Minimizing Stochastic Moiré Using Green-Noise Halftoning

Daniel L. Lau▲

Department of Electrical and Computer Engineering, University of Kentucky, Lexington Kentucky, USA

In amplitude modulated halftoning, Moiré refers to the low frequency textures created by superimposing the monochrome halftones of cyan, magenta, yellow, and black inks. Recently, a stochastic Moiré phenomenon has been discovered that accounts for the low frequency graininess created by superimposing frequency modulated (FM) halftones. A particularly interesting property of stochastic Moiré is that it is most visible when overlapping FM patterns are uncorrelated and are of equal intensity. So in a printer that cannot guarantee perfect alignment in the component screens, the only way to minimize the visibility of Moiré is to introduce clustering such that two overlapping dither patterns of equal intensity have different principal frequencies. In this article, we introduce an adaptive error diffusion halftoning algorithm that produces green-noise dither patterns with coarseness varying according to color content.

Journal of Imaging Science and Technology 47: 327-338 (2003)

Introduction

In the domain of monochrome imaging, digital halftoning refers to the process of converting a continuous tone image into a binary image for display on intensity challenged devices such as desktop ink jet and laser printers.[1] Based upon the resulting distribution of black and white pixels, algorithms are categorized as being either amplitude modulated (AM) or frequency modulated (FM) halftoning where AM halftoning creates the illusion of continuous tone by producing a regular array of printed dot clusters that vary in size according to tone. In this scheme, dark shades of gray are produced by large clusters and light shades by small. FM halftoning algorithms homogeneously distribute printed pixels in a random or stochastic fashion such that the spacing between isolated dots varies with tone with dark shades of gray produced by a tight packing of dots and light shades by a loose packing.

Because the isolated dots of FM halftoning are harder for a human observer to see than the clusters of AM halftoning and because the human visual system (HVS) is less sensitive to the textures of the random patterns produced by FM halftoning than to the regular grid textures of AM halftoning, FM halftoning creates the illusion of continuous tone with an apparent spatial resolution higher than that of AM halftoning. So FM prints appear to have better edge detail than AM, garnering the name, "high fidelity printing." But due to the fine dot structure, FM halftones exhibit much higher degrees of distortion in the printed output due to variations in the printed dot and have, therefore, been restricted to reliable printing devices such as ink jet printers.[2]

In electrophotographic or laser printers, isolated dots are printed with a high degree of variability from dot to dot and only reproduce reliably when minority pixels occur in clusters. These printing devices have, therefore, been restricted to AM halftoning, but in order to take advantage of the HVS's lack of sensitivity to stochastic dither patterns, a new category of AM-FM hybrid halftoning algorithms have become of interest. These produce the illusion of varying shades of gray by a pattern composed of randomly distributed, randomly shaped printed pixel clusters. These new halftoning algorithms are generally referred to as green-noise halftoning[3] due to the mid-frequency only content of the resulting patterns, and shown in Fig. 1 is a comparison of the three halftoning techniques on the image pillars.

In color printing, the illusion of continuous shades of chrominance and luminance is produced by superimposing the monochrome halftones of cyan, magenta, yellow, and black inks. For AM halftoning, the superposition of regular grids poses the added problem of Moiré, a visually disturbing low frequency texture created by the superposition of regular patterns that exist in neither of the component patterns.[4] This Moiré phenomenon is unavoidable, and its visibility in color prints can only be minimized. The minimization of Moiré is achieved by offsetting the screen angles of component halftones by 30° and is the chief reason why CMYK screens are set at screen angles of 15°, 75°, 0°, and 45° respectively.[5] While the 15° offset between cyan and yellow and between magenta and yellow would suggest a more highly visible Moiré pattern, the high luminance of yellow dots makes this Moiré difficult to discern on a white background.

Original manuscript received September 23, 2002

▲ IS&T Member

dllau@engr.uky.edu

Color Plate 6 is printed in the color plate section of this issue, pp. 372–376.

Figure 1. The image pillars halftoned using (left) Amplitude Modulated halftoning, (center) green-noise halftoning, and (right) Frequency Modulated halftoning.

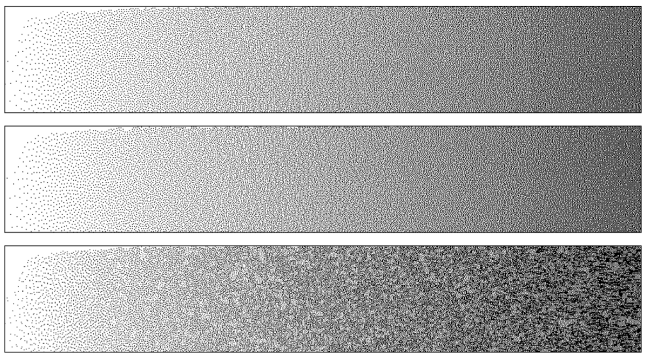

Figure 2. The (bottom) low frequency graininess produced by superimposing (top and middle) two FM halftone patterns.

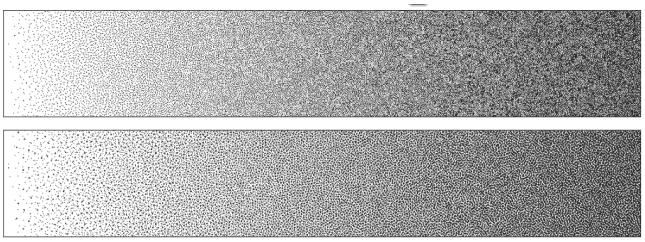

Figure 3. The composite halftone created by (top) two uncorrelated blue-noise masks and (bottom) an uncorrelated combination of a blue-noise and a green-noise mask.

In FM halftoning, the Moiré phenomenon was never thought to be of concern due to the lack of regular patterns, but Lau *et al.*[6] show that, in fact, the low frequency graininess produced by superimposing two uncorrelated FM dither patterns (Fig. 2) is an instance of stochastic Moiré. Prior to Lau *et al.*, the only significant research describing instances of the stochastic Moiré phenomenon was started by Glass[7,8] who studied the visual perception of order created by superimposing a random pattern of dots onto a geometrically distorted version of itself to create what are now generally referred to as *Glass Patterns*. From the visual perception of order created when a binary dot pattern is superimposed onto a rotated version of itself, Glass speculated that early steps of the human visual process may include the computation of local autocorrelations by excitation of line detectors and the subsequent averaging of local autocorrelations by collective excitation of the columns in the visual cortex.[7]

Like periodic Moiré in AM halftoning, stochastic Moiré cannot be eliminated but only have its visibility minimized.[9,10] This minimization is achieved by correlating the component dither patterns to form interlocking screens, whereby the overlapping of printed pixels between two patterns is minimized (dot-off-dot printing) or maximized (dot-on-dot printing). A good example of such a scheme with halftones in a correlated fashion is realized using jointly blue-noise masks as described by Sullivan *et al.*,[11] or even a DBS approach as developed by Lin and Allebach.[12] Because of the potentially disastrous impact misalignment would have on such halftones, the preferred approach to Moiré reduction is to halftone each color in an uncorrelated fashion and then to vary the statistical character of the overlapping screens as much as possible.[9] A possible scheme may involve using error diffusion with a Floyd and Steinberg[13] error filter for one color and a Stucki[14] error filter for another. Or, perhaps, a left to right serpentine raster scan is applied on one color and an up and down serpentine raster scan on the other.

In green-noise halftoning, stochastic Moiré has the same impact as in FM halftoning except that the lower frequency content of green-noise halftones results in lower frequency fluctuations in texture for stochastic Moiré.[15] The HVS being more sensitive to low frequency textures makes the stochastic Moiré of green-noise more noticeable and, therefore, of greater impact on the resulting print[10] (see Fig. 1). Like FM halftoning, stochastic Moiré visibility is minimized by varying the statistical character of overlapping patterns, but unlike FM halftones, the statistical nature of green-noise halftones can be manipulated by varying the coarseness of overlapping patterns such that the average size of clusters and the related spacing between clusters changes without modifying the pattern's intensity. So in green-noise halftoning, stochastic Moiré visibility is minimized by varying the coarseness between overlapping dither patterns.[10] Shown

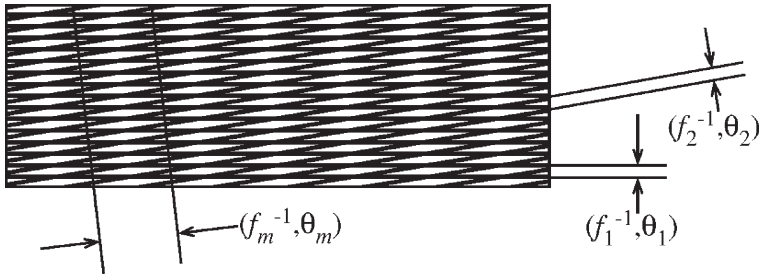

Figure 4. An instance of periodic Moiré where two overlapping line gratings with line frequencies f_1 and f_2 and screen angles θ_1 and θ_2 lead to a Moiré grating with line frequency f_m and screen angle θ_m.

