

Digital halftoning technique using a blue-noise mask

Theophano Mitsa

Department of Electrical and Computer Engineering, University of Iowa, Iowa City, Iowa 52242

Kevin J. Parker

Department of Electrical Engineering, University of Rochester, Rochester, New York 14627

Received December 23, 1991; revised manuscript received May 11, 1992; accepted May 12, 1992

A novel digital halftoning technique, by which the halftoning is achieved by a pixelwise comparison of the gray-scale image to an array (halftone screen), the blue-noise mask, is presented. This mask is designed so that the halftone image has blue-noise (high-frequency) characteristics in the frequency domain. The algorithm for the construction of the blue-noise mask and an algorithm for the construction of binary patterns with the same first-order but different second-order statistics are presented. Two psychovisual tests in which human subjects rated halftone patterns and images according to various criteria are also described.

1. INTRODUCTION

Many printing devices and displays, such as facsimile machines and laser printers, are bilevel and therefore are not capable of reproducing gray-scale images. Gray-scale images are converted to binary images by using halftone processes. The two most widely used halftone processes are ordered dither and error diffusion. Reviews of halftone processes can be found in Refs. 1 and 2.

For ordered dither the halftoning is achieved by a simple pixelwise comparison of the gray-scale image to a periodic array (halftone screen). Dispersed-dot ordered dither occurs when the halftone dots are of a fixed size, and clustered-dot ordered dither occurs when the halftone dots are of variable sizes. In conventional error diffusion³ every pixel of the gray-scale image is compared with a fixed threshold, and the resulting binary value is subtracted from the original value of that pixel. The resulting difference (error), after being multiplied by appropriate weights, is diffused into a given size neighborhood.⁴ Different error-diffusion schemes exist according to the position, number, and value of weights in the error filter.⁵⁻⁸ Error-diffusion techniques produce binary patterns with blue-noise (high-frequency) characteristics. In the image domain blue-noise characteristics correspond to visually pleasing, isotropic, uncorrelated patterns, without low-frequency graininess. Error-diffusion and dispersed-dot dither methods are suitable for printers that can accurately reproduce single black or white pixels, whereas, if single-dot reproduction is not possible, clustered-dot dither is the preferred choice.

In this paper we present a novel digital halftone process in which the halftoning is achieved by a pixelwise comparison of the gray-scale image with an array, the blue-noise mask. The blue-noise mask is constructed to have specific first and second-order properties.^{9,10} When the mask is thresholded at any level, for example at $A\%$ of the maximum level, exactly A of every 100 pixels will have a value greater than the threshold value. In addition, the spatial distribution of black versus white pixels will be arranged in such a manner as to form a blue-noise pattern.

In Section 2 the effect of halftoning with ordered dither and error diffusion in the Fourier domain is described, and the relationship of these techniques to the blue-noise mask is presented. In Section 3 we describe the algorithm for the construction of the blue-noise mask. In Section 4 a power-spectrum-matching algorithm that creates binary patterns with the same first-order but different second-order statistics is described. Section 4 also describes a visual study in which human subjects rated binary patterns created with the aforementioned power-spectrum-matching algorithm according to their isotropy. The purpose of this study was to determine how variations in the second-order statistics of a binary pattern affect its visual appearance. In another visual study, described in Section 5, the quality of images halftoned by using the blue-noise mask and also by using ordered dither and error diffusion was evaluated by human subjects. This psychovisual test indicates that halftoning using a blue-noise mask outperforms ordered-dither techniques and gives results comparable with those produced by halftoning with error diffusion.

2. PRINCIPLES OF THE BLUE-NOISE MASK AND CONNECTION WITH ERROR DIFFUSION AND ORDERED DITHER

Ordered-dither and error-diffusion techniques are of particular importance to the development of the blue-noise mask, since the blue-noise mask has the structure of an ordered-dither mask and the blue-noise-producing properties of error diffusion. Ordered dither and error diffusion are different in that, when it is applied to uniform (flat) images, ordered dither produces locally periodic patterns, whereas error diffusion, in general, produces locally aperiodic patterns. Because of the different natures of the patterns, the Fourier transform is being used to study the patterns produced by ordered dither,¹¹ and radially averaged power spectra are used to study the patterns produced by error diffusion.⁴ In Subsections 2.A and 2.B we give a brief overview of the Fourier-domain effects of

halftoning with error diffusion and ordered dither, and this overview enables us to establish in Subsection 2.C the connection between the blue-noise mask and its two parent halftone processes.

A. Ordered Dither

The binary pattern that results after the ordered-dither array is thresholded at a constant level is called the dot profile¹¹ for that level. Kermisch and Roetling¹² derived an expression for the Fourier transform of a halftone image in terms of the Fourier transform of the original image. Allebach¹¹ presented an analysis of the dot profile for the ordered-dither case and described its relationship to the Fourier transform of the halftoned image.

A halftone image $h(x, y)$ with size $K \times K$ can be written in terms of the samples f_{ij} of the input gray-scale image $f(x, y)$ as¹¹

$$h(x, y) = \sum_{i=0}^{K-1} \sum_{j=0}^{K-1} p[i, j, f_{ij}] \text{rect}[(x - iR)/R] \text{rect}[(y - jR)/R], \quad (1)$$

where R is the distance between addressable points on the display and $\text{rect}(x)$ is equal to 1 for $|x| \leq 1/2$ and 0 otherwise. The function $p(i, j, g)$ is the dot profile, i.e., the binary dot shape that represents level g on the display device. Here and in the remainder of this paper we assume that the gray level g is normalized between 0 and 1. For example, $p(110, 30, 0.2)$ gives the value of the dot profile for gray level $g = 0.2$ at location (110, 30). The dot profile is binary and periodic with period M , where M is the size of the basic period (halftone cell) of the ordered-dither array. As Allebach showed in Ref. 11, the Fourier spectrum of an image halftoned by using a mask can be written as follows:

$$H(u, v) = \text{sinc}(Ru) \text{sinc}(Rv) \times \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} F_{mn} \left(u - \frac{m}{MR}, v - \frac{n}{MR} \right), \quad (2)$$

where

$$F_{mn}(u, v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} P[m, n, f(x, y)] \times \exp[-i2\pi(ux + vy)] dx dy, \quad (3)$$

$$P(m, n, g) = \frac{1}{M^2} \sum_{k=-M/2}^{(M/2)-1} \sum_{l=-M/2}^{(M/2)-1} p(k, l, g) \times \exp\left(-i2\pi \frac{mk + nl}{M}\right). \quad (4)$$

The function $P(m, n, g)$ is the discrete Fourier transform of the dot profile $p(i, j, g)$ for a fixed level g . It can also be regarded as a nonlinear transformation of the original image $f(x, y)$ for fixed i and j . Thus the spectrum of the halftone image consists of the displaced spectra of nonlinearly transformed versions $P[m, n, f(x, y)]$ of the input gray-scale image.¹¹ The terms $F_{mn}(u, v)$ are referred to as spectral orders. The $F_{00}(u, v)$ spectral order contains the spectrum of the input gray-scale image, whereas the other nonzero spectral orders are responsible for the visible patterns and loss of detail in the halftone image.¹¹

