Dithering with Blue Noise

ROBERT A. ULICHNEY

Digital halftoning, also referred to as spatial dithering, is the method of rendering the illusion of continuous-tone pictures on displays that are capable of only producing binary picture elements. The concept of blue noise—high frequency white noise—is introduced and found to have desirable properties for haltoning. Efficient algorithms for dithering with blue noise are developed, based on perturbed error diffusion. The nature of dither patterns produced is extensively examined in the frequency domain. New metrics for analyzing the frequency content of aperiodic patterns for both rectangular and hexagonal grids are developed; blue-noise dithering is found to be ideally suited for rectangular grids. Several carefully selected digitally produced examples are included.

GLOSSARY OF SYMBOLS

Δ Width of the annuli over which the power spectrum estimate $\hat{P}(f)$ is averaged to form the radially averaged spectrum $P_r(f_r)$.

λ_g Principle radial wavelength in a homogeneously distributed field of binary pixels representing the constant gray level g.

σ_g^2 Variance of an individual output binary pixel resulting from halftoning homogeneous region of gray level g.

$e[n]$ Error filter. Governs how past quantization errors are negatively distributed or "diffused" into the yet to be quantized image in the error diffusion algorithm.

f Continuous-space frequency vector,

$$f = \begin{bmatrix} f_1 \\ f_2 \end{bmatrix}.$$

f_g Principal radial frequency in a field of homogeneously distributed binary pixels representing gray level g.

f_r Radial frequency. The scalar distance in frequency units from zero frequency in a two-dimensional Fourier transform.

g Gray level. It has a continuous value in the range 0 (white) to 1 (black) inclusive.

$I[n]$ Quantized discrete-space image. The output from a halftoning process.

$I[n; g]$ The binary output image resulting from halftoning an image consisting of a fixed gray level $J[n] = g$.

$J[n]$ Continuous amplitude, discrete-space image. Input to a halftoning process.

n Discrete-space spatial index vector,

$$n = \begin{bmatrix} n_1 \\ n_2 \end{bmatrix}.$$

$N_r(f_r)$ Number of discrete frequency samples in an annulus about radial frequency f_r.

$P(f)$ Power spectrum. In this report, only the power spectra of binary output of aperiodic halftone processes on a single-input gray level are considered.

$\hat{P}(f)$ Power spectrum estimate.

$P_r(f_r)$ Radially averaged power spectrum. Sample mean of the frequency samples of $\hat{P}(f)$ in the annulus, $\|f\| - f_r| < \Delta/2$, about f_r.

$s^2(f_r)$ Sample variance of the frequency samples of $\hat{P}(f)$ in the annulus, $\|f\| - f_r| < \Delta/2$, about f_r.

$s^2(f_r)/P_r^2(f_r)$ Anisotropy of $\hat{P}(f)$.

S Sample period.

u_1, u_2 Frequency baseband replication vectors.

U Frequency baseband replication matrix, $[u_1 : u_2]$.

v_1, v_2 Spatial sampling vectors.

V Spatial sampling matrix, $[v_1 : v_2]$.

I. INTRODUCTION

The nature of various types of noise is often described by a color name. The most well known example is "white noise," so named because its power spectrum is flat across all frequencies, much like the visible frequencies in white light. "Pink noise" is used to describe low-frequency white noise, the power spectrum of which is flat out to some finite high-frequency limit. There is even the curious case of "brown noise," named for the spectrum associated with Brownian motion [13]. Introduced in this paper is "blue noise," the high-frequency complement of pink noise, which is shown to be important in the generation of good quality digital halftones.

Digital Halftoning is the method of rendering the illusion

of continuous-tone pictures on displays capable of only producing binary picture elements. It is sometimes referred to as spatial dithering. Several approaches to this process are known: Overviews of existing halftone algorithms have been written by Jarvis *et al.* [17], Stoffel and Moreland [30], Stucki [31], and Ulichney [34]. When a rendering system, such as that of the offset printing process, cannot accommodate small isolated black or white pixels, a "clustered-dot" screen is needed. Otherwise "dispersed-dot" methods are preferred, and are the focus of this paper. Techniques can also be categorized by computational complexity as either point or neighborhood processes.

Ordered dither is a point process that produces output by comparing a single continuous-tone input value to a deterministic periodic array of thresholds. For images on square grids, the threshold arrays attributed to Bayer [5] have been in wide use. Homogeneous threshold arrays have also been developed for hexagonal and asymmetric grids, with the theoretical background for the study and design of these dither patterns available in [34, ch. 6-7].

Schemes which generate blue noise are neighborhood operations. As in two-dimensional white-noise patterns, blue-noise patterns are aperiodic and radially symmetric. Although white-noise patterns do not suffer from the correlated periodicity of ordered dither, the fact that they possess energy at very low frequencies result in a grainy appearance. Blue-noise patterns enjoy the benefits of aperiodic, uncorrelated structure without low-frequency graininess.

The important algorithm known as Error Diffusion is closely examined and with some variation is found to be a good blue noise generator. This algorithm is examined for both rectangular and hexagonal grids. After the success of hexagonal grids for the case of ordered dither, it is surprising to learn, as will be theoretically established, that a rectangular grid is the superior choice for dithering with blue noise.

The input to a halftone process is $J[n_1, n_2] = J[n]$, a continuous-tone discrete-space monochrome image with gray values between zero (white) and one (black). The output is a binary discrete-space image $I[n]$.

Of particular interest are the binary patterns, $I[n; g]$, resulting from dithering an input image of only one fixed gray value $J[n] = g$. The aperiodic patterns generated in this paper can be modeled as stochastic processes. The unconditional probability mass function of any individual binary output pixel, $I[n; g]$, is

$$p_I(I[n; g]) = \begin{cases} g, & \text{for } I[n; g] = 1 \\ (1 - g), & \text{for } I[n; g] = 0 \end{cases} \qquad (1)$$

that is, it can be modeled as a Bernoulli distribution. Since this is true for all n, $I[n]$ is a stationary random process with

$$E\{I[n; g]\} = g \qquad (2)$$

and

$$\text{var }\{I[n; g]\} = \sigma_g^2 = g(1 - g). \qquad (3)$$

The mean of $I[n; g]$ is exactly what is expected, since the gray level g is being represented. The variance of $I[n; g]$ varies with g, and has a maximum at $g = \frac{1}{2}$, midway between the extremes of zero variance at solid black and white.

Such patterns will be closely examined in the frequency domain. The symbols used in this paper are listed in the Glossary of Symbols.

II. FOURIER ANALYSIS

The Fourier transform has been employed to evaluate halftone images for the cases of ordered dither on square grids [2], [18], [25], and ordered dither on hexagonal grids [33], [14]. In this section, metrics for analyzing the nature of aperiodic binary dither patterns will be developed in the frequency domain.

The most characteristic feature of a halftone technique is the texture generated in areas of uniform gray. The rendition of high-frequency detail depends primarily on how sharp the image was (or to what extent high-pass filtering was performed) prior to halftoning. *Radially averaged power spectra* along with a measure of anisotropy provide a mechanism for studying aperiodic patterns.

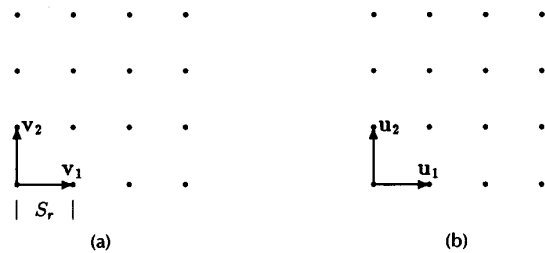Figs. 1 and 2 show square and hexagonal sampling grids

Fig. 1. Square grid. (a) The spatial-domain sample locations. (b) The frequency-domain sample locations.
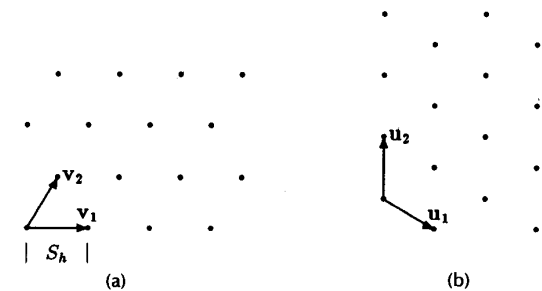

Fig. 2. Hexagonal grid. (a) The spatial-domain sample locations. (b) The frequency-domain sample locations.

or lattices. A periodic sampling grid can be defined by a spatial sampling matrix $V = [v_1 : v_2]$, composed of two linearly independent sampling vectors,

$$v_1 = \begin{bmatrix} v_{11} \\ v_{21} \end{bmatrix} \qquad v_2 = \begin{bmatrix} v_{12} \\ v_{22} \end{bmatrix}. \qquad (4)$$

The Fourier transform on such a grid will be repeated in the frequency domain as defined by the *frequency baseband replication matrix* $U = [u_1 : u_2]$.

These two matrices are reciprocally related by

$$V^T U = I$$

where I is the identity matrix. In this paper, frequency space has units of cycles per unit length. The details of computing the Fourier transform for this application are given in [34,

ch. 3]; the treatment of general periodic grids is also found in [8].

Both the rectangular and hexagonal grids to be considered are regular, that is, have pixel shapes tiling the plane that are regular polygons. In both cases, $|v| = |v_1| = |v_2| = S$. For the rectangular grid

$$V = \begin{bmatrix} S_r & 0 \\ 0 & S_r \end{bmatrix}$$

$$U = \begin{bmatrix} \dfrac{1}{S_r} & 0 \\ 0 & \dfrac{1}{S_r} \end{bmatrix}. \tag{5}$$

For the hexagonal grid

$$V = \begin{bmatrix} S_h & \dfrac{S_h}{2} \\ 0 & \dfrac{\sqrt{3}S_h}{2} \end{bmatrix}$$

$$U = \begin{bmatrix} \dfrac{1}{S_h} & 0 \\ \dfrac{-1}{\sqrt{3}S_h} & \dfrac{2}{\sqrt{3}S_h} \end{bmatrix}. \tag{6}$$

Sample periods for the rectangular and hexagonal case are denoted S_r and S_h. To maintain a constant number of samples per unit area in the examples that follow

$$\frac{S_r}{S_h} = \sqrt{\frac{\sqrt{3}}{2}}.$$

Note that the magnitude of the baseband replication vectors u_1 and u_2 are equal within each case. That is, $|u| = |u_1| = |u_2|$, and

$$|u| = \begin{cases} \dfrac{1}{S_r}, & \text{for rectangular grids} \\ \dfrac{2}{\sqrt{3}S_h}, & \text{for hexagonal grids.} \end{cases} \tag{7}$$

Halftone processes which do not produce output by thresholding with a deterministic, periodic threshold array will in general be aperiodic. As mentioned in the Introduction, such aperiodic dither patterns can be modeled as stochastic processes.

A. Estimating the Power Spectrum

The Fourier transform of the autocorrelation function of a stationary random process is the power spectrum $P(f)$. In most cases, the autocorrelation function of a given aperiodic halftone process will not be known, so a method of spectral estimation must be employed to produce an estimate $\hat{P}(f)$ of $P(f)$. The very simple technique known as Bartlett's Method [4] of averaging periodograms, named after the one who first suggested the technique for the one-dimensional case, is used in this study to produce $\hat{P}(f)$.

A periodogram is the magnitude squared of the Fourier transform of sample output divided by the sample size. All spectral estimates in this study were produced by averaging 10 periodograms of 256 × 256 output pixels from a given halftone rendering of a fixed gray level $I[n; g]$. Since some of the processes to be presented have transient behavior near edges or boundaries, the segments are cropped sufficiently far from output edges to avoid such artifacts; only the "steady-state" output is measured.

It can be shown [23] that a spectral estimate formed by averaging K periodograms has an expectation equal to $P(f)$ smoothed by convolution with the Fourier transform of a triangle function with a span equal to the size of the sample segments, and variance

$$\text{var}\{\hat{P}(f)\} \approx \frac{1}{K} P^2(f). \tag{8}$$

B. Radially Averaged Power Spectra and Anisotropy

A desirable attribute of a well-produced aperiodic halftone of a fixed gray level is radial symmetry; directional artifacts are perceptually disturbing. $\hat{P}(f)$ is a function of two dimensions. Although anisotropies in $I[n; g]$ can be qualitatively observed by studying three-dimensional plots of $\hat{P}(f)$, a more quantitative metric is proposed.

Fig. 3 shows how spectral estimates $\hat{P}(f)$ can be partitioned into annuli of width Δ for regular rectangular (Fig. 3(a)) and hexagonal (Fig. 3(b)) grids. Each annulus has a central radius f_r, the radial frequency, and $N_r(f_r)$ frequency samples.

Two useful one-dimensional statistics can be derived from averages within these annuli. The sample mean of the frequency samples of $\hat{P}(f)$ in the annulus, $\|f\| - f_r| < \Delta/2$ about f_r, is defined as the radially averaged power spectrum

$$P_r(f_r) = \frac{1}{N_r(f_r)} \sum_{i=1}^{N_r(f_r)} \hat{P}(f). \tag{9}$$

The sample variance of the same frequency samples is defined as

$$s^2(f_r) = \frac{1}{N_r(f_r) - 1} \sum_{i=1}^{N_r(f_r)} (\hat{P}(f) - P_r(f_r))^2. \tag{10}$$

Note that the sum is divided by $N_r(f_r) - 1$ and not $N_r(f_r)$, so as to yield an unbiased estimate of the variance (see [11]).

For each gray-level output of a halftone process to be analyzed, two plots are presented. First, the radially averaged power spectrum divided by σ_g^2 will be shown. Since spectral energy increases with σ_g^2 (see (3)), normalizing $P_r(f_r)$ by this amount will render all plots on the same relative scale. Because of the importance of σ_g^2, its relationship to gray level g is now shown in Fig. 4.

Secondly, the anisotropy of $\hat{P}(f)$ will be plotted. Anisotropy is defined in this paper as

$$\frac{s^2(f_r)}{P_r^2(f_r)} \tag{11}$$

a measure of the relative variance of frequency samples within a given annulus. Note that this measure can be described as the square of the coefficient of variation, or as a "noise-to-signal" ratio. Because the range of anisotropy values can be quite large, it will be plotted in decibels.

