Digital Color Halftoning with Generalized Error Diffusion and Multichannel Green-Noise Masks

Daniel L. Lau, Gonzalo R. Arce, *Senior Member, IEEE*, and Neal C. Gallagher, *Fellow, IEEE*

***Abstract*—In this paper, we introduce two novel techniques for digital color halftoning with green-noise—stochastic dither patterns generated by homogeneously distributing minority pixel clusters. The first technique employs error diffusion with output-dependent feedback where, unlike monochrome image halftoning, an interference term is added such that the overlapping of pixels of different colors can be regulated for increased color control. The second technique uses a green-noise mask, a dither array designed to create green-noise halftone patterns, which has been constructed to also regulate the overlapping of different colored pixels. As is the case with monochrome image halftoning, both techniques are tunable, allowing for large clusters in printers with high dot-gain characteristics, and small clusters in printers with low dot-gain characteristics.**

***Index Terms*—AM, color, dither techniques, FM, green-noise, halftoning.**

I. Introduction

DIGITAL halftoning is a technique used by binary display devices to create, within the human eye, the illusion of continuous tone. Designed to mimic analog techniques, dot-clustered ordered dithering or amplitude modulated (AM) halftoning produces this illusion by varying the size of round printed dots which are arranged along an ordered grid. When using AM halftoning, the parameters of particular importance are the lines-per-inch (lpi) or the number of rows/columns of the regular grid[1] and the screen angle or the orientation of the regular grid relative to the horizontal axis. Typically, monochrome screens have an angle of $45°$ as the human visual system is least sensitive to diagonal artifacts [1].

In color printers, the illusion of continuous shades of color is produced by superimposing the binary halftones of cyan, magenta, yellow, and black (CMYK) inks. As the dots of an AM halftone form a regular grid, clustered-dot dithering suffers from *moiré*—the secondary interference patterns created by superimposing two or more regular patterns. In order to minimize the appearance of moiré, the screens of cyan, magenta, yellow, and black are typically oriented at the angles of $15°$, $75°$, $45°$, and $0°$ to create a pleasant *rosette* pattern.

Manuscript received February 5, 1999; revised September 17, 1999. This research was supported in part by the National Science Foundation under Grant CDA-9703088 and by Lexmark International. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jan P. Allebach.

The authors are with the Department of Electrical and Computer Engineering, University of Delaware, Newark, DE 19716 USA (e-mail: lau@ece.udel.edu; arce@ece.udel.edu; gallaghe@ece.udel.edu).

Publisher Item Identifier S 1057-7149(00)03564-8.

[1]The highest quality AM halftones will have 150 lpi or more.

The problems of moiré and screen angles are avoided in frequency modulated halftoning where continuous tone is produced by varying the distance between printed dots and not varying the size. Typically, FM halftones are produced by the process of error diffusion which creates a stochastic arrangement of dots. Besides avoiding moiré, FM halftoning, by isolating minority pixels, maximizes the spatial resolution of the printed image relative to the printer [2], but this distribution also maximizes the perimeter-to-area ratio of printed dots [3]—making FM halftones more susceptible to printer distortions such as dot-gain, the increase in size of a printed dot. Whether a function of the printing process (mechanical dot-gain) or of the optical properties of the paper (optical dot-gain), dot-gain causes the printed halftone to appear darker than the original ratio of white-to-black pixels [4]. In printers with high dot-gain characteristics, AM halftoning, with its lower spatial resolution and moiré, may be the preferred technique, as its clustered-dots have the lower perimeter-to-area ratio.

An alternative to AM and FM halftoning, Levien's [5] error diffusion with output-dependent feedback is an AM-FM hybrid which creates the illusion of continuous tone by producing a stochastic patterning of dot clusters which vary in both their size and in their separation distance. The major advantage, of this new technique over prior error diffusion schemes, is that by adjusting a single parameter, the output is tunable—capable of creating halftones with large clusters in printers with high dot-gain characteristics and small clusters in printers with low dot-gain characteristics. Error diffusion with output-dependent feedback, therefore, can trade halftone visibility for printer robustness.

Studied by Lau *et al.* [2], Levien's technique creates patterns described in terms of their spectral content as green-noise—containing no low or high frequency spectral components. This green-noise model is presented in accordance with Ulichney's [6] blue-noise model which describes the spectral characteristics of the ideal error-diffused halftone patterns as having no low-frequency content. Furthermore, as Mitsa and Parker [7] used the spectral characteristics of blue-noise to generate the blue-noise mask, a binary dither array which greatly reduces the computational complexity associated with FM halftoning, Lau *et al.* [8], using the spatial and spectral characteristics of green-noise, have introduced the green-noise mask.

The problem yet to be addressed in the evolution of green-noise halftoning is its application to color. FM halftoning has been studied in great detail with respect to color printing. The techniques introduced range from simply halftoning each color independently to more complex model-based techniques which transform the CMYK color space to alternate spaces such as the

CIE La*b* space [9]. Yao and Parker [10] have even introduced the blue-noise mask to color halftoning.

This paper introduces, to color printing, green-noise halftoning. The first technique extends error diffusion with output-dependent feedback to not only cluster pixels of like color but also to regulate the clustering of pixels of different colors. That is, with this new technique, the halftoning of different colors can be correlated such that the superimposing of different inks can be increased or decreased. This is in direct contrast to independently halftoning each channel—offering far greater control of resulting halftone patterns.

A second technique, to be introduced in this paper, incorporates a desired correlation between colors to construct a multichannel green-noise mask—extending the capabilities of the prior work in [8] to include color halftoning. By design, this new mask maintains all the desirable attributes of the monochrome mask (isotropic, tunable coarseness) while also regulating the overlapping of pixels of different colors.

II. Halftone Statistics

Point process statistics have a long history in stochastic geometry [11]–[13] and were recently introduced to digital halftoning by Lau *et al.* [2] for the study of periodic dither patterns. In this framework, Φ is a stochastic model governing the location of points in $\Re^2$. ϕ, a sample of Φ, is the set $\{x_i\colon\ i = 1, 2, \cdots\}$ where x_i is a point in $\Re^2$, and the scalar quantity $\phi(B)$ represents the number of x_i's in B, a subset of $\Re^2$. In terms of a monochrome digital halftone pattern, a point is defined to be a minority pixel [a white (1) pixel for gray level $0 \le g < 1/2$ and a black (0) pixel for $1/2 \le g \le 1$].

So given I_g, a binary dither pattern representing a monochrome image of constant gray level g, $\phi[n] = 1$ indicates that the pixel $I_g[n]$ is a minority pixel; otherwise, $\phi[n] = 0$. As a random quantity, the first order moment or the expected value of $\phi[n]$ is its *intensity* $\mathcal{I}[n]$, which is the unconditional probability that $I_g[n]$ is a minority pixel with $\mathcal{I}[n] = g$ for $0 \le g < 1/2$ and $\mathcal{I}[n] = 1 - g$ for $1/2 \le g \le 1$. For a stationary point process where the statistical properties of Φ are independent of n, $\mathcal{I}[n] = \mathcal{I}$.

A second metric for characterizing the statistical properties of Φ is the *reduced second moment measure* $\mathcal{K}[n;\ m]$ defined as

$$\mathcal{K}[n;\ m] = \frac{\boldsymbol{E}\{\phi[n]|\phi[m] = 1\}}{\boldsymbol{E}\{\phi[n]\}} \tag{1}$$

the ratio of the probability that pixel $I_g[n]$ is a minority pixel under the condition that $I_g[m]$ is a minority pixel to the unconditional probability that $I_g[n]$ is a minority pixel. $\mathcal{K}[n;\ m]$ can be interpreted as a measure of the influence of a minority pixel at m on the pixel $I_g[n]$ with $\mathcal{K}[n;\ m] > 1$ indicating that $I_g[n]$ is more likely to be a minority pixel given $I_g[m]$ and $\mathcal{K}[n;\ m] < 1$ indicating that $I_g[n]$ is less likely to be a minority pixel. $\mathcal{K}[n;\ m] = 1$ indicates that $I_g[m]$ has no influence on $I_g[n]$, and in the case of a dither pattern constructed from uncorrelated (white) noise, $\mathcal{K}[n;\ m] = 1$ for all n and m.

For a stationary point process, $\mathcal{K}[n;\ m] = \mathcal{K}(r,\ \theta)$ where r is the distance from m to n and θ is the direction. For an isotropic point process, the statistical properties of Φ are invariant to rotation, and therefore, $\mathcal{K}(r,\ \theta)$ for an isotropic process is written as $\mathcal{R}(r)$ and is referred to as the *pair correlation*. $\mathcal{R}(r)$ is explicitly defined as

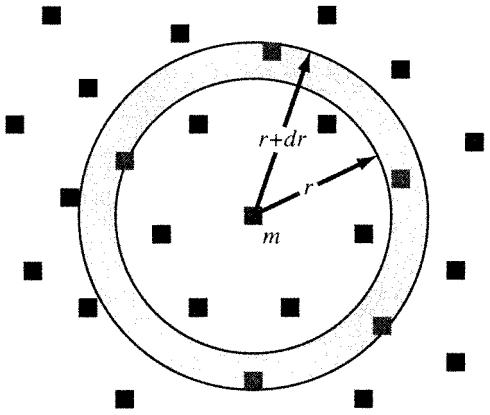

Fig. 1. Spatial ring $R_m(r) = \{n\colon\ r \le |n - m| < r + dr\}$ used to calculate the pair correlation $\mathcal{R}(r)$ for a binary dither pattern.

$$\mathcal{R}(r) = \frac{\boldsymbol{E}\{\phi(R_m(r))|\phi[m] = 1\}}{\boldsymbol{E}\{\phi(R_m(r))\}} \tag{2}$$

the ratio of the expected number of minority pixels located in the ring $R_m(r) = \{n\colon\ r < |n - m| \le r + dr\}$ (Fig. 1) under the condition that $I_g[m]$ is a minority pixel to the unconditional expected number of minority pixels located in $R_m(r)$. $\mathcal{R}(r)$ is also the average $\mathcal{K}[n;\ m]$ of all pixels in the set $R_m(r)$.

$\mathcal{R}(r)$ offers an especially useful tool for characterizing a periodic dither patterns as illustrated in Fig. 2 where three dither patterns representing gray level $g = 15/16$ are shown with their corresponding pair correlations (both the calculated and the ideal). The first pattern (left) is a white-noise dither pattern, and as such has a pair correlation $\mathcal{R}(r) = 1$ for all r as the value of any single pixel in I_g is independent of all other pixels. The name (white-noise) derives from the fact that the resulting power spectrum remains constant for all frequencies [6].

