

Blue-Noise Multitone Dithering

J. Bacca Rodríguez, *Member, IEEE*, G. R. Arce, *Fellow, IEEE*, and D. L. Lau, *Member, IEEE*

Abstract—The introduction of the blue-noise spectra—high-frequency white noise with minimal energy at low frequencies—has had a profound impact on digital halftoning for binary display devices, such as inkjet printers, because it represents an optimal distribution of black and white pixels producing the illusion of a given shade of gray. The blue-noise model, however, does not directly translate to printing with multiple ink intensities. New multilevel printing and display technologies require the development of corresponding quantization algorithms for continuous tone images, namely multitone. In order to define an optimal distribution of multitone pixels, this paper develops the theory and design of multitone, blue-noise dithering. Here, arbitrary multitone dot patterns are modeled as a layered superposition of stack-constrained binary patterns. Multitone blue-noise exhibits minimum energy at low frequencies and a staircase-like, ascending, spectral pattern at higher frequencies. The optimum spectral profile is described by a set of principal frequencies and amplitudes whose calculation requires the definition of a spectral coherence structure governing the interaction between patterns of dots of different intensities. Efficient algorithms for the generation of multitone, blue-noise dither patterns are also introduced.

Index Terms—Blue-noise dithering, digital halftoning, direct binary search (DBS), error diffusion, multitone.

I. INTRODUCTION

HALFTONING is the process of converting a continuous tone image into a pattern of black and white dots [1], where the illusion of continuous tone is the result of the low-pass characteristics of the human eye that make it unable to discriminate printed dots. Ulichney's introduction of the blue-noise model [2], later revised by Lau and Ulichney [3], has had a profound impact in halftoning, since it describes the spectral and spatial characteristics of a visually pleasant dither pattern. Blue-noise dithering is characterized by an arrangement where the minority pixels are spread as homogeneously as possible to create patterns that are aperiodic, isotropic, and do not contain low-frequency components.

Recent advances in printing technology now allow for the reproduction of dots of different intensities. Hardware implementations include the use of several different inks, different ink concentration, and/or variable dot sizes. The availability of these techniques poses the problem of reproducing a continuous tone image with dots of at least three intensities (black, white, and one or more intermediate gray level inks). Image rendering via multiple inks is known as multilevel halftoning or multi-

toning. Several algorithms for multitone have been proposed in the literature, mostly as extensions of previously developed halftoning algorithms. The results obtained with these methods are compared and evaluated subjectively since they rely on the blue-noise model for halftones and a comparable theory designed explicitly for multitones is not available. Such theory, the blue-noise model for multitone dithering, is developed in this paper.

Multitone, like halftoning, aims at generating images that are visually pleasant to the human eye. For this reason, the fundamental principles of blue-noise halftoning can be generalized to the multitone case. Characteristics like homogeneity and isotropy are desirable in multitone dither patterns. In the frequency domain, the radial symmetry and a low-frequency response close to zero are requirements imposed by the properties of the human eye and, in consequence, they must be considered in the development of the theory of blue noise for multitone. On the other hand, new challenges arise when dots of intermediate intensities are allowed, for example, dot patterns of different inks could interfere with each other creating variations in the intended average value of the picture or generating low-frequency noise.

To determine the spectral profile of multitones and the characteristics required for its optimality, threshold decomposition (TD) [4] is introduced as a tool for the analysis of multitone patterns. TD allows for the representation of multitones as the superposition of N spatially correlated halftone patterns, where N is the number of available inks. Each one of these patterns can be characterized itself as a blue-noise pattern. The spectral profile of blue-noise multitones is thus defined as the aggregation of the profiles of these halftones plus the cross spectra generated by the interaction between dots of different intensities. The spectral correlation between halftones composing a blue-noise multitone is characterized by means of their optimal spectral coherence.

