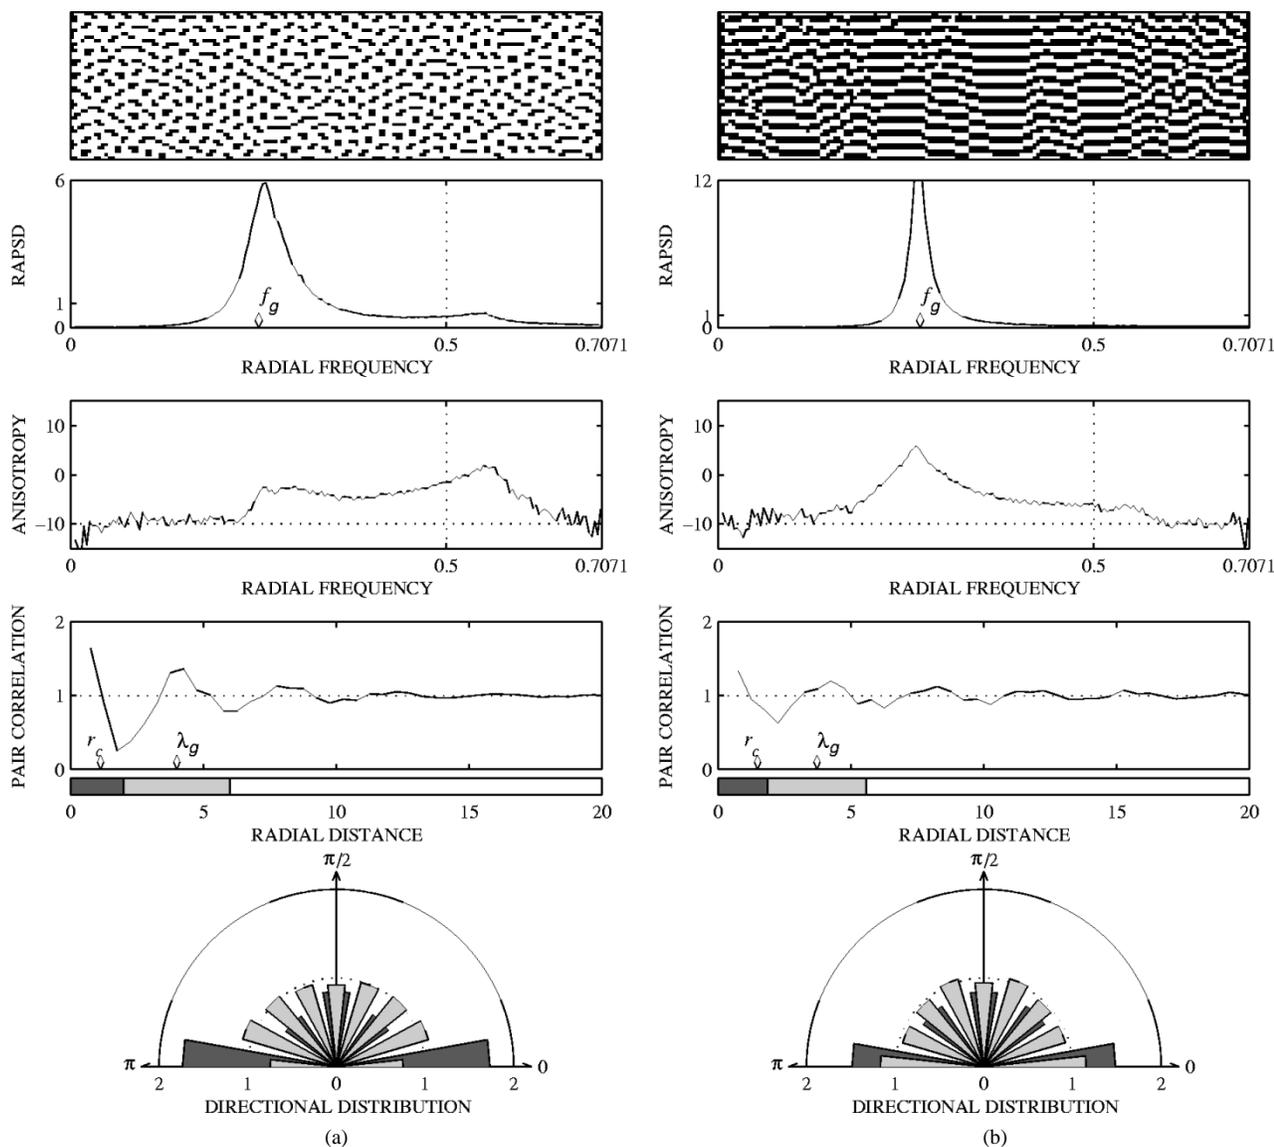


Fig. 26. Green-noise halftone statistics for: (a) $g = 1/4$ and (b) $g = 1/2$ using four hystereses and 12 error coefficients with $H = 1/2$.

reason that green-noise halftoning which clusters pixels is less susceptible to the effects of dot gain, and by varying the size of clusters, resulting patterns can be optimized for specific dot gain characteristics with large clusters reserved for printing processes with high variation from printed dot-to-printed dot and small clusters reserved for printing processes with low variation.