The second pattern (center) is composed of blue-noise which represents gray level g by distributing the minority pixels within I_g as homogeneously as possible—resulting in a dither pattern where minority pixels are placed, on average, a distance of λ_b apart where

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

and D is the minimum distance between addressable points on the display [6], [7]. Referred to as the *principle wavelength* of blue-noise, λ_b is illustrated in $\mathcal{R}(r)$ as series of peaks at integer multiples of λ_b. The term "blue-noise" denotes that the spectral components of a blue-noise dither pattern lie almost exclusively in the high (blue) frequency range.

The final pattern (right) is green-noise where gray level g is represented by homogeneously distributed minority pixel clusters. These clusters are separated, center-to-center, by an av-

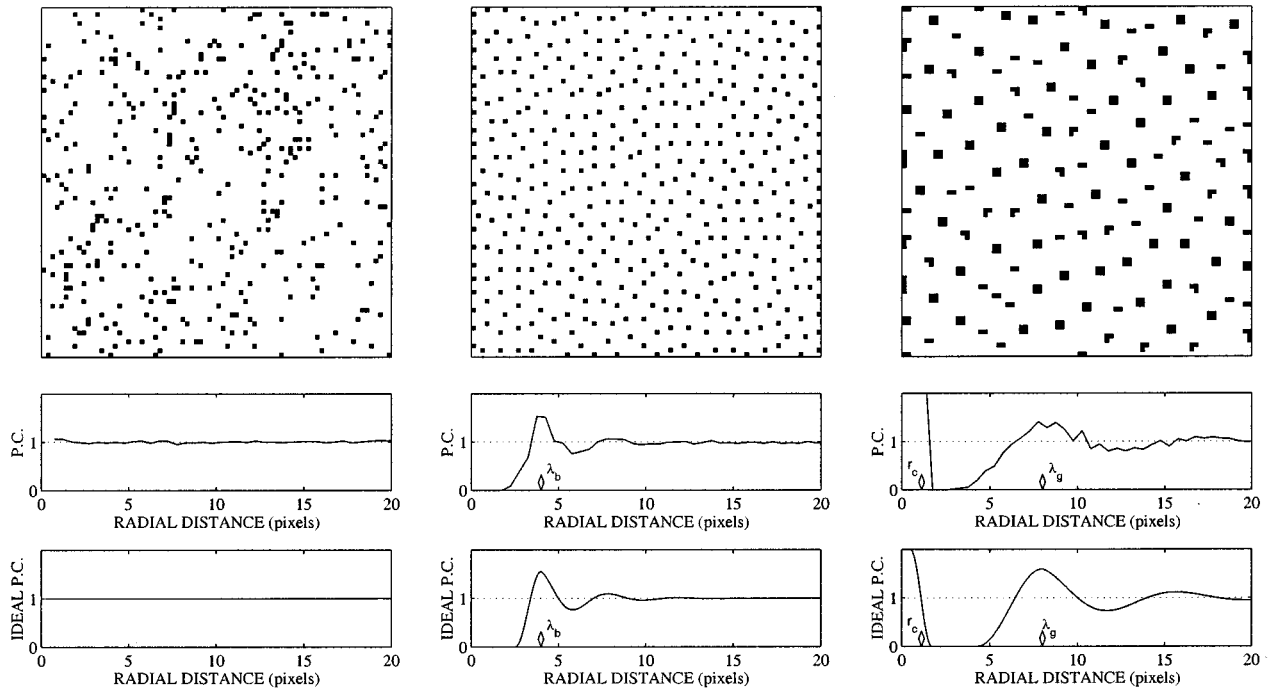

Fig. 2. Pair correlations, $\mathcal{R}(r)$, for (left) white-noise, (center) blue-noise and (right) green-noise dither patterns representing gray level $15/16$.

erage distance of λ_g, the *principle wavelength* of green-noise, where

$$\lambda_g = \begin{cases} D/\sqrt{g/\overline{M}}, & \text{for } 0 < g \leq 1/2 \\ D/\sqrt{(1-g)/\overline{M}}, & \text{for } 1/2 < g \leq 1 \end{cases} \tag{4}$$

and $\overline{M}$ is the average number of minority pixels per cluster. In $\mathcal{R}(r)$, it is the separation of clusters λ_g apart that leads to a series of peaks at integer multiples of λ_g; furthermore, it is the clustering of minority pixels that leads to a nonzero component for r near zero with $\mathcal{R}(r) > 1$ for $0 \leq r < r_c$. The parameter r_c is the *cluster radius* and is related to $\overline{M}$ as

$$\pi r_c^2 = \overline{M} \tag{5}$$

where πr_c^2 is the area covered by a circle with radius r_c. Lau *et al.* [2] note the λ_g is most apparent in $\mathcal{R}(r)$ when the variation in cluster size is small as increasing the variation leads to a "whitened" dither pattern where the peaks and valleys of $\mathcal{R}(r)$ become blurred. The term "green" refers to the resulting patterns' predominantly mid-frequency spectral components with the blue-noise model a limiting case ($\overline{M} = 1$).

A. Color Halftoning

In the case of a color halftone, the monochrome model I_g must be revised as a dither pattern is now composed of C colors where, for generality, the quantity C is an arbitrary integer. For RGB and CMYK, where images are composed of the additive colors red, green and blue or the subtractive colors cyan, magenta, yellow, and black, $C = 3$ and 4, respectively. So for color images, the halftone pattern I_g is now composed of the monochrome binary dither patterns $I_{g_1}, I_{g_2}, \cdots, I_{g_C}$ where g_i is the gray level of pattern I_{g_i} and ϕ_i is the corresponding point process.

In this new framework, the quantity $\mathcal{K}_{i,j}[n;\ m]$ is the reduced second moment measure between colors such that

$$\mathcal{K}_{i,j}[n;\ m] = \frac{\boldsymbol{E}\{\phi_{g_i}[n] | \phi_{g_j}[m] = 1\}}{\boldsymbol{E}\{\phi_{g_i}[n]\}} \tag{6}$$

is the ratio of the conditional probability that $\phi_{g_i}[n]$ is a minority pixel given that a minority pixel exists at sample m of ϕ_{g_j} to the unconditional probability that $\phi_{g_i}[n]$ is a minority pixel [8]. Similar to $\mathcal{K}[n;\ m]$, $\mathcal{K}_{i,j}[n;\ m] = 1$ indicates that the location of minority pixels in colors i and j are uncorrelated. The pair correlation between colors g_i and g_j follows as

$$\mathcal{R}_{i,j}(r) = \frac{\boldsymbol{E}\{\phi_{g_i}(R_m(r)) | \phi_{g_j}[m] = 1\}}{\boldsymbol{E}\{\phi_{g_i}(R_m(r))\}} \tag{7}$$

the ratio of the expected number of minority pixels of color g_i located in the ring $R_m(r) = \{n\colon\ r < |n-m| \leq r + dr\}$ under the condition that $\phi_{g_j}[m]$ is a minority pixel to the unconditional expected number of minority pixels with color g_i located in $R_m(r)$.

III. Generalized Error Diffusion

A. Monochrome

In error diffusion (Fig. 3), the output pixel $y[n]$ is determined by adjusting and thresholding the input pixel $x[n]$ such that

$$y[n] = \begin{cases} 1, & \text{if } (x[n] + x^e[n]) \geq 0 \\ 0, & \text{else} \end{cases} \tag{8}$$

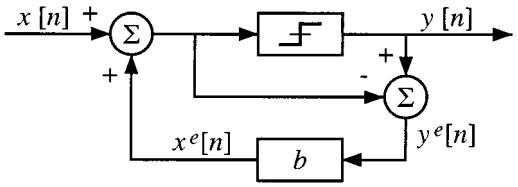

Fig. 3. Error diffusion algorithm as introduced by Floyd–Steinberg [14].

where $x^e[n]$ is the diffused quantization error accumulated during previous iterations as

$$x^e[n] = \sum_{i=1}^{M} b_i \cdot y^e[n-i] \tag{9}$$

with $y^e[n] = y[n] - (x[n] + x^e[n])$ and $\sum_{i=1}^{M} b_i = 1$. Using vector notation, (9) becomes

$$x^e[n] = \boldsymbol{b}^T \boldsymbol{y}^e[n] \tag{10}$$

where $\boldsymbol{b} = [b_1, b_2, \cdots, b_M]^T$ and $\boldsymbol{y}^e[n] = [y^e[n-1], y^e[n-2], \cdots, y^e[n-M]]^T$.

In [5], Levien adds an output-dependent feedback term (Fig. 4) to (8) such that

$$y[n] = \begin{cases} 1, & (x[n] + x^e[n] + x^h[n]) \geq 0 \\ 0, & \text{else} \end{cases} \tag{11}$$

where $x^h[n]$ is the hysteresis or feedback term defined as

$$x^h[n] = h \sum_{i=1}^{N} a_i \cdot y[n-i] \tag{12}$$

with $\sum_{i=0}^{N} a_i = 1$ and h is an arbitrary constant. Referred to as the *hysteresis constant*, h acts as a tuning parameter with larger h leading to coarser output textures [5] as h increases ($h > 0$) or decreases ($h < 0$) the likelihood of a minority pixel if the previous outputs were also minority pixels. Equation (12) can also be written in vector notation as

$$x^h[n] = h\boldsymbol{a}^T \boldsymbol{y}[n] \tag{13}$$

where $\boldsymbol{a} = [a_1, a_2, \cdots, a_N]^T$ and $\boldsymbol{y}[n] = [y[n-1], y[n-2], \cdots, y[n-N]]^T$. The calculation of the parameters $\boldsymbol{y}^e[n]$ and $x^e[n]$ remains unchanged in Levien's approach. So in summary of Levien's error diffusion with output-dependent feedback, the binary output pixel $y[n]$ is determined as

$$y[n] = \begin{cases} 1, & \text{if } (x[n] + \boldsymbol{b}^T \boldsymbol{y}^e[n] + h\boldsymbol{a}^T \boldsymbol{y}[n]) \geq 0 \\ 0, & \text{else} \end{cases} \tag{14}$$

where $\boldsymbol{y}^e[n] = [y^e[n-1], y^e[n-2], \cdots, y^e[n-N]]^T$ such that $y^e[n] = y[n] - (x[n] + x^e[n])$ and $\boldsymbol{y}[n] = [y[n-1], y[n-2], \cdots, y[n-N]]^T$.

B. Color

Now consider the C-channel case where an output pixel is not the binary pixel $y[n]$ but the C-dimensional vector $\vec{y}[n]$ such that

$$\vec{y}[n] = \begin{bmatrix} y_1[n] \\ y_2[n] \\ \vdots \\ y_C[n] \end{bmatrix} \tag{15}$$

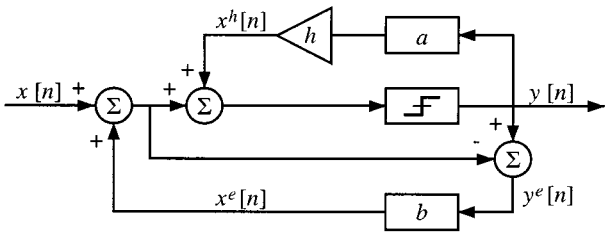

Fig. 4. Error diffusion with output-dependent feedback algorithm as introduced by Levien [5].