Given the spatial and spectral characterization of optimal blue-noise multitones, the need for algorithms that generate such multitones arises. To this end, threshold decomposition is extended so it can be applied to continuous tone images. This representation can be used in combination with a given blue-noise halftoning algorithm to generate multitone dither patterns that show the spectral characteristics of blue-noise multitones previously defined. In particular, the application of such scheme to error diffusion and DBS is shown.

The paper is divided as follows. Section II summarizes the blue-noise theory for halftones as presented by Ulichney [2] and Lau and Ulichney [3]. Section III shows the development of the blue-noise model for multitones, including the new blue-noise multitone structure. Section IV shows examples of the multitones obtained with such methods together with their spectral analysis, and Section V is dedicated to the conclusions.

Manuscript received September 18, 2007; revised April 16, 2008. First published June 24, 2008; last published July 11, 2008 (projected). The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Brian L. Evans.

J. Bacca Rodríguez and G. R. Arce are with the Department of Electrical and Computer Engineering, University of Delaware, Newark, DE 19716 USA.

D. L. Lau is with the Department of Electrical and Computer Engineering, University of Kentucky, Lexington, KY 40506 USA.

Digital Object Identifier 10.1109/TIP.2008.926145

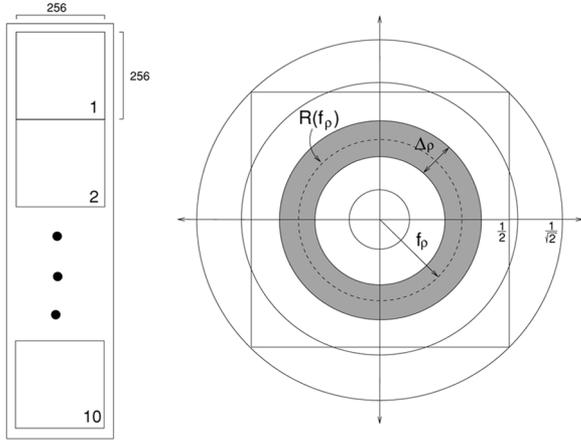


Fig. 1. Calculation of the RAPSD of a dither pattern. Ten sections of 256×256 pixels are extracted from a large dither pattern of the desired gray level, the periodogram of each pattern is calculated and $\hat{P}(\mathbf{f})$ is calculated as their average. To obtain the RAPSD, the average of $\hat{P}(\mathbf{f})$ is taken over annuli of width $\Delta\rho$ as indicated.

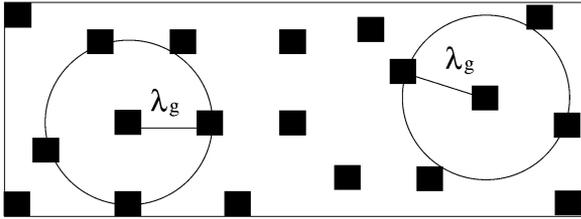


Fig. 2. Average distance between pixels in a blue-noise halftone pattern. In areas of constant intensity, minority pixels tend to spread apart an average distance λ_g in blue-noise dithering.

II. BLUE-NOISE MODEL FOR BINARY DITHER PATTERNS

A. Spectral Statistics of Halftones

Ulichney [2] characterized well formed dither patterns in the Fourier domain by means of the radially average power spectrum density (RAPSD) and anisotropy measures. He focused on binary patterns resulting from the dithering of an input image composed of pixels of the same intensity g . These patterns can be characterized as Bernoulli processes with a probability density function

$$P(H[\mathbf{n}]) = \begin{cases} g, & \text{for } H[\mathbf{n}] = 1 \\ 1 - g, & \text{for } H[\mathbf{n}] = 0 \end{cases} \quad (1)$$

with a second moment (variance) equal to $\sigma^2 = g(1 - g)$. The characteristics of a dither pattern in the frequency domain can be studied using its power spectrum $P(\mathbf{f})$. Ulichney estimated the power spectrum ($\hat{P}(\mathbf{f})$) as the average of ten periodograms obtained from squared dither patterns of 256×256 pixels, cropped from larger patterns as indicated in Fig. 1. The end result is a 2-D estimate that can be partitioned into annuli of width $\Delta\rho$. The RAPSD is the radial average of $\hat{P}(\mathbf{f})$ on this annuli calculated as