Green noise offers a wide range of halftone renditions; therefore, selecting the optimal hysteresis constant for error diffusion with output-dependent feedback is an important problem that must be addressed based on specific printer traits. As a demonstration of the advantages of using green noise instead of blue, Fig. 29 shows the resulting images using the circular dot-overlap model where the ratio of the diameter of a printed dot to the minimum distance between samples is 10:6. Note that the blue-noise image [Fig. 29(a)] appears much darker than the green-

noise image [Fig. 29(b)]. Also note that for blue noise, isolated white pixels are nearly erased as neighboring black dots almost completely overlap the white pixels.

VI. CONCLUSION

In addition to introducing green noise, this paper has introduced the use of the two statistics for the analysis of digital halftoning techniques: 1) the pair correlation and 2) the directional distribution function. Being spatial-domain statistics, both offer a more intuitive understanding of the underlying point process compared to the conventional spectral-domain statistics: the radially averaged power spectrum and the anisotropy. Also in this paper, we have presented models in the spatial and spectral domain of both blue-noise and green-noise dither patterns, and using such, we have analyzed several variations to error-diffusion with output-dependent feedback—remarking on the attributes of

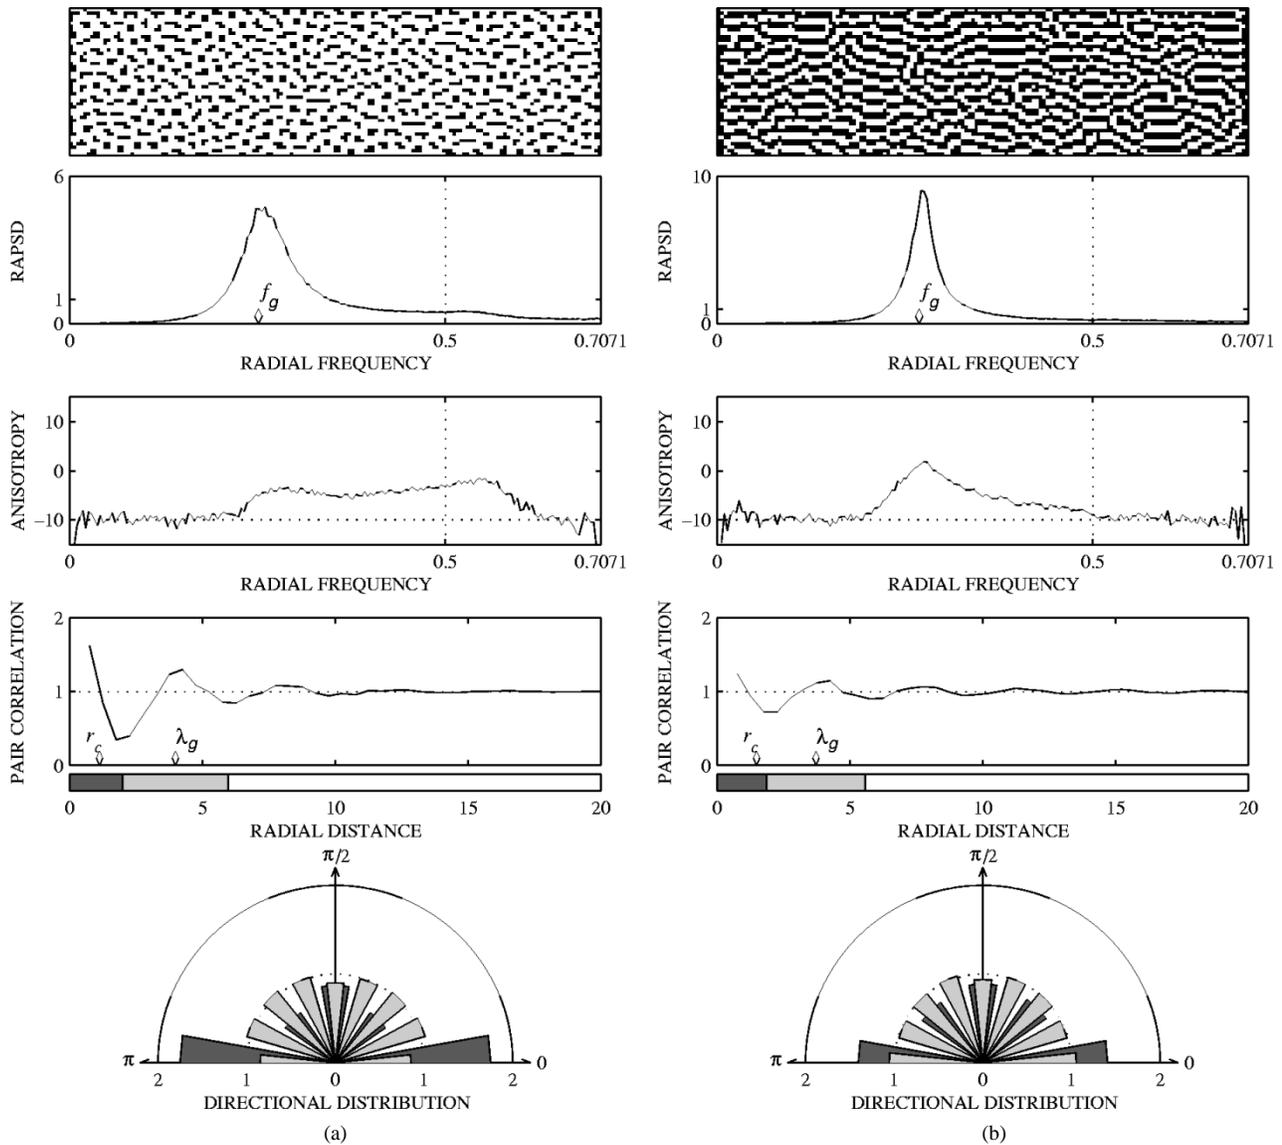


Fig. 27. Green-noise halftone statistics for: (a) $g = 1/4$ and (b) $g = 1/2$ using four hystereses and 12 error coefficients with $H = 1/2$ and 30% random error filter weights and 30% random hysteresis filter weights.