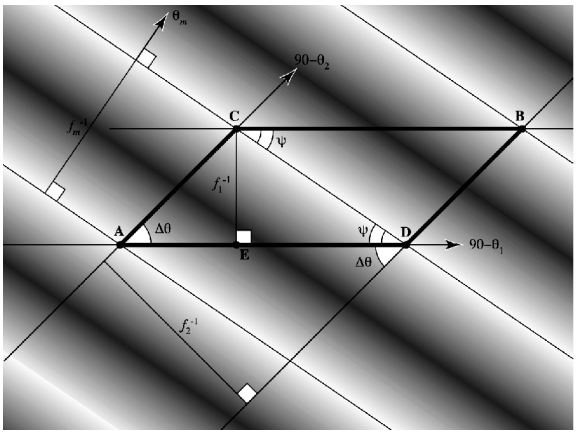

$$\Delta\theta=\theta_2-\theta_1$$
$$|\mathbf{AD}|=f_2^{-1}/\sin(\Delta\theta)$$
$$|\mathbf{AE}|=f_1^{-1}/\tan(\Delta\theta)$$
$$\psi=\tan^{-1}(f_1^{-1}/(|\mathbf{AD}|-|\mathbf{AE}|))$$
$$f_m^{-1}=|\mathbf{AD}|\sin(\psi)$$
$$\theta_m=180-\theta_1-\psi$$

Figure 5. Diagram showing the derivation of the periodic Moiré line frequency, f_m, and screen angle, θ_m.

in Fig. 3 is proof of this concept where Fig. 3 (top) shows the overlapping of two uncorrelated halftone gradients produced using blue-noise masks and pattern Fig. 3 (bottom) shows the overlapping of uncorrelated gradients using a blue-noise mask for one pattern and a green-noise mask for the other.

The drawback of using increased pattern coarseness to minimize stochastic Moiré is that in areas of the digital print where the halftone is composed of dots from a single color, the high coarseness this pattern may have, in the hopes of minimizing stochastic Moiré in other areas, has the undesirable effect of creating a visually apparent texture due to increased visibility of printed dots that now occur in clusters.[3] So in general, minimizing stochastic Moiré by means of green-noise is a balancing act of reducing Moiré visibility at the cost of low frequency graininess introduced by coarse halftones, with the printer manufacturer choosing the optimal ratio for their device.[9] But what would clearly help is a halftoning scheme that adaptively modifies coarseness according to the color composition of the original image at each pixel. In this study, we introduce such a technique by employing Levien's[16] error diffusion with output dependent feedback for two-color halftones, where coarseness is determined according to the difference in tone at each pixel, such that as the difference in tone gets smaller, the difference in coarseness gets larger.

This article has been divided into five sections with the following section reviewing the general theory of periodic Moiré analysis as introduced by Amidror *et al.*[4] Following that, this study reviews the stochastic Moiré analysis introduced by Lau *et al.*,[9,10] but unlike those earlier works, this study also demonstrates the proposed stochastic Moiré analysis applied to AM screens. From stochastic Moiré analysis, this study reviews the green-noise halftoning model and shows how Levien's error diffusion with output dependent feedback[16] can be modified for two-color halftones with an adaptive hysteresis parameter for minimizing stochastic Moiré visibility. In Ref. 10, minimizing stochastic Moiré visibility is done using static blue and green-noise masks where pattern coarseness is fixed according to the mask. In concluding, this article discusses some of the results of several visual perception studies involving Glass Patterns and how these results may impact Moiré analysis when taking color into account.

Periodic Moiré Analysis

The periodic Moiré phenomenon in color AM halftoning has been studied in great detail by Amidror *et al.*[4] who, in an effort to develop a spectral model, looks initially at the Moiré patterns produced by overlapping two line gratings as depicted in Fig. 4. In the spatial domain, Moiré is the repeating pattern of light and dark regions where light regions are areas where the black lines of the two gratings overlap and the dark regions are the areas where the black lines of the two gratings are perfectly non-overlapping. From the geometric analysis of Fig. 5, the overlapping of two line gratings, with line frequencies of f_1 and f_2 and screen angles of θ_1 and θ_2,

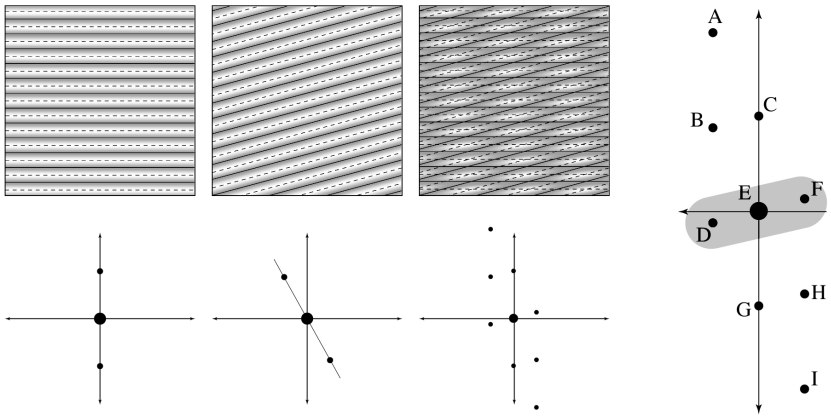

Figure 6. Amidror *et al.'s* periodic Moiré analysis using cosinusoidal gratings in place of binary halftone patterns with the impulses labeled **D** and **F** corresponding to the Moiré line grating with frequency f_m and screen angle θ_m. The peaks and valleys of the two cosinusoidal gratings are indicated by dashed and solid lines respectively. Portions of this figure originally appeared in Ref. 10 and are used with permission from the Optical Society of America.

results in a Moiré line grating with line frequency f_m and screen angle θ_m such that:

$$f_m^{-1} = |\mathbf{AD}|\sin(\psi), \text{ and} \tag{1}$$

$$\theta_m = 180° - \theta_1 - \psi \tag{2}$$

In Amidror's spectral analysis, the binary line gratings are replaced with cosinusoidal gradients that range in value from zero to one, such that their Fourier representations are each composed of three impulses arranged at the screen angle of the gradient (Fig. 6). Amidror simplifies the problem further by treating the superposition of gratings as a multiplicative relationship with the resulting surface's Fourier representation acquired by convolution of the component patterns, leading to the nine impulses of Fig. 6, labeled **A** through **I**. The four impulses (**A**, **D**, **F**, and **I**) of this resulting surface that do not directly correspond to the impulses of either component gradient (**B**, **C**, **E**, **G**, and **H**) are the Moiré artifacts with the impulses labeled **D** and **F** corresponding to the Moiré line grating with frequency f_m and screen angle θ_m, as derived by geometric analysis.

Noting the HVS's increased sensitivity to low frequency artifacts, the visibility of Moiré is decreased by moving the Moiré impulses as far from DC as possible, resulting in a screen angle offset of 90°. In the special instance where $f_1 = f_2$, it is possible to align the line gratings such that lines perfectly overlap everywhere. Such an alignment would result in a Moiré line grating with frequency $f_m = 0$. So in this special case, Moiré visibility can also be minimized by moving the Moiré impulses to DC. The problem that this line-on-line scheme has is that with even a slight misalignment of either screen angle, the Moiré impulses would move off of DC and result in very visible artifacts. Amidror refers to this approach to Moiré visibility minimization as being *unstable* due to the wildly varying visibility of Moiré derived from a small change in screen orientation.