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

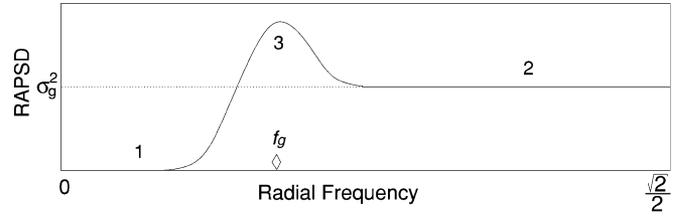


Fig. 3. Ideal radial average of the power spectrum of a blue-noise halftone pattern illustrating its three main characteristics: Low-frequency response close to zero (1), flat high-frequency region (2), and a peak at the principal frequency of the pattern (3).

where f_ρ is the central radius and $N(R(f_\rho))$ the number of samples in the annuli. To measure radial symmetry, Ulichney defined a spectral measure called anisotropy [2]. In general, the radial symmetry of the halftoning algorithms used in this paper has already been proven, making a further analysis of the anisotropy unnecessary.

B. Blue-Noise Spectra

Ulichney stated that the optimal dither patterns is such that the average distance between nearest-neighboring minority pixels is defined according to

$$\lambda_g = \begin{cases} \frac{S}{\sqrt{g}}, & \text{for } g \leq \frac{1}{2} \\ \frac{S}{\sqrt{1-g}}, & \text{for } g > \frac{1}{2} \end{cases} \quad (3)$$

where S is the minimum distance between addressable pixels. λ_g is referred to as the *principal wavelength* of the pattern. In the Fourier domain, a well formed dither pattern is characterized by a low-frequency response close to zero, a flat high-frequency region, and a peak at the so-called *principal frequency* of the pattern. A plot of an ideal RAPSD of a blue-noise dither pattern is shown in Fig. 3. The principal frequency f_g is the inverse of the principal wavelength

$$f_g = \begin{cases} S^{-1}\sqrt{g}, & \text{for } g \leq \frac{1}{2} \\ S^{-1}\sqrt{1-g}, & \text{for } g > \frac{1}{2} \end{cases} \quad (4)$$

When analyzing the RAPSD of blue-noise patterns of intensity near $g = (1/2)$, Lau and Ulichney [3] noticed that by forcing a principal frequency greater than $(1/2)$, a halftoning algorithm was sacrificing radial symmetry as the sampling grid constrained the placement of dots along the diagonals. Lau and Ulichney argued that, for levels $(1/4) < g < (3/4)$, the grid-defiance illusion of the patterns was lost due to the added diagonal correlation. So, in order to maintain a continuous wavefront inside the baseband, some clustering should be allowed in dither patterns representing these gray levels. Equations (3) and (4) were, therefore, modified. The principal frequency is redefined as

$$f_g = \begin{cases} S^{-1}\sqrt{g}, & \text{for } g \leq \frac{1}{4} \\ \frac{S^{-1}}{2}, & \text{for } \frac{1}{4} < g < \frac{3}{4} \\ S^{-1}\sqrt{1-g}, & \text{for } g \geq \frac{3}{4} \end{cases} \quad (5)$$

and the principal wavelength is defined as its inverse.