each variation. In effect, we have demonstrated a wide range of modifications that may be used to optimize a green-noise processes for specific printer traits.

Fig. 30 shows a comparison of five major classes of dither patterns: 1) white-noise; 2) blue-noise; 3) green-noise; 4) dispersed ordered dither; and 5) clustered-dot dither. Ulichney [12], who makes a similar comparison, writes that blue noise is pleasant because it does not add a structure of its own to an image as do dispersed-order dither and clustered-dot dither, both being arranged to form a regular pattern. Blue noise also does not look too “noisy” or uncorrelated as does white noise.

Green noise, being a stochastic patterning of dots, also does not add structure to an image as does dispersed-ordered dither or clustered-dot dither. It also does not appear noisy or uncorrelated as does white, but green

filter coefficient arrangement:	2 hysteresis and 2 error weights				4 hysteresis and 12 error weights				
	hysteresis: H=1/2		H=1		H=3/2		H=1/2		
error and feedback filter perturbation:	0%	50% error only	0% hyst only	75% hyst only	0%	50% hyst only	0%	30% error & hyst	
$g=1/16$	f_g	0.25	0.24	0.18	0.18	0.16	0.13	0.18	0.18
	\bar{M}	1.00	1.05	1.86	1.88	2.58	3.52	1.95	1.89
$g=1/8$	f_g	0.35	0.33	0.26	0.24	0.20	0.18	0.26	0.23
	\bar{M}	1.00	1.15	1.90	2.12	3.24	4.03	1.90	2.31
$g=1/4$	f_g	0.40	0.35	0.30	0.26	0.21	0.22	0.25	0.25
	\bar{M}	1.53	2.07	2.80	3.58	5.43	5.41	3.99	3.89
$g=1/2$	f_g	0.38	0.35	0.30	0.28	0.25	0.24	0.27	0.27
	\bar{M}	3.43	4.03	5.62	6.47	7.95	8.89	6.92	6.92

Fig. 28. Table showing the principle frequency f_g and the average number of pixels per cluster \bar{M} for several configurations of error diffusion with output-dependent feedback employing a serpentine raster scan.