B. Error Diffusion and Blue Noise

The binary patterns that result from error-diffusion techniques can be modeled as a stationary random process.² Therefore the power spectrum $P(u, v)$ will be used to study the properties of these patterns in the frequency domain. Homogeneous binary patterns, such as the ones produced by error diffusion, can be approximated for many gray levels as being radially symmetric in the Fourier transform domain, given the low anisotropy of their power spectrum as described by Ulichney.² Thus, instead of studying two-dimensional (2-D) plots of $P(u, v)$, we use a simpler one-dimensional (1-D) measure,² the radially averaged power spectrum. For computing the radially averaged power spectrum, the spectral estimate $\hat{P}(u, v)$ is segmented into annuli of width Δ . The width Δ is chosen such that for radial frequency f_r it contains the pixels that satisfy

$$\text{int}(u^2 + v^2)^{1/2} = f_r, \quad (5)$$

where the $\text{int}()$ operation denotes integer truncation. The sample mean of the frequency samples within the annulus $\text{int}(u^2 + v^2)^{1/2} = f_r$ is defined as the radially averaged power spectrum

$$P_r(f_r) = \frac{1}{N_r(f_r)} \sum_{\text{int}(u^2+v^2)^{1/2}=f_r} \hat{P}(u, v), \quad (6)$$

where $N_r(f_r)$ is the number of frequency samples within an annulus.

Ulichney² was the first to analyze the visually pleasing properties of error diffusion mathematically and to attribute them to blue noise. Blue noise is high-frequency noise, and its visually pleasing properties arise from the fact that the human visual system is more sensitive to low frequencies than to high frequencies.¹³

The radially averaged power spectrum of a blue-noise binary pattern (for level g) is shown in Fig. 1. The cutoff frequency f_g is known as the principal frequency, and it assumes its highest value for level $g = 1/2$, since at this level the populations of black and white dots are equal and therefore only high-frequency components appear in the binary image. The dependence of the principal frequency on the gray level g can be expressed as follows:

$$f_g = \begin{cases} \sqrt{g}/R & \text{for } g \leq 0.5, \\ \sqrt{1-g}/R & \text{for } g > 0.5, \end{cases} \quad (7)$$

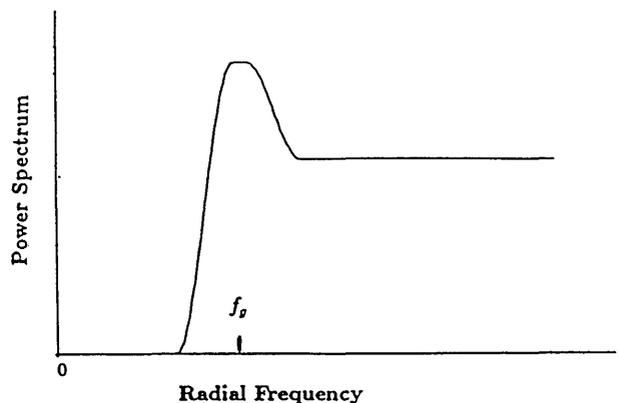


Fig. 1. Radially averaged power spectrum of a blue-noise binary pattern.

where R is again the distance between addressable points on the display.

C. Properties of the Blue-Noise Mask

In this section we examine the properties of the blue-noise mask and establish its connection with error diffusion and ordered dither. Our motivation for building this mask was to combine the advantages of error diffusion and ordered dither, i.e., to combine the speed of ordered dither with the high-quality results of error diffusion. This was done by building our mask as a 2-D single-valued function, i.e., giving it the structure of an ordered-dither mask and also giving it the blue-noise-producing properties of error diffusion. It should be noted here that blue noise refers to the spectral shape of the resulting halftone image and not of the mask itself.

More specifically the blue-noise mask is built such that when it is thresholded at any level, the resulting binary pattern has the correct first-order statistics and the power spectrum has blue-noise characteristics. It is important to note here that the blue-noise mask we constructed is 256×256 pixels, which is significantly larger than conventional ordered-dither masks. The blue-noise mask is built with wraparound properties such that a smaller blue-noise mask can be used to halftone a larger image by tiling the mask over an appropriate number of periods.

The blue-noise mask is a 2-D single-valued function, and it is completely independent of the image to be halftoned. The only connection required between the blue-noise mask and the image to be halftoned is that they must have the same dynamic range (B bits). When the mask is thresholded at any level g , the resulting binary pattern $p(i, j, g)$ is called the dot profile for that level, as it is for the ordered-dither techniques. The value of the dot profile at a particular location (i, j) and for level g is given by

$$p(i, j, g) = \begin{cases} 1 & g > m_{ij} \\ 0 & g \leq m_{ij} \end{cases} \quad (8)$$

where m_{ij} is the corresponding pixel of the blue-noise mask.

Since the blue-noise mask is a 2-D single-valued function, the dot profiles for the different levels are not independent. The dependence of the dot profiles can be written as¹¹

$$\text{if } g_2 > g_1 \cap p(i, j, g_1) = 1 \Rightarrow p(i, j, g_2) = 1. \quad (9)$$

This restriction is true for all halftone techniques that are implemented with a halftone screen. It is this restriction that is responsible for the superiority of error-diffusion techniques over conventional ordered dither, since the dot profiles for error diffusion are independent and can be optimum for many gray levels.

An important property of the blue-noise mask is that, when it is thresholded at any level, the resulting binary pattern has blue-noise properties. The motivation behind that particular design of dot profiles can be found in Eq. (2). According to this equation, the spectrum of any halftone image consists of the spectrum of the original image at the $F_{00}(u, v)$ (zeroth) spectral order plus spectra of distorted versions of the original image at the nonzero spectral orders. Our goal is to minimize the energy of the nonzero spectral orders that are close to the zeroth

order because these spectral orders are responsible for most of the artifacts in the halftone image.¹¹

As Eq. (2) shows, energy at spatial-frequency location (m, n) in $P(m, n, g)$, the discrete Fourier transform of the dot profile, corresponds to energy in the $F_{mn}(u, v)$ spectral order in the Fourier transform $H(u, v)$ of the halftone image. Therefore by minimizing the low frequencies in $P(m, n, g)$ we minimize the spectral orders that are close to the zeroth order and in that way achieve reduction of the visually annoying artifacts in the halftone image. Thus the introduction of blue noise in the dot profiles is intended to reduce the interference from nonzero spectral orders and to create unstructured patterns that are free of periodic artifacts.

Blue noise is the connection between the blue-noise mask and error diffusion. However, the dot profiles of the blue-noise mask are designed to have a principal frequency different from that of the corresponding patterns of error diffusion. The scaling factor $1/\sqrt{2}$ is introduced for purposes that are explained in Sections 3 and 4. Also, because halftoning by using a blue-noise mask requires only a simple pixelwise comparison, it is a faster technique than error diffusion. The speed of error diffusion depends on the particular implementation; for example, error diffusion schemes with fixed weights are faster to implement than are perturbed-weights schemes.

3. CONSTRUCTION OF THE BLUE-NOISE MASK

In this section we describe the algorithm for the construction of the blue-noise mask. We refer to this algorithm as ACBNOM (algorithm for the construction of the blue-noise mask). As previously, the gray level g is normalized between 0 and 1, i.e., $0 < g < 1$.