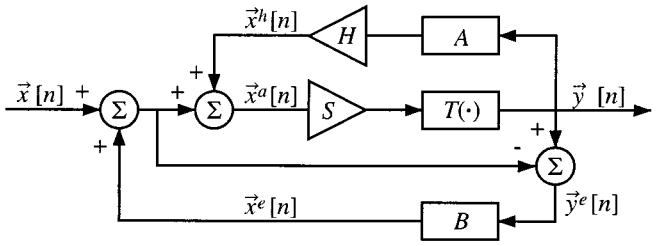

Fig. 5. Generalized error diffusion algorithm.

where $y_i[n]$ is the binary output pixel of color i. Assuming all C channels are halftoned independently, the binary output pixel $y_i[n]$ is determined as

$$y_i[n] = \begin{cases} 1, & \text{if } (x_i[n] + x_i^e[n] + x_i^h[n]) \geq 0 \\ 0, & \text{else} \end{cases} \tag{16}$$

where $x_i^e[n]$ and $x_i^h[n]$ are the error and hysteresis terms, respectively, for the ith color. The error term, being a vector, is calculated as

$$\vec{x^e}[n] = \begin{bmatrix} \boldsymbol{b}_1^T & 0 & \cdots & 0 \\ 0 & \boldsymbol{b}_2^T & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \boldsymbol{b}_C^T \end{bmatrix} \begin{bmatrix} \boldsymbol{y}_1^e[n] \\ \boldsymbol{y}_2^e[n] \\ \vdots \\ \boldsymbol{y}_C^e[n] \end{bmatrix}$$
$$\vec{x^e}[n] = \boldsymbol{B}\boldsymbol{Y}^e[n] \tag{17}$$

where $\boldsymbol{b}_i$ are the filter weights regulating the diffusion of error in the ith channel and $\boldsymbol{y}_i^e[n]$ is the vector $[y_i^e[n-1], y_i^e[n-2], \cdots, y_i^e[n-N]]^T$ composed exclusively from errors in channel i such that $y_i^e[n] = y_i[n] - (x_i[n] + x_i^e[n])$. The hysteresis term $\vec{x^h}[n]$, also a vector, is calculated as

$$\vec{x^h}[n] = \begin{bmatrix} h_1 & 0 & \cdots & 0 \\ 0 & h_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & h_C \end{bmatrix} \begin{bmatrix} \boldsymbol{a}_1^T & 0 & \cdots & 0 \\ 0 & \boldsymbol{a}_2^T & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \boldsymbol{a}_C^T \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{y}_1[n] \\ \boldsymbol{y}_2[n] \\ \vdots \\ \boldsymbol{y}_C[n] \end{bmatrix}$$
$$\vec{x^h}[n] = \boldsymbol{H}\boldsymbol{A}\boldsymbol{Y}[n] \tag{18}$$

where $\boldsymbol{a}_i$ are the filter weights and h_i is the hysteresis constant that regulates the diffusion of feedback in the ith channel.

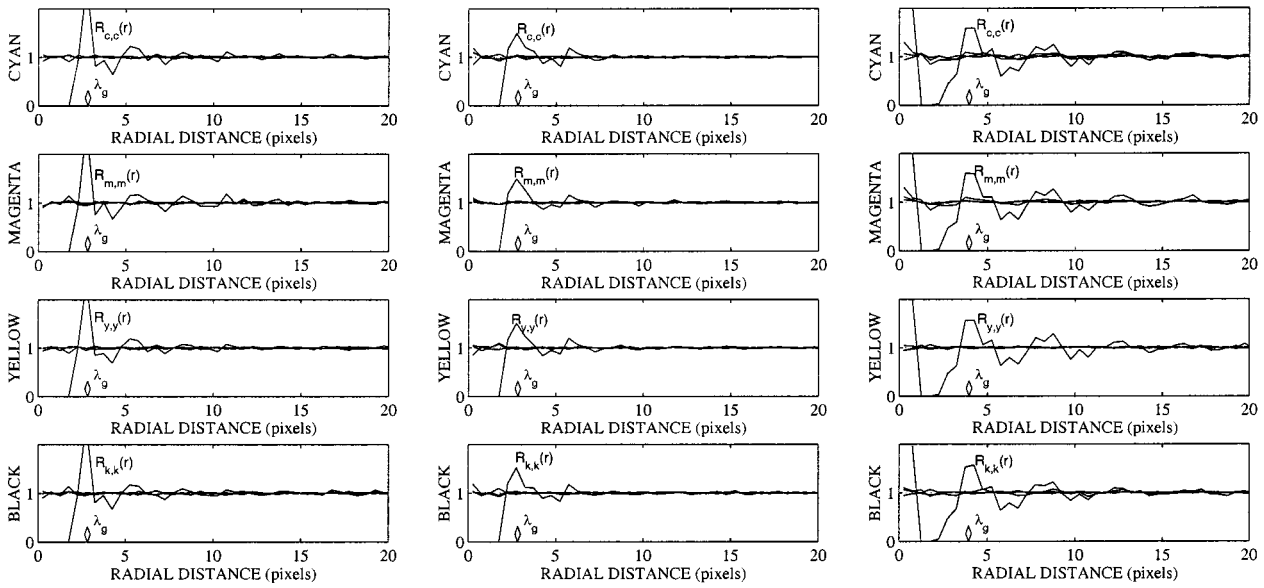

Fig. 6. Pair correlations for CMYK halftone patterns with no diffusion between colors using: (left) Floyd–Steinberg error diffusion weights and no hysteresis, (center) Levien diffusion with small hysteresis constant ($h = 0.5$), and (right) Levien diffusion with medium hysteresis constant ($h = 1.0$).

Generalized even further, (17) becomes

$$\vec{x^e}[n] = \begin{bmatrix} \boldsymbol{b}_{1,1}^T & \boldsymbol{b}_{1,2}^T & \cdots & \boldsymbol{b}_{1,C}^T \\ \boldsymbol{b}_{2,1}^T & \boldsymbol{b}_{2,2}^T & \cdots & \boldsymbol{b}_{2,C}^T \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{b}_{C,1}^T & \boldsymbol{b}_{C,2}^T & \cdots & \boldsymbol{b}_{C,C}^T \end{bmatrix} \begin{bmatrix} \boldsymbol{y}_1^e[n] \\ \boldsymbol{y}_2^e[n] \\ \vdots \\ \boldsymbol{y}_C^e[n] \end{bmatrix} \tag{19}$$

where quantization error can now be diffused between channels through $\boldsymbol{b}_{i,j}$, the error filter weights which regulate the diffusion of error from channel j to channel i; furthermore, (18) becomes:

$$\vec{x^h}[n] = \begin{bmatrix} h_{1,1} & h_{1,2} & \cdots & h_{1,C} \\ h_{2,1} & h_{2,2} & \cdots & h_{2,C} \\ \vdots & \vdots & \ddots & \vdots \\ h_{C,1} & h_{C,2} & \cdots & h_{C,C} \end{bmatrix} \cdot \begin{bmatrix} \boldsymbol{a}_{1,1}^T & \boldsymbol{a}_{1,2}^T & \cdots & \boldsymbol{a}_{1,C}^T \\ \boldsymbol{a}_{2,1}^T & \boldsymbol{a}_{2,2}^T & \cdots & \boldsymbol{a}_{2,C}^T \\ \vdots & \vdots & \ddots & \vdots \\ \boldsymbol{a}_{C,1}^T & \boldsymbol{a}_{C,2}^T & \cdots & \boldsymbol{a}_{C,C}^T \end{bmatrix} \begin{bmatrix} \boldsymbol{y}_1[n] \\ \boldsymbol{y}_2[n] \\ \vdots \\ \boldsymbol{y}_C[n] \end{bmatrix} \tag{20}$$

where the previous outputs of channel j can impact all other channels where $h_{i,j} \neq 0$ for $i \neq j$ as $\boldsymbol{a}_{i,j}$ and $h_{i,j}$ regulate the diffusion of feedback from channel j to i.

Before concluding this section, we make one last, but significant, modification to error diffusion (Fig. 5) by first defining the thresholding function $T(\cdot)$ as

$$\begin{aligned} \vec{y}[n] &= T(\vec{x}[n] + \vec{x^e}[n] + \vec{x^h}[n]) \\ &= T(\vec{x^a}[n]) \end{aligned} \tag{21}$$

where $\vec{x^a}[n] = \vec{x}[n] + \vec{x^e}[n] + \vec{x^h}[n]$ is the *accumulated* input pixel. The $C \times C$ *interference* matrix $\boldsymbol{S}$ is added to (21) as

$$\vec{y}[n] = T(\boldsymbol{S}\vec{x^a}[n]) \tag{22}$$

such that $S_{i,j}$ is the influence on the thresholding function of color i by the accumulated input of color j. The effect of $S_{i,j}$ is to increase ($S_{i,j} > 0$) or decrease ($S_{i,j} < 0$) the likelihood of a minority pixel at $y_i[n]$ based on the *likelihood* of a minority pixel at $y_j[n]$. Finally, we summarize error diffusion by the *generalized error diffusion equation*

$$\vec{y}[n] = T(\boldsymbol{S}(\vec{x}[n] + \boldsymbol{B}\boldsymbol{Y}^e[n] + \boldsymbol{H}\boldsymbol{A}\boldsymbol{Y}[n])). \tag{23}$$

C. Simulations

Shown in Fig. 12(a) is the resulting CMYK dither patterns created by halftoning a 96×96 pixel color image of constant color value $\vec{x}[n] = [7/8, 7/8, 7/8, 7/8]^T$ using the Floyd–Steinberg [14] error filter weights with no hysteresis and no dependencies between colors. Before halftoning, low-level white-noise was added to the first scan line of the original color image in order to minimize edge effects and also to unsynchronize the resulting dither patterns. Since in this configuration where all colors are halftoned exactly the same way and with each color of the original image identical to the other colors, the resulting pattern of each color will also be identical (synchronized) to the other patterns. So by adding a single or even a few lines of low-level noise eliminates this synchronization between colors; furthermore, adding several columns of white-noise also minimizes edge effects. In this paper, dither patterns created by error diffusion are the result of using a serpentine (left-to-right and then right-to-left) raster scan on a continuous-tone image where low-level white-noise ($mean = 0$, $var = 0.1$) has been added for the sole purpose of unsynchronizing each channel to both the edge rows and columns. The halftoned images are then cropped to exclude those same rows and columns.