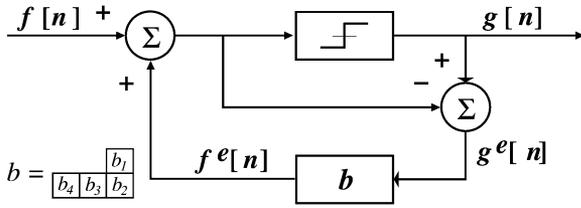


Fig. 4. Error diffusion halftoning. The pixel $f[n]$ is passed through a quantizer to obtain the corresponding pixel of the halftone $g[n]$, the difference between these two is diffused to the neighboring pixels by means of the filter b .

C. Blue-Noise Halftoning

Since the introduction of the blue-noise model, many halftoning algorithms have been proposed aiming to produce patterns with such characteristics.¹ The simplest are screening algorithms, implemented by means of a thresholding operation, like Bayer's dither array [6]. The use of this kind of array resulted in the introduction of periodic artifacts. Several algorithms have been proposed to design dither arrays with better characteristics (Sullivan *et al.* [7], Chu [8], Mitsa and Parker [9], Ulichney [10], and Allebach and Lin [11]). A summary of some of these methods and others applied to multitone and color printing can be found in [12].

Alternative halftoning methods affect not only the pixel being quantized but also its vicinity, resulting in a higher computational complexity. The earliest of these techniques is the error diffusion algorithm (Floyd and Steinberg [13]), summarized in the block diagram in Fig. 4: A quantizer takes the value of the input pixel and compares it with a threshold to decide the value of the corresponding pixel in the halftone. The error introduced is calculated and diffused to the neighboring, soon-to-be-processed pixels using a filter $\mathbf{b} = [(7/16), (1/16), (5/16), (3/16)]$. The process is repeated with the following pixel until the whole image has been halftoned. The picture is processed on a left-to-right and top-to-bottom raster scan. The output of this algorithm when applied to a gray-scale ramp is shown in Fig. 5. The checkerboard patterns observed around gray level (1/2) indicate the diagonal correlation introduced by a high cut-off frequency. Other than that, worms and other geometric artifacts can be appreciated in different regions of the picture. Fig. 6 shows the RAPSD of patterns of intensities (1/16), (1/8), (1/4), and (1/2) halftoned with Floyd and Steinberg's algorithm. The geometric artifacts in the gray-scale ramp reflect here as spectral peaks at the principal frequency of the pattern or its multiples.

Several variations of this algorithm have been proposed over the years, including the use of different filter shapes and number of weights as proposed by Jarvis *et al.* [14], Stucki [15] (12 weights), or Shiao and Fan [16] (six weights). The scanning path can be modified from a traditional raster scan to a serpentine or others like the Peano [17] and Hilbert paths [18], or paths dictated by a matrix as in dot-diffusion (Knuth [19]). The threshold of error diffusion can also be modified based on previous outputs [20] or the intensity of the current pixel as indicated by Eschbach and Knox [21]. Other authors suggest to make the shape

¹A summary of the blue-noise principles and some of these methods can be found in [5]

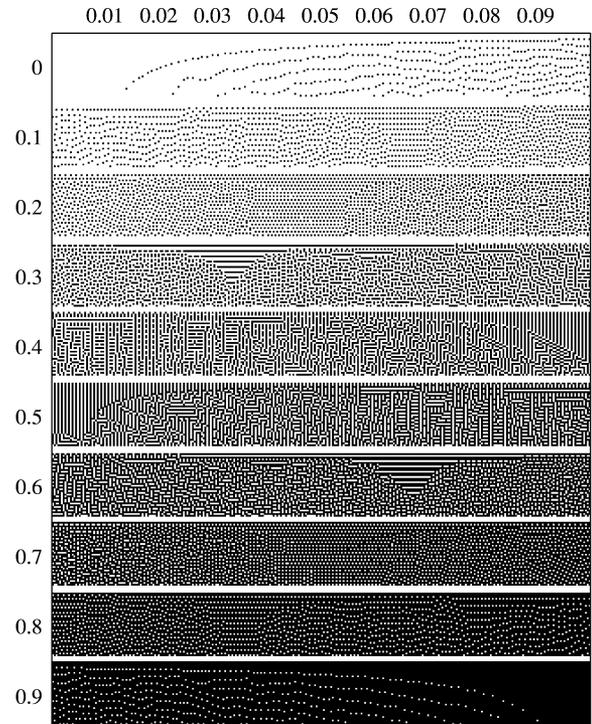


Fig. 5. Halftone of a gray-scale ramp generated with Floyd–Steinberg error diffusion.