The dot profile for each level g , i.e., the binary pattern $p(i, j, g)$ that results when the blue-noise mask is thresholded at constant level g , is constructed such that it has blue-noise characteristics in the Fourier domain and correct first-order statistics in the image domain, i.e., the mean of the dot profile $p(i, j, g)$ is g .

In building a single-valued function, one constructs the dot profiles sequentially, i.e., the dot profile for level $g_i + \Delta g$ is built from the dot profile for level g_i by converting the value of a given number of pixels. If the next dot profile corresponds to a higher gray level than does the current dot profile (upward construction), then during conversion zeros are replaced by ones; otherwise, ones are replaced by zeros (downward construction). In what follows, when we discuss the construction of the dot profiles we will imply the upward construction unless we state otherwise. Note that the words "zeros" and "ones" are used as abbreviations for pixels that contain a zero or a one, respectively.

For a $M \times N$ B -bit mask we design 2^B dot profiles such that all available gray levels are represented uniquely. As the dot profile $p(i, j, g_i + \Delta g)$ is constructed from $p(i, j, g_i)$, the number of zeros that will change to ones is equal to $X = (M \times N)\Delta g$, where usually $\Delta g = 1/2^B$. The blue-noise mask has the property that, when it is thresholded at any constant gray level g , the resulting binary pattern is exactly the dot profile we designed for that level. This is achieved by building the blue-noise mask while con-

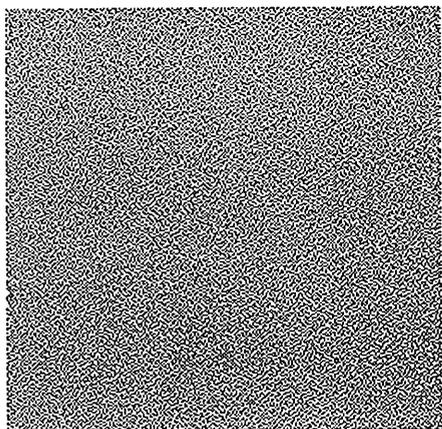


Fig. 2. Starting pattern for the creation of the blue-noise mask.

structuring the dot profiles in such a way that the mask keeps track of the pixels that change value at each iteration after the implementation of the constraints. More specifically, the pixel-value conversion that occurs at location (i, j) of the current dot profile is spatially encoded in another array as a decrement or increment of its pixel value at location (i, j) . This array, which performs the spatial encoding of the dot-profile changes, will be referred to as the cumulative array. The cumulative array is a single-valued three-dimensional (3-D) function of spatial location and gray level and will be denoted by $c(i, j, g)$. When the building of all dot profiles is completed, the cumulative array becomes the blue-noise mask $m(i, j)$.

We start building the dot profiles from level $g = 0.5$, and then we create the dot profiles for levels $g > 0.5$. The dot profile for level $g = 0.5$, $p(i, j, 0.5)$, is chosen to be a suitable blue-noise pattern and is shown in Fig. 2. The cumulative array is also first created at level $g = 0.5$, and it is initialized such that when thresholded at gray level $g = 0.5$, the resulting binary pattern is the dot profile $p(i, j, 0.5)$ we designed for that level. Therefore we assign the value 2^{B-1} to every pixel in the cumulative array that corresponds to a zero in the dot profile $p(i, j, 0.5)$; otherwise, we assign the value $2^{B-1} - 1$:

$$\forall p(i, j, 0.5) = 0 \Rightarrow c(i, j, 0.5) = 2^{B-1}, \quad (10)$$

$$\forall p(i, j, 0.5) = 1 \Rightarrow c(i, j, 0.5) = 2^{B-1} - 1. \quad (11)$$

Below we describe the algorithm we developed in order to create the dot profiles of an $M \times N$ mask for the upward-construction case ($g > 0.5$). At the l th iteration this algorithm creates the dot profile $p(i, j, g_l + \Delta g)$ from $p(i, j, g_l)$ by the following steps:

1. Compute the power spectrum $P(u, v, g_l)$ of the dot profile $p(i, j, g_l)$.
2. Compute the radially averaged power spectrum $P_r(f_r, g_l)$ of $P(u, v, g_l)$:

$$P_r(f_r, g_l) = \frac{1}{N_r(f_r)} \sum_{\text{int}(u^2+v^2)^{1/2}=f_r} P(u, v, g_l). \quad (12)$$

3. Design a 1-D blue-noise filter $D_r(f_r, g_l + \Delta g)$ by dividing the desirable radially averaged power spectrum $P_r'(f_r, g_l + \Delta g)$ by the existing radially averaged power

spectrum $P_r(f_r, g_l)$:

$$|D_r(f_r, g_l + \Delta g)|^2 = \frac{P_r'(f_r, g_l + \Delta g)}{P_r(f_r, g_l)}. \quad (13)$$

As will also be explained in Section 4, the inverse-filtering operation is well conditioned because of the unavoidable low-frequency leakage that prevents $P_r(f_r, g_l)$ from becoming zero. $P_r'(f_r, g_l + \Delta g)$ is designed such that the principal frequency is given by

$$f_g = \begin{cases} \sqrt{g}/\sqrt{2}R & \text{for } g \leq 0.5 \\ \sqrt{1-g}/\sqrt{2}R & \text{for } g > 0.5 \end{cases}, \quad (14)$$

where $g = g_l + \Delta g$. One reason for introducing the scaling factor $K = 1/\sqrt{2}$ is to satisfy design constraints. Specifically, the blue-noise filter is designed in 1-D space and then is replicated over all angles to fill the entire 2-D space. If $K = 1$, then the highest spatial frequency is $f_r = 1/\sqrt{2}$, which corresponds to location A in Fig. 3. Thus designing $D_r(f_r, g_l + \Delta g)$ with $K = 1$ and replicating it (thin line and corresponding arrow line in Fig. 3) will cause truncation of the filter at all angles except those that correspond to points A, B, D, and E in Fig. 3. On the other hand, if $D_r(f_r, g_l + \Delta g)$ is designed with $K = 1/\sqrt{2}$, then the highest spatial frequency is $f_r = 1/2$ (location C in Fig. 3) and the replication of $D_r(f_r, g_l + \Delta g)$ causes no truncation effects (thick line in Fig. 3). Another reason for introducing the scaling factor $K = 1/\sqrt{2}$ will be discussed in Section 4 and is based on the fact that the introduction of this scaling factor produces patterns with high isotropy.