(a)



(b)

Fig. 29. The resulting images using: (a) blue noise and (b) green noise using a circular dot-overlap model such that the ratio of the diameter of a printed dot to the minimum distance between samples is 10:6.

noise, being composed of pixel clusters, does increase the visibility of the halftone pattern. Given ideal printing conditions, blue noise is, therefore, the obvious choice for

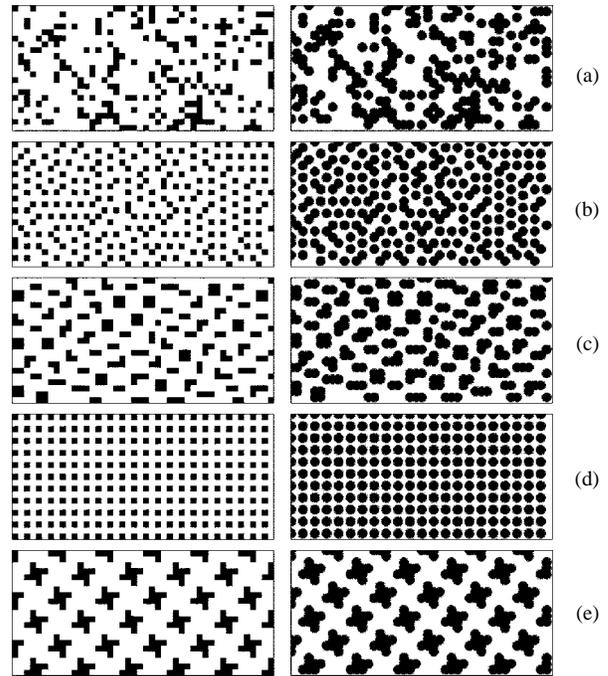


Fig. 30. Sample halftone patterns of: (a) white-noise; (b) blue-noise; (c) green-noise; (d) dispersed-dot ordered; and (e) clustered-dot ordered dithering under (left) ideal printing conditions and (right) with circular dot-overlap model.

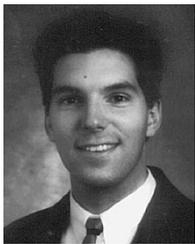
reproducing continuous tone images due to the increased spatial resolution which it achieves, but when printing conditions are not ideal, printing distortions such as dot-gain make clustered-dithering techniques such as green noise preferable to blue.

In closing, we add that the impact of blue noise on digital halftoning has been wide spread and has triggered many advancements. In particular, the blue-noise mask [21] has greatly decreased the computational complexity of dithering with blue noise. Equivalent developments can be made for green noise, as is noted in [22] for the green-noise mask.

REFERENCES

- [1] V. Ostromoukhov, P. Emmel, N. Rudaz, I. Amidror, and R. D. Hersch, "Dither algorithms for variable dot-size printers," in *Proc. Int. Conf. Image Processing*, 1996, pp. 553–556.
- [2] I. Amidror, R. D. Hersch, and V. Ostromoukhov, "Spectral analysis and minimization of moiré patterns in color separation," *J. Electron. Imaging*, vol. 3, pp. 295–317, July 1994.
- [3] P. Fink, *PostScript Screening: Adobe Accurate Screens*. Mountain View, CA: Adobe Press, 1992.
- [4] M. Rodriguez, "Graphic arts perspective on digital halftoning," *Proc. SPIE, Human Vision, Vision Processing and Digital Display V*, vol. 2179, pp. 144–149, Feb. 1994.
- [5] D. Blatner and S. Roth, *Real World Scanning and Halftones*. Berkeley, CA: Addison-Wesley, 1993.
- [6] T. P. Blumer, "Image display apparatus and method using halftone super cells," U.S. Patent 5 287 195, 1994.
- [7] M. A. Coudray, "Causes and corrections of dot gain on press," *Screen Printing: J. Technol. Manage.*, vol. 86, pp. 18–26, Aug. 1996.
- [8] R. A. Ulichney, *Digital Halftoning*. Cambridge, MA: MIT Press, 1987.
- [9] L. Velho and J. M. Gomes, "Digital halftoning with space filling curves," *Comput. Graphics*, vol. 25, pp. 81–90, July 1991.
- [10] T. Scheermesser and O. Bryngdahl, "Control of texture in image halftoning," *J. Opt. Soc. Amer. A*, vol. 13, pp. 1645–1652, Aug. 1996.