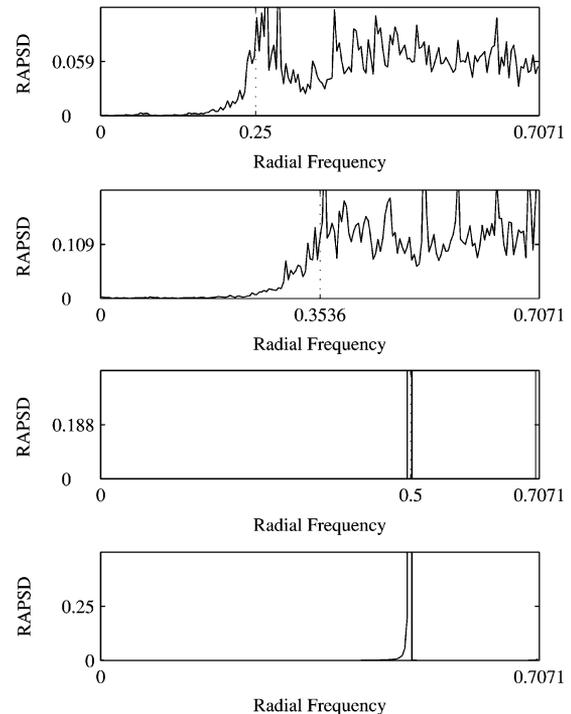


Fig. 6. RAPSD of halftones generated with Floyd–Steinberg error diffusion for gray levels (1/16)(1/8)(1/4), and (1/2).

of the filter, as well as the weights, input dependant as shown by Eschbach [22] and Ostromoukhov [23]. A joint optimization of thresholds and weights based on a model for the human visual system is presented by Li and Allebach [24]. Error diffusion has even been modified to generate green-noise as in [25] and [26].

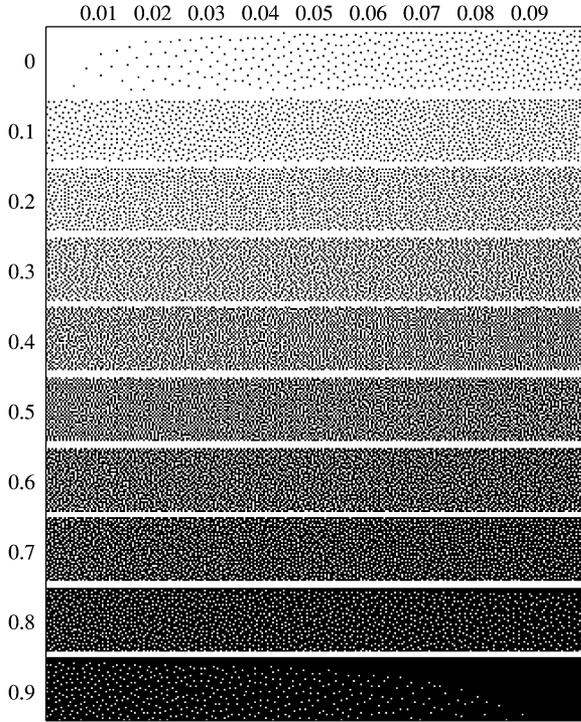


Fig. 7. Halftone of a gray-scale ramp generated with Ulichney's error diffusion.