For a statistical analysis of the resulting dither pattern, Fig. 6 (left) shows four plots labeled cyan, magenta, yellow, and black corresponding to the CMYK dither pattern in Fig. 12(a). Shown in the first plot (labeled cyan) is the pair correlation between colors cyan versus cyan [$\mathcal{R}_{c,c}(r)$], cyan versus magenta [$\mathcal{R}_{c,m}(r)$], cyan versus yellow [$\mathcal{R}_{c,y}(r)$] and cyan versus black [$\mathcal{R}_{c,k}(r)$]. The small diamonds placed along the horizontal axis indicate the principle wavelengths and cluster radii for

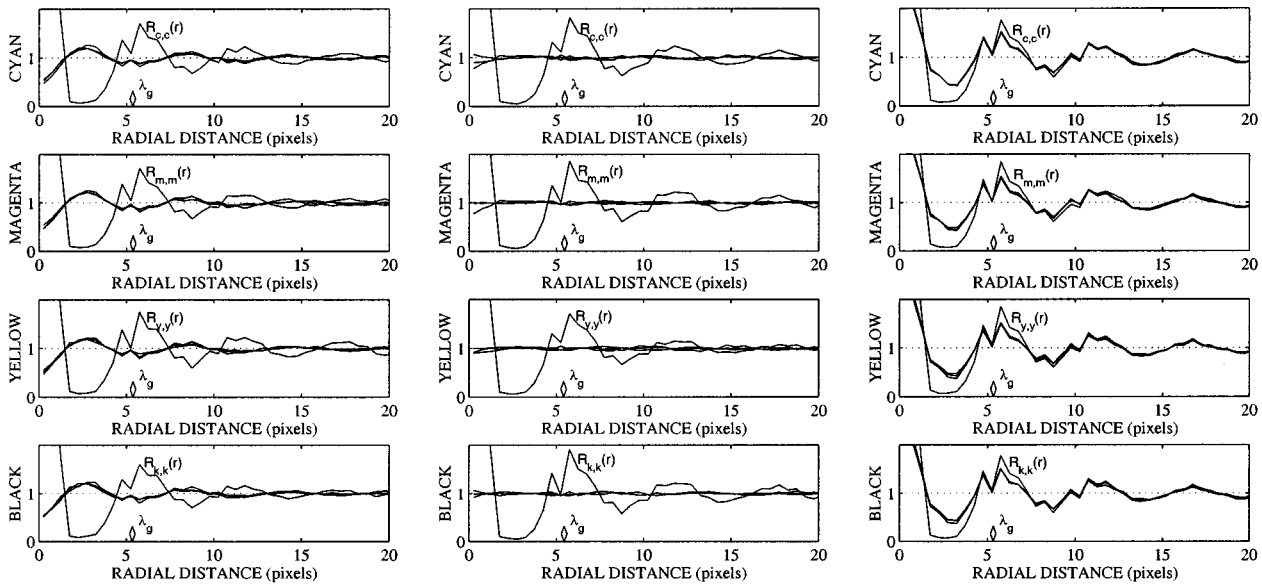

Fig. 7. Pair correlations for CMYK halftone patterns using Levien diffusion with high hysteresis constant ($h = 1.5$) with: (left) a negative interference term ($S_{i \neq j} = -0.2$), (center) no interference term ($S_{i \neq j} = 0$), and (right) a positive interference term ($S_{i \neq j} = +0.2$).

$\mathcal{R}_{c,c}(r)$. As would be expected for a monochrome image, the pair correlation $\mathcal{R}_{c,c}(r)$ exhibits blue-noise characteristics as the pair correlation shows

1) $\mathcal{R}_{c,c}(r) < 1$ for r near zero;
2) frequent occurrence of the interpoint distance λ_g;
3) decreasing influence with increasing r.

Having no diffusion between colors and zero interference ($\boldsymbol{S} = \boldsymbol{I}$, the identity matrix), the pair correlations between channels are predominantly flat as minority pixels of color i have no influence on minority pixels of color j. The remaining plots show similar relationships for colors magenta, yellow, and black.

Shown in Fig. 6 (center) are the resulting pair correlations to Fig. 12(b) where Levien's error diffusion scheme has been implemented with a low hysteresis constant $h = 0.5$, no diffusion between colors and zero interference. With a low hysteresis constant, this scheme generates blue-noise patterns very similar to that generated using the Floyd-Steinberg filter weights. Fig. 6 (right) [Fig. 12(c)] shows Levien's error diffusion scheme with a medium hysteresis constant $h = 1.0$ where the patterns begin to exhibit clustering as the average size of a minority pixel cluster is 1.95 pixels. In each color, the pair correlation exhibits strong green-noise characteristics as each plot shows

1) clustering as indicated by $\mathcal{R}_{i,i}(r) \geq 1$ for $r \leq r_c$ ($r_c = 0.79$);
2) frequent occurrence of the intercluster distance $\lambda_g = 3.95$;
3) decreasing influence with increasing r. As before with zero influence and no diffusion between colors, the pair correlations between colors remains predominantly flat for all r.

The dither patterns of Fig. 7 illustrate the effects of $\boldsymbol{S}$, the interference matrix, with Fig. 7 (left) [Fig. 12(d)] showing the case where $\boldsymbol{S}$ is the matrix defined by $S_{i,j} = 1$ for $i = j$ and $S_{i,j} = -0.2$ for $i \neq j$. In all instances where $S_{i,j} < 0$, $\boldsymbol{S}$ has the effect of reducing the superposition of minority pixels of different colors with lesser $S_{i,j}$ leading to lesser overlap. That is, given that a cyan pixel is very likely to be printed, minority pixels for magenta, yellow and black are less likely to be printed at that same pixel location. This behavior is well illustrated in the pair correlations where $\mathcal{R}_{i,j} < 1$ for $\mathcal{R}_{i,i} > 1$. For comparison, Fig. 7 (center) [Fig. 12(e)] shows the case where $\boldsymbol{S}$ is the identity matrix (no interference) with a flat pair correlation between minority pixels of different colors. For further comparison, Fig. 7 (right) [Fig. 12(f)] shows the case where $\boldsymbol{S}$ is the matrix defined by $S_{i,j} = 1$ for $i = j$ and $S_{i,j} = +0.2$ for $i \neq j$. Here, the effect of $\boldsymbol{S}$ is to increase the superposition of minority pixels such that a minority pixel of color i with a high likelihood of being printed making a minority pixel of any color j more likely.

IV. Multi-Channel Green-Noise Masks

The green-noise mask is a novel approach to dither array screening where a continuous-tone image is converted to a binary halftone image by performing a pixel-wise comparison between the original and the dither array or mask. Previously, halftoning with green-noise has implied error-diffusion based methods which although are tunable (capable of creating halftone patterns with large clusters for printers with high dot-gain characteristics and small clusters for printers with low dot-gain characteristics) carry a high computational cost. Now through the use of a green-noise mask, halftoning can create a stochastic patterning of dots with adjustable coarseness but with the same computational freedom as ordered-dither halftoning schemes—an advantage that, for many printing devices, overcomes the drawbacks of distortions inherent to dither array halftoning such as tiling artifacts. Many such drawbacks, though, can be minimized and sometimes visually eliminated using device dependent compensation techniques.

Introduced in [8] for monochrome images, the green-noise mask is defined by the set, $\{\phi_g\colon\ 0 \leq g \leq 1\}$, of $M \times N$

binary green-noise dither patterns with one pattern, ϕ_g, corresponding to each possible discrete gray level g (256 patterns for 8-bit gray-scale images). This set satisfies the *stacking constraint* that for any two gray-levels k and g with $k < g$, $\phi_k \subset \phi_g$ (if $\phi_k[m, n] = 1$ then $\phi_g[m, n] = 1$). As a consequence, a pixel of the $M \times N$ dither array or mask $DA[m, n]$ is defined simply as the minimum g for which $\phi_g[m, n] = 1$. The size parameters M and N are arbitrary integers with larger masks constructed by tiling edge-to-edge the original $M \times N$ mask such that the output pixel, $y[m, n]$, after halftoning the input pixel, $x[m, n]$, is defined as

$$y[m, n] = \boldsymbol{T}(x[m, n] - DA[\mathrm{mod}_M(m), \mathrm{mod}_N(n)]) \quad (24)$$

where T is the thresholding function of (21).

For color halftoning, the multichannel green-noise mask is defined by the set $\{\phi_{c,g}\colon c = 1, 2, \cdots, C \text{ and } 0 \le g \le 1\}$ where $\phi_{c,g}$ is the binary green-noise dither pattern for color c and intensity level g (for 24-bit RGB color this corresponds to 256 patterns per channel or 3×256 total patterns). Like the monochrome set, this set must also satisfy the stacking constraint but only within a given color such that $\phi_{c,g}[m, n] = 1$ if $\phi_{c,k}[m, n] = 1$ for color c and intensity levels k and g with $k < g$. A pixel, $\vec{DA}[m, n]$, of the multichannel green-noise mask is therefore defined as:

$$\vec{DA}[m, n] = \begin{bmatrix} \min\{g\colon \phi_{1,g}[m, n] = 1\} \\ \min\{g\colon \phi_{2,g}[m, n] = 1\} \\ \vdots \\ \min\{g\colon \phi_{C,g}[m, n] = 1\} \end{bmatrix} \quad (25)$$

where the output after halftoning is defined by

$$\vec{y}[m, n] = \boldsymbol{T}(\vec{x}[m, n] - \vec{DA}[\mathrm{mod}_M(m), \mathrm{mod}_N(n)]). \quad (26)$$

A. Monochrome BIPPCCA

The physical construction of binary dither patterns for the monochrome green-noise mask is done through BIPPCCA (the BInary Pattern Pair Correlation Construction Algorithm). The basic premise of BIPPCCA is to take an empty array (containing no minority pixels) and assign, to each element, a probability of that element becoming a minority pixel. BIPPCCA will then convert the most likely elements to minority pixels, one at a time, until the ratio of black to white pixels is g, the desired gray level. The most likely element is the majority pixel with the highest probability during the current iteration, and in order for the resulting dither pattern to have desired statistical properties (i.e., a desired pair correlation), BIPPCCA will adjust, at each iteration, the probability of each majority pixel in the array according to the current set of minority pixels.

In BIPPCCA, the initial assignment of probabilities is done in an uncorrelated manner, but as each new minority pixel is added, the probabilities of all neighboring majority pixels are adjusted according to the desired pair correlation of the resulting pattern. As $\mathcal{R}(r)$ is a function of the radial distance between pixels, a majority pixel's probability is increased if its radial distance from the newest minority pixel, r, corresponds to $\mathcal{R}(r) > 1$ and decreased if r corresponds to $\mathcal{R}(r) < 1$. As an example, consider using BIPPCCA to construct a blue-noise pattern where the pair correlation is zero for r near zero. This feature of $\mathcal{R}(r)$ is achieved in BIPPCCA if with each new minority pixel, the probability of every element directly adjacent is set to zero; furthermore, as $\mathcal{R}(r)$ has a peak at λ_b, the blue-noise principle wavelength, all elements a distance λ_b from each new minority pixel should be increased to ensure a peak exists in the pair correlation of the final pattern.