- [11] R. Levien, "Output dependant feedback in error diffusion halftoning," in *Proc. IS&T's 8th Int. Congr. Advances in Non-Impact Printing Technologies*, Oct. 1992, pp. 280-282.
- [12] R. A. Ulichney, "Dithering with blue noise," *Proc. IEEE*, vol. 76, pp. 56-79, Jan. 1988.
- [13] N. A. C. Cressie, *Statistics for Spatial Data*. New York: Wiley, 1983.
- [14] D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic Geometry and Its Applications*. New York: Wiley, 1987.
- [15] T. Mitsa and K. J. Parker, "Digital halftoning technique using a blue noise mask," *J. Opt. Soc. Amer.*, vol. 9, pp. 1920-1929, Aug. 1992.
- [16] R. W. Floyd and I. Steinberg, "An adaptive algorithm for spatial gray-scale," *Proc. SID*, vol. 17, no. 2, pp. 75-78, 1976.
- [17] P. Hall, *Introduction to the Theory of Coverage Processes*. New York: Wiley, 1988.
- [18] P. Stucki, "Mecca-a multiple-error correcting computation algorithm for bilevel image hardcopy reproduction," IBM Res. Lab., Zurich, Switzerland, Tech. Rep. RZ1060, 1981.
- [19] T. N. Pappas and D. L. Neuhoff, "Printer models and error diffusion," *IEEE Trans. Image Processing*, vol. 4, pp. 66-79, Jan. 1995.
- [20] T. N. Pappas, C. Dong, and D. L. Neuhoff, "Measurement of printer parameters for model-based halftoning," *J. Electron. Imaging*, vol. 2, pp. 193-204, July 1993.
- [21] K. J. Parker and T. Mitsa, "Method and apparatus for halftone rendering of a gray scale image using a blue noise mask," U.S. Patent 5 341 228, 1994.
- [22] G. R. Arce and D. L. Lau, "Method and apparatus for producing halftone images using a green noise mask," U.S. Patent pending.



Daniel L. Lau received the B.S.E.E. degree with highest distinction from Purdue University, West Lafayette, IN, in 1995. He is currently working towards the Ph.D. degree in electrical engineering at the University of Delaware, Newark.

He has worked in image processing at the Lawrence Livermore National Laboratory. His research interests include image processing, digital halftoning, multimedia, nonlinear filters, and art conservation. He has written several articles

on nonlinear filters for signal and image processing, and he has consulted with industry on halftoning and digital printing. His most recent work has focused on accurate color reproduction.



Gonzalo R. Arce (S'82-M'82-SM'93) received the B.S.E.E. degree with highest honors from the University of Arkansas, Fayetteville, in 1979 and the M.S. and Ph.D. degrees in electrical engineering from Purdue University, West Lafayette, IN, in 1980 and 1982, respectively.

Since 1982, he has been with the Department of Electrical and Computer Engineering at the University of Delaware, Newark, where he is currently Professor and Associate Chair. He frequently serves as consultant to industry and government in the areas of digital printing, image processing, and digital communications. His research interests include robust and nonlinear signal processing and its applications, wireless communications, image and video signal processing, secure multimedia communications, and halftoning. He holds two U.S. Patents.

Dr. Arce has served as Associate Editor of the IEEE TRANSACTIONS ON SIGNAL PROCESSING, Guest Editor of IEEE TRANSACTIONS ON IMAGE PROCESSING, and Guest Editor for *Optics Express*. He received an NSF Research Initiation Award in 1983 and the Best Paper award from the U.S. Army's ATIRP Federated Laboratory Consortium in 1997. He is a Fellow of the Center for Advanced Studies at the University of Delaware. He is a member of the Digital Signal Processing Technical Committee of the Circuits and Systems Society and a Member of the Board of Nonlinear Signal and Image Processing Society.



Neal C. Gallagher (S'72-M'75-SM'85-F'87) received the Ph.D. degree from Princeton University, Princeton, NJ, in 1974.

He is the Charles Evans Black Professor and Chair of Electrical Engineering at the University of Delaware, Newark. Currently he is working on the use of finite element methods to solve electromagnetic problems, the use of rank order filtering techniques to improve the performance of phase locked loops, the modeling of cellular automata using wave theory, and the autofocus

problem in SAR. Before joining the University of Delaware he was on the faculty at Purdue University for 18 years. Prior to that he was a Faculty Member at Case Western Reserve University. Over the years has published in the areas of stochastic processes, quantization and source coding, electromagnetic theory, and microwave components and optics.

Dr. Gallagher is presently Editor of *Applied Optics Information Processing*. He is a Fellow of the IEEE for his work in nonlinear digital signal processing, and he is a Fellow of the Optical Society of America (OSA) for his work on diffractive optics. He has held various editorial and committee positions in the IEEE, OSA, and SPIE.