To improve the radial symmetry and the cut-off frequency of error-diffusion dither patterns, Ulichney used a serpentine scan and introduced randomness in the weights of the error filter. The weights used by Ulichney were calculated as $[b_1 + R_1, b_2 - R_2, b_3 - R_1, b_4 + R_2]$, where $R_1 = (5/16)U[-1, 1]$, $R_2 = (1/16)U[-1, 1]$, and $U[m, n]$ represents a uniformly distributed random variable in the interval $[m, n]$. The results obtained applying this algorithm are illustrated in Figs. 7 and 8. The randomness in the weights breaks out most of the geometric patterns present in the halftone in Fig. 5, including the checkerboards near $g = (1/2)$. The RAPSD plots in Fig. 8 show that this improvement is linked to a frequency content more evenly distributed in the band above the principal frequency. Floyd–Steinberg's original technique is an implementation of the original blue-noise model, whereas Ulichney's is a realization of the model by Lau and Ulichney.

Further upgrades in the quality of halftones can be obtained using algorithms that perform an iterative search whose final goal is to obtain an improved version of an initial halftone. These algorithms are the most complex and, even though they are not practical for implementation, their results serve as benchmarks and optimization references for other algorithms. Among them is Analoui and Allebach's direct binary search (DBS) [27]. DBS improves the quality of a halftone under an error measure that can include models of the human visual system (HVS) and the printing device, and it is summarized in the following.

Assume a digital image and its corresponding halftone are represented by $f[\mathbf{m}]$ and $g[\mathbf{m}]$ respectively. The HVS model is represented as a linear filter named $h(\mathbf{x})$, and the printer model by $p(\mathbf{x})$. The perceived printed image and halftone, represented

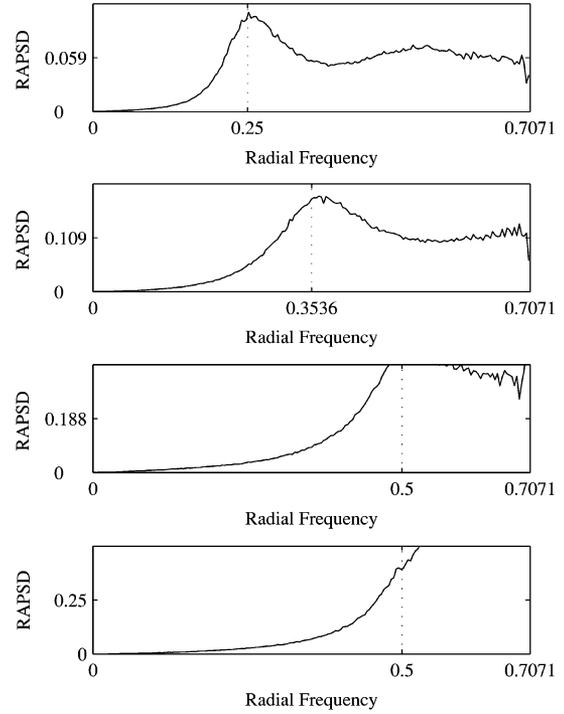


Fig. 8. RAPSD of halftones generated with Ulichney's error diffusion for gray levels $(1/16)(1/8)(1/4)$, and $(1/2)$.

by \tilde{f} and \tilde{g} , are obtained by filtering with a linear filter that comprises the effect of the printing process and the HVS ($\tilde{p} = p * h$, where “*” represents convolution). The perceived error between the image and the halftone is given by $\tilde{e}(\mathbf{x}) = \tilde{g}(\mathbf{x}) - \tilde{f}(\mathbf{x})$, with the total squared error defined as

$$E = \int |\tilde{e}(\mathbf{x})|^2 dx. \quad (6)$$

Consider a trial change of the pixel $g[\mathbf{m}_0]$. This pixel could be swapped with one of its 8 nearest neighbors, say $g[\mathbf{m}_1]$, or toggled to a different intensity (1 or 0). The effect of such pixel change is