The zero frequency term is proportional to g^2 in all cases; the spike at this frequency will not be shown since it does not contribute to the structure of the dither pattern.
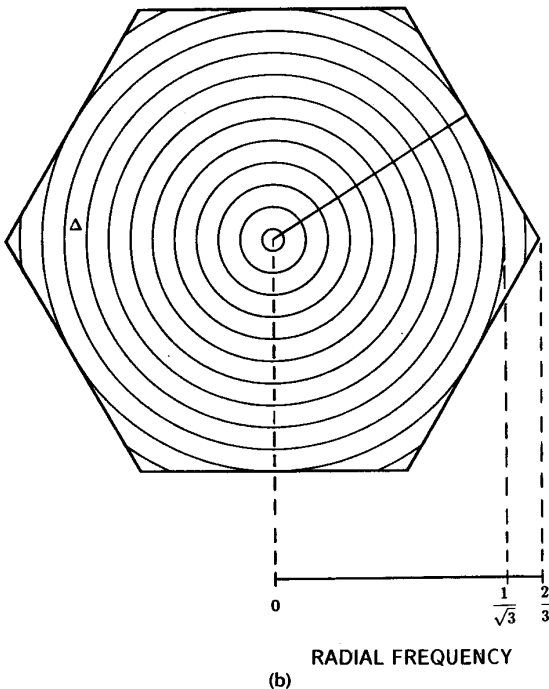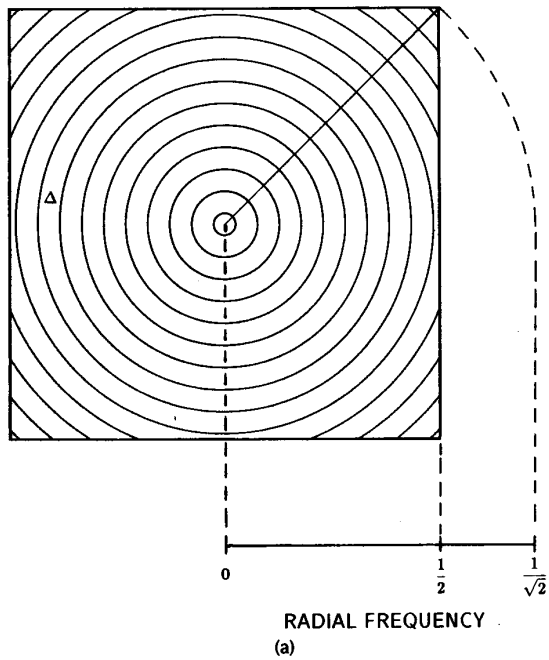
RADIAL FREQUENCY
(a)

RADIAL FREQUENCY
(b)

Fig. 3. Segmenting the spectral estimate into concentric annuli. The radial frequency f_r is in units of S^{-1}. (a) Regular rectangular (square) grid, 181 annuli actually used. (b) Regular hexagonal grid, 148 annuli actually used.

1) Quality of Measurement: To what extent will $P_r(f_r)$ and $s^2(f_r)/P_r^2(f_r)$ be meaningful metrics? From (8) and the fact that $K = 10$ segments are used in the estimate $\hat{P}(f)$

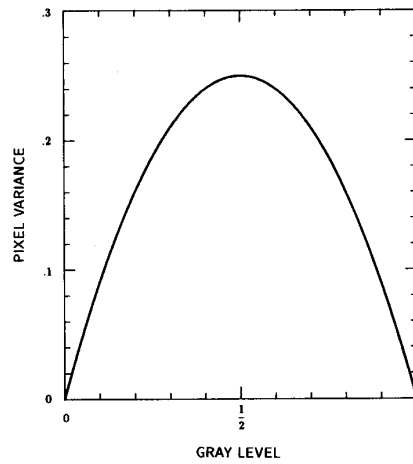$$\frac{\text{var } \{\hat{P}(f)\}}{P^2(f)} \approx \frac{1}{10}. \tag{12}$$

Fig. 4. The dependence of pixel variance σ_g^2 on gray level g.

If $P(f)$ is perfectly radially symmetric, the measure of anisotropy $s^2(f_r)/P_r^2(f_r)$ is merely an estimate of the above ratio. Therefore, an anisotropy of $\frac{1}{10}$ or -10 dB should be considered "background noise," and a reference line at this level will appear in each plot.

Also, if anisotropy is low, that is, close to -10 dB, indicating good radial symmetry, then $\hat{P}(f)$ is effectively a function of one independent variable f_r instead of two variables f. The variance of $P_r(f_r)$ is that of (12) divided by $N_r(f_r)$, assuming that each of the $N_r(f_r)$ samples are independent. This reduction in variance as $N_r(f_r)$ increases is indeed observed in the experimental data to be presented.

$N_r(f_r)$ depends on the width of the annuli Δ. As indicated earlier, in this paper all estimates $\hat{P}(f)$ will consist of 256^2 frequency samples. The size of Δ was chosen so that exactly one sample along each frequency axis fell into each annulus. A plot of $N_r(f_r)$ for (a) rectangular and (b) hexagonal grids is plotted in Fig. 5.
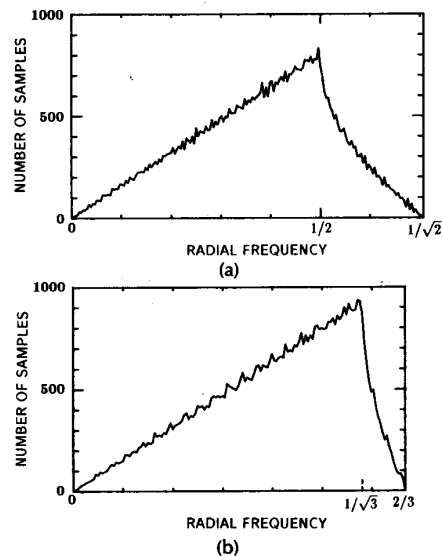

Fig. 5. Number of frequency samples $N_r(f_r)$ within each annulus centered on f_r. (a) Regular rectangular (square) grid. (b) Regular hexagonal grid.

The irregularities in the shape of these plots are a consequence of rectangular and hexagonal grids not being perfectly radially symmetric. The number of grid points that fall into a particular annulus essentially increases linearly, as one would expect, up to the largest annulus that will completely fit within the shape of the baseband; this occurs at $f_r/S^{-1} = \frac{1}{2}$ for rectangular grids and $f_r/S^{-1} = 1/\sqrt{3}$ for hexagonal grids.

III. WHITE NOISE DITHERING

In this section, the process of creating a dispersed-dot halftone by the point process of thresholding an input image with uniformly distributed, uncorrelated (white) noise is investigated. The quality of output from this method does not deserve consideration for practical use; the well-known method of dispersed-dot ordered dither performs much better than this one, taking no more computational effort.

So why should a section be devoted to this so called technique of "random dither"? There are two reasons.

The first is historical. The idea was the first used to exploit the fact that electronic displays can have independently addressable dots. Goodall [15] in 1951 and Roberts [24] in 1962 demonstrated how contouring due to insufficient gray levels can be corrected by adding noise of this type. This is perhaps the first technique that comes to mind to correct the shortcomings of using a fixed threshold, and in the early days of digital halftoning it was always referenced for comparison; in fact the name "ordered dither" was meant to contrast random dither.

Secondly, investigating the output of single gray levels $I[n; g]$, dithered in this way, provides a means to check the validity of the newly introduced metrics of radially averaged power spectrum and anisotropy. Since $I[n; g]$ is white noise, it has a known autocorrelation function, namely an impulse at the origin with area σ_g^2. So, the power spectrum should be radially symmetric with fixed amplitude σ_g^2. Such radial symmetry has been observed optically in Fraunhofer diffraction patterns of randomly distributed apertures [16].

The random numbers used in this section are, strictly speaking, pseudo-random. They are produced by means of a multiplicative congruential random number generator [19] available with many programming libraries. This is a very efficient scheme requiring only one multiplication and division per random number, and in the case used in this study, has a repeat period of 2^{32}.

The effects of dithering with white noise on regular rectangular and hexagonal grids will now be considered separately.

A. Rectangular Grids

A random dithered gray-scale ramp is shown in Fig. 6. An example of the effect of random dither on a scanned image [21] is given in Fig. 7. They suffer from a grainy appearance. This is the case at any displayed resolution because of the presence of long wavelengths (low frequencies) at all gray levels. Both the rectangularly sampled gray-scale ramp and scanned image contain 300 by 400 pixels.

The radially averaged power spectrum and anisotropy are essentially the same for all gray levels, and is shown in Fig. 8(a) for one particular gray level $g = \frac{1}{8}$. The well-behaved nature of these plots validate four things:
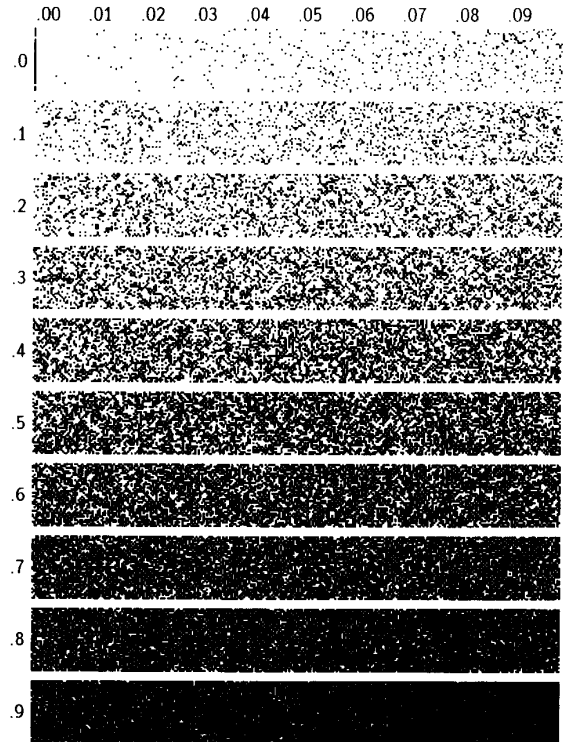

Fig. 6. Rectangular random dither of a gray-scale ramp.

Fig. 7. Rectangular random dither of a scanned picture.

1) The amplitude of $P_r(f_r)$ is correct, that is, flat, as expected for white noise.
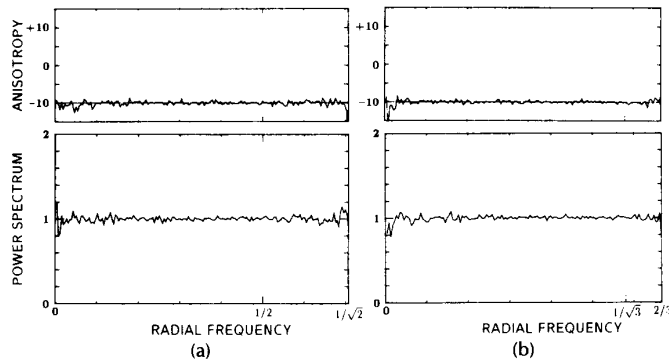2) The values of σ_g^2 are as predicted by (3), Fig. 4.

Fig. 8. Analysis of random dither. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f_r (in units of S^{-1}) for the fixed gray level $g = \frac{1}{8}$, with $\sigma_g^2 \approx 0.1094$. (a) Rectangular case. (b) Hexagonal case.

3) The apparent variance of $P_r(f_r)$ decreases with $N_r(f_r)$ (Fig. 5) with a minimum at $\frac{1}{2}$, (1/ √3for the hexagonal case).

4) The anisotropy measure is correct. White noise is radially symmetric so the anisotropy should be at the "background noise" minimum of −10 dB.

The characteristic features seen here will serve as a reference for the other plots to be studied.

B. Hexagonal Grids

Hexagonal radially averaged power spectra exhibit the same well-behaved features as in the rectangular case. A random dithered gray-scale ramp and scanned image are shown in Figs. 9 and 10. The continuous-tone source images used were the same as those used in the rectangular case,
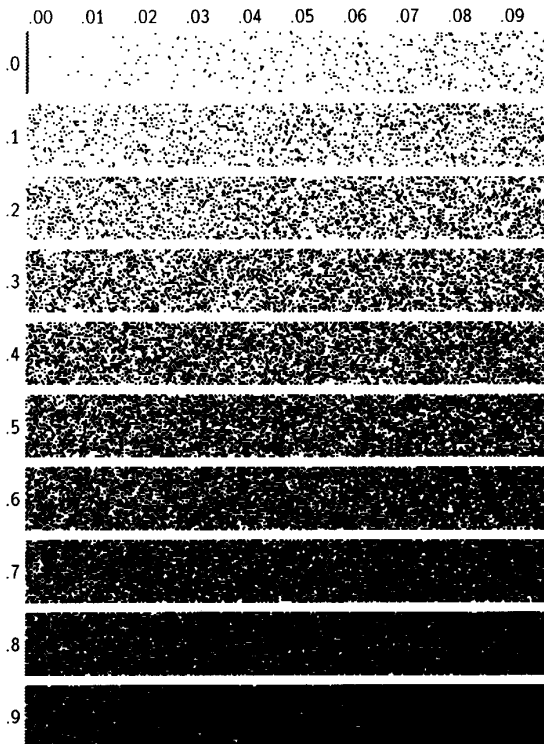

Fig. 9. Hexagonal random dither of a gray-scale ramp.

Fig. 10. Hexagonal random dither of a scanned picture.

and displayed with the same number of samples per unit area. These hexagonally sampled images contain 273 by 440 pixels, with alternate rows offset by one half horizontal period.

Again, the radially average power spectrum and anisotropy are essentially the same for all gray levels. One representative sample at $g = \frac{1}{8}$ is shown in Fig. 8(b).

IV. Principle Wavelength

Consider the problem of rendering a fixed gray level g with binary pixels on regular rectangular or hexagonal grids. For gray levels between $g = 0$ (white) and $g = \frac{1}{2}$, the black pixels are in the minority and it is their spatial distribution on a background of white that determine the perceived pat-

terns. The patterns for dark gray levels, between $g = 1$ (black) and $g = \frac{1}{2}$, are perceived as white pixels on a black background.

The goal is to distribute the binary pixels as homogeneously as possible. The pixels, if so arranged, would be separated by an average distance in two dimensions. This distance is called the *Principal Wavelength*, and would have the value

$$\lambda_g = \begin{cases} |v|/\sqrt{g}, & 0 \le g \le \frac{1}{2} \\ |v|/\sqrt{1-g}, & \frac{1}{2} < g \le 1 \end{cases} \tag{13}$$

where $|v| = S$ as described in Section II. Several intuitive properties justify the relation in (13):

1) As the gray level approaches perfect black and white the wavelength approaches infinity.

2) Wavelength decreases symmetrically with equal deviations from black and white toward middle gray.

3) The square of the wavelength is inversely proportional to the number of minority pixels per unit area. Note that the number of minority pixels per unit area is proportional to g for light gray levels, and $(1 - g)$ for dark gray levels. As property 3) defines the equation within a proportionality constant, the equality is complete by determining the wavelength for a known homogeneous distribution at a particular gray level. That constant is $|v|$ for both rectangular and hexagonal grids.

Since the distribution is assumed to be homogeneous, the corresponding power spectrum would be radially symmetric. The principal wavelength would be manifested as the *Principal Frequency*

$$f_g = \begin{cases} \sqrt{g}\,|u|, & 0 \le g \le \frac{1}{2} \\ \sqrt{1-g}\,|u|, & \frac{1}{2} < g \le 1 \end{cases} \tag{14}$$

where $|u|$ is as defined in (7).

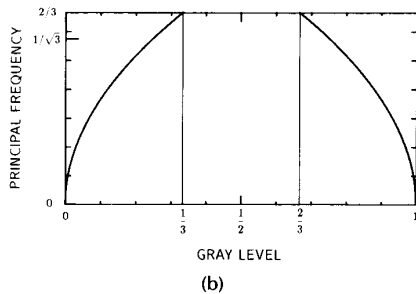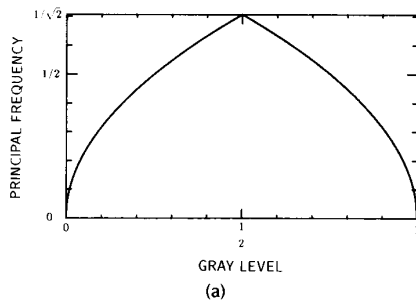f_g is ploted in Fig. 11 as a function of g for (a) rectangular

Fig. 11. Principal frequency f_g (in units of S^{-1}) as a function of gray level g. (a) Regular rectangular (square) grid. (b) Regular hexagonal grid.

and (b) hexagonal grids. These plots reveal an amazing shortcoming of hexagonal grids; they cannot support a principal frequency for $\frac{1}{3} < g < \frac{2}{3}$! Fig. 12(b) illustrates the highest frequency which can exist on a hexagonal grid, $f_r/S^{-1} = \frac{2}{3}$, for either a pattern for $g = \frac{1}{3}$ or its complement $g = \frac{2}{3}$.