4. Produce a 2-D real and even, radially symmetric filter $D(u, v, g_l + \Delta g)$ by computing the square root of $|D_r(f_r, g_l)|^2$ and replicating it for all angles in the Fourier domain such that the entire 2-D Fourier space is filled:

$$D(u, v, g_l + \Delta g) = [|D_r(f_r, g_l + \Delta g)|^2]^{1/2}, \quad \text{int}(u^2 + v^2)^{1/2} = f_r. \quad (15)$$

5. Apply the blue-noise filter $D(u, v, g_l + \Delta g)$ to the Fourier transform $P_F(u, v, g_l)$ of the current dot profile $p(i, j, g_l)$:

$$P'(u, v, g_l + \Delta g) = P_F(u, v, g_l) \times D(u, v, g_l + \Delta g). \quad (16)$$

6. Compute the inverse Fourier transform (IFT) of $P'(u, v, g_l + \Delta g)$ to obtain $p'(i, j, g_l + \Delta g)$:

$$p'(i, j, g_l + \Delta g) = \text{IFT}[P'(u, v, g_l + \Delta g)], \quad (17)$$

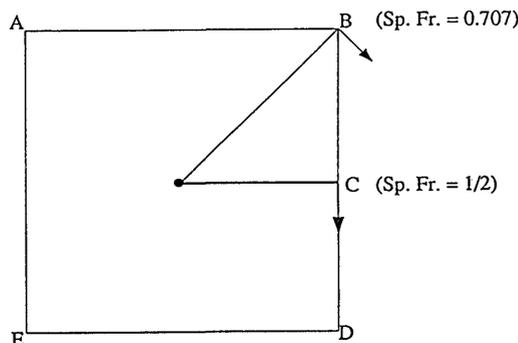


Fig. 3. Replication of the 1-D filter. Sp. Fr., spatial frequency.

where $p'[i, j, g_l + \Delta g]$ is a modified pattern that has the desired blue-noise characteristics in the Fourier domain for level $g_l + \Delta g$. However, $p'(i, j, g_l + \Delta g)$ is no longer binary.

7. To enforce the binary state and to implement the correct first-order statistics for level $g_l + \Delta g$, we form a difference, or error, array:

$$e(i, j, g_l + \Delta g) = p'(i, j, g_l + \Delta g) - p(i, j, g_l). \quad (18)$$

8. Rank order the errors of all pixels that contain a zero.

9. Replace $(M \times N)\Delta g$ zeros with ones in $p(i, j, g_l)$. The zeros that will be replaced with ones are those with the highest error. It is important to note here that before converting the value of a pixel, we check its neighborhood mean to ensure that no clumps will be caused by the conversion. The new pattern that results after the replacement of the zeros, $p(i, j, g_l + \Delta g)$, is the dot profile for level $g_l + \Delta g$ because it is binary and has the desired first- and second-order statistics for level $g_l + \Delta g$. The binary pattern $p(i, j, g_l + \Delta g)$ will also be used as the starting pattern for the next $(l + 1)$ iteration.

10. Update the cumulative array by adding one only to those pixels that still correspond to a zero in the dot profile $p(i, j, g_l + \Delta g)$:

$$c(i, j, g_l + \Delta g) = c(i, j, g_l) + \overline{p(i, j, g_l + \Delta g)}, \quad g_l \geq 0.5, \quad (19)$$

where the overbar indicates a logical NOT operation that changes zeros to ones and vice versa. In this fashion, when the blue-noise mask is thresholded at constant level $g_l + \Delta g$, the resulting binary pattern is the dot profile $p(i, j, g_l + \Delta g)$.

After having built the dot profiles for all levels above 0.5, we proceed to the construction of the levels for which $0 < g < 0.5$ (downward construction). The procedure here is the same as described in the above steps, except that to construct the next dot profile we replace ones with zeros. Also, in the case of downward construction the cumulative array is updated as follows:

$$c(i, j, g_l - \Delta g) = c(i, j, g_l) - p(i, j, g_l - \Delta g), \quad g_l \leq 0.5. \quad (20)$$

After the dot profiles for all gray levels have been constructed, the cumulative array contains a uniform distribution of values between 0 and $2^B - 1$ and becomes the blue-noise mask. Thus the elements of the blue-noise mask can be written in terms of the dot profiles¹¹ as follows:

$$m_{ij} = \min[g : p(i, j, g) = 1]. \quad (21)$$

Equation (21) shows that the value of the blue-noise mask at location (i, j) is equal to gray level g , where g is the first level whose dot profile has a black-to-white (zero-to-one) conversion at location (i, j) .

As an example, a constant-gray-scale image ($g = 0.89$) is halftoned by using the 256×256 blue-noise mask and a perturbed-weights error-diffusion scheme.² The resulting patterns are shown in Figs. 4 and 5, respectively. The corresponding radially averaged power spectra are shown

in Fig. 6. As expected, the principal frequency peak appears at a lower frequency for the blue-noise mask than for the error-diffusion method.

4. BINARY-PATTERN-POWER-SPECTRUM MATCHING ALGORITHM

A. Creation of the Starting Pattern

As was mentioned in Section 3, the creation of the blue-noise mask starts at $g = 0.5$ with a suitable blue-noise pattern. The starting pattern ($g = 0.5$) is of great importance to the overall construction of the blue-noise mask because the dot profiles for all levels are built from this pattern and therefore clumps or other types of inhomogeneities that might appear in the starting pattern will also propagate to other levels. An initial approach would be to use an error-diffusion pattern at $g = 0.5$ for a starting pattern. However, this would be inappropriate because of the inhomogeneous nature of error-diffusion patterns at level $g = 0.5$.

The starting pattern is created with an algorithm similar to the one used for the creation of the blue-noise mask. This algorithm will be called BIPPSMA, which stands for binary-pattern-power-spectrum matching algorithm. The main difference between ACBNOM and BIPPSMA is that at every level BIPPSMA changes the second-order statis-

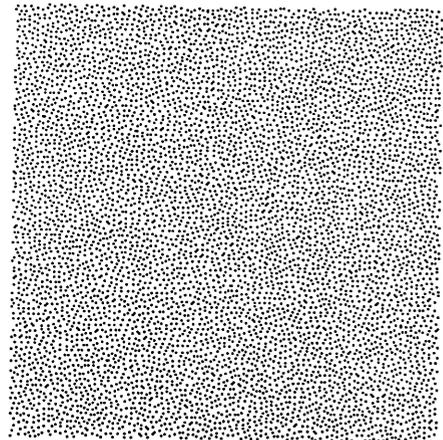


Fig. 4. Constant gray-scale image ($g = 0.89$) halftoned by using the blue-noise mask.

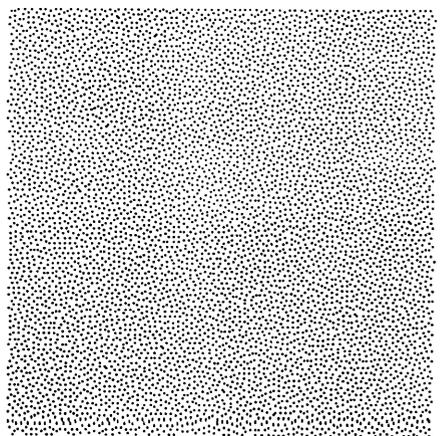


Fig. 5. Constant gray-scale image ($g = 0.89$) halftoned by using error diffusion.