For two dimensional line gratings, the Fourier representations of the overlapping patterns are composed of nine impulses each with convolution leading to the resulting superimposed surface. Again, any impulses not corresponding directly to an impulse of the component gratings are Moiré components, and minimizing Moiré visibility is achieved by choosing a screen angle offset that maximizes the minimum distance between the Moiré impulses and the spectral origin (DC). Assuming that the horizontal and vertical line frequencies are equal between the two component patterns, Moiré visibility can also be minimized by perfectly aligning the patterns to produce the familiar dot-on-dot or dot-off-dot pattern.[17]

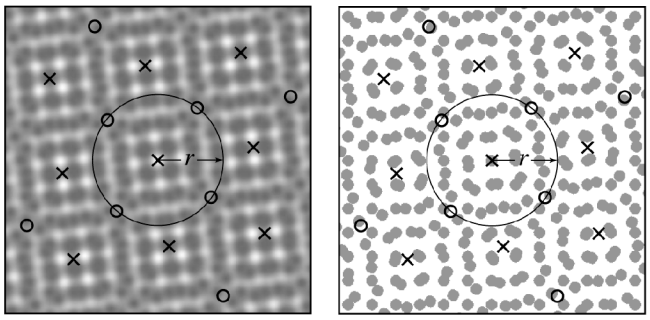

Figure 7. The in and out-of-phase points for 2-D gratings and for corresponding AM halftone pattern. This figure originally appeared in Ref. 10 and is used with permission from the Optical Society of America.

Returning to the spatial domain, the repeating pattern of light and dark regions produced by superimposing cosinusoidal gradients can be described in terms of phase, with Moiré being a repeating pattern of in phase and 180° out of phase points.[6] In Fig. 6, the locations where the two gradients are in phase are the points where the two peaks (solid lines) intersect or where two valleys (dashed lines) intersect. Out of phase points are the locations where a peak of one pattern intersects a valley of the other. The distance, r, between these in phase and out of phase points between two cosinusoidal gradients is defined as:

$$r = \frac{\lambda}{2\cos(\theta)}, \tag{3}$$

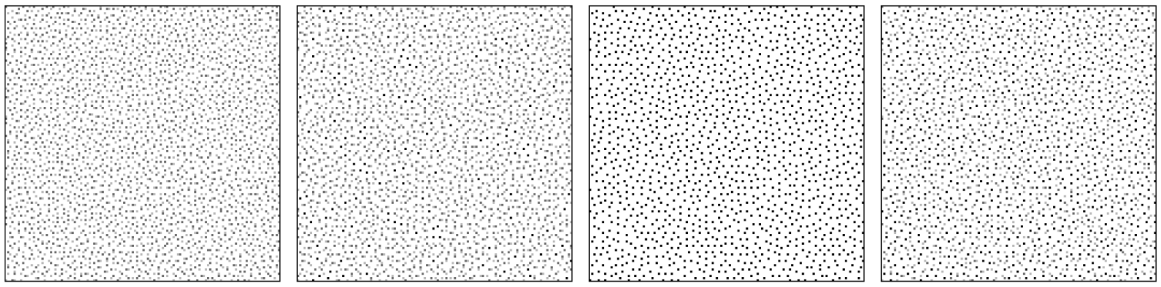

Figure 8. Four instances of two-color FM halftones using 10% cyan and 10% magenta.

where θ represents the relative screen angle between the two screens and λ is the inverse of the screen frequency. Note the similarity with Eq. (1). Using 2-D gratings, Fig. 7 shows the locations of the in (X's) and out of phase (O's) points.

As before, there are two approaches to minimizing the visibility of Moiré. The first is that given the low pass nature of the HVS, minimizing r ($\theta = 90°$) would create a pattern that maximizes the frequency of the fluctuations in phase alignment. Maximizing these fluctuations moves the spectral content of the Moiré pattern to where the HVS is least sensitive, creating an apparent image composed exclusively of a DC color/texture. The second approach to Moiré minimization involves maximizing r ($\theta = 90°$) such that the Moiré pattern fluctuates too slowly per unit length to be noticed by the HVS. In this case, Moiré is purely DC with even slight misalignment of a screen leading to highly visible Moiré artifacts.[10]

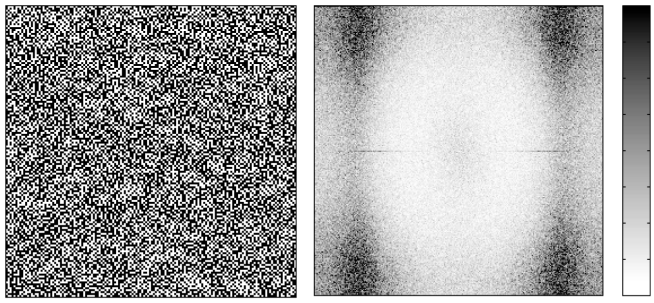

Figure 9. The dither pattern created by superimposing two FM dither patterns ($g_a = g_b = 1/3$) along with its corresponding power spectrum with the specrtral origin (DC point) located at the center of the plot.

Stochastic Moiré Analysis

For a concise definition of stochastic Moiré, Fig. 8 shows black-and-white versions of several stochastic halftone patterns with each composed of 10% cyan coverage and 10% magenta coverage. Pattern **A** shows the instance where minority pixels of cyan and magenta are homogeneously distributed such that magenta pixels are placed midway between cyan pixels. Shown in pattern **C** is the case where magenta pixels are placed directly on top of all cyan pixels. While it may be argued by a given observer that pattern **A** may create a hue different from that of pattern **C**, we ignore those differences and focus on the fact that pattern **A** certainly portrays a different texture from that of **C**. Any differences in hue are assumed to be correctable by manipulating the percentage of coverage of the various inks.[9] Furthermore while **A** and **C** have different textures, the stochastic Moiré model makes no determination as to whether **A** or **C** is a *better* texture.

Now in pattern **B**, the magenta and cyan dither patterns are completely uncorrelated with some minority pixels (of both patterns) overlapping and some falling directly in between those of the other pattern. All other minority pixels fall somewhere in between these two extremes. Stochastic Moiré refers to the change in texture that occurs from one point to another within B. Like periodic Moiré, the optimal stochastic halftone is the one that either minimizes or maximizes the amount of fluctuation in this texture per unit area.[10] In the case of pattern **A** and **C**, the amount of fluctuation is minimized with patterns **A** and **C** offering equivalent optimality regarding Moiré. Pattern **D** offers the other Moiré extreme by maximizing the amount of fluctuation per unit area.

Looking at the Fourier representation of stochastic Moiré in Fig. 9, as suggested by Amidror, we see that little information is revealed in the spectral domain regarding Moiré visibility, and we, therefore, need to derive some type of spatial analysis. Relying on phase information, we consider two stochastic dither patterns to be in phase when the minority pixels of one pattern directly overlap minority pixels of the other. They are considered perfectly out of phase when minority pixels of one pattern fall directly in between the minority pixels of the other. For minority pixels that correspond to neither of these extremes, phase can be measured as being proportional to the Euclidean distance from a minority pixel of one pattern to the nearest minority pixel of the other.[9] In this framework, characterizing stochastic Moiré becomes a problem of characterizing the change in phase moving between minority pixels within one of the two patterns.[6]

Formally, Lau[9] states that the dither patterns $\phi_A = \{a_i : i = 1, 2, \ldots, N_A\}$ and $\phi_B = \{b_j : j = 1, 2, \ldots, N_B\}$ define a discrete-space, 2-D function $D[n]$ such that:

$$D[n] = \frac{1}{\lambda_b} \sum_{i=1}^{N_A} d_i \, \delta[n - a_1], \tag{4}$$

where

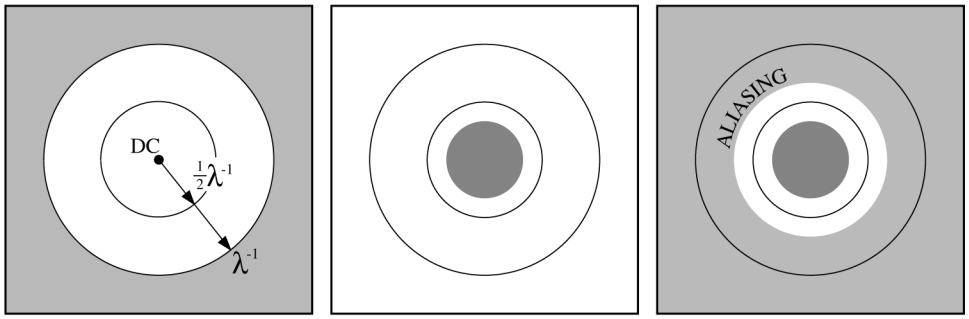

(a) Poisson point process (b) stochastic moiré surface (c) sampled surface without aliasing

Figure 10. Diagram of the spectral domain representations of (a) the Poisson point process defining the sampling grid; (b) the circular band-limited stochastic Moiré surface with bandwidth less than half the principal frequency; and (c) the power spectrum of the sampled surface with no aliasing. This figure originally appeared in Ref. 10 and is used with permission from the Optical Society of America.