On a rectangular grid (Fig. 12(a)), the checkerboard pattern for gray level $g = \frac{1}{2}$, and frequency $f_r/S^{-1} = 1/\sqrt{2}$, can
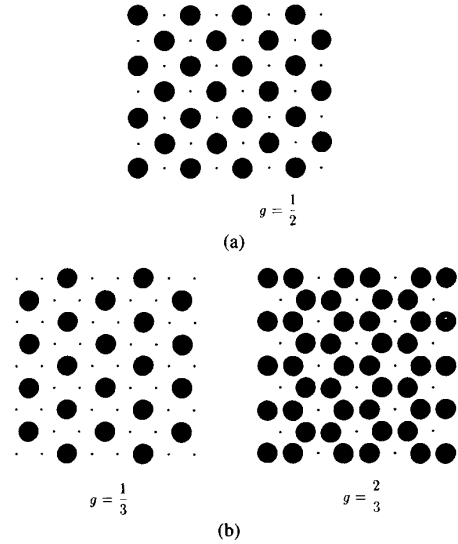

Fig. 12. Patterns with the highest possible spatial frequency (corresponding to the corners of the spectral baseband). (a) Rectangular grid, $f_g/S^{-1} = 1/\sqrt{2}$, $g = \frac{1}{2}$. (b) Hexagonal grid, $f_g/S^{-1} = \frac{2}{3}$, $g = \frac{1}{3}$, and $g = \frac{2}{3}$.

most definitely be supported. These spatial patterns correspond to the high-frequency corners of the frequency baseband.

For gray levels in the range $\frac{1}{3} < g < \frac{2}{3}$ on a hexagonal grid, pixels must be grouped together resulting in frequencies lower than $f_r/S^{-1} = \frac{2}{3}$. A principal frequency exists for all gray levels on rectangular grids.

A well-formed binary dither pattern rendering of a fixed gray level should consist of an isotropic field of binary pixels with an average separation of λ_g. This average separation should vary in an uncorrelated manner, but the wavelengths of this variation must not be significantly longer that λ_g. Recall that the failure of dithering with white noise was due to the presence of long wavelengths.

Fig. 13 depicts these well-formed dither pattern characteristics in the frequency domain. The radially averaged power spectrum of a fixed gray level g has three important features. First, its peak should be at the principal frequency for that gray level f_g. This frequency marks a sharp transition region below which little or no energy exists. And finally, the uncorrelated high-frequency fluctuations are characterized by high-frequency white noise, or "blue noise."

The visually pleasing nature of some of the patterns generated by the error diffusion algorithm to be presented in the next section can be attributed to a spectral signature as just described. However, several shortcomings exist in this algorithm. In Section VI, error diffusion enhanced with certain stochastic perturbations is found to be a good blue noise generator.
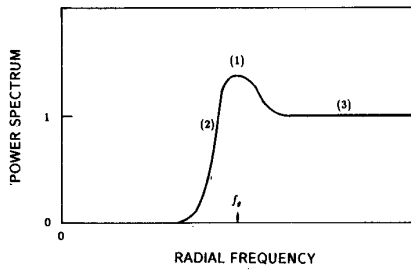
Fig. 13. Spectral characteristics of a well-formed dither pattern: (1) Low-frequency cutoff at principal frequency. (2) Sharp transition region. (3) Flat high-frequency "blue noise" region.

The presence of significant low-frequency energy is responsible for the visibility of disturbing artifacts in half-toning patterns. For dispersed-dot ordered dither, half of the total possible gray-level patterns available for a given threshold array have low-frequency components which correspond to wavelengths of the size of the threshold period [34]. For good blue noise processing, the lowest frequency is essentially f_g. The negative feedback of error diffusion acts as a low frequency inhibitor.

V. ERROR DIFFUSION ALGORITHM

The error diffusion algorithm, first introduced by Floyd and Steinberg in 1975 [9], [10], requires neighborhood operations and is thus more computationally intensive than point-process halftone methods. It is currently the most popular neighborhood halftoning process and has received considerable attention in spite of some shortcomings. A generic form of this algorithm is graphically illustrated in Fig. 14.
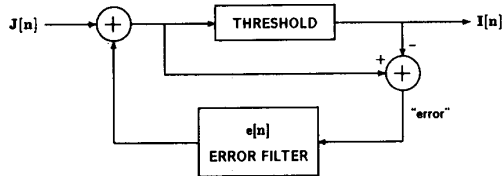

Fig. 14. The error diffusion algorithm.

The threshold in this case is fixed at $\frac{1}{2}$ where the input $J[n]$ varies as usual from $g = 0$ (white) to $g = 1$ (black). The resulting binary output value of 0 or 1 is compared with the original gray level value. The difference is suitably called the "error" for location n. The signal consisting of past error values is passed through an error filter $e[n]$ to produce a correction factor to be added to future input values. Errors are thus "diffused" over a weighted neighborhood determined by $e[n]$.

Fig. 15 summarizes error filter impulse responses promoted in the literature. Note that in all cases the values are deterministic and sum to 1 so that errors are neither amplified nor reduced. The first three listed are designed for rectangular grids; the error filter in Fig. 15(d) is intended for use on a hexagonal grid. In this section, their effect on resulting dither patterns will be examined in both the spatial and frequency domains.

(a)

(b)

(c)

(d)

Fig. 15. Error filters reported in the literature. (a) Floyd and Steinberg (1975) [9], [10]. (b) Jarvis, Judice, and Ninke (1976) [17]. (c) Stucki (1981) [32]. (d) Stevenson and Arce (1985) [29]. (a), (b) and (c) are for rectangular grids; (d) is for hexagonal grids. ("•" represents the origin.)

A. The Floyd and Steinberg Filter

The original error filter or set of "weights" suggested by Floyd and Steinberg is shown in Fig. 15(a). They argued that a filter with four elements was the smallest number that could produce good results. The values were chosen to particularly assure the checkerboard pattern at middle gray.

Fig. 16 shows the effect of error diffusion with the Floyd and Steinberg filter for the gray scale ramp (Fig. 16(a)) and, scanned picture (Fig. 16(b)). The reason for the popularity of this algorithm is clear; several gray levels are represented by pleasingly isotropic, structureless distributions of dots. However, some shortcomings are also apparent:

1) Correlated artifacts in many of the gray level patterns. This can best be seen in the gray scale ramp.

2) Directional hysteresis due to the raster order of processing. This artifact is most apparent in very light and very dark patterns. Note the light regions of the sky in the scanned picture.

3) Transient behavior near abrupt edges or boundaries.

The radially averaged power spectrum $P_r(f_r)$ and anisotropy measure $s^2(f_r)/P_r^2(f_r)$ are plotted for several gray levels in Fig. 17. In each case here and throughout this paper, the principal frequency f_g is marked with a small diamond on the frequency axis of the power spectrum.

The lack of symmetry in these patterns is strongly acknowledged in the anisotropy measure, especially for Fig. 17(d) with $g = \frac{1}{4}$ and Fig. 17(e) with $g = \frac{1}{2}$. Recall that the "background noise limit" due to the spectral estimation method used here is -10 dB (12) indicated by a reference line at the level. Any measure greater than 0 dB at any frequency indicates an especially anisotropic pattern; at such a level, the sample variance is greater than the square of the average.

B. Filters with 12 Weights

In 1976 Jarvis, Judice, and Ninke [17] reported an error filter with the 12 elements shown in Fig. 15(b). Although
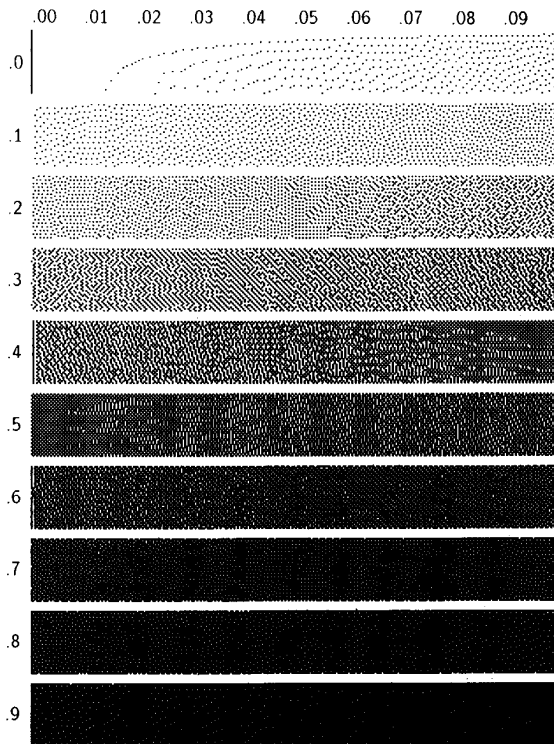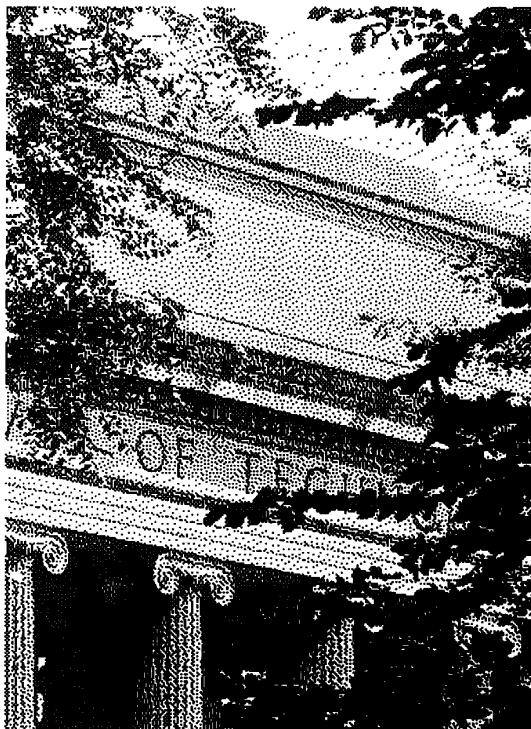
Fig. 16. Error diffusion with the Floyd and Steinberg filter. (a) Gray-scale ramp. (b) Scanned picture.

their technique was identical to error diffusion, they called it the "Minimum Average Error" algorithm. Output from error diffusion with this filter is shown in Fig. 18. The larger filter size does reduce some of the artifacts seen with the four-element filter of Floyd and Steinberg, but directional hysteresis in the very dark and light regions has increased, and pixels are clustered together more in the middle gray regions.

It also sharpens the picture more. Pictures usually look better with some degree of sharpening; in the case of Fig. 18(b), the enhancement due to sharpening somewhat masks the problems with patterns in uniform gray regions. Besides compounding the comparison process, this sharpening feature is less than ideal for two reasons. Edges are amplified along one direction only and the degree of this amplification may or may not be to the degree desired. In fact, Stucki [32] argues that it is objectionable and uses an additional filter to *inhibit* this edge amplification. For any half-toning technique, a precise degree of *symmetric* sharpening can easily be controlled by first high-pass filtering the image. The effect of prefiltering with a digital Laplacian is demonstrated in Section VII.

The 12-element error filter used by Stucki is shown in Fig. 15(c). For computational efficiency, he selected values that are all powers of 2. The effect of error diffusion with this error filter is found to be quite similar to the images of Fig. 18.

The patterns generated for fixed gray levels with the Jarvis *et al.* filter is examined in the frequency domain in Fig. 19. While all gray levels shown still suffer some anisotropy $g = \frac{1}{32}, \frac{1}{16}$, and $\frac{1}{8}$ show a stronger concentration of energy at the principal frequency f_g than those of Fig. 17. The clustering of pixels at $g = \frac{1}{4}$ and $\frac{1}{2}$ results in significant energy at frequencies lower than f_g. The radial spectra for dither patterns generated on rectangular grids are generally very similar for gray levels g and $1 - g$; the similar patterns exist at $g = \frac{3}{4}$ and $g = \frac{1}{4}$, $g = \frac{7}{8}$ and $g = \frac{1}{8}$, and so on. So for this and other rectangular cases gray levels greater than $\frac{1}{2}$ are not shown.

C. Hexagonal Case

In Section IV it was argued that hexagonal grids were inferior to rectangular grids as far as generating blue noise. This is based on the inability of a hexagonal grid to support a principal wavelength λ_g for $\frac{1}{3} < g < \frac{2}{3}$ (see Fig. 11(b)). In spite of this deficiency, the other features of hexagonal grids, particularly its superior covering efficiency (see [34, ch. 2]), are still reason to devote attention.

The only reported attempt at performing error diffusion on a hexagonal grid was by Stevenson and Arce [29], whose error filter was given in Fig. 15(d). They stated that this is the filter which gave the "highest image quality," but admitted that no optimization was done. The effect of hexagonal error diffusion with this filter is shown in Fig. 20. The gray scale-ramp reveals many disturbing texture patterns.