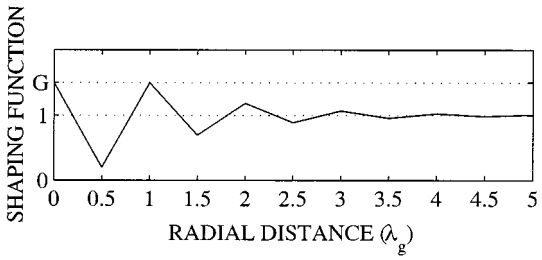

Fig. 8. Pair correlation shaping function, $\tilde{\mathcal{R}}(r)$, used to construct green-noise with principle wavelength λ_g.

In practice, how much to increase or decrease a given probability, in BIPPCCA, is defined according to $\tilde{\mathcal{R}}(r)$, the *pair correlation shaping function*. $\tilde{\mathcal{R}}(r)$ is a user-defined function based on the desired pair correlation with increasing $\tilde{\mathcal{R}}(r)$ leading to stronger correlations and decreasing $\tilde{\mathcal{R}}(r)$ leading to reduced. At $\tilde{\mathcal{R}}(r) = 0$, minority pixels are completely inhibited. In designing $\tilde{\mathcal{R}}(r)$, it is important to note that $\tilde{\mathcal{R}}(r)$ does not have absolute control over the resulting $\mathcal{R}(r)$, but with careful tuning, $\mathcal{R}(r)$ will approximate the shape of $\tilde{\mathcal{R}}(r)$. Shown in Fig. 8 is the shaping function used by Lau *et al.* [8] to construct green-noise dither patterns. This function has peaks at integer multiples of λ_g, the green-noise principle wavelength, and valleys mid-way between. The parameter G is a tuning parameter and is shown in [8] to create visually pleasing patterns when $G = 1.10$. Being piecewise linear, this pair correlation shaping function is an especially simple approximation of the pair correlation of the ideal green-noise pattern for a given gray level and cluster size, but by itself, resulting dither patterns tend to look noisy and nonstationary.

Because stationarity is a necessary property for digital halftoning [6], the most likely pixel will no longer be the majority pixel with the highest probability, but instead be the majority pixel with the highest product $U[m, n] \times CM[m, n]$ where $U[m, n]$ is the probability of a given pixel and $CM[m, n]$ is a function of the density of minority pixels within the surrounding area. Referred to as the *concentration matrix*, CM makes majority pixels more likely to become minority pixels in areas of low minority pixel concentration and less likely in areas of high.

In BIPPCCA, the concentration of minority pixels is measured as the output after applying a low-pass filter, H_{LP}, using circular convolution. In [8], Lau *et al.* construct green-noise patterns using the Gaussian filter, H_{LP}, defined as

$$H_{LP}(r) = \exp\left(\frac{-r^2}{\lambda_g^2}\right) \quad (27)$$

where H_{LP} has a wide-spread impulse response for large λ_g where clusters are far apart and a narrow-spread impulse response for small λ_g where clusters are close together. How much to increase or decrease a probability according to the minority pixel concentration is then determined by the user through a mapping of the filtered output to the concentration matrix. Fig. 9

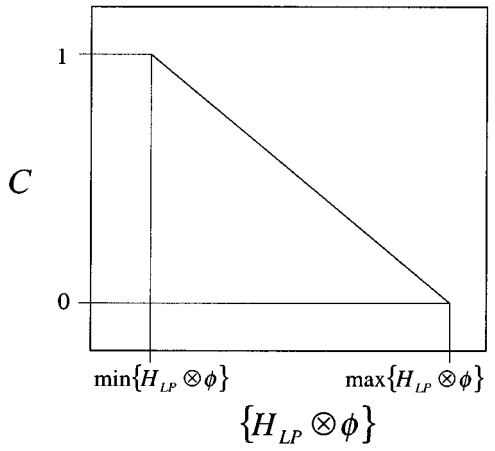

Fig. 9. Mapping function used to construct the concentration matrix C from the output after filtering ϕ with the low-pass filter H_{LP} using circular convolution.

shows the mapping of concentration values used in [8] to determine CM where $\{H_{LP} \otimes \phi_g\}$ represents the output after filtering the binary dither pattern of the current iteration, ϕ_g, with the low-pass filter H_{LP}. In this mapping, values of $\{H_{LP} \otimes \phi_g\}$ are scaled in a linear fashion such that $\max\{H_{LP} \otimes \phi_g\} \rightarrow 0$ and $\min\{H_{LP} \otimes \phi_g\} \rightarrow 1$.

In summary, the steps for monochrome BIPPCCA are performed as follows where ϕ_g is initially an $M \times N$ array with no minority pixels.

1) Create an $M \times N$ array, U, of uniformly distributed random numbers such that $U[m, n] \in (0, 1]$ is the probability that $\phi_g[m, n]$ will become a minority pixel.
2) Construct the concentration matrix CM using a user-defined mapping of $\{H_{LP} \otimes \phi_g\}$, the output after filtering ϕ_g with the low-pass filter H_{LP} using circular convolution, and then locate the majority pixel in ϕ_g with the highest modified probability (the majority pixel $\phi_g[m, n]$ such that $U[m, n] \times CM[m, n] > U[o, p] \times CM[o, p]$ for all $1 \leq o \leq M$ and $1 \leq p \leq N$ where $\phi_g[o, p]$ is also a majority pixel). Replace that pixel, $\phi_g[m, n]$, with a minority pixel.
3) Given the new minority pixel, $\phi_g[m, n]$, adjust the probability of each and every majority pixel, $\phi_g[o, p]$, such that:

$$(U[o, p])_{\text{new}} = (U[o, p])_{\text{old}} \times \tilde{\mathcal{R}}(r) \tag{28}$$

where r is the minimum wrap-around distance from the majority pixel $\phi_g[o, p]$ to the new minority pixel $\phi_g[m, n]$ defined for an $M \times N$ array as shown in (29) at the bottom of the page.
4) If the ratio of the total number of minority pixels to the total number of pixels in ϕ_g is equal to g, the desired gray-level, then quit with the desired dither pattern given by ϕ_g; otherwise, continue at step 2.

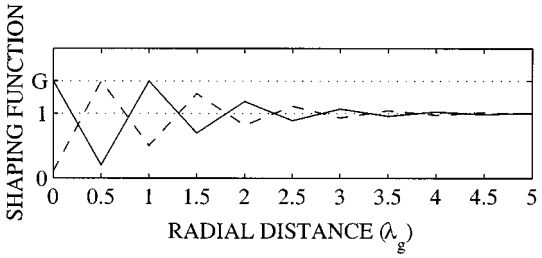

Fig. 10. Pair correlation shaping functions, $\bar{\mathcal{R}}_{i,i}(r)$ (solid line) and $\bar{\mathcal{R}}_{i,j}(r)$ (dashed line), used to construct color green-noise with principle wavelength λ_g and decreased overlap of minority pixels of different colors.

As described, the above algorithm is not suited to the design of green-noise masks as the stacking constraint will not be satisfied for all patterns. BIPPCCA must, therefore, be constrained to create a pattern, ϕ_g, which is constructed to have, as a subset, all constructed patterns ϕ_k for which $k < g$. ϕ_g must also be a subset of all constructed patterns ϕ_l for which $l > g$.

In order to constrain BIPPCCA, assume first that $g < 1/2$ and that $\phi_k \subset \phi_g \subset \phi_l$. The first step in then to initialize ϕ_g to ϕ_k, where $\phi_k[n] = 1$ is a minority pixel, instead of an all majority pixel array. Step 3 is then applied for each and every minority pixel in ϕ_g to the probability matrix, U, of uniformly distributed random numbers. BIPPCCA then continues at step 2 where, in order to satisfy the constraint that $\phi_g \subset \phi_l$, only those majority pixels ($\phi_g[n] = 0$) for which $\phi_l[n] = 1$ are considered for swapping. BIPPCCA then continues using these modified steps until a sufficient number of minority pixels in exist in ϕ_g. For the case of $g > 1/2$, the same modifications as above are used except that ϕ_g is initialized to ϕ_l where minority pixels are equal to 0. Step 2 is then constrained by ϕ_k and not ϕ_l.

ϕ_g to ϕ_k and applying step 3 for each and every minority pixel in ϕ_k to the probability matrix, U, of uniformly distributed random numbers. BIPPCCA then continues at step 4. The second constraint is satisfied in step 2 of BIPPCCA when locating the maximum likely majority pixel by considering only those majority pixels in ϕ_g which correspond to minority pixels in the constructed patterns, ϕ_l, for which $l > g$.

In applying this *constrained* BIPPCCA to mask design, note that the patterns composing the set $\{\phi_g\colon 0 \leq g \leq 1\}$ can be constructed in any order and that order does have an impact on the construction of each pattern as ϕ_g is constrained by the constructed pattern corresponding to the maximum gray level that is less than g and constrained by the constructed pattern corresponding to the minimum gray level which is greater than g. While no criteria for choosing an optimal ordering, or even an initial gray level, has been offered, generating patterns in a random order may offer better results than by constructing patterns according to consecutive gray levels. As an example, when constructing green-noise masks [8], use the interleaved ordering $\{0, 1, 1/4, 1-(1/4), 1/4+\Delta, 1-(1/4+\Delta), 1/4+2\Delta, 1-(1/4+2\Delta), \cdots, 1/2, 1/4-\Delta, 1-(1/4-\Delta), 1/4-2\Delta, 1-(1/4-2\Delta), \cdots, \Delta, 1-\Delta\}$ where $\Delta = 1/255$.

$$r = \sqrt{\min\{|m-o|,\, M-|m-o|\}^2 + \min\{|n-p|,\, N-|n-p|\}^2} \tag{29}$$

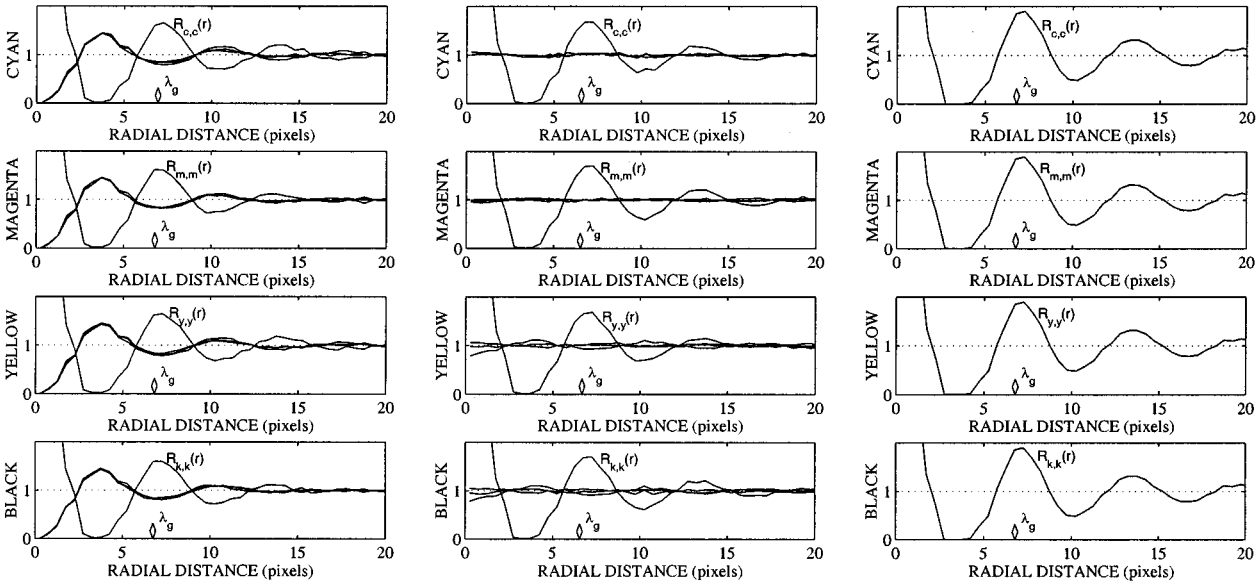

Fig. 11. CMYK green-noise dither patterns created using the VBIPPCCA algorithm with: (left) decreased overlapping, (center) uncorrelated overlapping, and (right) increased overlapping.