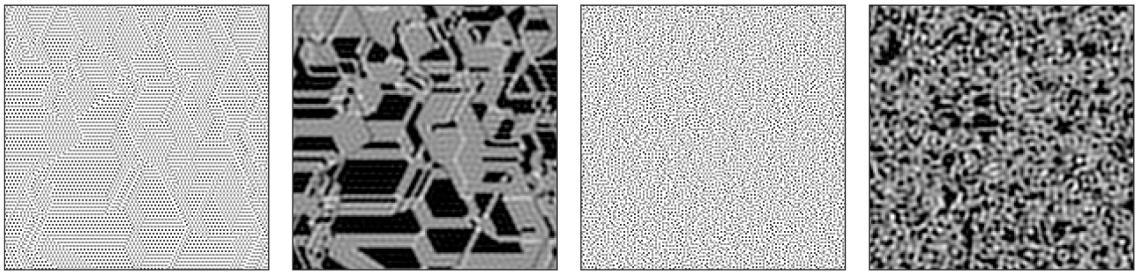

Figure 11. Two examples of stochastic Moiré and the corresponding Moiré surfaces produced by superimposing two FM halftone patterns.

$$d_i = \min_j \left\| a_i - b_j \right\|_2. \tag{5}$$

In its Eq. (4) form, $D[n]$ represents a continuous space signal $S(x)$, an image referred to as the *stochastic Moiré surface*, that has been sampled on a stochastic sampling grid defined by ϕ_A. If the sampling grid has a Poisson distribution, as would be the case of a blue-noise dither pattern with principle wavelength λ_a, we can obtain $S(x)$ from $D[n]$ using an ideal low pass filter with cutoff frequency $0.5\lambda^{-1}{}_a$ and magnitude $\lambda^2{}_a$, Fig. 10. The magnitude of the low pass filter is derived from Papoulis[18] who shows that if $x[n]$ is the sampled signal of $x(t)$ defined by a Poisson sampling process with average density g, then the sum

$$\tilde{X}(f) = \frac{1}{g} \sum_n x[n] e^{-j2\pi fn} \tag{6}$$

is an unbiased estimate of $X(f)$, the Fourier transform of $x(t)$, given that the aliases are nulled out. Hence the Fourier transform of $D[n]$ is scaled by $1/g$ or equivalently by $\lambda^2{}_a$ to obtain an unbiased estimate of $S(f)$.

Shown in Fig. 11 are the stochastic Moiré surfaces corresponding to two instances of two-color FM halftoning with characterizing stochastic Moiré now reduced to the problem of characterizing the spatial fluctuations in the continuous space, 2-D surface $S(x)$. How we measure the spatial fluctuations in these new monochrome images can be done by a weighted average of the spectral power in $S(x)$ such that spectral components, corresponding to frequencies where the HVS is particularly sensitive, weighted the highest. Lau[9] suggests measuring the visual cost of the stochastic Moiré surface as:

$$VC\left(\phi_A, \phi_B\right) = \text{Average AC Power}\left(S(f) \times HVS(f)\right) \tag{7}$$

where $HVS(f)$ is the spectral, low pass filter model of the human visual system used by Sullivan *et al.*[11] Assuming a 300 dpi printer and 20 inch viewing distance, the stochastic Moiré surfaces of Fig. 11 have visual cost measures of (left) 0.0484 and (right) 0.0311, suggesting an improvement in color rendition going from (left) to (right), a result with which visual inspection would certainly agree.

This new stochastic Moiré analysis can be applied to the periodic Moiré found in AM halftoning where every cluster is replaced with a single minority pixel at the cluster's centroid. Assume that the reference pattern that defines the sampling grid will have a screen angle of 45° and a periodicity of 8 pixels on a 1200 dpi halftone (150 lines per inch); the low pass filter will have a cutoff frequency of $\pi/8$ radians (75 Hz). Shown in Fig. 12 are the surfaces corresponding to the periodic Moiré created by screen angle offsets of 5°, 15°, 30°, and 45°. For a measure of the visual cost of each surface, Fig. 12 also shows the same surfaces but after applying the chosen HVS model. The average AC powers of these surfaces are 0.0220, 0.0134, 0.0063, and 0.0193 respectively. The higher visual cost for 45° over 15° derives mainly from the fact that while the amplitude of fluctuations seems larger for 15° in Fig. 12, their diagonal orientation results in less visual cost compared to the horizontal and vertical components for a 45° screen angle offset.

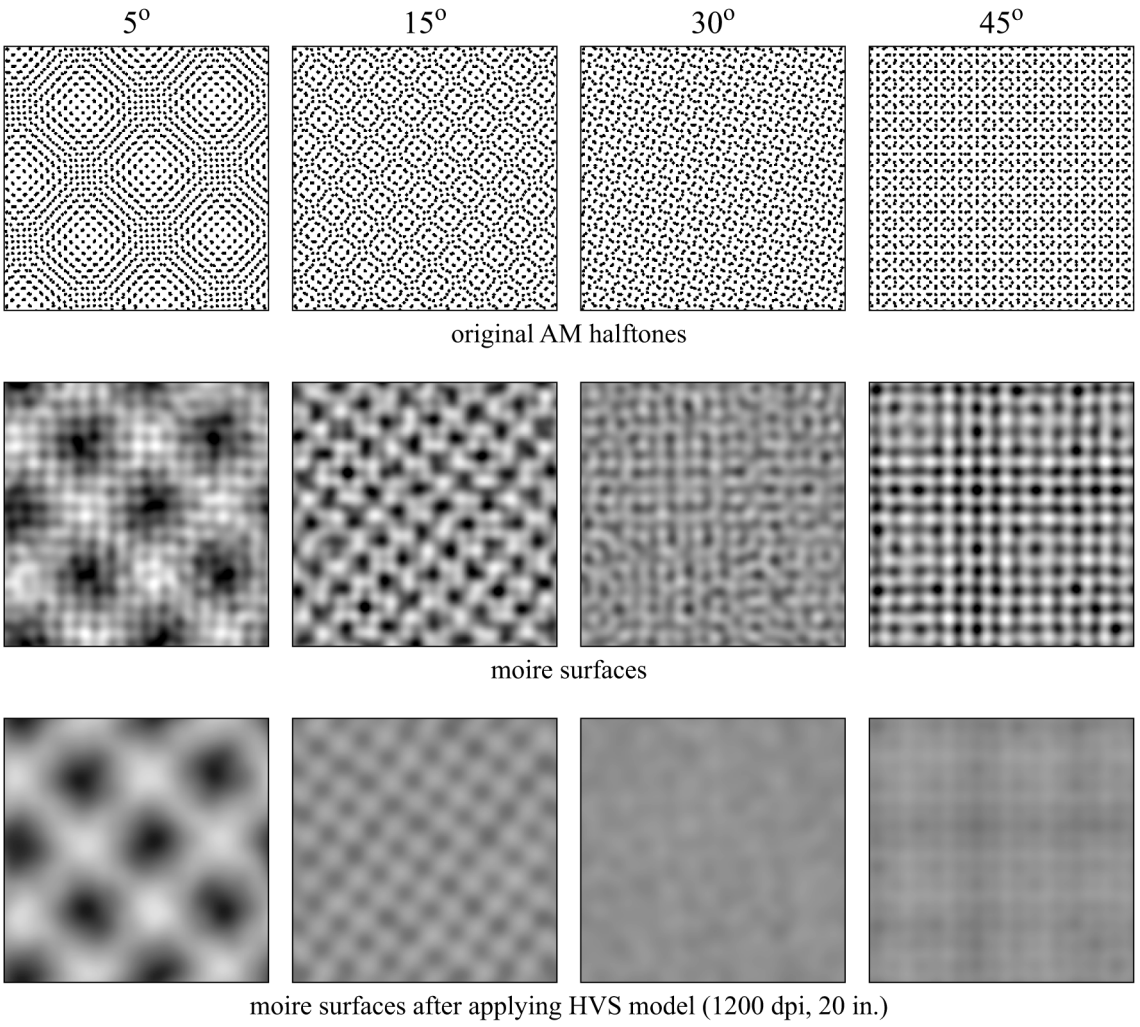

Figure 12. The Moiré surfaces and their corresponding modeled visual responses for 1200 dpi, 150 lpi AM halftoning with screens angles 5°, 15°, 30°, and 45°.