A close look at the radially averaged power spectrum in Fig. 21 reflects the large measure of anisotropy in seven selected patterns. Perhaps the reason for the bizarre shape of many of the anisotropy plots is due to the stable texture patterns that begin to "grow" in regions of constant gray producing localized spikes in the power spectra.
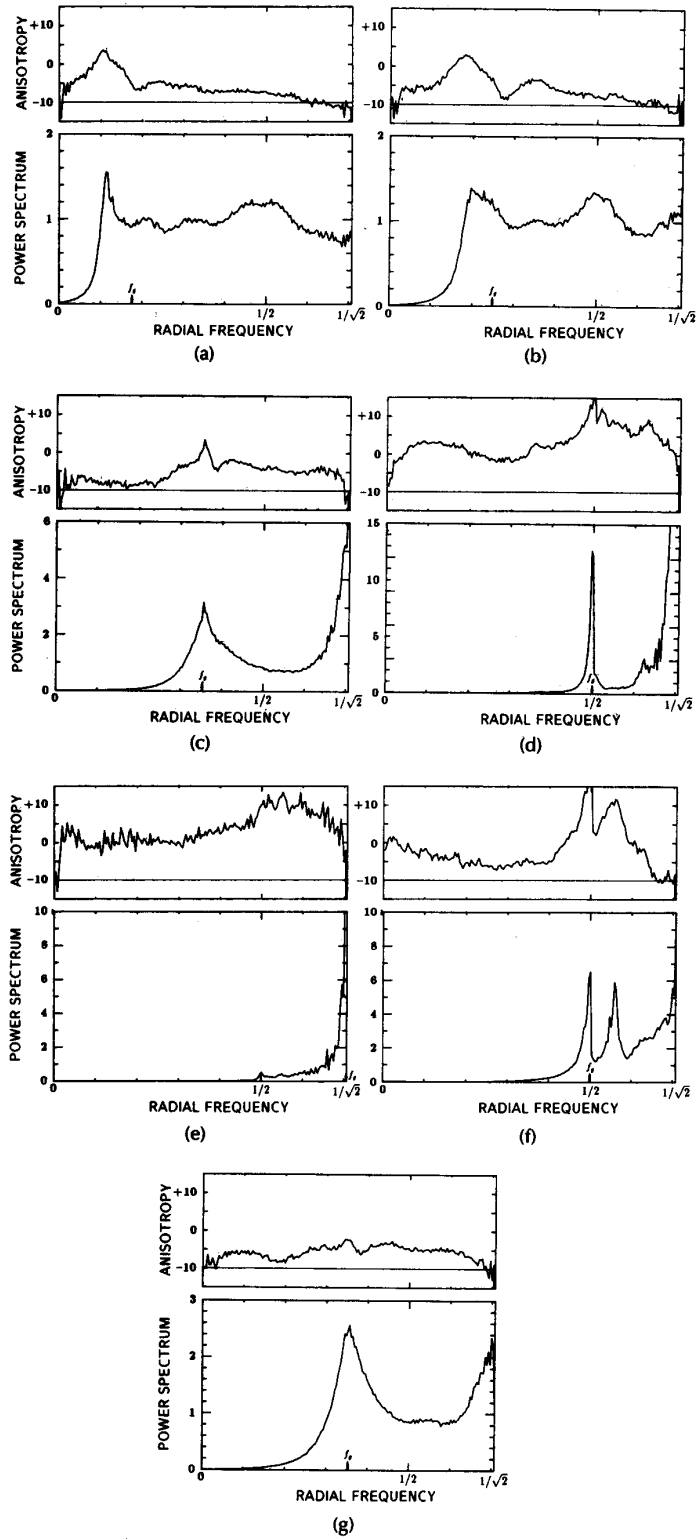
Fig. 17. Analysis of the Floyd and Steinberg filter. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f, (in units of S^{-1}) for selected fixed gray levels: (a) $g = \frac{1}{32}$, $f_g/S^{-1} \approx 0.1768$, $\sigma_g^2 \approx 0.0303$. (b) $g = \frac{1}{16}$, $f_g/S^{-1} = 0.25$, $\sigma_g^2 \approx 0.0586$. (c) $g = \frac{1}{8}$, $f_g/S^{-1} \approx 0.3495$, $\sigma_g^2 \approx 0.1094$. (d) $g = \frac{1}{4}$, $f_g/S^{-1} = 0.5$, $\sigma_g^2 = 0.1875$. (e) $g = \frac{1}{2}$, $f_g/S^{-1} \approx 0.7071$, $\sigma_g^2 = 0.25$. (f) $g = \frac{3}{4}$, $f_g/S^{-1} = 0.5$, $\sigma_g^2 = 0.1875$. (g) $g = \frac{7}{8}$, $f_g/S^{-1} \approx 0.3536$, $\sigma_g^2 \approx 0.1094$.
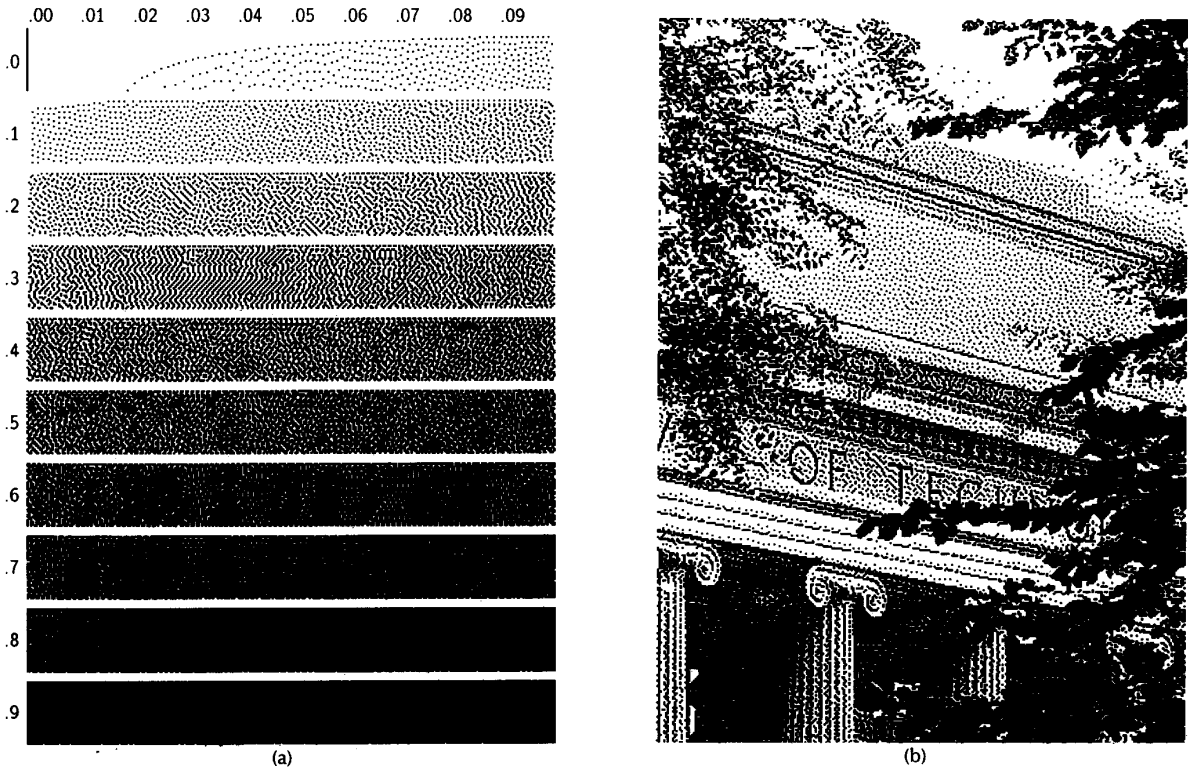
Fig. 18. Error diffusion with the Jarvis *et al.* Filter. (a) Gray-scale ramp. (b) Scanned picture.
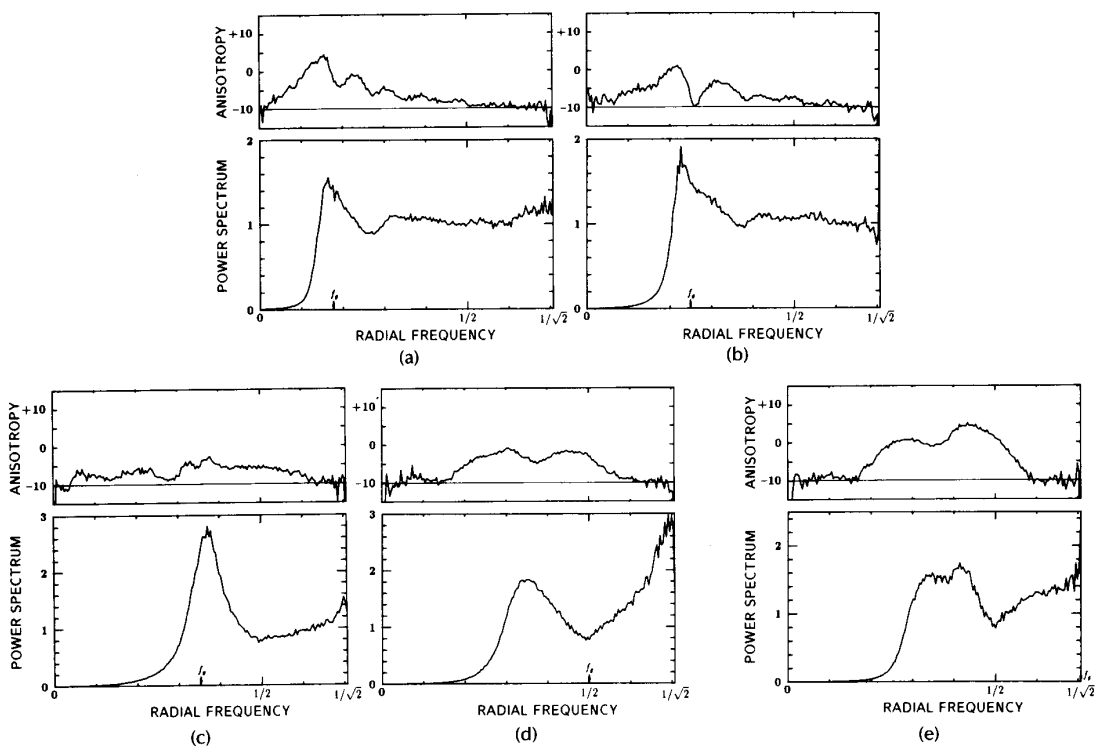
Fig. 19. Analysis of the Jarvis *et al.* filter. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f_r (in units of S^{-1}) for selected fixed gray levels. (a) $g = \frac{1}{32}$, $f_g/S^{-1} \approx 0.1768$, $\sigma_g^2 \approx 0.0303$. (b) $g = \frac{1}{16}$, $f_g/S^{-1} = 0.25$, $\sigma_g^2 \approx 0.0586$. (c) $g = \frac{1}{8}$, $f_g/S^{-1} \approx 0.3495$, $\sigma_g^2 \approx 0.1094$. (d) $g = \frac{1}{4}$, $f_g/S^{-1} = 0.5$, $\sigma_g^2 = 0.1875$. (e) $g = \frac{1}{2}$, $f_g/S^{-1} \approx 0.7071$, $\sigma_g^2 = 0.25$.

Fig. 20. Error diffusion with the Stevenson and Arce filter. (a) Gray-scale ramp. (b) Scanned picture.

VI. Error Diffusion with Perturbation

The idea of "dithering" or perturbing a method in image processing to defeat visual artifacts of a regular and deterministic nature has been used before. Randomizing sampling grids [7] is one method used to reduce the aliasing effects of undersampled images. Allebach improved the classical screen by randomizing the centers of dot clusters [1], [3] which eliminated the occurrence of morié patterns. And in the color printing industry, arbitrary or "irrational" clustered-dot screen angles have been digitally produced by employing random perturbations [12], [26].

Proposed here are several modifications to the basic error diffusion algorithm graphically depicted in Fig. 14. They are categorized in the following four areas.

Choice of Error Filter: An enormous number of choices are available for this, the deterministic part of the algorithm. An error filter can consist of weights of any 1) number, 2) position, and 3) value. For computational efficiency, as small a filter as possible is preferred.

Threshold Perturbation: In 1983, Billotet-Hoffman and Bryngdahl [6] suggested using an ordered dither threshold array in place of the fixed threshold used in error diffusion. However, the resulting halftoned output differs little from conventional ordered dither. A modification to this idea would be to perturb a fixed threshold within a given maximum percentage with ordered dither and/or white noise. So, additional parameters include:

1) choice of period size of ordered dither,
2) magnitude of ordered dither perturbation,
3) magnitude of white noise perturbation.

Raster Direction: The directional artifacts seen in the examples of error diffusion are due largely to the traditional raster order of processing. Many choices of space-filling curves to define the order of processing are possible. Although they did not call it error diffusion, Witten and Neal [38] demonstrated fairly good results by essentially using an error filter with *one* deterministic weight and processing all of the two-dimensional image data on a Peano curve (a type of fractal). While this particular approach imposes heavy demands on memory, the idea of using nonstandard raster ordering should be tried.

One idea that breaks up the directionality of a normal raster without the expense of a full two-dimensional buffer is to process along a serpentine raster (see Fig. 22). Neighborhood operations in image processing hardware or software buffer image data in full lines. So the choice of serpentine raster processing does not require any memory increase over a normal raster.

Stochastic Filter Perturbations: Along with threshold perturbations, random noise can be added to the elements or weights of the error filter. This idea was proposed by Schreiber of MIT and demonstrated by Woo [39], but only on the 12-element filter of Jarvis *et al.*

The magnitude or range of additive noise can be adjusted for each element. The sum of all of the weights in the resulting stochastic filter should still be unity at all times. This condition can be met by pairing weights of comparable value. For each pair of weights a scaled random value is added to one and subtracted from the other.