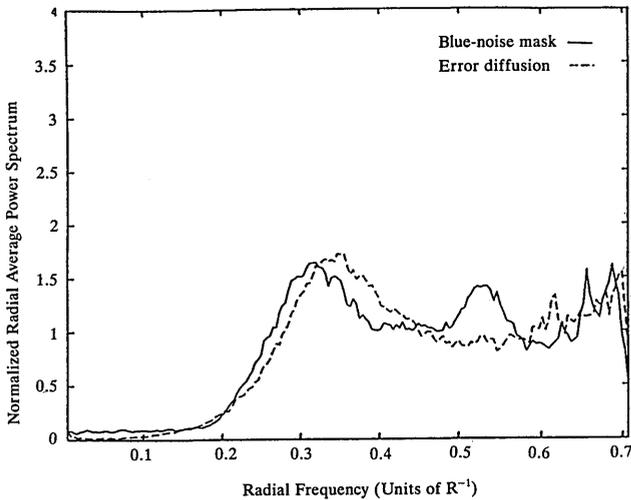


Fig. 6. Radially averaged power spectra of a constant gray-scale image ($g = 0.89$) halftoned by using the blue-noise mask and error diffusion.

tics of the binary pattern while preserving the first-order statistics. BIPPSMA starts with a binary pattern (the seed pattern) that results from thresholding an $M \times N$ white-noise array at level $g = 0.5$. The steps of BIPPSMA for the l th iteration and input pattern $h_l(i, j)$ are summarized as follows:

1. Compute the power spectrum $P_l(u, v)$ of $h_l(i, j)$.
2. Compute the radially averaged power spectrum $P_{lr}(f_r)$.
3. Design the 1-D filter: $|D_{lr}(f_r)|^2 = P_{lr}'(f_r)/P_{lr}(f_r)$.
4. Design the 2-D filter: $D_l(u, v) = (|D_{lr}(f_r)|^2)^{1/2}$.
5. Filter $H_l(u, v)$ [the Fourier transform of $h_l(i, j)$]: $H_l'(u, v) = H_l(u, v) \times D_l(u, v)$.
6. Compute the inverse Fourier transform of $H_l'(u, v)$: $h_l'(i, j) = \text{IFT}[H_l'(u, v)]$.
7. Form the error array $e_l(i, j) = h_l'(i, j) - h_l(i, j)$.
8. Rank order the errors for the zeros and ones.
9. Convert the pairs that have the highest error, and create $h_{l+1}(i, j)$.
10. Compute the mean-square error between $h_{l+1}(i, j)$ and $h_l'(i, j)$.

Steps 1–7 in BIPPSMA above are the same as the corresponding steps in ACBNOM. However, in step 8 all pixels (both zeros and ones) are ranked according to their error, and then in step 9 the zeros and ones with the highest errors exchange positions. In this way the second-order statistics of the binary pattern change with the implementation of the blue-noise filter in step 3 (the same as in ACBNOM), whereas the binary state and the mean value are preserved. Note that in step 3 the principal frequency of the desired radially averaged power spectrum $P_{lr}'(f_r)$ is computed from Eq. (14) for $g = 0.5$. The algorithm iterates until a mean-square-error criterion (between the desired and current pattern) is satisfied. This mean-square error is computed between the binary pattern $h_{l+1}(i, j)$ that results after the pairwise exchange of zeros and ones and the continuous-tone filtered pattern $h_l'(i, j)$. If the mean-square error at the l th iteration is smaller than the mean-square error of the previous iteration, then the algorithm repeats with $h_{l+1}(i, j)$ as the input

pattern. Otherwise, the algorithm is terminated, and $h_l(i, j)$ is chosen as the binary pattern that best satisfies the specified constraints. In a typical run on a 256×256 array at level $g = 0.87$, the algorithm converges, i.e., the mean-square-error value is stabilized, in approximately 20 iterations. The number of pixels that are replaced at each iteration N_p was chosen to be 256. Different values of N_p were tried, and for smaller values the number of iterations increased, whereas for bigger values clumping problems occurred.

Apart from providing a starting pattern for ACBNOM, BIPPSMA is also useful in helping us understand the effect of different second-order statistics on the visual appearance of a binary pattern. Indeed, it was the results of a psychovisual test performed on binary patterns created with BIPPSMA that led us to the introduction of a scaling factor $K = 1/\sqrt{2}$ in Eq. (14) for the creation of the dot profiles of the blue-noise mask. Specifically, as will be explained in Section 5, it was found that power-spectrum-matching algorithms such as BIPPSMA and ACBNOM create the most homogeneous patterns when a scaling factor $K = 1/\sqrt{2}$ is introduced in the formula for the computation of the principal frequency.

B. Creation and Comparison of Binary Patterns with the Same First-Order but Different Second-Order Statistics

In the halftoning literature today the phrase “blue-noise pattern” implies a binary pattern that is of the form shown in Fig. 1 and has the following characteristics:

1. In the Fourier domain most of the energy is concentrated in the high frequencies, whereas the low frequencies have negligible energy.² An important note here is that the terms “low frequencies” and “high frequencies” are used in the halftoning literature to denote the relative positions of the frequencies with respect to the principal-frequency location, i.e., frequencies below the principal frequency are termed “low” and frequencies above the principal frequency are termed “high.”
2. In the Fourier domain the principal frequency is located at the position that corresponds to the average distance of the minority pixels in the binary pattern and is given by Eq. 7.
3. In the image domain the binary pattern is isotropic, i.e., there are no preferred directions in the distribution of black and white dots.

However, this definition is somewhat vague because it does not specify the exact amount of energy that must be contained in the principal-frequency peak in order for one to have a blue-noise pattern. Another interesting issue is how perturbations of the principal-frequency position in the Fourier transform affect the appearance of the binary pattern in the image domain. We address these issues here and examine (a) the strictness of the definition of the principal frequency as the cutoff frequency given by Eq. (7) and the effect of a scaling factor introduced into this equation and (b) the effect of different sizes of the principal-frequency peak on the visual appearance of the binary pattern. For the latter purpose we designed with BIPPSMA binary patterns that had the same first-order statistics but had spectra that differed in the location, height, and width of the principal-frequency peak.

BIPPSMA was applied to binary patterns with dimensions 256×256 at two different constant gray levels, $g_1 = 0.87$ and $g_2 = 0.95$. For both gray levels the starting pattern was white-noise thresholded at levels g_1 and g_2 . We experimented with different variations in the power spectrum of the desired signal. To examine the significance of the principal-frequency location, we introduced a scaling factor K in Eq. (7) such that

$$f_g = \begin{cases} K\sqrt{g}/R & \text{for } g \leq 0.5 \\ K\sqrt{1-g}/R & \text{for } g > 0.5 \end{cases} \quad (22)$$

We experimented with three scaling factors, $K1 = 0.543$, $K2 = 1/\sqrt{2}$, and $K3 = 1$. The third scaling factor corresponds to the standard definition of the principal frequency in Eq. (7).

To examine the effect of different principal-frequency energies on the appearance of the binary patterns in the image domain, we experimented with different heights and widths of the principal-frequency peak. Specifically, we tried two different values for the width of the principal-frequency peak, $0.015R^{-1}$ and $0.0429R^{-1}$, respectively, where R is the distance between addressable points in the spatial domain. For the height we tried the ratios of 1.5, 1.0 (no peak), and 2.5 between the maximum of the principal-frequency peak and the high-frequency baseline. This ratio will be referred to as the high-to-low ratio. In Figs. 7, 8, and 9 we show binary patterns at gray level $g = 0.87$ that correspond to scaling factors $K = 0.543$, $K = 1/\sqrt{2}$, and $K = 1$, respectively.

It is interesting to compare the spectra of these patterns (Fig. 10) with the desired spectra for each pattern, shown in Fig. 11. Note that the dc term is not shown. As can be seen from Fig. 10, BIPPSMA successfully creates binary patterns with power spectra that match the desired ones, with the exception of some low-frequency leakage and some additional energy at the principal-frequency location for $K = 1$. Low-frequency leakage is inevitable when one uses this particular algorithm, for the following reason: BIPPSMA replaces continuous, filtered values of $h_i'(i, j)$ with binary values in $h_{i+1}(i, j)$ according to the error array. This substitution of binary values for continuous values can be viewed as a quantization procedure, which classically is treated as a white-noise process. A degree of low-frequency leakage is actually a stabilizing feature in BIPPSMA because in this way we avoid the problem of dividing by zero in step 3 (in both BIPPSMA and ACBNOM).