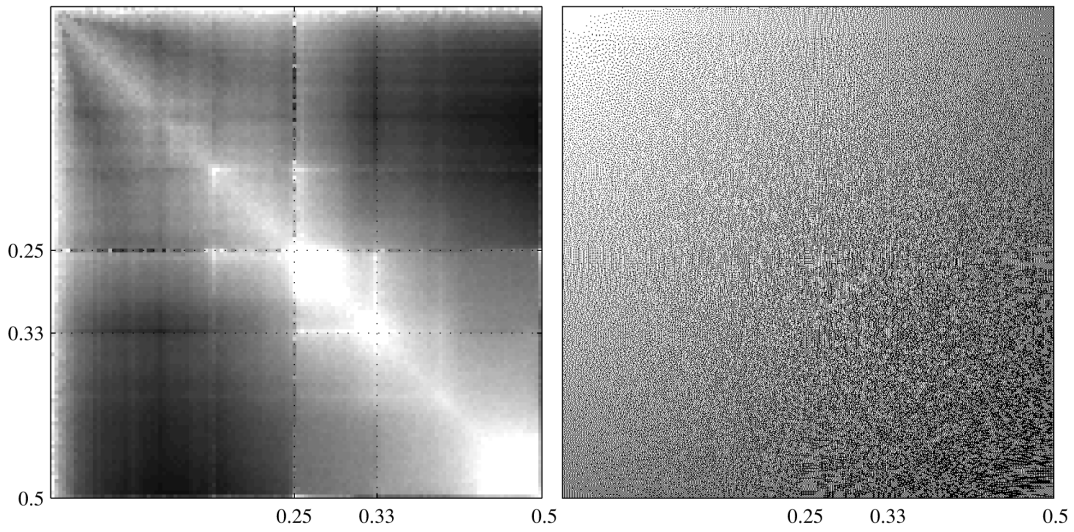

Figure 13. The 2D plot (left) of stochastic Moiré visibility (black = 0, white = 0.035) for g_a and $g_b \in (0, 1/2]$, for Floyd and Steinberg's error diffusion, and the corresponding dither pattern (right) produced by combining two linear gradient halftones, arranged perpendicular by using a pixel-wise logical OR operation (black = 1, white = 0). Portions of this figure originally appeared in Ref. 10 and are used with permission from the Optical Society of America.

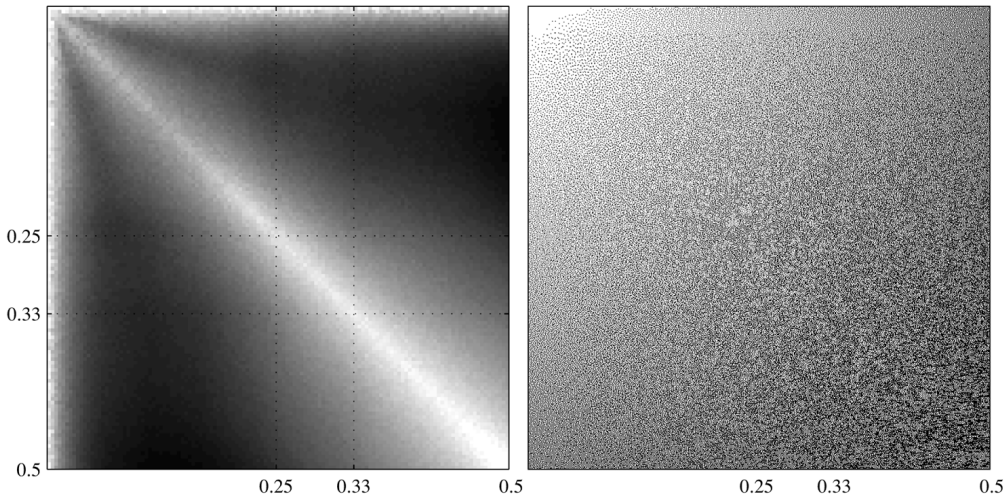

Figure 14. The 2D plot (left) of stochastic Moiré visibility (black = 0, white = 0.035) for g_a and $g_b \in (0, 1/2]$, for Ulichney's perturbed filter weight error diffusion, and the corresponding dither pattern (right) produced by combining two linear gradient halftones, arranged perpendicular by using a pixel-wise logical OR operation (black = 1, white = 0). Portions of this figure originally appeared in Ref. [10] and are used with permission from the Optical Society of America.

Moving to error diffusion, we can analyze this adaptive halftoning technique by looking at the dither patterns produced by superimposing two uncorrelated halftones, such that the resulting pattern is produced using a pixel-wise logical OR function (black = 1) as demonstrated in Figs. 13 and 14.[10] Here in Figs. 13 and 14 (right), two halftone gradients ranging from 0 to 0.5 are shown superimposed with perpendicular orientations. Corresponding to these dither patterns are the visual costs of stochastic Moiré, assuming a 300 dpi printer and 20 inch viewing distance, shown in Figs. 13 and 14 (left) for all intensity combinations of g_a and $g_b \in [0, 0.5]$.

In Fig. 13, halftoning is done using error diffusion with Floyd and Steinberg's original four weight error filter on a serpentine raster scan. What we see, in Fig. 13, is that while Floyd and Steinberg selected their filter weights to produce the classic checkerboard $g = 1/2$ pattern, they were not concerned about the noticeable artifacts that would appear in two-color halftones as Floyd and Steinberg's error diffusion suffers from correlated artifacts at gray levels 1/4, 1/3, and 1/2 that lead to large patches of fixed phase. These large patches lead to large low frequency components in the power spectra of the Moiré surfaces and, therefore, large visibility measures. Looking at both the visibility measure and the superimposed dither pattern, we see a direct and obvious correlation where the strong visibility component at $g_a = g_b = 1/4$ identifies the fluctuations in texture along the $g_a = 1/4$ and the $g_b = 1/4$ lines in the superimposed dither pattern as being so uncharacteristic of the rest of the pattern that visual observation shows an apparent "cross" through the center of the image.

In Fig. 14, halftoning is done using error diffusion with Ulichney's perturbed filter weights on a serpentine raster scan. Through the perturbation of error filter coeffcients, Ulichney's error diffusion breaks up periodic textures in the dither pattern, creating a far more pleasant blue-noise halftone.[1] What we see in Ulichney's error diffusion is a halftoning algorithm with consistent behavior across all gray levels and without the outliers or spikes that exist in the visibility measure of Floyd and Steinberg's error diffusion. The lack of any wildly varying or abrupt changes in visual cost means that smooth color gradients will show a more consistent behavior in texture.[10] We see this clearly in the superimposed dither pattern where while Floyd and Steinberg have problems maintaining a smooth appearance with the large visibility measure along the $g_a = 1/4$ and $g_b = 1/4$ lines clearly apparent, Ulichney's is a halftone with no dramatic shifts in texture with the same visible impression of noise across the entire pattern.

From visual inspection of the stochastic Moiré visibility plots of Figs. 13 and 14, an obvious feature is the increased visibility along the 1:1 ratio line with this increase along the $g_a = g_b$ line a natural property of bluenoise. This can be seen if we note that bluenoise halftoning produces a Poisson point process and that maximizing spatial fluctuations in texture is accomplished by breaking up clusters of in phase or clusters of out of phase pixels in the halftone. Looking at the 1-D case given that $\phi_A = \{a_m : m = 1, 2, \ldots\}$ and $\phi_B = \{b_n; n = 1, 2, \ldots\}$ represent two Poisson point processes with an average separation between consecutive points of λ_A and λ_B, suppose that $a_i \in \phi_A$ and $b_j \in \phi_B$ are such that $a_i = b_j$, indicating that ϕ_A and ϕ_B are in phase at a_i. In order to minimize the fluctuation in texture between ϕ_A and ϕ_B near a_i, we need to maximize the probability that a_{i+1} overlaps b_{j+1}, such that the texture remains constant in the region $[a_i, a_{i+1}]$, where:

$$\mathbf{Pr}\left(a_{i+1} = b_{j+1}\right) = \int_{x=a_i}^{+\infty} \mathbf{Pr}\left(a_{i+1} = b_{j+1} \mid b_{j+1} = x\right) \cdot \mathbf{Pr}\left(b_{j+1} = x\right) dx \qquad (8)$$

For a fixed λ_B, the probability that a_{i+1} overlaps b_{j+1} is maximized when $\lambda_A = \lambda_B$ for uncorrelated ϕ_A and ϕ_B. Since λ_A and λ_B are functions of their corresponding minority pixel intensities, fluctuations in texture are minimized when the intensities of the two component dither patterns are equal.