A random value χ is generated by the method described in Section III with the adjusted uniform probability density
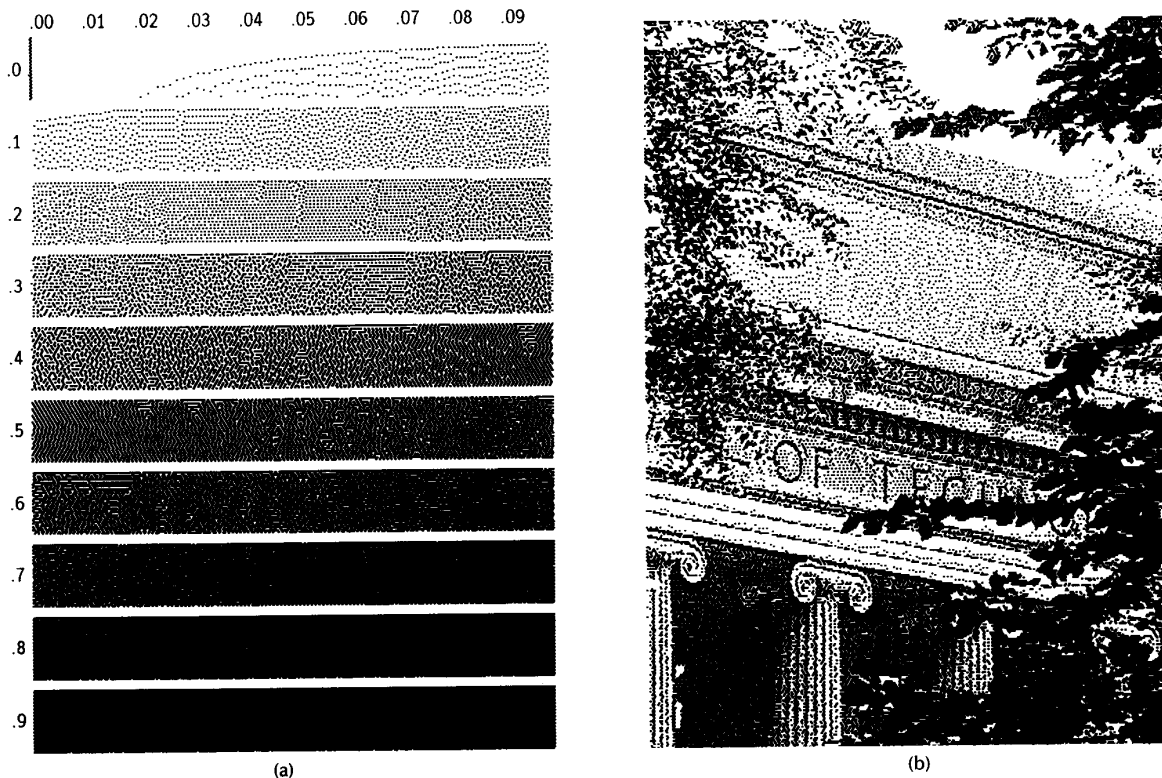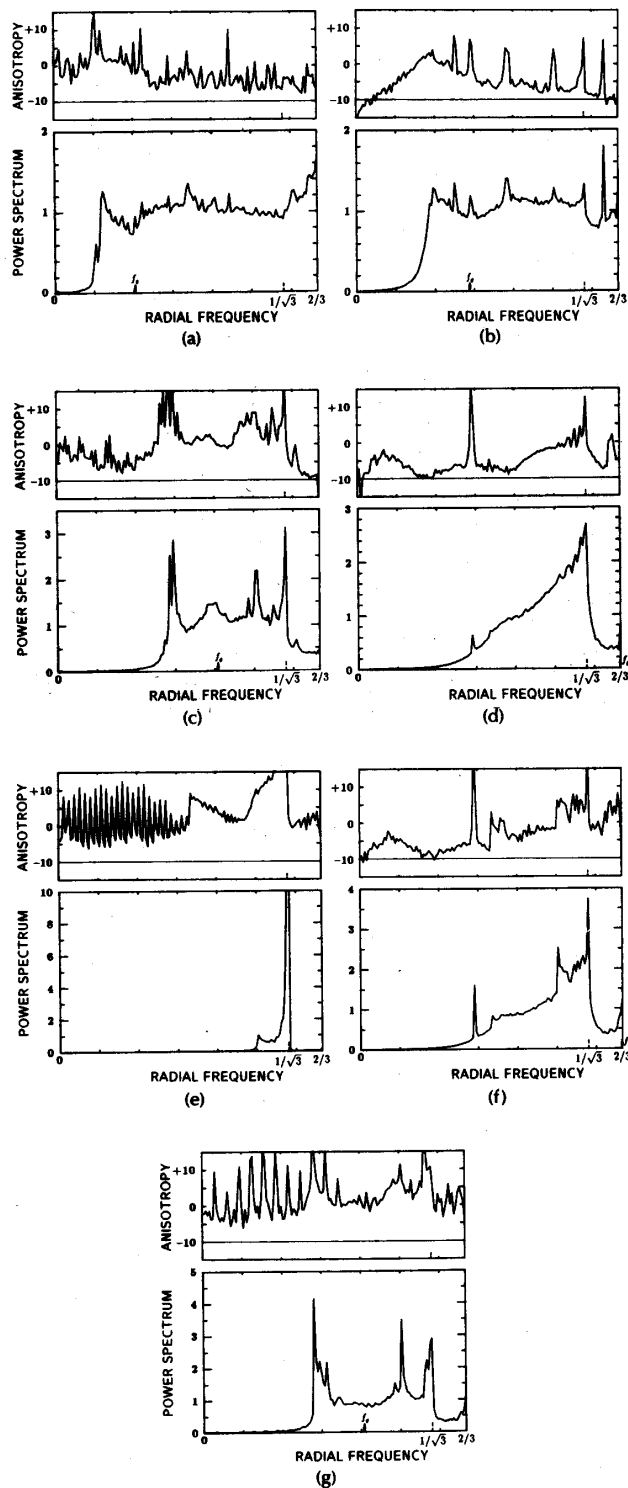
Fig. 21. Analysis of the Stevenson and Arce filter. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f_r (in units of S^{-1}) for selected fixed gray levels. (a) $g = \frac{1}{32}$, $f_g/S^{-1} \approx 0.2041$, $\sigma_g^2 \approx 0.0303$. (b) $g = \frac{1}{16}$, $f_g/S^{-1} \approx 0.2887$, $\sigma_g^2 \approx 0.586$. (c) $g = \frac{1}{8}$, $f_g/S^{-1} \approx 0.4082$, $\sigma_g^2 \approx 0.1094$. (d) $g = \frac{1}{3}$, $f_g/S^{-1} = \frac{2}{3}$, $\sigma_g^2 \approx 0.1111$. (e) $g = \frac{1}{2}$ (f_g is not supported), $\sigma_g^2 = 0.25$. (f) $g = \frac{2}{3}$, $f_g/S^{-1} = \frac{2}{3}$, $\sigma_g^2 \approx 0.1111$. (g) $g = \frac{7}{8}$, $f_g/S^{-1} \approx 0.4082$, $\sigma_g^2 \approx 0.1094$.
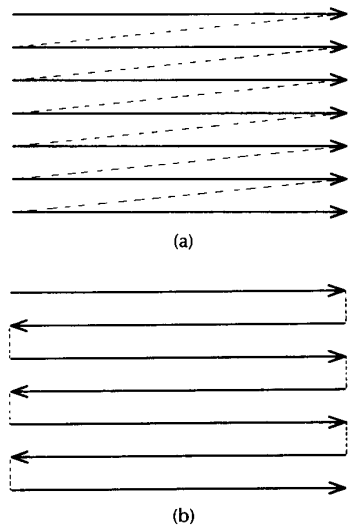
Fig. 22. Two processing path options. (a) Normal raster. (b) Serpentine raster.

function

$$p_x(\chi) = \begin{cases} \dfrac{a_i}{2}, & \text{for } -a_i < \chi < a_i \\ 0, & \text{otherwise} \end{cases} \qquad (15)$$

for each pair of weights i in the filter. Then at every image location for each pair of weights $a_i\chi$ is added to the first weight, and subtracted from the second. The a_i's expressed as a percentage of the smaller weight in the pair, are yet other adjustable parameters.

The error filter can also be perturbed by randomizing the *positions* of the weights.

The number of adjustable parameters available to modify the basic error diffusion algorithm is tremendously large. Much is to be learned about the effect of each parameter used independently and in combination with others. Over a hundred combinations of these parameters were experimented with in this investigation; the examples that follow are a carefully selected representative sample.

A. One Weight

The most computationally inexpensive form of the error diffusion algorithm is one implemented with an error filter with one weight. In this case all of the error is diffused to only one location; no multiplication is necessary.

If the one weight is fixed to some predetermined location, the resulting patterns fail in a big way. This is demonstrated in Fig. 23 where the position of the weight was fixed diagonally adjacent to the origin. Failure is at least as great for any other location choice, or when a serpentine raster is used.

However, when the *position* of the one weight is randomly determined over some finite set of candidate positions, a much more acceptable result emerges. The images of Fig. 24 where produced with the position of the one weight selected with equal probability between only two candidate locations, immediately below and preceding the filter origin.
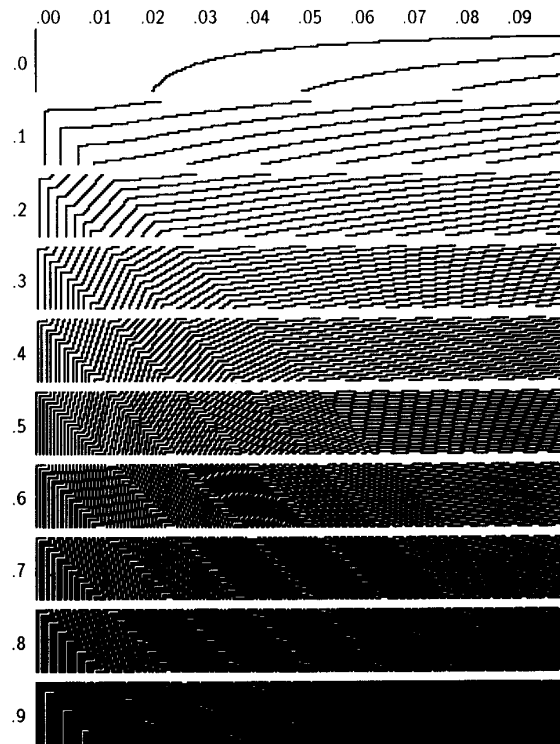

Fig. 23. Failure of one deterministic weight location.

This example represents a broad class of parameter combinations. Effectively the same output results for any local neighborhood of candidate locations, as well as for a broad range of probability mass functions governing the selection of a location. Several combinations were tried without change in output, including the four Floyd and Steinberg locations using the values of their weights as the probability mass function for position selection. Even when two or three weights were randomly selected, no significant difference was seen.

The radially averaged power spectrum is displayed in Fig. 25. The spectra for the various gray levels all reveal the desirable properties of

1) very low anisotropy,
2) flat blue noise region,
3) cutoff at f_g.

The only feature that the spectra fall short of is that of a sharp transition region. This is most pronounced at $g = \frac{1}{2}$. The low-frequency leakage is responsible for the grainy texture; the associated patterns can be called "light-blue noise." The suppression of frequencies below f_g is sufficient, however, to produce a vast improvement over the white noise images of Section III.

B. Two Weights

In this section, variations on the error filter shown in Fig. 26 will be considered. As might be expected, using this filter unperturbed yields unacceptable results as seen in Fig. 27. The strong diagonal texture patterns result in an extremely anisotropic power spectrum evidenced by Fig. 28.
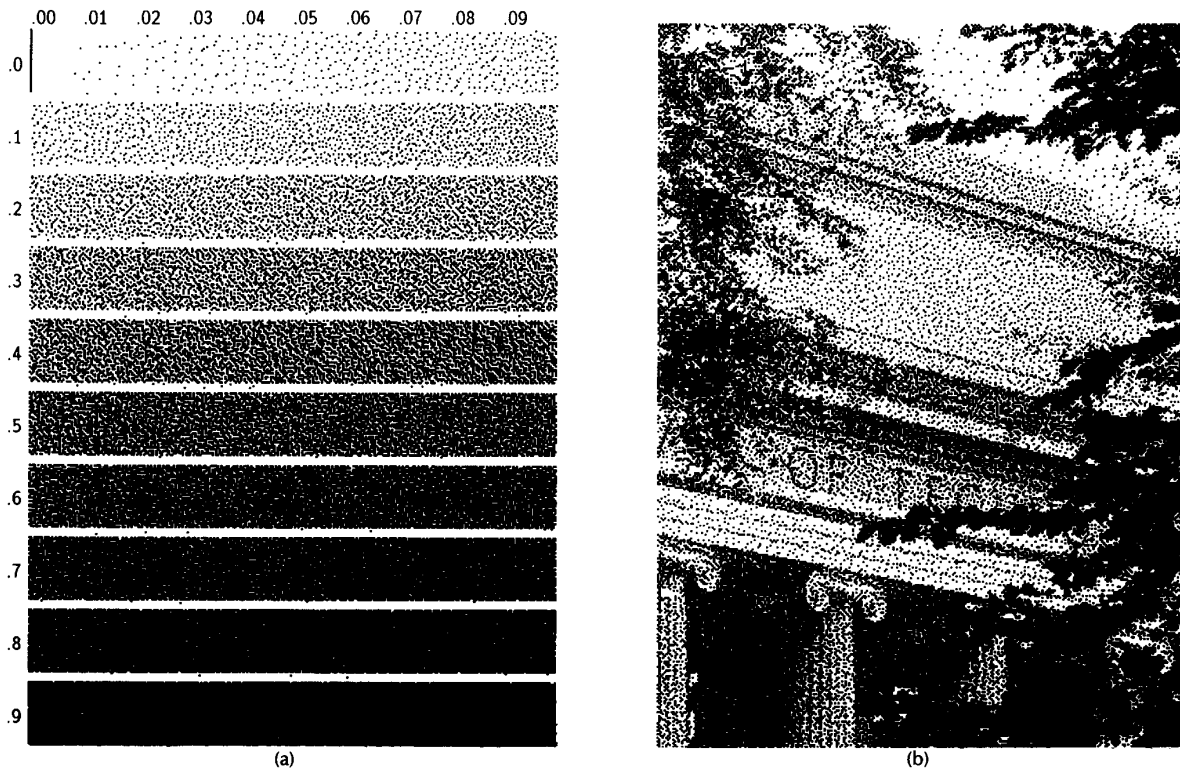
Fig. 24. Effect of one randomly positioned weight. (a) Gray-scale ramp. (b) Scanned picture.
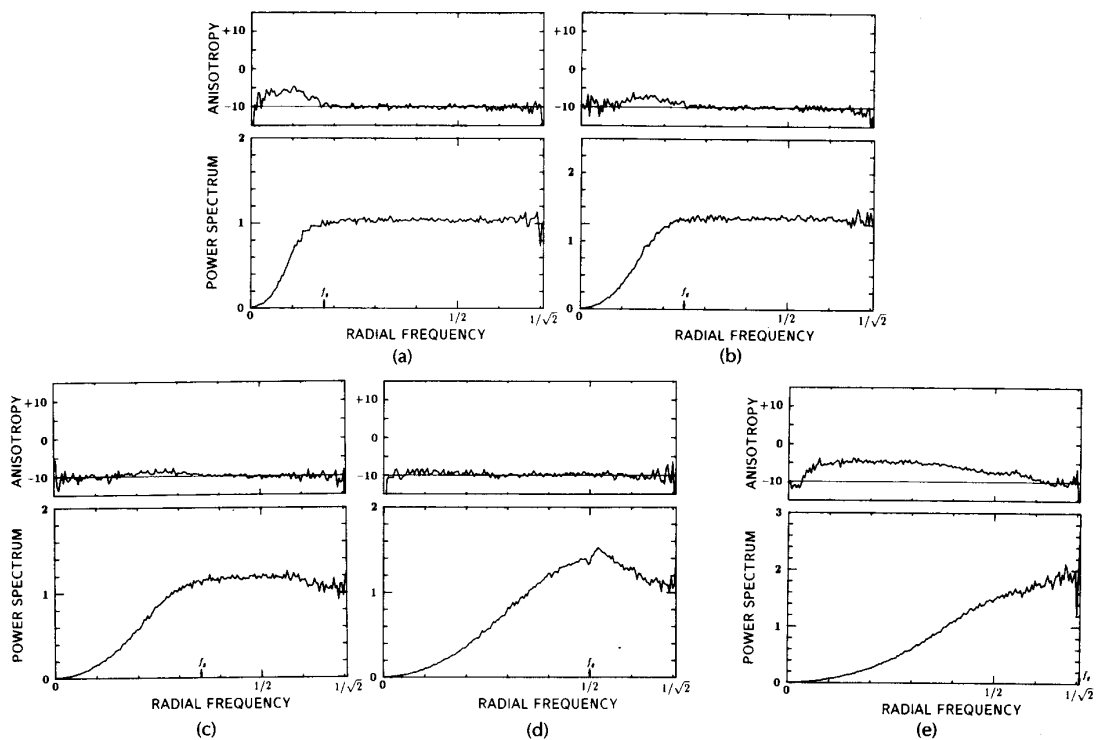
Fig. 25. Analysis of one randomly positioned weight. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f_r (in units of S^{-1}) for selected fixed gray levels. (a) $g = \frac{1}{32}$, $f_g/S^{-1} \approx 0.1768$, $\sigma_g^2 \approx 0.0303$. (b) $g = \frac{1}{16}$, $f_g/S^{-1} = 0.25$, $\sigma_g^2 \approx 0.0586$. (c) $g = \frac{1}{8}$, $f_g/S^{-1} \approx 0.3495$, $\sigma_g^2 \approx 0.1094$. (d) $g = \frac{1}{4}$, $f_g/S^{-1} = 0.5$, $\sigma_g^2 = 0.1875$. (e) $g = \frac{1}{2}$, $f_g/S^{-1} \approx 0.7071$, $\sigma_g^2 = 0.25$.

$$\left(\frac{1}{2}\times\right) \qquad \bullet \quad 1$$
$$1$$

Fig. 26. Deterministic part of a two-weight error filter.