B. Color BIPPCCA

The physical construction of binary dither patterns for the multichannel green-noise mask is done through MBIPPCCA (the Multichannel BInary Pair Correlation Construction Algorithm). In color, a binary dither pattern representing the color $\vec{g} = [g_1, g_2, \cdots, g_C]^T$ is defined by the set of monochrome images $\{\phi_{1,g_1}, \phi_{2,g_2}, \cdots, \phi_{C,g_C}\}$ where ϕ_{i,g_i} is the binary dither pattern corresponding to the ith color with intensity g_i. MBIPPCCA constructs these monochrome images according to the previous algorithm, but unlike BIPPCCA, when a minority pixel is added to the ith color, the probabilities corresponding to majority pixels of color j are adjusted according to $\mathcal{R}_{j,i}(r)$, the desired pair correlation between minority pixels of colors j and i. So for a CMYK dither pattern, for each minority cyan pixel added, MBIPPCCA will make use of the user-defined shaping functions $\tilde{\mathcal{R}}_{c,c}(r)$, $\tilde{\mathcal{R}}_{m,c}(r)$, $\tilde{\mathcal{R}}_{y,c}(r)$, and $\tilde{\mathcal{R}}_{k,c}(r)$ to adjust the probabilities of majority pixels in the cyan, magenta, yellow, and black colors, respectively.

Because stationarity is also a desired property for digital color halftoning, MBIPPCCA will apply CM just as in the monochrome case with each color filtered independently of the others. Returning to the CMYK case, this implies that the maximum likely majority pixel of the cyan color is the majority pixel of the cyan color with the highest product $U_c[m, n] \times CM_c[m, n]$ where $U_c[m, n]$ is the probability array for cyan pixels and $CM_c[m, n]$ is the concentration matrix formed by applying a user-defined mapping to the concentration of minority cyan pixels.

In summary, the steps for MBIPPCCA are performed as follows where $\{\phi_{i,g_i}: i = 1, 2, \cdots, C\}$ is the initial set of empty $M \times N$ arrays.

1) Create a set of $M \times N$ arrays, $\{U_1, U_2, \cdots, U_C\}$, of uniformly distributed random numbers such that $U_i[m, n] \in (0, 1]$ is the probability that $\phi_{i,g_i}[m, n]$ will become a minority pixel.
2) For $i = 1, 2, \cdots, C$, if the ratio of the total number of minority pixels to the total number of pixels in ϕ_{i,g_i} is less than g_i, then
   a) Construct the concentration matrix CM_i using a user-defined mapping of $\{H_{LP} \otimes \phi_{i,g_i}\}$, the output after filtering ϕ_{i,g_i} with the low-pass filter H_{LP} using circular convolution.
   b) Locate the majority pixel in ϕ_{i,g_i} with the highest modified probability (the majority pixel $\phi_{i,g_i}[m, n]$ such that $U_i[m, n] \times CM_i[m, n] > U_i[o, p] \times CM_i[o, p]$ for all $1 \le o \le M$ and $1 \le p \le N$ where $\phi_{i,g_i}[o, p]$ is also a majority pixel), and replace that pixel, $\phi_{i,g_i}[m, n]$, with a minority pixel.
   c) Given the new minority pixel, $\phi_{i,g_i}[m, n]$, adjust the probability of each and every majority pixel, $\phi_{l,g_l}[o, p]$ for $l = 1, 2, \cdots, C$, such that
   $$(U_l[o, p])_{\text{new}} = (U_l[o, p])_{\text{old}} \times \tilde{\mathcal{R}}_{l,i}(r) \qquad (30)$$
   where r is the minimum wrap-around distance from the majority pixel $\phi_{l,g_l}[o, p]$ to the new minority pixel $\phi_{i,g_i}[m, n]$.
3) If for all colors i, the ratio of the total number of minority pixels to the total number of pixels in ϕ_{i,g_i} is equal to g_i, the desired intensity of color i, then quit with the desired color dither pattern given by the set $\{\phi_{i,g_i}: i = 1, 2, \cdots, C\}$; otherwise, continue at step 2b.

Like BIPPCCA, the above algorithm is not suited to the design of multichannel green-noise masks as the stacking constraint will not be satisfied for all patterns. MBIPPCCA must, therefore, satisfy the same constraints as BIPPCCA in order to be used for mask construction. The first of these two constraints, $\phi_{i,g_i} \subset \phi_{i,k_i}$, is satisfied by first initializing ϕ_{i,g_i} to ϕ_{i,k_i} and applying step 2c for each and every minority pixel in ϕ_{i,k_i}. MBIPPCCA can then continue at step 3. The second constraint is satisfied in step 2b of MBIPPCCA when locating the maximum likely majority pixel in color i by considering only those

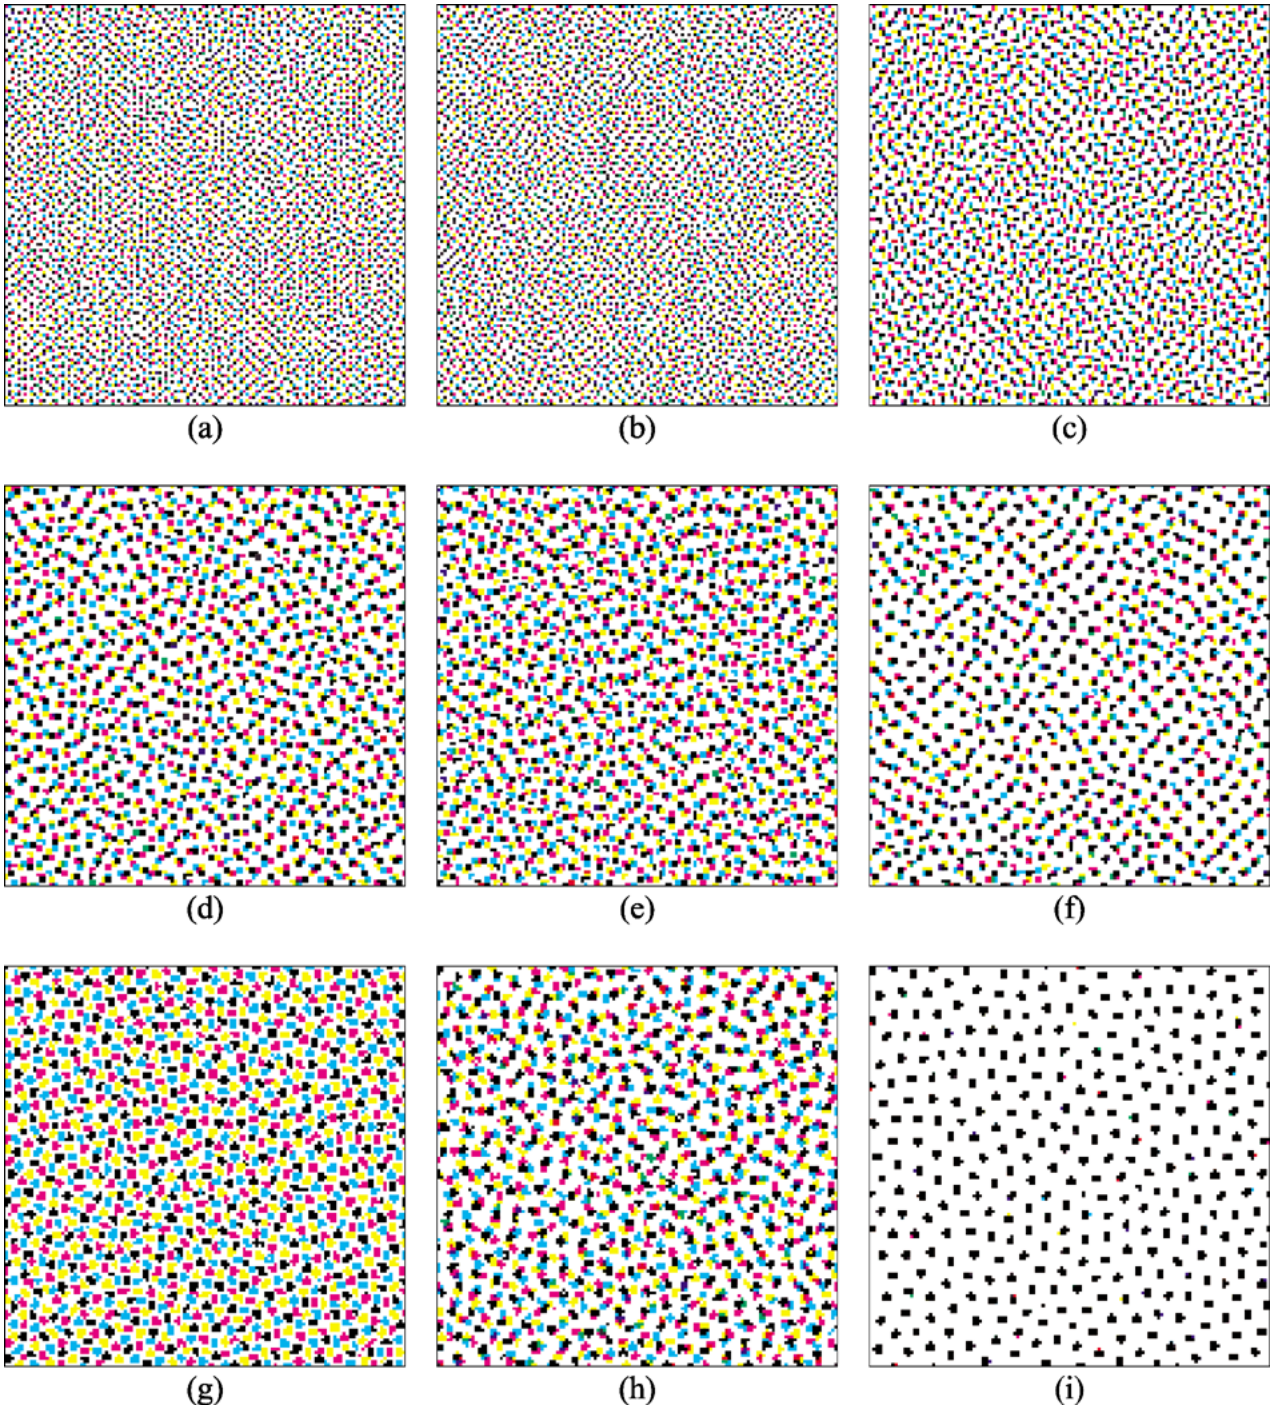

Fig. 12. *Color plate 1.* CMYK dither patterns with: (a) created via Floyd-Steinberg error diffusion, (b)–(f) created via generalized error diffusion, and (g)–(i) created via VBIPPCCA.

majority pixels in ϕ_{i,g_i} that correspond to minority pixels in the constructed patterns, ϕ_{i,l_i}, for which $l > g$. Note that with these constraints, patterns of the multichannel green-noise mask can be constructed in any order, and that the order, to which patterns of any color i are constructed, need not be the same as any other color j.