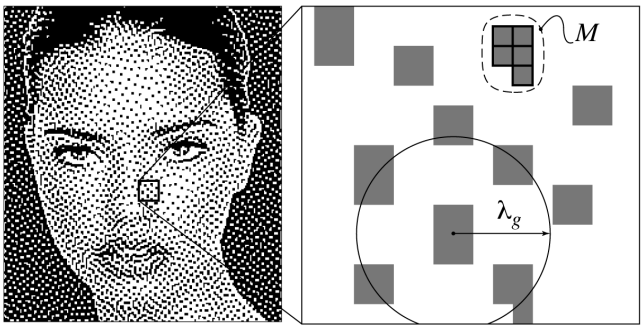

Figure 15. Diagram showing the relationship between the average spacing between clusters, λ_g, and the average size of clusters, $\overline{M}$, in a green-noise dither pattern.

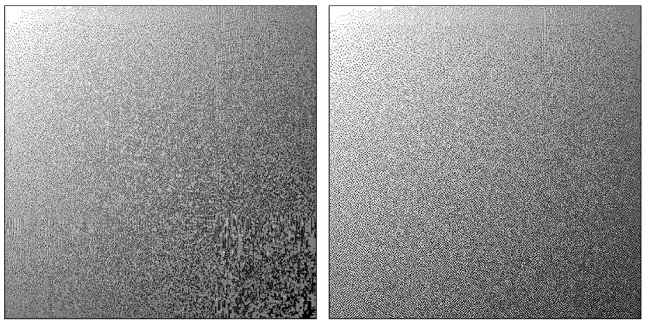

Figure 16. The composite halftones generated using error diffusion with output dependent feedback with (left) $h_b[n] = 0$ and (right) $h_b[n] = 1.0$.

Green-Noise Halftoning

In green-noise halftoning, the spatial arrangement of pixels is such that minority pixels form clusters which are themselves spread as homogeneously as possible – leading to an arrangement where clusters are separated (center-to-center) by an average distance of λ_g.[3] The parameter λ_g is the principle wavelength of green-noise, and given that the square of this wavelength is inversely proportional to the gray level divided by $\overline{M}$, the average number of pixels per cluster, λ_g is such that:

$$\lambda_g = \begin{cases} 1/\sqrt{g/\overline{M}}, & \text{for } 0 < g \le 1/2 \\ 1/\sqrt{(1-g)/\overline{M}}, & \text{for } 1/2 < g \le 1 \end{cases} \tag{9}$$

where λ_g is in units of pixels on the display (see Fig. 15).

Lau *et al.*[3] note that in green-noise, the cluster size/shape greatly influences the separation between clusters with large variation in cluster size/shape leading to large variations in separation between clusters. Just as with blue-noise, variations significantly longer than λ_g lead to a dither pattern which may appear too white relative to the optimal green-noise pattern.

For deriving the stochastic Moiré surface from two overlapping green-noise dither patterns, Lau *et al.*[10] rely on the parent process of a green-noise pattern where each minority pixel cluster is replaced with a single minority pixel at the cluster's centroid. This new pattern has statistical characteristics identical to a blue-noise dither pattern and, therefore, acts like a Poisson point sampling grid. So given λ_g's dependence on $\overline{M}$ in addition to g, it is, therefore, possible to modify the average spacing between parent process points (the underlying blue-noise process) without modifying g. In terms of stochastic Moiré visibility, modifying $\overline{M}$ allows green-noise patterns to move off the 1:1 line in a 2-color halftone without changing hue. So combining a blue-noise and a green-noise dither pattern allows us to avoid the worst case scenario of stochastic Moiré that occurs when overlapping patterns have equal spacing between their parent points. As was demonstrated previously in Fig. 3, blue and green-noise combinations will show a dramatic improvement in pattern homogeneity in the distribution of black pixels.

The drawback to increasing pattern coarseness, for the purpose of minimizing stochastic Moiré visibility, is the increased low frequency graininess of the component dither patterns caused by the increased visibility of the clusters. So given the detrimental impact of employing larger average cluster sizes, a printer manufacturer has the task of choosing the optimal trade-off between halftone visibility and stochastic Moiré visibility keeping in mind the additional impact of printer distortion.[10] In the case of halftoning by means of blue or green-noise masks, halftone coarseness is fixed at the time of mask construction and is independent of the image content. But using an adaptive scheme like Levien's error diffusion with output dependent feedback, halftone coarseness could be a function of the local color content of the image.

In Levien's algorithm, the quantized output, $y[n]$, is defined as:

$$y[n] = \begin{cases} 1, & \text{if } \left(x[n] + x_e[n] + x_h[n]\right) \ge 0 \\ 0, & \text{else} \end{cases} \tag{10}$$

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

$$x_h[n] = h\sum_{i=1}^{N} a_i \cdot y[n-i] \tag{11}$$

with $\sum_{i=0}^{N} a_i = 1$

and h an arbitrary constant. Referred to as the *hysteresis constant, h* acts as a tuning parameter with larger h leading to coarser output textures[16] as h increases ($h > 0$) or decreases ($h < 0$) the likelihood of a minority pixel if the previous outputs were also minority pixels. The error term, $x_e[n]$, is the diffused quantization error accumulated during previous iterations as:

$$x_e[n] = \sum_{i=1}^{M} b_i \cdot y_e[n-i] \tag{12}$$

with $y_e[n] = y[n] - (x[n] + x_e[n])$ and $\sum_{i=1}^{M} b_i = 1$.

For demonstration of stochastic Moiré using error diffusion with output dependent feedback, Fig. 16

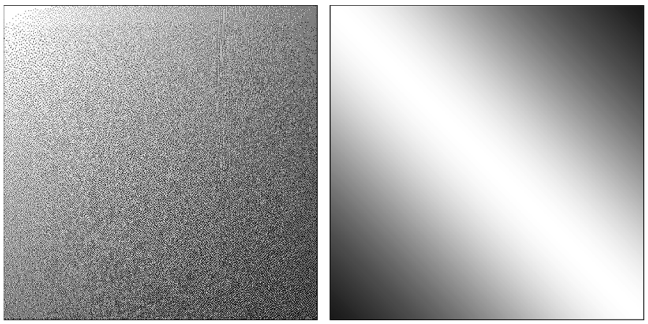

Figure 17. The (left) composite halftone generated using error diffusion with output dependent feedback where $h_A[n]$ was set to 0 for all n, while $h_B[n]$ was set equal to the image shown on the (right) derived from Eq. (13) with C_1 = 1.0 and C_2 = 10.0.

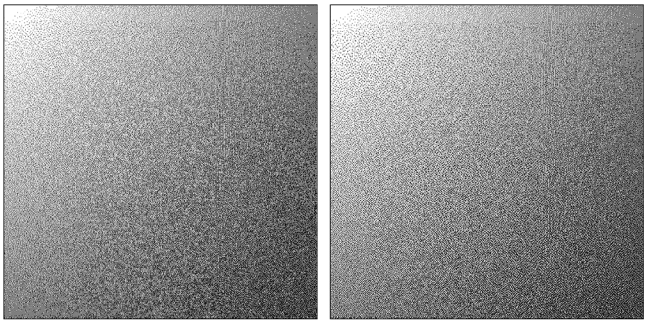

Figure 18. The composite halftones generated using error diffusion with output dependent feedback where (left) (C_1, C_2) = (0.6, 10.0) and (right) (C_1, C_2) = (1.0, 4.0).

(left) is the composite halftone pattern created by superimposing two halftone gradients both generated with h = 0.0. In Fig. 16 (right), h = 1.0 for the second gradient. Clearly, setting h = 1.0 for pattern **B** has greatly improved the homogeneity of black pixels across the image, but at the same time, has greatly increased the coarseness of the pattern in the lower-left corner where g_a = 0 and g_b = 0.5.