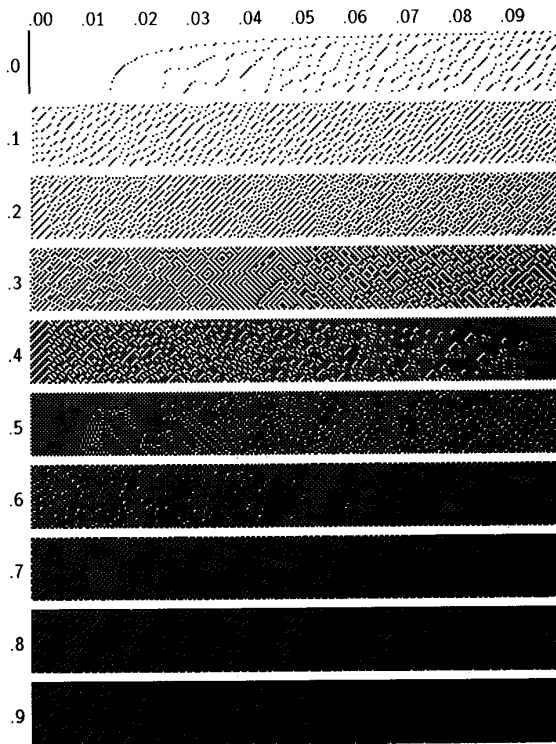**Fig. 27.** Effect of two deterministic weights.
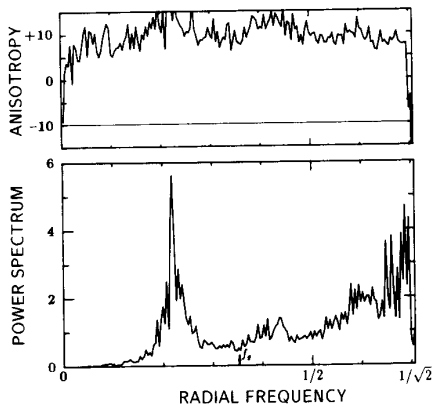
Fig. 28. Analysis of two deterministic weights. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f_r (in units of S^{-1}) for the fixed gray level $g = \frac{1}{8}$. $f_g/S^{-1} \approx 0.3495$, $\sigma_g^2 \approx 0.1094$.

The use of a serpentine raster corrects the directionality of these textures but still leaves many undesirable patterns. This is shown in Fig. 29.

These patterns can be largely reduced by adding the perturbation of 100-percent randomness to the two weights; a case that is particularly interesting because it simply requires the selection of a random number distributed
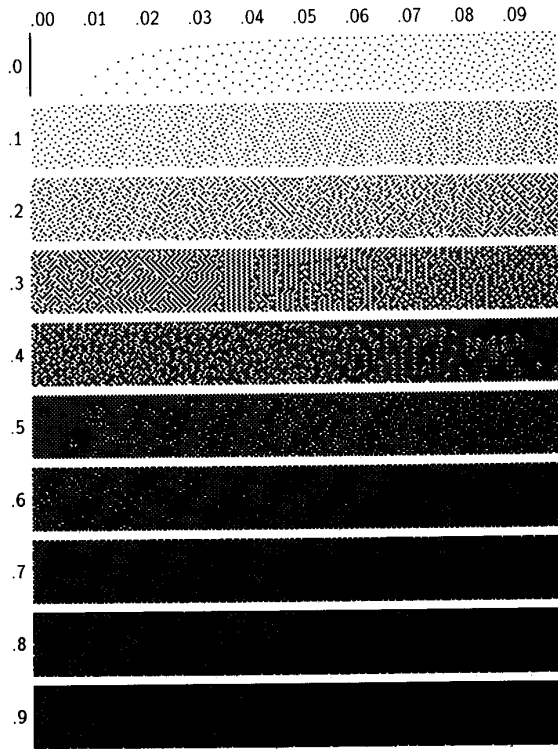
(a)

(b)

Fig. 29. Effect of two deterministic weights on a serpentine raster. (a) Gray-scale ramp. (b) Scanned picture.

between 0 and 1 for one weight, and its two's complement for the other. The very random nature of this approach, while making the patterns more isotropic, passed a bit too

```
     .00  .01  .02  .03  .04  .05  .06  .07  .08  .09
 .0
 .1
 .2
 .3
 .4
 .5
 .6
 .7
 .8
 .9
```
(a)

(b)

Fig. 30. Effect of two 50-percent random weights on a serpentine raster. (a) Gray-scale ramp. (b) Scanned picture.

much low-frequency content. A compromise is seen in Fig. 30 where the weights were perturbed with 50-percent noise. The well-behaved radially averaged power spectra results are shown in Fig. 31.

C. Four Weights

Although better than that which can be achieved with one weight, the transition regions for the two-weight case just examined were still not as steep as desired. Experiments with various choices of three weights did not produce a significant improvement.

In trying several combinations of deterministic values in a four-element error filter, none proved better than the famed filter of Floyd and Steinberg. In this section, two variations on this basic filter are presented, both processed with serpentine rasters. As in the two-weight case, processing on a serpentine raster without random perturbations (a case not shown for economy of space) will correct directional artifacts, but still produces some stable asymmetric structures. Fig. 32 shows the results of error diffusion holding the filter values constant but perturbing the threshold by 30 percent with white noise. The noisy threshold breaks up most remaining stable texture patterns yielding good radial symmetry at the expense of adding some low-frequency energy. The noise value of 30 percent in this example appeared to be the best compromise between low frequency graininess (100-percent threshold noise), and stable texture patterns (0-percent threshold noise).

In Fig. 33, instead of perturbing the threshold, noise was added to the weights. For this purpose, the two larger weights ($\frac{7}{16}$ and $\frac{5}{16}$) and the two smaller weights ($\frac{3}{16}$ and $\frac{1}{16}$) were paired together. To prevent weights with negative values, the maximum noise amplitude (100 percent) is the value of the smaller weight in each pair. In this case, the value of 50-percent noise added to each pair appeared to optimize the tradeoff between graininess and stable textures. Adding noise in this way succeeded at balancing this compromise better than the noisy threshold method.

The radially averaged power spectra in Fig. 34 reveal a good blue noise process. Along with low anisotropy and flat high-frequency regions, the extra number of weights provided additional low-frequency inhibition and steeper transition regions.

Here, as in many of the other cases considered, the most disturbing patterns evolve near $g = \frac{1}{2}$. Unless the power spectra contain impulses at the corners of the frequency baseband, perturbations to this perfect spatial checkerboard pattern, no matter how slight, become very visible. Rendering gray levels near this value is perhaps the greatest weakness of the error diffusion algorithm and variations on it.

If plots of $P_r(f_r)$ at several other gray levels could be added to Fig. 34 and shown in sequence, the form of the plot would look like a wave beginning at $f_r = 0$ at $g = 0$ moving to the right as g increased until it hit the "high-frequency wall" at $g = \frac{1}{2}$. Then as g increased from $\frac{1}{2}$ to 1, the "wave" would retreat in a symmetric fashion.

D. Hexagonal Case

After much experiment it was found that the most successful methods of halftoning by error diffusion on rectangular grids work quite well on hexagonal grids after adjusting the error filters slightly. Fig. 35 displays the deterministic part of the two- and four-element filter that will be demonstrated in this section.

As in the rectangular case, processing on a serpentine raster corrects directional artifacts, but some random perturbation is needed to break up those stable patterns that

(a)

(b)

(c)

(d)

(e)

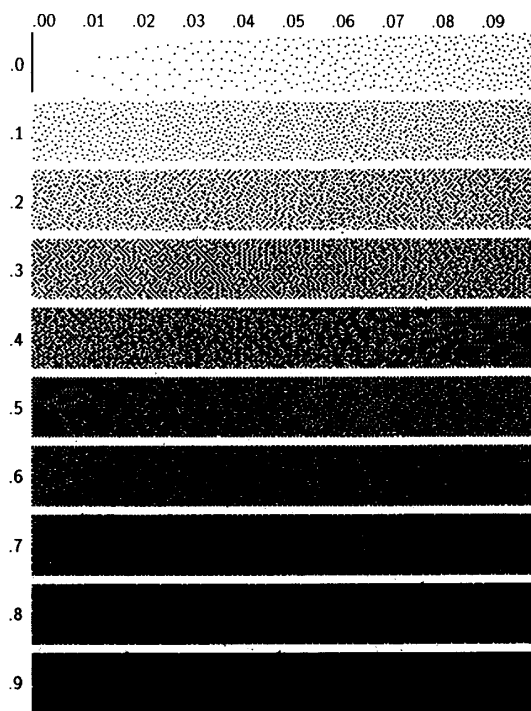**Fig. 31.** Analysis of two 50-percent random weights on a serpentine raster. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f_r (in units of S^{-1}) for selected fixed gray levels. (a) $g = \frac{1}{32}$, $f_g/S^{-1} \approx 0.1768$, $\sigma_g^2 \approx 0.0303$. (b) $g = \frac{1}{16}$, $f_g/S^{-1} = 0.25$, $\sigma_g^2 \approx 0.586$. (c) $g = \frac{1}{8}$, $f_g/S^{-1} \approx 0.3495$, $\sigma_g^2 \approx 0.1094$. (d) $g = \frac{1}{4}$, $f_g/S^{-1} = 0.5$, $\sigma_g^2 = 0.1875$. (e) $g = \frac{1}{2}$, $f_g/S^{-1} \approx 0.7071$, $\sigma_g^2 = 0.25$.

(a)

(b)
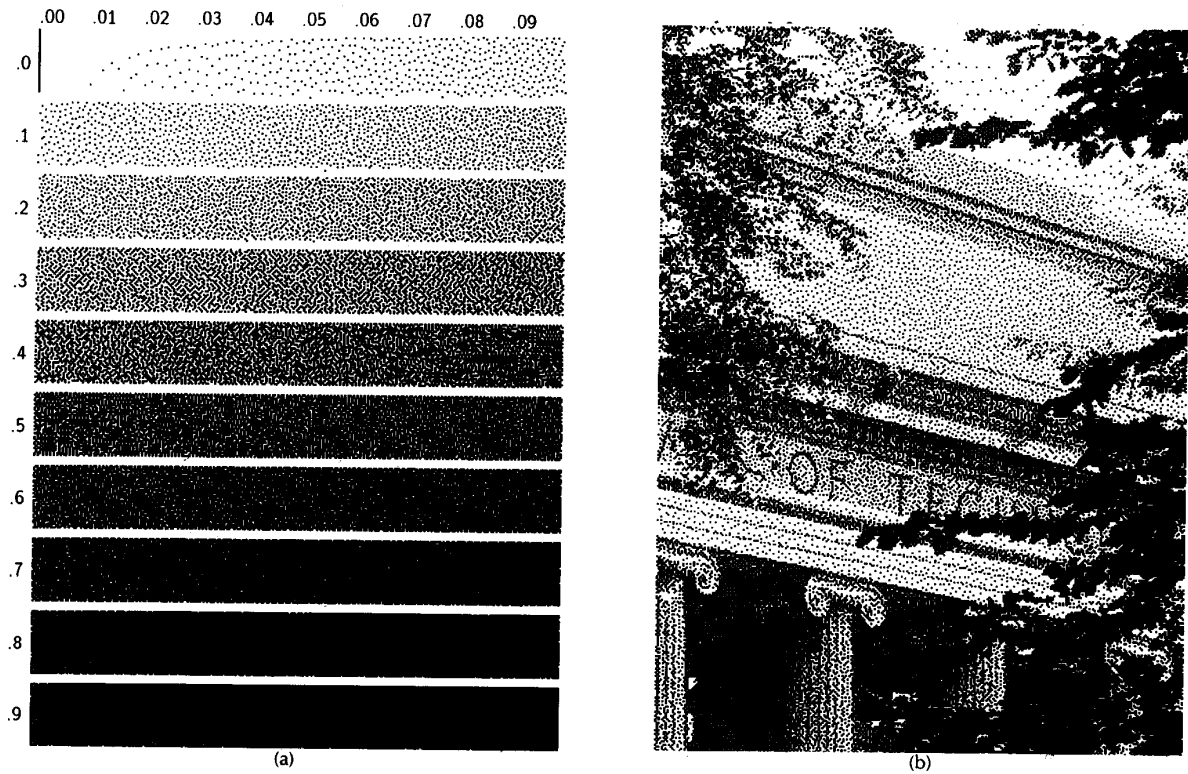
Fig. 32. Floyd and Steinberg filter with a 30-percent random threshold processed on a serpentine raster. (a) Gray-scale ramp. (b) Scanned picture.
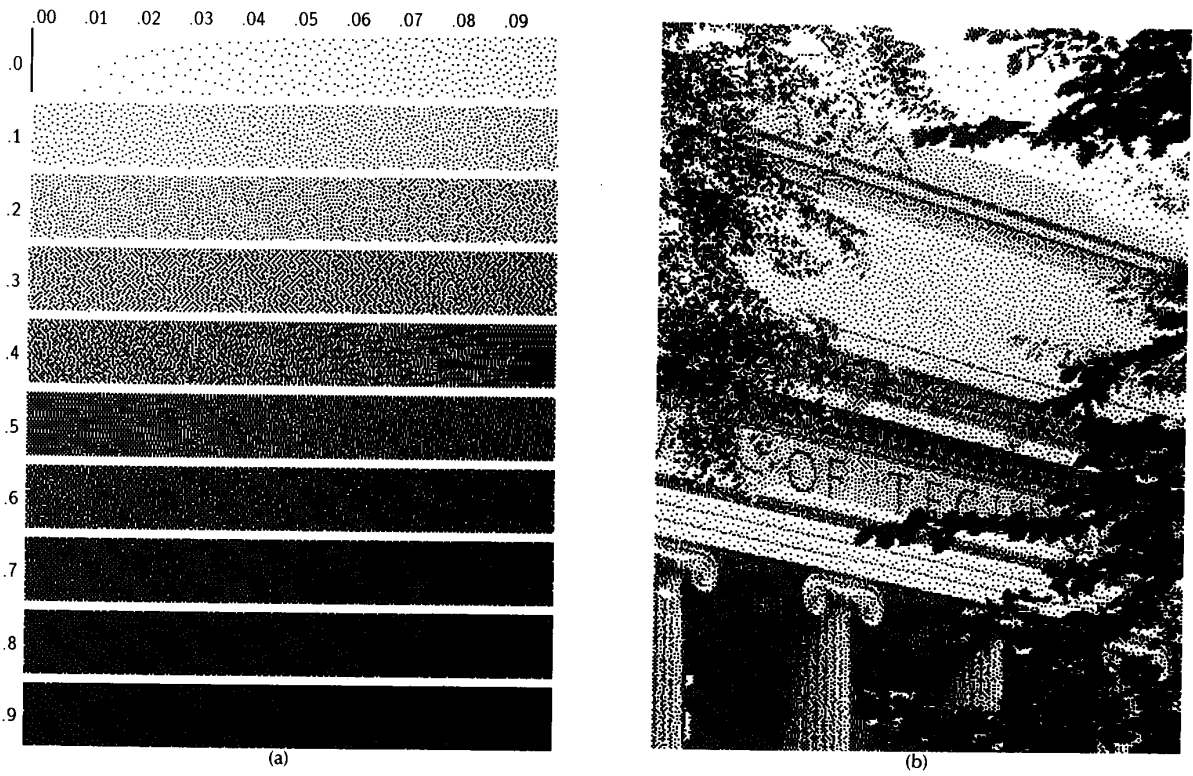
Fig. 33. Floyd and Steinberg filter with 50-percent random weights processed on a serpentine raster. (a) Gray-scale ramp. (b) Scanned picture.
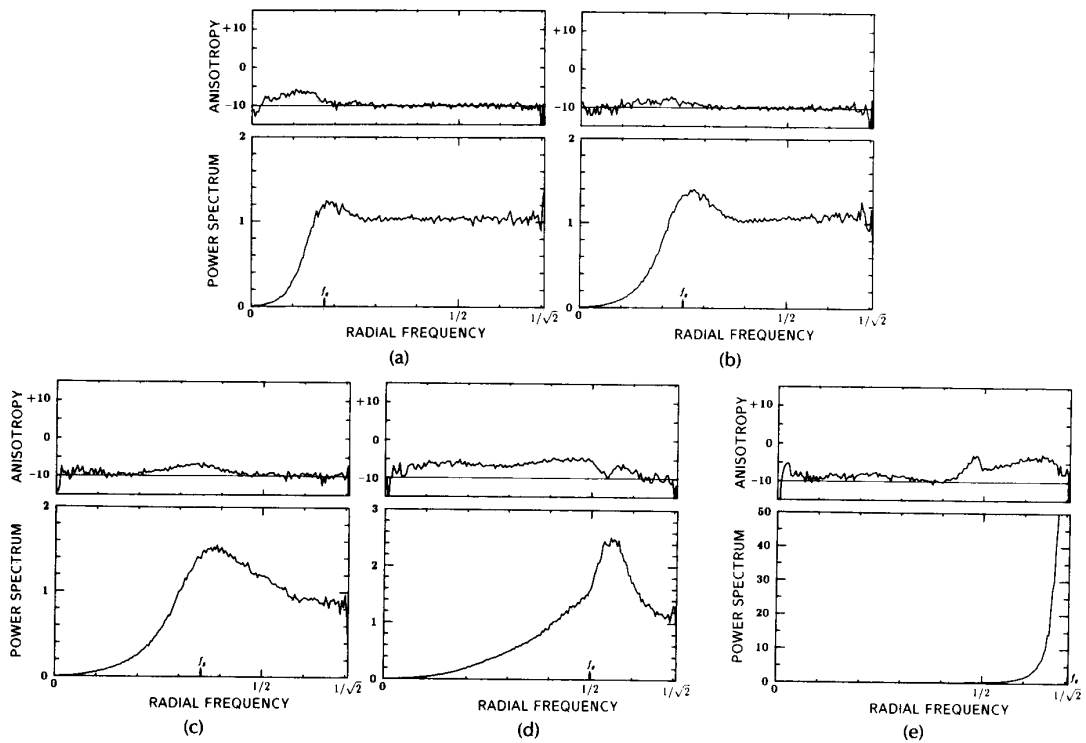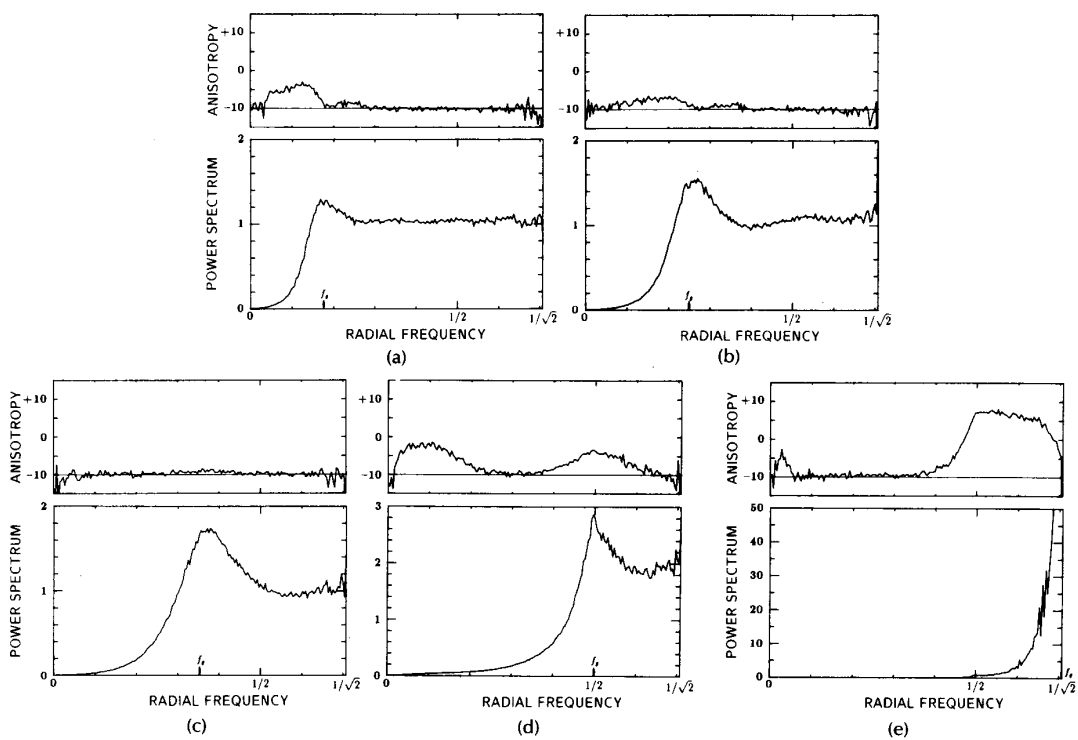