C. Simulations

Before constructing masks, Fig. 10 shows a set of pair correlation shaping functions where the function $\tilde{\mathcal{R}}_{i,i}(r)$ shapes the pair correlation between pixels of the same color and the function $\tilde{\mathcal{R}}_{i,j}(r)$ shapes the pair correlation between pixels of different colors. This pair of shaping functions is used to reduce the amount of overlap between pixels of different colors [Fig. 12(g)]. Using this same pair of shaping functions, but with $\tilde{\mathcal{R}}_{i,j}(r) = 1$ for all r, patterns with no correlation between channels [Fig. 12(h)] can be constructed. To create a pattern where the overlapping of pixels of different colors is increased, the function $\tilde{\mathcal{R}}_{i,j}(r)$ is set to have the same shape as $\tilde{\mathcal{R}}_{i,i}(r)$ [Fig. 12(i)].

These patterns of Fig. 12(g)–(i), generated by MBIPPCCA, were constructed to represent a 91×91 pixel input image of

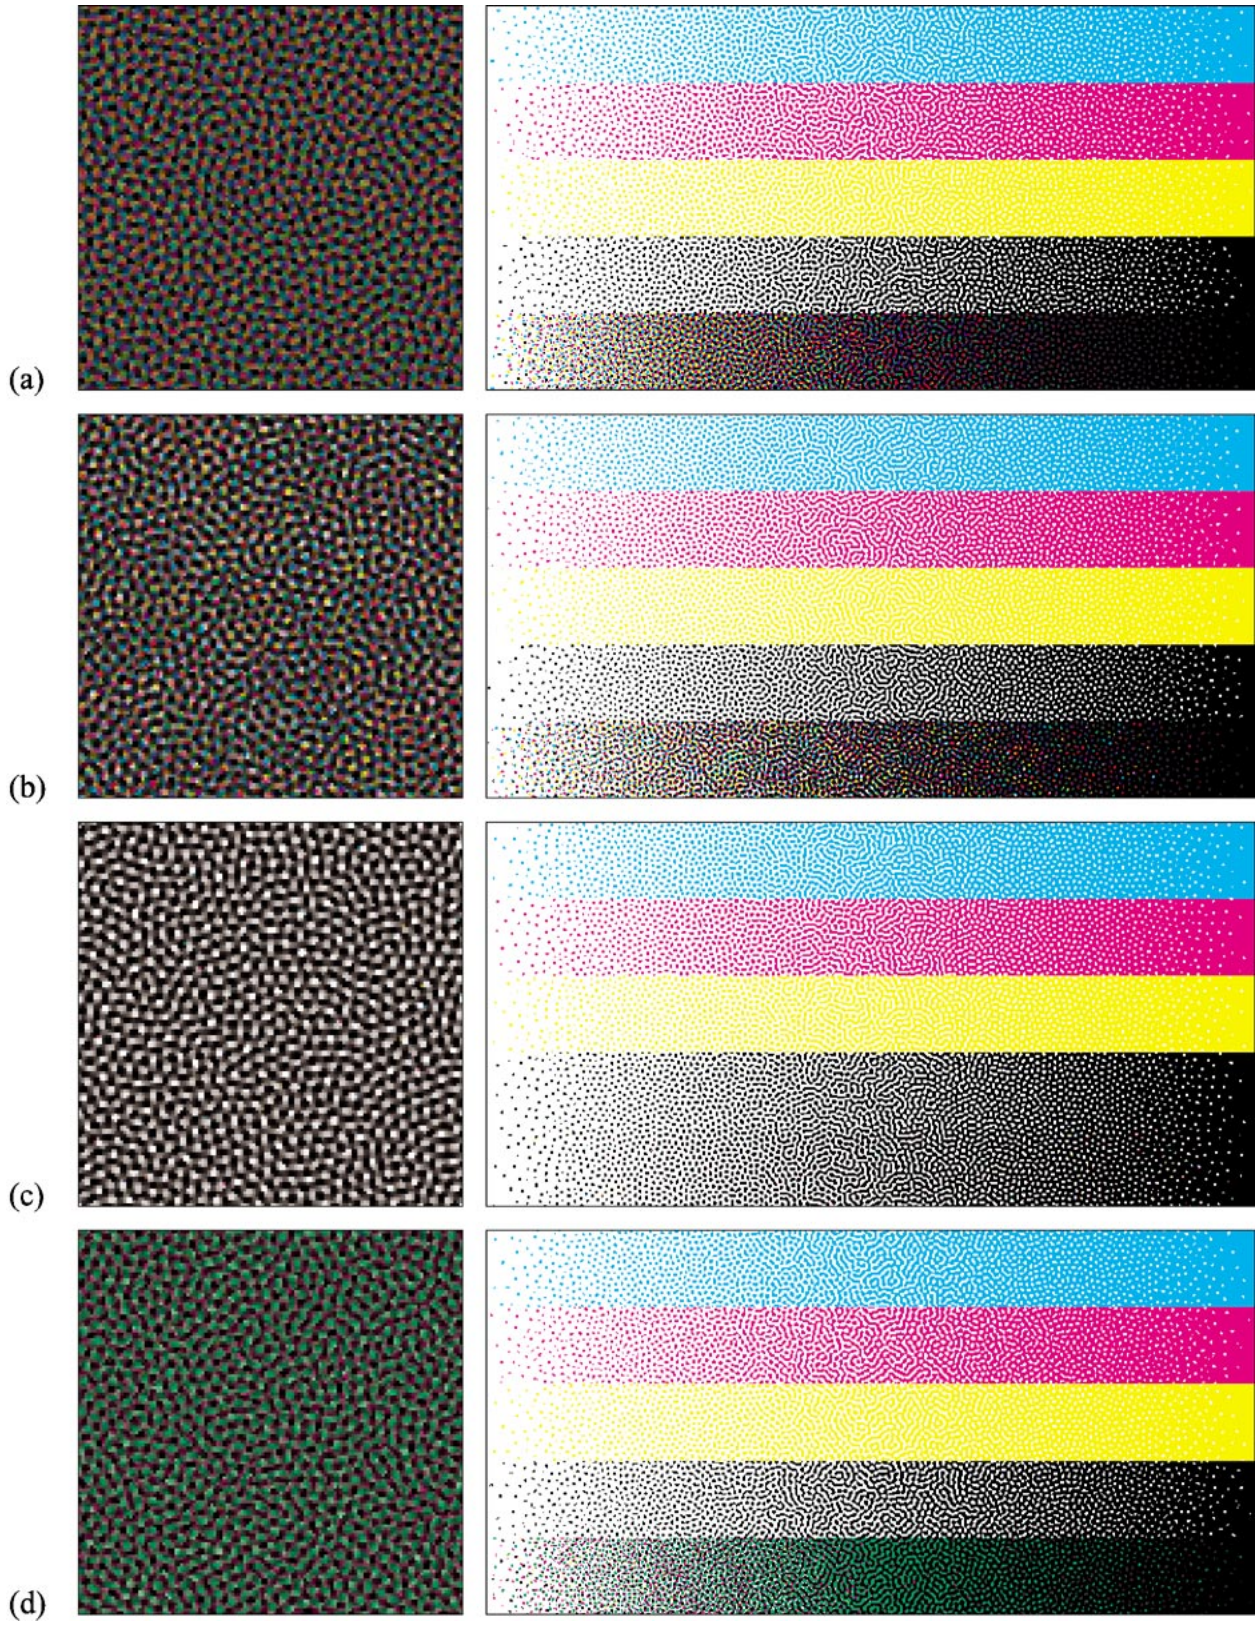

Fig. 13. *Color plate 2.* CMYK green-noise masks constructed from VBIPPCCA such that: (a) minimizes dot overlap, (b) has uncorrelated overlap, (c) maximizes overlap, and (d) maximizes overlap for colors cyan and yellow.

constant color $\vec{x}[n] = [7/8, 7/8, 7/8, 7/8]^{\boldsymbol{T}}$ with an average of 5 pixels per cluster ($\lambda_g = 6.32$ pixels). The statistical measures of the spatial relationships between pixels for these three patterns are shown in Fig. 11. The results, shown here, demonstrate MBIPPCCA's ability to capture the same spatial relationships between pixels as those created in Fig. 12(d)–(f) via generalized error diffusion. The key is in the shaping functions, and through these shaping functions, the same relationships between