To derive a conclusion about the significance of the principal-frequency location and the energy of the principal-frequency peak in the visual appearance of a binary pattern, we performed a psychovisual test. For that purpose 15 binary patterns (256×256) with different locations, heights, and widths of the principal-frequency peak were created for level $g_1 = 0.87$. Also, nine images with different locations and heights of the principal-frequency peak were created for $g_2 = 0.95$. We did not vary the width for $g_2 = 0.95$ because of the small effect of different widths on the visual appearance of a binary pattern. The binary patterns were printed at 70 dots/in. (27.6 dots/cm) on an Apple Laserwriter, and these patterns were presented in random order to 10 subjects under identical conditions. The criterion according to which the binary patterns were rated was isotropy, i.e., the

lack of preferred directions in the distribution of the black and white dots. It is important to note here that the resolution of 70 dots/in. (27.6 dots/cm) was not chosen for reasons of printing optimization; on the contrary, our purpose was to increase the detectability of inhomogeneities in the binary patterns. The viewing distance was approximately 10 in. (25.4 cm). Our rating scale is a frequently used scale¹⁴ from 5 (best) to 1 (worst). Our conclusions based on the psychovisual test results are as follows:

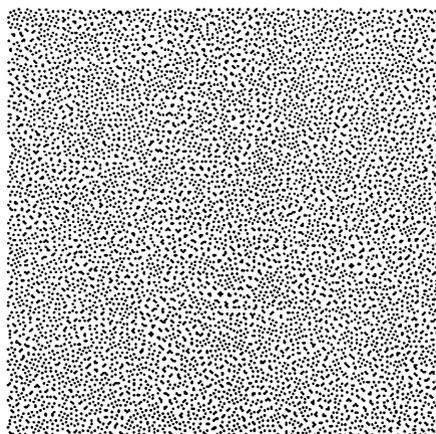


Fig. 7. Binary pattern ($g = 0.87$) with $K = 0.543$.

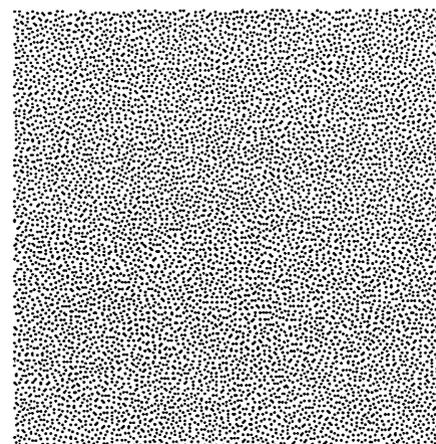


Fig. 8. Binary pattern ($g = 0.87$) with $K = 1/\sqrt{2}$.

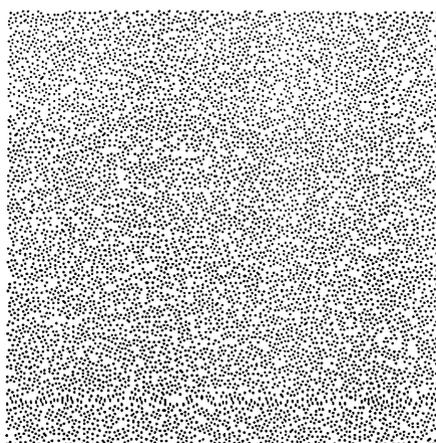


Fig. 9. Binary pattern ($g = 0.87$) with $K = 1$.

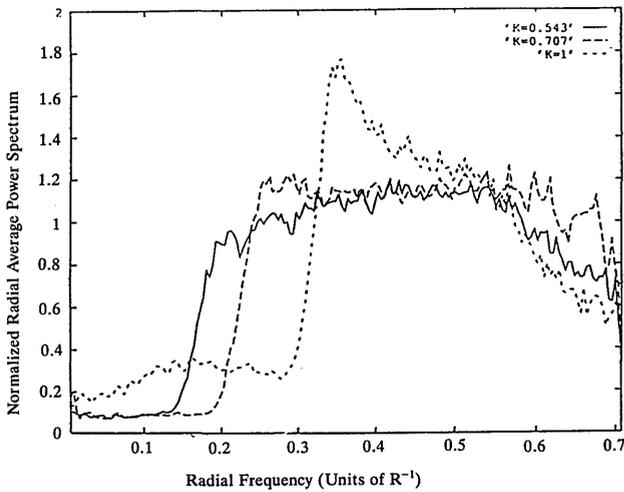


Fig. 10. Radially averaged power spectra of binary patterns at $g = 0.87$ with $K = 0.543$ (solid curve), $K = 1/\sqrt{2}$ (dashed curve), and $K = 1$ (dotted curve).

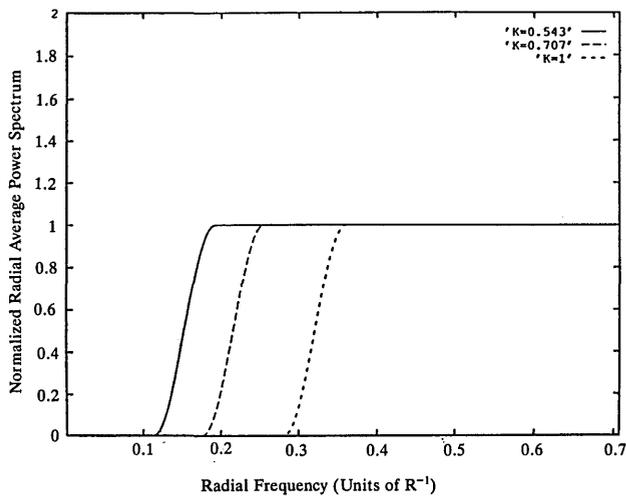


Fig. 11. Desired radially averaged power spectra for binary patterns at $g = 0.87$ with $K = 0.543$ (solid curve), $K = 1/\sqrt{2}$ (dashed curve), and $K = 1$ (dotted curve).

1. When a power-spectrum-matching algorithm (BIPPSMA) was used, the best-rated patterns were those with $K = 1/\sqrt{2}$. Thus the definition of the principal frequency as the cutoff frequency given by Eq. (7) can be broadened to include a scaling factor $K = 1/\sqrt{2}$. It is important to note here that the conclusion that patterns with $K = 1/\sqrt{2}$ are rated better than those with $K = 1$ is not a general conclusion but is restricted to power-spectrum-matching algorithms of the type of BIPPSMA and ACBNOM.

2. Binary patterns with the same location but different energies of the principal-frequency peak were rated closely; therefore the size (energy) of the principal-frequency peak is not important in the visual appearance of a binary pattern.

The results of the psychovisual test for level $g = 0.87$ are shown in Fig. 12. As can be seen from this figure, binary patterns that corresponded to a scaling factor of $K = 1/\sqrt{2}$ received significantly higher rating than patterns

that corresponded to scaling factors $K = 1$ and $K = 0.543$, whereas different high-to-low ratios had little effect on the rating of images with the same K .