Fig. 34. Analysis of the Floyd and Steinberg filter with 50-percent random weights processed on a serpentine raster. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f_r (in units of S^{-1}) for selected fixed gray levels. Processed on a serpentine raster. (a) $g = \frac{1}{32}$, $f_g/S^{-1} \approx 0.1768$, $\sigma_g^2 \approx 0.0303$. (b) $g = \frac{1}{16}$, $f_g/S^{-1} = 0.25$, $\sigma_g^2 \approx 0.0586$. (c) $g = \frac{1}{8}$, $f_g/S^{-1} \approx 0.3495$, $\sigma_g^2 \approx 0.1094$. (d) $g = \frac{1}{4}$, $f_g/S^{-1} = 0.5$, $\sigma_g^2 = 0.1875$. (e) $g = \frac{1}{2}$, $f_g/S^{-1} \approx 0.7071$, $\sigma_g^2 = 0.25$.

$$\left(\frac{1}{2}\times\right) \qquad\qquad \bullet \;\; 1$$
$$\qquad\qquad\qquad\qquad\qquad 1$$

(a)

$$\left(\frac{1}{16}\times\right) \qquad\qquad \bullet \;\; 7$$
$$\qquad\qquad\qquad\qquad 3 \;\; 5 \;\; 1$$

(b)

Fig. 35. Two hexagonal error filters to be perturbed. (a) Two weights. (b) Variant of Floyd and Steinberg filter. ("•" represents the origin.)

remain. Again, for economy of space (a concern that may at this point seem moot given the length of this paper) only those examples which illustrate the most successful compromise between low-frequency graininess and stable texture patterns are shown.

Parallel to the methods used for Figs. 30 and 33, the hexagonal filters will have 50-percent noise added to their weights, and processing will be on a serpentine raster. The result of error diffusion with the two-element filter is illustrated in Fig. 36; the result of the four-element filter is shown in Fig. 37.

Both render gray levels in a much more isotropic manner than that of Fig. 20. The large 12-element filter did sharpen the picture more, however. Again, as was argued in the rectangular case, sharpening should not be compounded with halftoning. In the next section, an example is given of how precisely controlled sharpening can be achieved in a separate operation for both rectangular and hexagonal grids.

The radially averaged power spectrum of both the two- and four-weight stochastic hexagonal filters are well behaved with the four-weight case having sharper transition regions. So this case is included in Fig. 38. Notice how well the peak of these spectra follows the principal frequency (marked as usual with a diamond on the frequency axis). The important exception is for $g = \frac{1}{2}$ which is in the region where the hexagonal grid cannot support a principal frequency.

Two other particularly interesting cases are those for $g = \frac{1}{3}$ and $g = \frac{2}{3}$. The principal frequency for these gray values is at the high-frequency limit for a hexagonal grid $f_g = \frac{2}{3}$. These cases are similar to the $g = \frac{1}{2}$ case for rectangular grids.

In balance, anisotropy is generally low for the stochastic hexagonal filter, not nearly as wild as in Fig. 21. Outside of the range $\frac{1}{3} < g < \frac{2}{3}$, such filters are also good blue noise generators.

The wave analogy given for rectangular grid blue-noise spectra can also be used for hexagonal grids. A sequence of plots of $P_r(f_r)$ for several gray levels would show a "wave" beginning at $f_r = 0$ at $g = 0$ moving to the right as g increased until it hit the "high-frequency wall" at $g = \frac{1}{3}$. Skipping to $g = \frac{2}{3}$, the "wave" would retreat in a symmetric fashion as g increased from $\frac{2}{3}$ to 1.

VII. Concluding Remarks

A. Sharpening

The improved output perceived from a halftoning method that intrinsically sharpens can misleadingly outweigh other shortcomings in its ability to render gray levels accurately and without algorithmic artifacts. The virtues of a halftoning scheme should be decoupled from its ability to sharpen.

Sharpening does improve, or at least defeat, unsharp-

ening degradations that halftoning imparts. The proper degree of sharpening is a subjective quality and can easily be controlled independently of halftoning. Sharpening can be combined with interpolation for the case of digital enlarging. The sharpened Gaussian [27], [28] is often used for such a dual purpose.

When sharpening is not combined with resampling, a separate high-pass filter operation is needed. Perhaps the most popular high-pass filter used for the purpose of sharpening is the digital Laplacian. The Laplacian ∇^2 is an operator which produces the second spatial derivative

$$\nabla^2 J(x) = \frac{\partial^2 J(x)}{\partial x_1^2} + \frac{\partial^2 J(x)}{\partial x_2^2}. \tag{16}$$

When applied to an image, it produces large amplitudes at edge locations, and zero in constant or uniformly varying regions (regions where the zeroth or first derivative is zero).

This operator can be described as convolution with the filter $\nabla^2 \delta(x)$. The discrete-space analog to this on a rectangular grid is the five-element filter shown in Fig. 39(a), although the nine-element variety achieves a similar effect. A hexagonal version is shown in Fig. 39(b).

Denoting a digital Laplacian filter as $\Psi[n]$, sharpening is achieved by subtracting the Laplacian filtered image from the original image

$$J_{\text{sharp}}[n] = J[n] - \beta\Psi[n] * J[n]. \tag{17}$$

The amount of sharpening is controlled by the value of $\beta \geq 0$. A tradeoff must be made between the accentuation of edge detail and amplification of noise.

An example of the effect of presharpening in this way on a rectangular grid with $\beta = 2.0$ is displayed in Fig. 40. It should be compared to that in Fig. 33(b) which was identically halftoned without presharpening.

Adaptive sharpening techniques exist which are not as sensitive to noise but are, as one might expect, more compute-intensive.

B. Blue Noise is Pleasant

Fig. 41 collectively compares greatly enlarged portions of the four major classes of dither patterns arranged in order of increasing correlation (decreasing entropy). All patterns are representations of a fixed gray level $g = \frac{1}{9}$, and thus all have roughly the same number of black pixels.

While white noise appears too random or "noisy," ordered dither appears "structured." The purpose of a dither pattern is to represent a continuous-tone level. It therefore should not have any form or structure of its own; a pattern succeeds when it is innocuous. Blue noise is visually pleasant because it does not clash with the structure of an image by adding one of its own or degrade it by being too "noisy" or uncorrelated.

Blue noise even defies the structure of the underlying grid. Even though the dots in Fig. 41(b) are perfect squares, each precisely aligned to a given position on a rectangular grid, the collective ensemble tends to destroy this rigorous alignment creating what can be called a *grid defiance illusion*.

For many years, noise with $1/f$ power spectrum distributions have been known to exist in electrical systems. But recently, discoveries have repeatedly confirmed the existence of a $1/f$ power spectrum in almost every aspect of

Fig. 36. Effect of two 50-percent random weights on a hexagonal grid processed on a serpentine raster. (a) Gray-scale ramp. (b) Scanned picture.
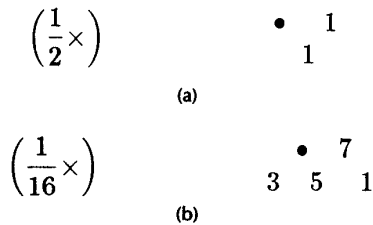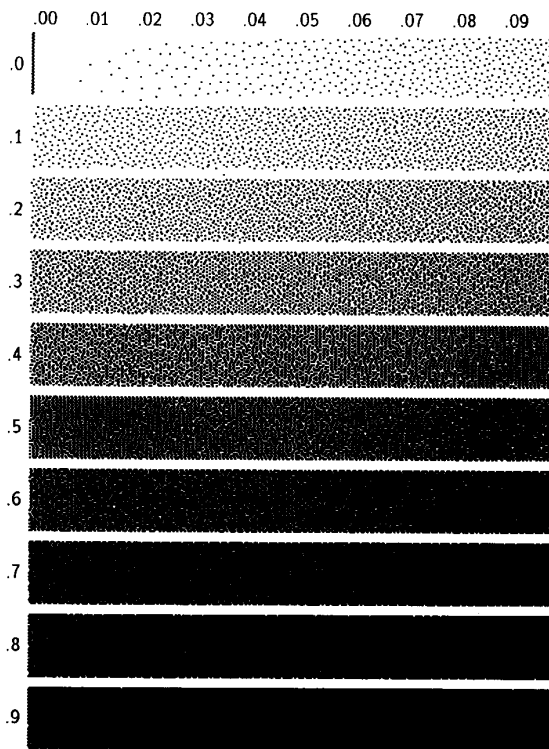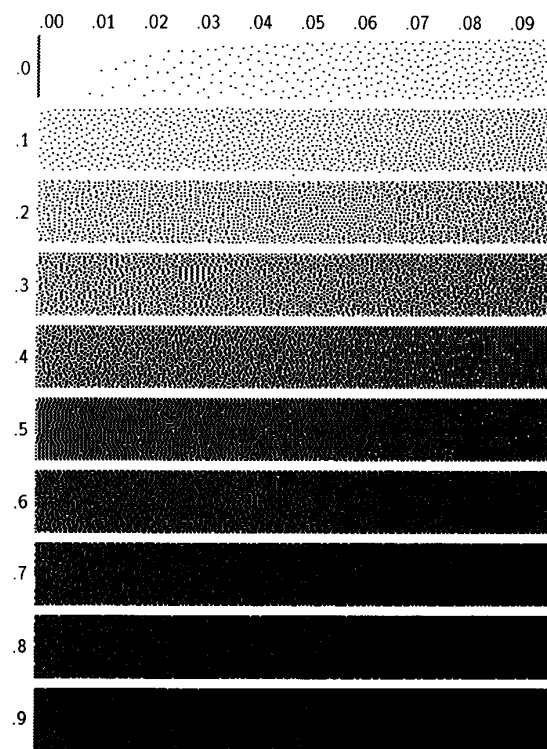