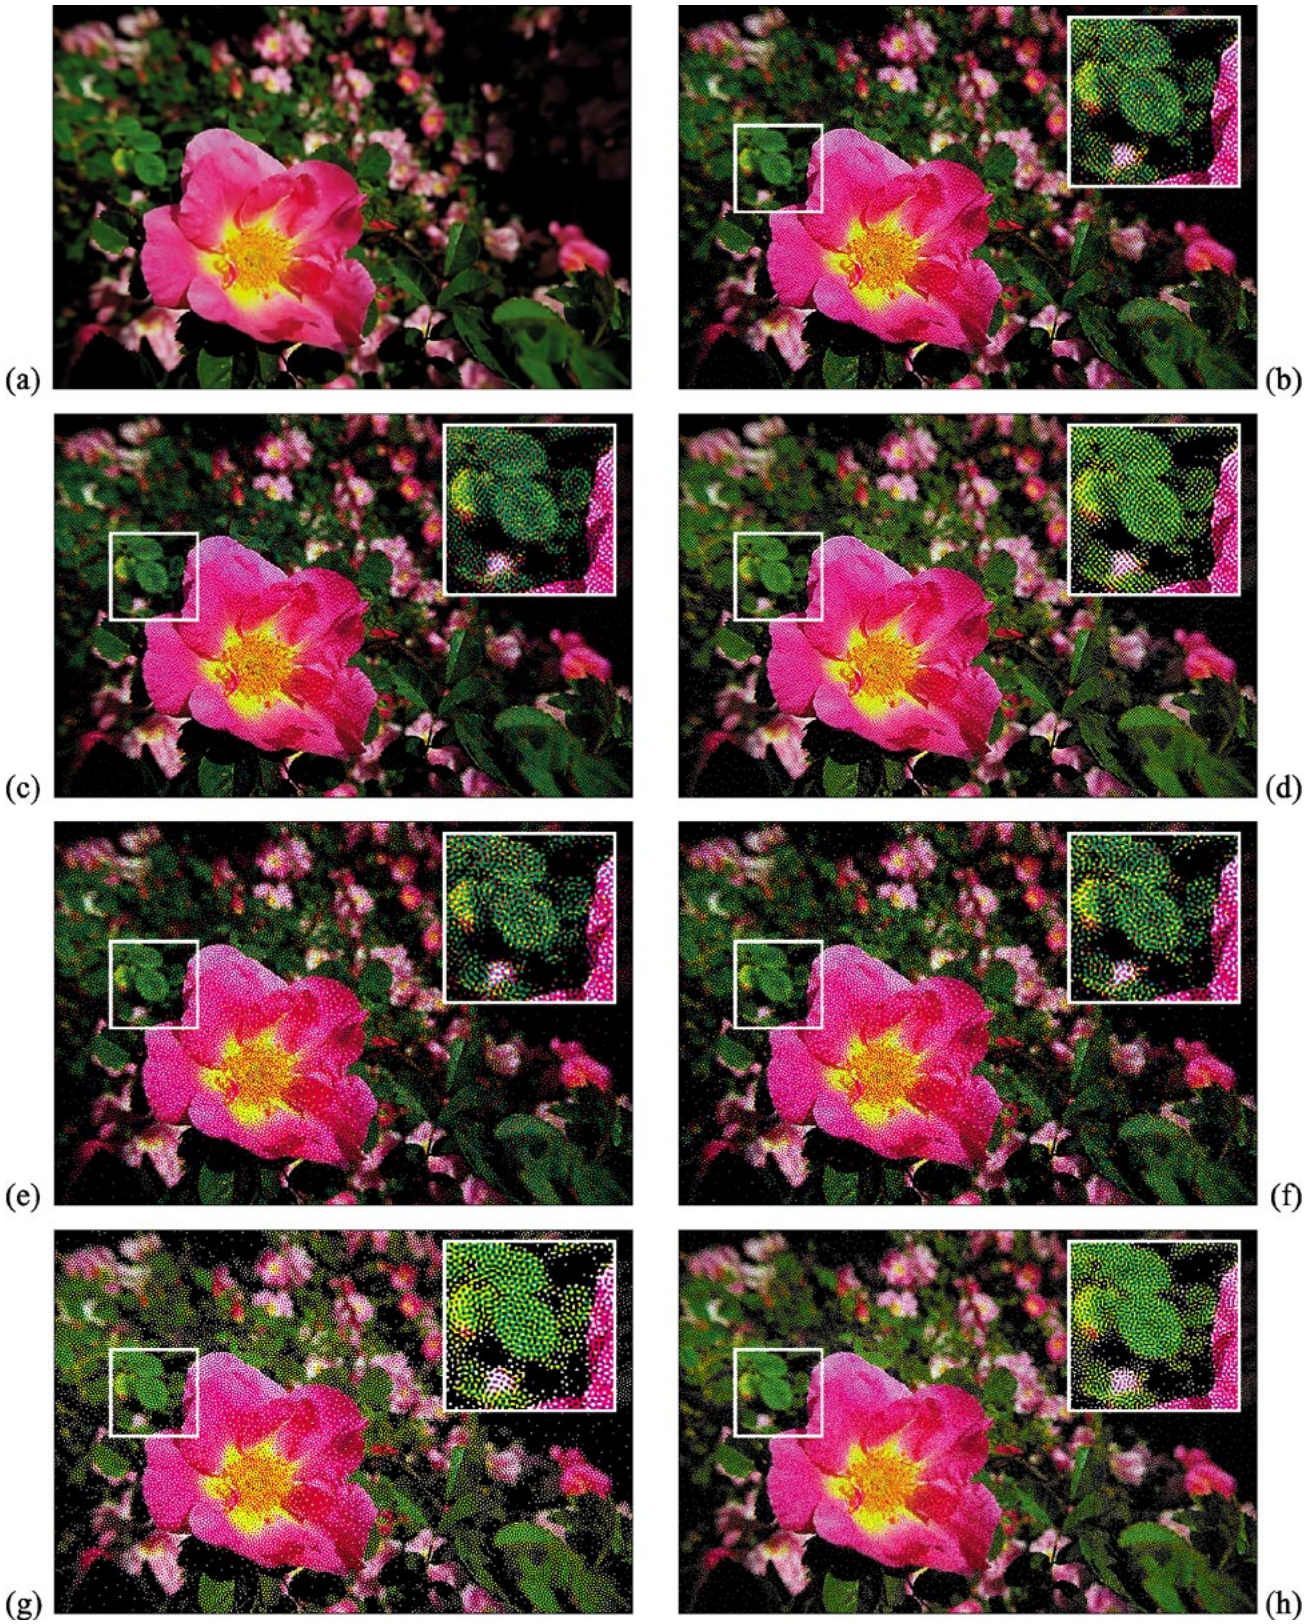

Fig. 14. *Color plate 3.* CMYK halftoned images of the: (a) original using, (b–d) generalized error diffusion, and (e–h) multichannel green-noise masks.

minority pixels can be encouraged in the design of green-noise masks.

Mask design can be seen in Fig. 13 (left) where the three design criteria: 1) decreased; 2) uncorrelated; and 3) increased pixel overlap are employed in masks (a), (b) and (c), respectively. Mask (d) is a special mask designed more to demonstrate the range of possibilities for dither array generation. In this instance, the colors cyan and yellow are designed to overlap

while not overlapping with black or magenta. The colors black and magenta are uncorrelated with respect to each other. The CMYK-color scales shown in Fig. 13 (right) are given to further illustrate the clustering behavior of each mask. By design, each mask has an average cluster size of 2 pixels at extreme gray levels ($g = 0, 1$) and an average cluster size of 12 pixels at $g = 1/2$.

V. Conclusions

In summarizing this paper, it is important to note that this paper does not present a process of optimal color reproduction but instead offers two new techniques for getting there. Future work will look at optimizing the parameters of green-noise for specific output devices as, again, the advantage to using green-noise is that it is tunable—allowing for various cluster sizes for various dot-gain characteristics. Previously, techniques such as error diffusion with output-dependent feedback and the green-noise mask could only be optimized or tuned within a given channel/color. Now both can consider the interactions of the component colors.

Noting Fig. 14(a) where the continuous tone CMYK image flowers is shown with its corresponding halftone reproductions, the generalized error diffusion scheme gives its best reproduction in Fig. 14(d) where the amount of overlap is increased relative to the uncorrelated overlap of Fig. 14(c) and the decreased overlap of Fig. 14(b). In these three instances, the configurations of parameters (A, B, H, and S) are exactly the same as those of Section III-C, Fig. 7.

In Fig. 14(e)–(g), the same comparison of overlap is made using the multichannel green-noise masks of Fig. 13(a)–(c), respectively, where increased overlap gives the best color reproduction. Although these patterns appear "grainy" relative to their error diffused counterparts, this shortcoming is not a function of masks in general but is a shortcoming of the design criteria used in the construction of these specific masks. That is, these masks are composed of clusters which are too large for the quality of printer being used (Tektronix Phaser 440 dye-sublimation)—resulting in halftone patterns with visually disturbing artifacts such as the annoying cyan clusters in the predominantly magenta flower pedals.

A mask such as that (not pictured) used in Fig. 14(h) is much better suited to this printer as the spatial relationship between minority pixels is closer to blue-noise; furthermore, this mask also takes into account the improved color reproduction achieved by increasing the overlap of minority pixels of different colors. As this mask makes good use of the printers ability to print individual pixels, it is a clear example of the tunability of green-noise for color halftoning.

References

[1] J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Trans. Syst., Man, Cybern.*, vol. 21, pp. 33–38, Jan./Feb. 1991.

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

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

[4] T. N. Pappas and D. L. Neuhoff, "Printer models and error diffusion," *IEEE Trans. Image Processing*, vol. 4, pp. 66–79, Jan. 1995.

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

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

[7] 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.

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

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

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

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

[12] P. J. Diggle, *Statistical Analysis of Spatial Point Patterns*. London, U.K.: Academic, 1983.

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

[14] R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial gray-scale," in *Proc. Society Information Display*, vol. 17, 1976, pp. 75–78.

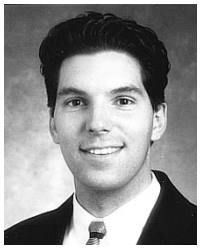

Daniel L. Lau received the B.S.E.E. degree with highest distinction from Purdue University, West Lafayette, IN, in May 1995, and the the Ph.D. degree in electrical engineering from the University of Delaware, Newark, in May 1999.

He spent a year signal and image processing at the Lawrence Livermore National Laboratory, Livermore, CA. His research interests include image processing, digital halftoning, multimedia, and nonlinear filters. He is especially interested in applying digital image processing techniques to applications in art conservation. He has written several papers on nonlinear filters for signal processing, and has consulted with industry on halftoning and digital printing.

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, University of Delaware, Newark, where he is currently a Professor and Associate Chair. He frequently serves as a Consultant to industry and government in the areas of digital printing, image processing, communications, and tomography. He has 15 years of research experience in the areas of halftoning and digital printing. He has published over a dozen journal and conference papers in the area of halftoning. His other research interests include robust signal processing and its applications, communication theory, image processing, and secure multimedia communications.

Neal C. Gallagher (S'72–M'75–SM'85–F'87) is the Dean of Engineering at Colorado State University. Over the years, he has published in the areas of stochastic processes, quantization and source coding, electromagnetic theory and microwave components, and optics. He has held various editorial and committee positions in the IEEE, OSA, and SPIE. He is currently the Chair of the Information Processing, Data Storage, and Holography Technical Group of the OSA's Technical Committee and thereby a member of its executive committee.

Dr. Gallagher is a Fellow of the IEEE for his work in nonlinear digital signal processing and a Fellow of the Optical Society of America for his work on diffractive optics.