5. PSYCHOVISUAL TEST

In this section we present the results of a psychovisual test that was performed to assess and compare the quality of images that were halftoned by using the blue-noise mask, error diffusion (a perturbed-weights scheme²), Bayer's dispersed-dot dither, and the 45° classical screen clustered-dot dither. These techniques were chosen for comparison with the blue-noise mask because they represent the three most important halftoning categories and because they rank among the best in their respective categories.

Eleven images of various content were halftoned by using the above techniques, and the resulting images were printed at 150 dots/in. (59.1 dots/cm) on an Apple Laser-writer Plus. As before, the printing resolution was chosen not for printing optimization reasons but to increase the visibility of the artifacts introduced by the different halftoning methods. The different types of image included portraits, geometric shapes, natural scenes, and a medical image. One of these images, halftoned with the aforementioned techniques, is shown in Figs. 13–16. Twelve subjects rated the images according to their simi-

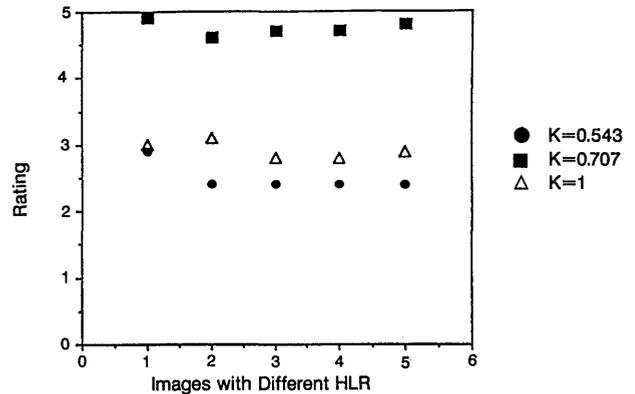


Fig. 12. Psychovisual test results ($g = 0.87$).

Table 1. Rating Scale

Rating	Rendition of Original Image
5	Excellent
4	Good
3	Fair
2	Poor
1	Bad

Table 2. Psychovisual Test Results

Halftone Method	Mean Opinion Score
Error diffusion	4.3
Blue-noise mask	3.9
Dispersed-dot dither	2.7
Clustered-dot dither	1.6

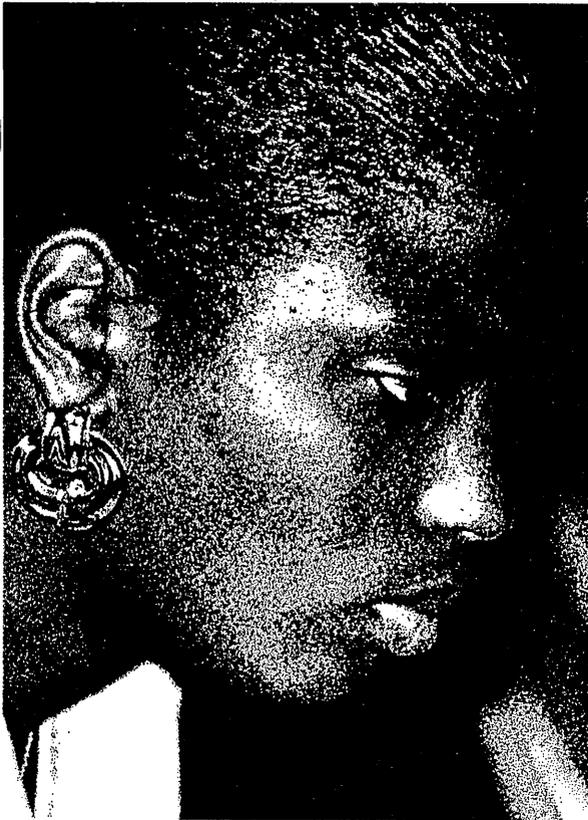


Fig. 13. Gray-scale image halftoned by using the blue-noise mask.

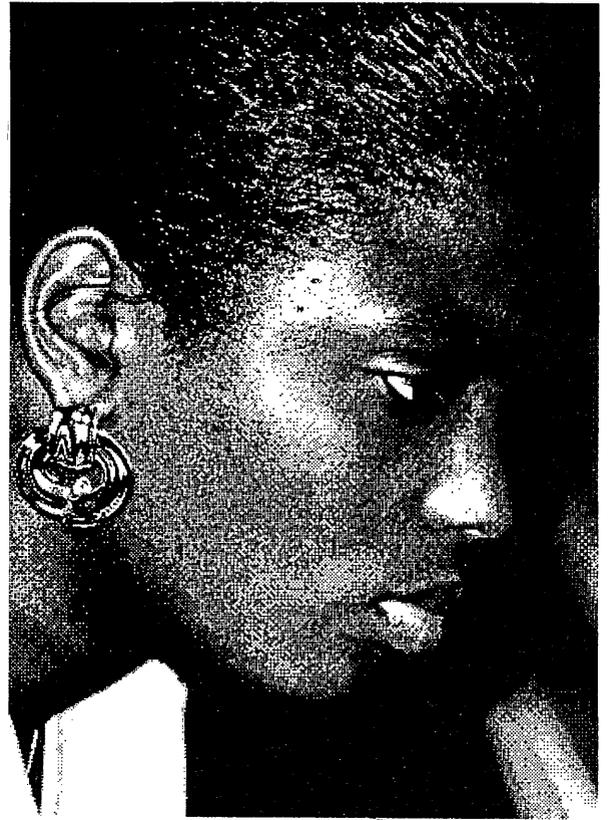


Fig. 15. Gray-scale image halftoned by using ordered dither.

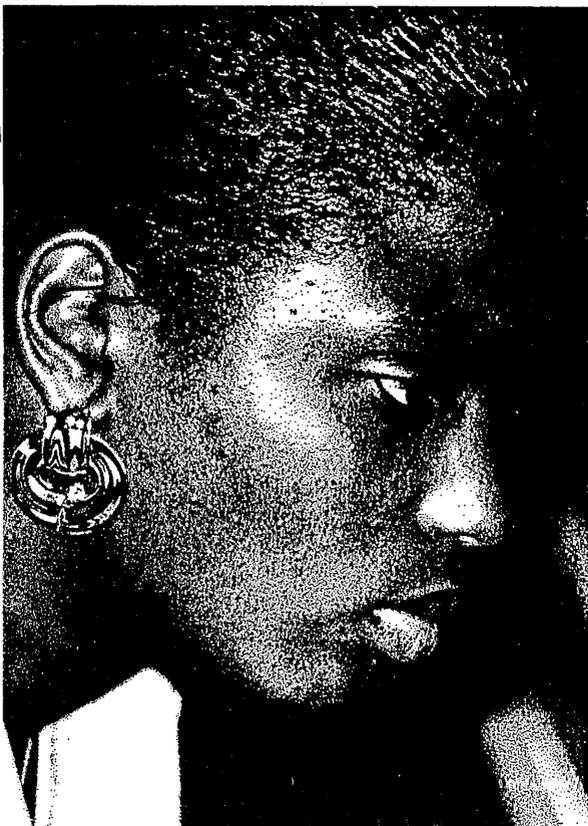


Fig. 14. Gray-scale image halftoned by using error diffusion.