Fig. 37. Effect of four 50-percent random weights on a hexagonal grid processed on a serpentine raster. (a) Gray-scale ramp. (b) Scanned picture.
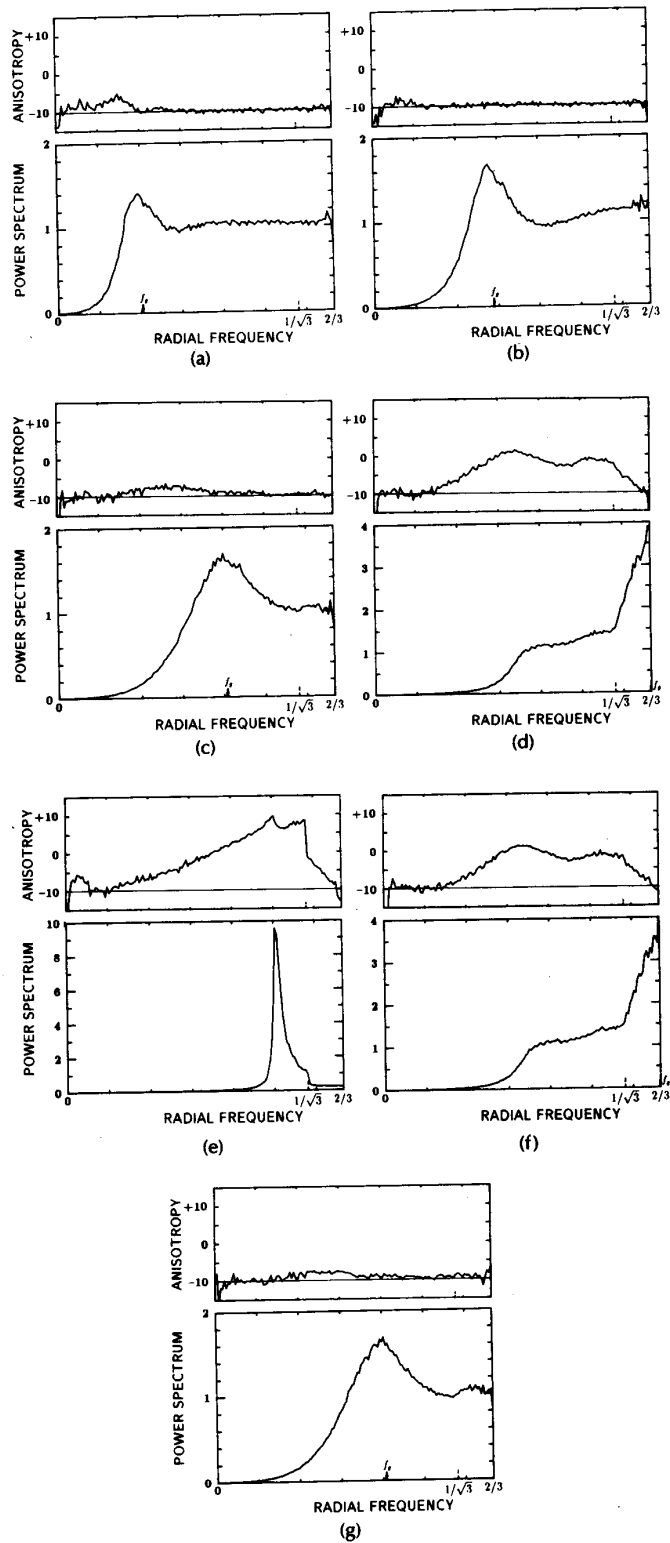
Fig. 38. Analysis of the stochastic hexagonal filter. Radial spectra (in units of σ_g^2), and anisotropy (in dB) as a function f_r (in units of S^{-1}) for selected fixed gray levels. (a) $g = \frac{1}{32}$, $f_g/S^{-1} \approx 0.2041$, $\sigma_g^2 \approx 0.0303$. (b) $g = \frac{1}{16}$, $f_g/S^{-1} \approx 0.2887$, $\sigma_g^2 \approx 0.0586$. (c) $g = \frac{1}{8}$, $f_g/S^{-1} \approx 0.4082$, $\sigma_g^2 \approx 0.1094$. (d) $g = \frac{1}{3}$, $f_g/S^{-1} = \frac{2}{3}$, $\sigma_g^2 \approx 0.1111$. (e) $g = \frac{1}{2}$ (f_g is not supported), $\sigma_g^2 = 0.25$. (f) $g = \frac{2}{3}$, $f_g/S^{-1} = \frac{2}{3}$, $\sigma_g^2 \approx 0.1111$. (g) $g = \frac{7}{8}$, $f_g/S^{-1} \approx 0.4082$, $\sigma_g^2 \approx 0.1094$.

$$\begin{array}{ccc} & \frac{1}{4} & \\ \frac{1}{4} & -1 & \frac{1}{4} \\ & \frac{1}{4} & \end{array} \quad \text{or} \quad \begin{array}{ccc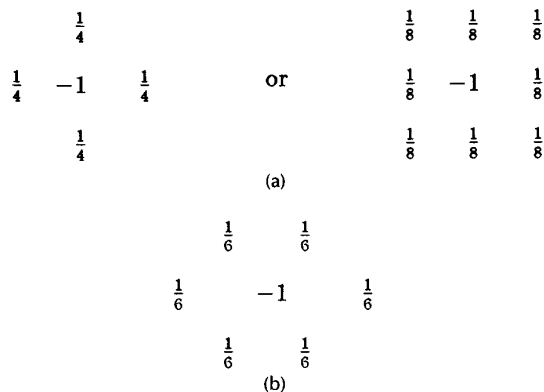} \frac{1}{8} & \frac{1}{8} & \frac{1}{8} \\ \frac{1}{8} & -1 & \frac{1}{8} \\ \frac{1}{8} & \frac{1}{8} & \frac{1}{8} \end{array}$$

(a)

$$\begin{array}{ccc} \frac{1}{6} & & \frac{1}{6} \\ \frac{1}{6} & -1 & \frac{1}{6} \\ \frac{1}{6} & & \frac{1}{6} \end{array}$$

(b)

Fig. 39. Digital "Laplacian" filters $\Psi[n]$. (a) Rectangular. (b) Hexagonal.

Fig. 40. Sharpening with a rectangular Laplacian filter. Laplacian amplitude $\beta = 2.0$.

(a)

(b)

(c)

(d)

Fig. 41. Comparison of dither patterns for a fixed gray level at $g = \frac{1}{9}$, arranged in order of decreasing entropy. (a) White noise. (b) Blue noise. (c) Dispersed-dot ordered dither. (d) Clustered-dot ordered dither.

nature [13], [20], [37] including such things as variations in sunspots, wobbling of the earth's axis, and flood levels of the River Nile. Evidence of $1/f$ fluctuations in human biological systems [22] has also been found; a $1/f$ spectrum was found in electroencephalogram (brain wave) measurements when subjects were exposed to "pleasing" stimuli.

A study by Richard Voss [35], [36] has found that practically all forms of music possess $1/f$ noise. Experiments with stochastic music composition revealed that listeners found $1/f$ music far more interesting than white ($1/f^0$) music, described as "too random," or brown ($1/f^2$) music, described as "too correlated."

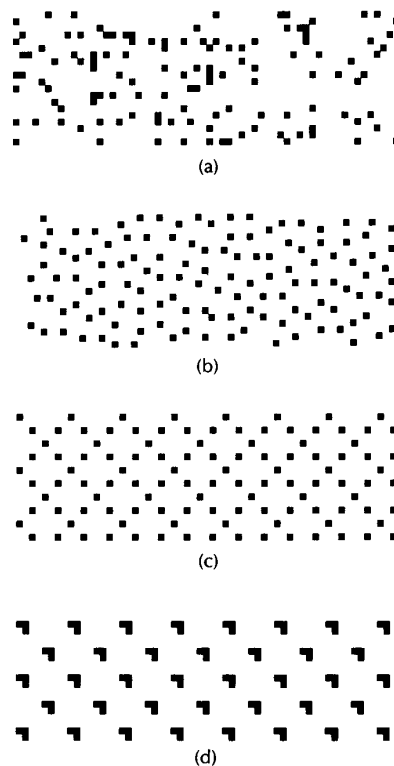Blue noise can be described as the "pleasing" complement of $1/f$ noise. The dominance of low frequencies in

$1/f$ phenomenon is responsible for its interesting and natural structure. Blue noise, by contrast, is not "interesting"; nor is it annoying. Being devoid of low frequencies and localized concentrations of spikes in the frequency domain, it has no structure and thus does not interfere with the interesting features of that which it is representing.

C. Summary

It was shown that well-formed dither patterns have radially symmetric power spectra that look like white noise at high frequencies, with minimal energy at low frequencies. This type of spectra are referred to as "blue noise." Such patterns appear random without low-frequency texture; in fact, upon examining them, dots are perceptually free of any alignment or structure in what can be called a grid defiance illusion.

From the concept of a principal wavelength, it was argued that rectangular grids are superior to hexagonal grids for generating blue-noise patterns; while hexagonal grids are more radially symmetric, square grids can support higher spatial frequencies.

New metrics for analyzing aperiodic dither patterns were introduced: the radially averaged power spectra for fixed gray levels and a measure of its anisotropy. It was shown that fairly good blue-noise patterns can be generated with versions of the perturbed error diffusion algorithm with small error filters. Perhaps other methods can be found that exhibit even better spectral characteristics for blue noise dithering.

REFERENCES

[1] J. P. Allebach and B. Liu, "Random quasi-periodic halftone process," *J. Opt. Soc. Amer.*, vol. 66, pp. 909-917, 1976.

[2] ——, "Analysis of halftone dot profile and aliasing in the discrete binary representation on images," *J. Opt. Soc. Amer.*, vol. 67, pp. 1147-1154, 1977.

[3] J. P. Allebach, "Random nucleated halftone screening," *Photogr. Sci. Eng.*, vol. 22, no. 2, pp. 89-89, 1978.

[4] M. S. Bartlett, *An Introduction to Stochastic Processes with Special Reference to Methods and Applications.* New York, NY: Cambridge Univ. Press, 1955, pp. 274-284.

[5] B. E. Bayer, "An optimum method for two level rendition of continuous-tone pictures," in *Proc. IEEE Int. Conf. on Communications, Conf. Rec.*, pp. (26-11)-(26-15), 1973.

[6] C. Billotet-Hoffman and O. Bryngdahl, "On the error diffusion technique for electronic halftoning," *Proc. SID*, vol. 24, pp. 253-258, 1983.

[7] M. A. Dippe and E. H. Wold, "Antialiasing through stochastic sampling," *Computer Graphics* (AMC SIGGRAPH'85 Conf. Proc.), vol. 19, no. 3, pp. 69-78, 1985.

[8] D. E. Dudgeon and R. M. Mersereau, *Multidimensional Digital Signal Processing.* Englewood Cliffs, NJ: Prentice-Hall, 1984, pp. 39-41.

[9] R. W. Floyd and L. Steinberg, "Adaptive algorithm for spatial grey scale," in *SID Int. Sym. Dig. Tech. Papers*, pp. 36-37, 1975.

[10] ——, "An adaptive algorithm for spatial greyscale," *Proc. SID*, vol. 17, no. 2, pp. 75-77, 1976.

[11] J. E. Freund, *Mathematical Statistics.* Englewood, NJ: Prentice-Hall, 1971.

[12] W. Gall, K. Wellendorf, and K. Kiel, "Production of screen printing blocks," U.S. Patent 4 499 489, 1985.

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

[14] T. Gardos, G. R. Arce, and J. P. Allebach, "The optimal ordered dither cell for the binary representation of continuous tone images on a hexagonal lattice," in *Proc. of the 1986 Conf. on Information Sciences and Systems* (Princeton Univ., Princeton, NJ, Mar. 1986).

[15] W. M. Goodall, "Television by pulse code modulation," *Bell Syst. Tech. J.*, vol. 30, pp. 33-49, 1951.

[16] E. Hecht and A. Zajac, *Optics.* Reading, MA: Addison-Wesley, 1974, pp. 361-363.

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

[18] D. Kermisch and P. G. Roetling, "Fourier spectrum of halftone images," *J. Opt. Soc. Amer.*, vol. 65, pp. 716-723, 1975.

[19] D. E. Knuth, *The Art of Computer Programming*, vol. 2. Reading, MA: Addison-Wesley, 1981, ch. 3, "Random numbers."

[20] B. Mandelbrot and R. Voss, "Why is nature fractal and when should noises be scaling?" *Noise in Physical Systems and 1/f Noise.* New York, NY: North-Holland Physics Publ., 1983, pp. 31-39.

[21] MIT Museum, Photo: "Building 10," contact print from 5×7 glass negative, c. 1929.

[22] T. Musha, "1/f fluctuations in biological systems," in *Proc. 3rd Symp. on 1/f Fluctuations*, pp. 143-146, 1981.

[23] A. V. Oppenheim and R. W. Schafer, *Digital Signal Processing.* Englewood Cliffs, NJ: Prentice-Hall, 1975, pp. 548-549.

[24] L. G. Roberts, "Picture coding using pseudo-random noise," *IRE Trans. Informat. Theory*, vol. IT-8, pp. 145-154, 1962.

[25] A. H. Robinson, "Multidimensional Fourier transforms and image processing with finite scanning apertures," *Appl. Opt.*, vol. 12, no. 10, pp. 2344-2352, 1973.

[26] G. Rosenfeld, "Screened image reproduction," U.S. Patent 4 456 924, 1984.

[27] W. F. Schreiber and D. E. Troxel, "Transformation between continuous and discrete representation of images: A perceptual approach," *IEEE Trans. Pattern Anal.*, vol. PAMI-7, no. 2, pp. 178-186, 1985.

[28] W. F. Schreiber, *Fundamentals of Electronic Imaging Systems: Some Aspects of Image Processing.* New York: NY: Springer-Verlag, 1986.

[29] R. L. Stevenson and G. R. Arce, "Binary display of hexagonally sampled continuous-tone images," *J. Opt. Soc. Amer. A*, vol. 2, no. 7, pp. 1009-1013, 1985.

[30] J. C. Stoffel and J. F. Moreland, "A survey of electronic techniques for pictorial reproduction," *IEEE Trans. Commun.*, vol. COM-29, pp. 1898-1925, 1981.

[31] P. Stucki, "Image processing for document reproduction," in *Advances in Digital Image Processing.* New York, NY: Plenum, 1979, pp. 177-218.

[32] ——, "MECCA—A multiple-error correcting computation algorithm for bilevel image hardcopy reproduction," Res. Rep. RZ1060, IBM Res. Lab., Zurich, Switzerland, 1981.

[33] R. Ulichney, "Generalized ordered dither," MIT, Tech. Rep. ATRP-T51, 1985. Also Digital Equipment Corporation, Maynard, MA, Tech. Rep. DEC-TR-412, 1985.

[34] ——, *Digital Halftoning.* Cambridge, MA: MIT Press, 1987.

[35] R. F. Voss and J. Clarke, "'1/f noise' in music and speech," *Nature*, vol. 258, no. 5533, pp. 317-318, Nov. 27, 1975.

[36] ——, "'1/f noise' in music: music from 1/f noise," *J. Acoust. Soc. Amer.*, vol. 63, no. 1, pp. 258-263, 1978.

[37] R. F. Voss "1/f (flicker) noise: A brief review," in *Proc. 33rd Annu. Symp. on Frequency Control*, pp. 40-46 (May 30-June 1, 1979).

[38] I. H. Witten and M. Neal, "Using peano curves for bilevel display of continuous-tone images," in *Proc. IEEE CG&A*, pp. 47-52, May 1982.

[39] B. Woo, "A survey of halftoning algorithms and investigation of the error diffusion technique," MIT, S.B. thesis, 1984.

Robert A. Ulichney received the B.S. degree in physics and computer science from the University of Dayton, Ohio, in 1976 and the M.S., E.E., and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, in electrical engineering and computer science, in 1979, 1984, and 1986, respectively.

For nine years, he has been conducting image processing research for Digital Equipment Corporation. He has written a book and several papers in this area, and serves as the technical sponsor for related projects at the University of Rochester, Rochester, NY, and the University of California, Davis. He is currently with DEC's Image Processing Research Group in Hudson, MA, and a Research Affiliate in the Research Laboratory of Electronics at MIT.