Following Lau's[19] suggestion of an image dependent hysteresis term, we can define $h[n]$ for the second of two, **A** and **B**, patterns as:

$$h_B[n] = h_A[n] + \frac{C_1}{1 + C_2 \cdot |g_a - g_b|^2} \tag{13}$$

where the scalar constants C_1 and C_2 now regulating coarseness with $h_B[n]$ ranging from $h_A + C_1/1+C_2$ when $|g_a - g_b| = 1$ and $h_A + C_1$ when $|g_a - g_b| = 0$. Using the two error and two hysteresis weight filter originally specified by Levien and setting C_1 = 1.0, C_2 = 10.0, and $h_A[n]$ = 0, Fig. 17 (right) shows the hysteresis value for the combined gradient pattern of Fig. 17 (left). Shown in Fig. 18 are further demonstrations of varying $h_b[n]$ with color composition where (C_1,C_2) = (0.6, 10.0) in Fig. 18 (left) and (C_1,C_2) = (1.0, 4.0) in Fig. 18 (right). Visual inspection reveals the dramatic improvement in homogeneity along the $g_a = g_b$ line, where the difference in coarseness is maximized, as well as the reduction in coarseness near (g_a, g_b) = (0.5, 0) in the lower-left corner of the pattern relative to the patterns of Fig. 16.

Moving now to more than two color images, we believe[10] that minimizing stochastic Moiré visibility between any two colors minimizes Moiré for all colors with coarseness between patterns ranked according to luminance. In the case of CMYK screens, this would imply a blue-noise screen for black, with a coarser pattern for magenta, a still coarser pattern for cyan, and with the coarsest pattern for yellow. For a demonstration, **Color Plate 6,** p. 375, compares the use of uncorrelated blue-noise dithering versus green-noise with varying coarseness between colors where, for a pixel in channel **K**, $h_K[n]$ is set constant at 0.3 due to 0.3's visually desirable attributes using two hysteresis and two error filter weights in monochrome halftones. In order to minimize stochastic Moiré visibility between magenta and black, $h_M[n]$ is set according to Eq. (13) for channel **M** such that:

$$h_M[n] = h_K[n] + \frac{C_1}{1 + C_2 \cdot |g_M - g_K|^2} \tag{14}$$

where C_1 = 0.6 and C_2 = 4.0. For cyan, $h_C[n]$ is set to minimize Moiré visibility between cyan and black as well as cyan and magenta; therefore, $h_C[n]$ is set such that:

$$\begin{aligned} h_C[n] &= h_M[n] + \max\left(h_{CK}[n], h_{CM}[n]\right), \text{where} \\ h_{CK}[n] &= \frac{C_1}{1 + C_2 \cdot |g_C - g_K|^2} \text{ and} \\ h_{CM}[n] &= \frac{C_1}{1 + C_2 \cdot |g_C - g_M|^2}. \end{aligned} \tag{15}$$

The same approach is applied to selecting the hysteresis parameter for channel **Y** such that:

$$\begin{aligned} h_Y[n] &= h_C[n] + \max\left(h_{YK}[n], h_{YM}[n], h_{YC}[n]\right), \text{where} \\ h_{YK}[n] &= \frac{C_1}{1 + C_2 \cdot |g_Y - g_K|^2} \\ h_{YM}[n] &= \frac{C_1}{1 + C_2 \cdot |g_Y - g_M|^2} \text{ and} \\ h_{YC}[n] &= \frac{C_1}{1 + C_2 \cdot |g_Y - g_C|^2}. \end{aligned} \tag{16}$$

Looking at the resulting CMYK halftones in **Color Plate 6**, we see that the difference in quality is most apparent in the fluctuations between red and yellow patches, assuming the print page is held at an appropriate distance from the observer's eye. Please note that if viewed from too close a distance, the graininess of the green-noise obscures the improvement in stochastic Moiré. From too far a distance, the two patterns blur to the point where fluctuations in color become indistinguishable. From a 20 inch viewing distance, the stochastic Moiré seems to become most apparent, overtaking the graininess of green-noise as the predominant halftone artifact.

The viewing distance and the visibility of Moiré varying with this distance brings up a very important issue regarding the use of green-noise as a means of minimizing Moiré visibility. That issue is how much of a difference in coarseness between colors should one have. Ultimately, coarseness should be a function of several factors including the printer manufacturer's personal interests to balance halftone visibility with stochastic Moiré visibility, and the device's ability to print dots reliably. A third factor is the HVS's sensitivity to stochastic Moiré visibility between certain colors, e.g., between cyan and magenta, as opposed to black and magenta or black and cyan. As described, the existing stochastic Moiré framework ignores these issues relating to color. But it should be known by the reader that the stochastic Moiré framework reviewed here was specifically designed to avoid such issues simply because of the many difficult questions that color assumes, aside from those relating to Moiré. On this same issue, Amidror *et al.*[4] simplified their analysis by restricting the problem of periodic Moiré for color halftones to the luminance channel such that the strength of the spectral impulses was proportional to the luminance of the corresponding color.

What Lau *et al.*[10] assumed is that by minimizing Moiré between any two channels, Moiré will be minimized across any three channels, and thus across the entire four channel image. And while this assumption may seem short sighted, studies of visual perception using Glass Patterns suggests very similar conclusions. In particular, Earle[20] shows that given a Glass Pattern composed of dots of three different luminance values, the HVS is more sensitive to correlation between the two most similar luminance values (cyan and magenta) and not necessarily the two lowest luminance values (black and magenta). So for CMYK screens, this work implies that the difference in coarseness between magenta and cyan should be larger than the difference between black and magenta, and that given a pattern composed of the three, the HVS is unlikely, or at least less likely, to notice the stochastic Moiré formed between cyan and black.

In our above derivation of the hysteresis values for the four channels, Eqs. (14-16) do not guarantee, with the exception of black and magenta, that two channels will not have the same principle wavelength. It is possible for cyan to be so coarse that even though it has a different intensity from black, it has a principal wavelength equal to that of black. Consider a situation where the intensity of black is low while both cyan and magenta have similar and high intensity values. In this situation, the cyan would be made coarse due to its similarity with magenta, but this coarseness comes not as a result of the low intensity of black such that cyan and black have the same principal wavelength. (Now as you may recall, similarity in principal wavelengths is the one indicator of a visually apparent stochastic Moiré.)

If we wanted to modify Eqs. (14-16) in order to prevent any two colors from having the same principal wavelength, we would need to measure the principal wavelength versus intensity versus hysteresis value, h, for a monochrome halftone and then use these data to identify the optimal hysteresis value for magenta, as being the h that adequately separates magenta's principal wavelength from that of black. For cyan, we would need to find the h that adequately separates its wavelength from both that of black and magenta, and lastly, repeat the process for yellow constrained by the three prior colors. But this may be unnecessary given studies by Prazdny[21] who shows that the perception of order in Glass Patterns is also affected by the similarity in dot sizes where differences in dot size erode the HVS's ability to detect correlations between points. What this means for stochastic Moiré is that the difference in cluster size between cyan and black acts as an additional deterrent to stochastic Moiré visibility.

What Eqs. (14-16) do guarantee is that colors with neighboring luminance values will have different principal wavelengths, with magenta always more coarse than black, cyan always more coarse than magenta, and yellow always more coarse than cyan. And this dissimilarity is, as Earle's and Prazdny's Glass Pattern studies would suggest, adequate for minimizing stochastic Moiré in four color prints. But clearly, psycho-visual evaluations of stochastic Moiré are a necessary and immediate concern for future stochastic Moiré research.

Conclusions

In closing, the reader should note that stochastic Moiré is a geometric phenomenon resulting from the superposition of two or more stochastic patterns, and it should be noted that Lau *et al.*[6] theorize that it is this stochastic Moiré phenomenon that leads to low frequency graininess in color FM halftones. The framework that followed the work of Lau *et al.* for measuring the visibility of stochastic Moiré should not be confused with stochastic Moiré itself, and it should be noted that there may exist other methods by which Moiré visibility can be analyzed. If nothing else, the framework used here supports Lau *et al.'s* assertion of the relationship between low frequency graininess and Moiré visibility. Furthermore, it justifies the combination of using blue and green-noise halftoning as a means of minimizing low frequency graininess between colors. It also adds more justification to the use of jointly bluenoise masks[11] and other halftoning algorithms[12] that closely correlate the component dither patterns of CMYK halftones.