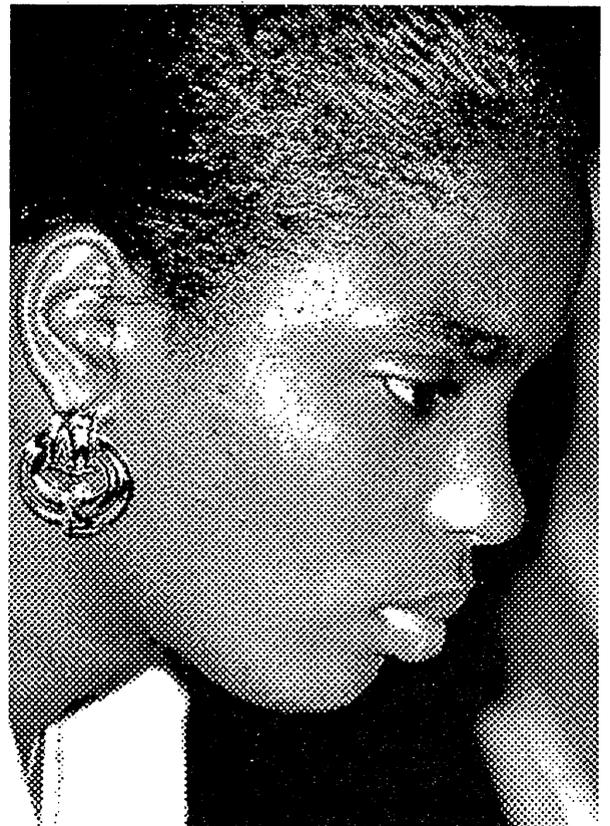


Fig. 16. Gray-scale image halftoned by using clustered-dot dither.

larity to the original gray-scale image at a viewing distance of 10 in. (25.4 cm). The viewers had no *a priori* knowledge about the halftone techniques that were used, and the images were presented to them in a random order under identical conditions. We used a standard rating scale from 5 (excellent) to 1 (bad).¹⁴ The entire scale is shown in Table 1. The mean opinion score for each halftoning method is shown in Table 2.

As can be seen from Table 2, error diffusion was rated best (4.3) with the blue-noise mask rated slightly lower (3.9), whereas the dispersed-dot and clustered-dot dither methods received lower ratings, 2.7 and 1.6, respectively. The blue-noise mask was rated slightly lower than error diffusion for the following reasons:

1. The blue-noise mask has correlated dot profiles. This is a disadvantage in the low-contrast areas.
2. The halftones that were created by using the blue-noise mask appear more grainy than those created by using error diffusion because of the small, but nonzero, leakage of low-frequency energy.
3. Error diffusion has intrinsic sharpening properties. The error-diffusion images appear sharper because of the edge-enhancement properties of error diffusion.¹⁵

However, error diffusion is not completely free of artifacts. Texture-shift artifacts can be seen in Fig. 14 (neck area), and they are due to the dependence of the binary value of a pixel on the previous binary assignments. The dispersed-dot and clustered-dot dither methods were rated poorly because of their periodic artifacts that give an artificial texture to the halftone images. Another type of artifact associated with ordered-dither techniques is the Moiré patterns that occur when the gray-scale image contains frequencies close to the frequency of the halftoning mask.

6. CONCLUSION

In recent years halftoning has become an essential part in binary output devices, such as facsimile machines, laser printers, and desktop-publishing systems. The most widely used halftone techniques today are ordered dither and error diffusion, with error diffusion being the preferred choice for printers that can accurately reproduce single black and white pixels. In this paper we introduced digital halftoning using a blue-noise mask, which is a novel (dispersed-dot) halftone algorithm that combines the structure of an ordered-dither mask and the blue-noise-producing properties of error diffusion.

We described the algorithm for the construction of the blue-noise mask (ACBNOM) and demonstrated how this algorithm builds the blue-noise mask by creating the binary patterns for each level and then combining them in a single-valued function. We also described the algorithm

for the construction of the starting pattern (BIPPSMA). This algorithm was also used for the construction of binary patterns with the same first-order but different second-order statistics. These patterns were rated in a psychovisual test according to their isotropy, and from the results of the test it was concluded that the introduction of a scaling factor $K = 1/\sqrt{2}$ in the computation of the principal frequency in power-spectrum-matching algorithms (such as BIPPSMA and ACBNOM) results in binary patterns with the highest isotropy. To compare the performance of the blue-noise mask with other halftone methods, we performed a psychovisual test in which 11 gray-scale images were halftoned with four halftone techniques, including the blue-noise mask, and then were presented to 12 subjects for evaluation. According to the results of the psychovisual test, the blue-noise mask performs significantly better than do ordered-dither techniques, and it gives results comparable with those of error-diffusion techniques.

REFERENCES

1. J. C. Stoffel and J. F. Moreland, "A survey of electronic techniques for pictorial reproduction," *IEEE Trans. Commun. COM-29*, 1898-1925 (1981).
2. R. A. Ulichney, *Digital Halftoning* (MIT Press, Cambridge, Mass., 1987).
3. R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial grayscale," *Proc. Soc. Inf. Displ.* **17**, 75-77 (1976).
4. R. A. Ulichney, "Dithering with blue noise," *Proc. IEEE* **76**, 56-79 (1988).
5. C. Billotet-Hoffman and O. Bryngdahl, "On the error diffusion technique for electronic halftoning," *Proc. Soc. Inf. Displ.* **24**, 253-258 (1983).
6. J. F. Jarvis, C. N. Judice, and W. H. Ninke, "A survey of techniques for the display of continuous-tone pictures on bilevel displays," *Comput. Graphics Image Process.* **5**, 13-40 (1976).
7. P. Stucki, "Image processing for document reproduction," in *Advances in Digital Image Processing*, P. Stucki, ed. (Plenum, New York, 1979), pp. 177-218.
8. R. L. Stevenson and G. R. Arce, "Binary display of hexagonally sampled continuous-tone images," *J. Opt. Soc. Am. A* **2**, 1009-1013 (1985).
9. T. Mitsa, "Digital halftoning using a blue-noise mask," Ph.D. dissertation (University of Rochester, Rochester, New York, 1991).
10. T. Mitsa and K. J. Parker, "Digital halftoning using a blue-noise mask," in *Image Processing Algorithms and Techniques II*, M. R. Civanlar, S. K. Mitra, and B. P. Mathur, eds., *Proc. Soc. Photo-Opt. Instrum. Eng.* **1452**, 47-56 (1991).
11. J. P. Allebach, "Analysis of halftone dot profile and aliasing in the discrete binary representation of images," *J. Opt. Soc. Am. A* **67**, 1147-1154 (1977).
12. D. Kermisch and P. Roetling, "Fourier spectrum of halftone images," *J. Opt. Soc. Am. A* **65**, 716-723 (1975).
13. F. W. Campbell and D. G. Green, "Optical and retinal factors affecting visual resolution," *J. Phys.* **181**, 576-593 (1965).
14. J. O. Limb, "Distortion criteria of the human viewer," *IEEE Trans. Syst. Man Cybern.* **SMC-9**, 778-793 (1979).
15. R. Eschbach and K. T. Knox, "Error-diffusion algorithm with edge enhancement," *J. Opt. Soc. Am. A* **8**, 1844-1850 (1991).