The contribution that this paper has made is to extend the idea of mixing blue and green-noise halftones to an adaptive halftoning algorithm; whereas, Lau *et al.*[10] relied upon masks where the coarseness of component patterns are fixed. While the error diffusion with output dependent feedback scheme used here is not optimal, this study uses it as a means of demonstrating to the reader how one might go about optimizing a specific adaptive halftoning scheme, and what issues should be taken into account in doing so. Many of these issues are avoided by using masks, and they include the reliability/repeatability of the printing device, pattern visibility of the component green-noise patterns, the visibility of Moiré between patterns of varying coarseness, and the visibility between patterns of different luminance. Future research with stochastic Moiré will have to address the issue of superimposing more than two dither patterns and also how to account for color in the component screens. ▲

References

1. R. A. Ulichney, *Digital Halftoning*, MIT Press, Cambridge, MA, 1987.
2. D. L. Lau and G. R. Arce, *Modern Digital Halftoning*, Marcel Dekker Inc, New York, NY, 2001.
3. D. L. Lau, G. R. Arce and N. C. Gallagher, Green-noise digital halftoning, *Proc. IEEE* **86,** 2424–2444 (1998).
4. I. Amidror, R. D. Hersch, and V. Ostromoukhov, Spectral analysis and minimization of Moiré patterns in color separation, *J. Electronic Imaging* **3,** 295–317 (1994).
5. D. Blatner and S. Roth, *Real World Scanning and Halftones*, Addison-Wesley Publishing Company, Berkeley, CA, 1993.
6. D. L. Lau, A. M. Khan and G. R. Arce, Stochastic Moiré, in *IS&T's*

PICS 2001: Image Processing, Image Quality, Image Capture Systems Conf., IS&T, Springfield, VA, 2001, pp. 96–100.
7. L. Glass, Moiré effect from random dots, *Nature* **223,** 578–580 (1969).
8. L. Glass, Perception of random dot interference patterns, *Nature* **246,** 360–362 (1973).
9. D. L. Lau, Stochastic Moiré II, in *IS&T's PICS 2002: Image Processing, Image Quality, Image Capture Systems Conf.*, IS&T, Springfield, VA, 2002, pp. 250–255.
10. D. L. Lau, A. M. Khan and G. R. Arce, Minimizing stochastic Moiré by means of green-noise masks, *J. Opt. Soc. Amer. A*, **19,** 2203–2217 (2002).
11. J. Sullivan, L. Ray and R. Miller, Design of minimum visual modulation halftone patterns, *IEEE Trans. Systems, Man, and Cybernetics* **21,** 33–38 (1991).
12. Q. Lin and J. Allebach, Color FM screen design using the DBS algorithm, *Proc. SPIE* **3300,** 353–361, (1998).
13. R. W. Floyd and L. Steinberg, An adaptive algorithm for spatial gray-scale, *Proc. SID* **17,** 75–78 (1976).
14. P. Stucki, Mecca-a multiple-error correcting computation algorithm for bilevel image hardcopy reproduction, *Tech. Rep. RZ1060*, IBM Research Laboratory, Zürich, Switzerland, 1981.
15. D. L. Lau, G. R. Arce and N. C. Gallagher, Digital color halftoning via generalized error diffusion and multi-channel green-noise masks, *IEEE Trans. Image Processing*, **9,** 923–935 (2000).
16. R. Levien, Output dependant feedback in error diffusion halftoning, in *IS&T's Eighth International Congress on Advances in Non-Impact Printing Technologies*, IS&T, Springfield, VA, pp. 280–282.
17. H. R. Kang, *Digital Color Halftoning*, SPIE Press, Bellingham, WA, 1999.
18. A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill Companies, New York, NY, 1994.
19. D. L. Lau and G. R. Arce, Robust halftoning with green-noise, in *IS&T's PICS 1999: Image Processing, Image Quality, Image Capture Systems Conference*, Springfield, VA, 1999.
20. D. C. Earle, Glass patterns: grouping by contrast similarity, *Perception* **28,** 1373–1382 (1999).
21. K. Prazdny, On the perception of glass patterns, *Perception* **13,** 469–478 (1984).

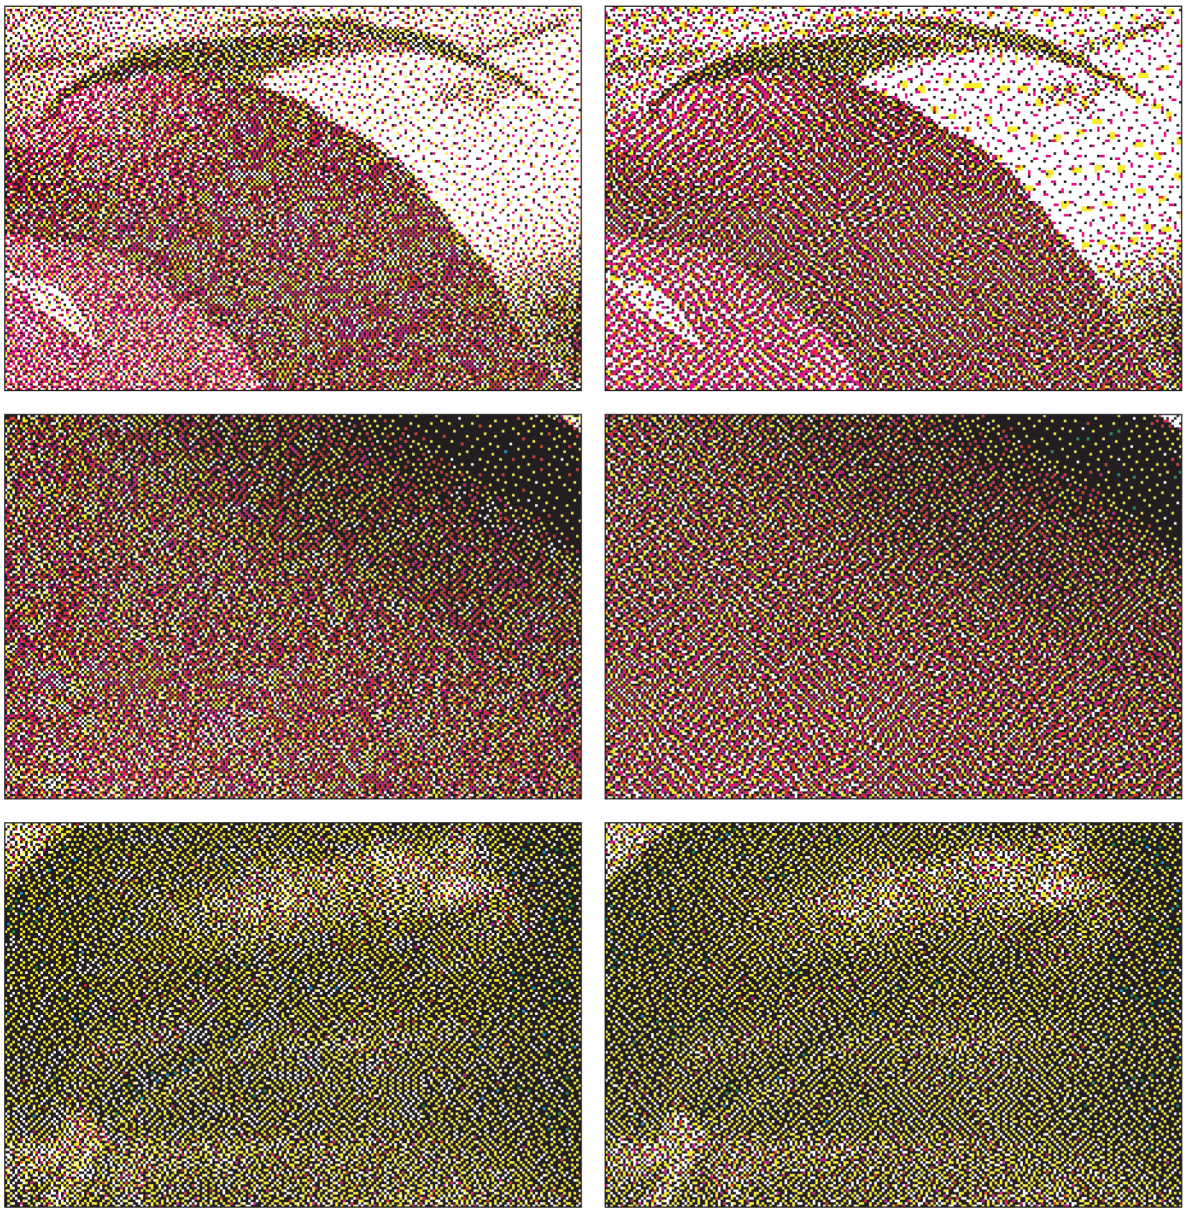

Color Plate 6. Four-color halftones the improvement in stochastic Moiré from (left) to (right) using varying levels of coarseness when patterns are uncorrelated. Holding these images 20 to 30 inches from the eye maximizes the visual differences between halftones. (*Lau,* pp. *327–338)*