$$\tilde{g}'[\mathbf{m}] = \tilde{g}[\mathbf{m}] + a_0 \tilde{p}[\mathbf{m} - \mathbf{m}_1] - a_1 \tilde{p}[\mathbf{m} - \mathbf{m}_0] \quad (7)$$

with a_0 and a_1 defined by

$$a_0 = \begin{cases} g[\mathbf{m}_0] - g[\mathbf{m}_1], & \text{for a swap} \\ -1, & \text{for a toggle if } g[\mathbf{m}_0] = 0 \\ 1, & \text{for a toggle if } g[\mathbf{m}_0] = 1 \end{cases} \quad (8)$$

$$a_1 = \begin{cases} -a_0, & \text{for a swap} \\ 0, & \text{for a toggle.} \end{cases}$$

The change in the error measure defined in (6) is given by

$$\Delta E = (a_0^2 + a_1^2) c_{\tilde{p}\tilde{p}}[\mathbf{0}] + 2(a_0 c_{\tilde{p}\tilde{e}}[\mathbf{m}_0] + a_1 c_{\tilde{p}\tilde{e}}[\mathbf{m}_1] + a_0 a_1 c_{\tilde{p}\tilde{p}}[\mathbf{m}_0 - \mathbf{m}_1]) \quad (9)$$

where $c_{\tilde{p}\tilde{p}}$ is the autocorrelation function of \tilde{p} and $c_{\tilde{p}\tilde{e}}$ is the cross correlation between \tilde{p} and \tilde{e} . The former remains constant

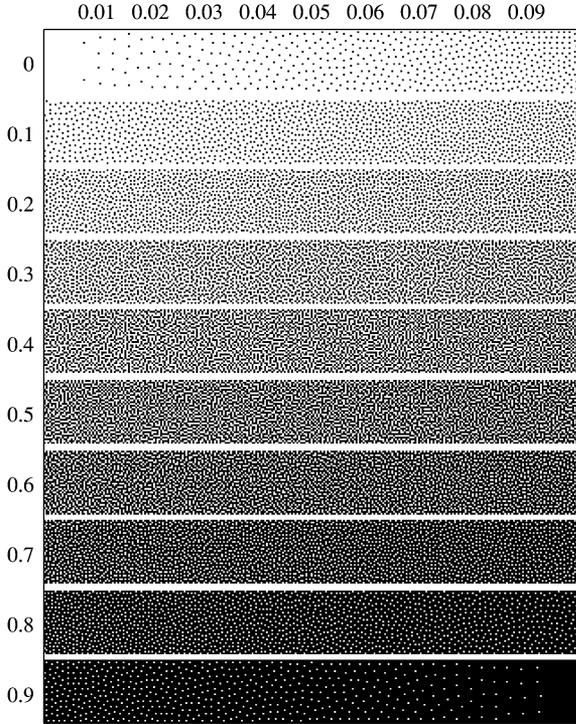


Fig. 9. DBS halftoning of a grayscale ramp.

during all the optimization process. The latter has to be updated every time the halftone is changed according to

$$c'_{\tilde{p}\tilde{e}}[\mathbf{m}] = c_{\tilde{p}\tilde{e}}[\mathbf{m}] + a_0 c_{\tilde{p}\tilde{p}}[\mathbf{m} - \mathbf{m}_0] + a_1 c_{\tilde{p}\tilde{p}}[\mathbf{m} - \mathbf{m}_1]. \quad (10)$$

Fig. 9 shows a halftone of the gray scale ramp produced using DBS. This halftone lacks the geometric artifacts present in Fig. 5 and looks smoother when compared with Fig. 7. Fig. 10 shows the RAPSD of DBS patterns for gray levels $(1/16)$, $(1/8)$, $(1/4)$, and $(1/2)$. This plots resemble the best the ideal characteristics of a blue-noise dither pattern in Fig. 3.

To appreciate the differences between some of the halftoning algorithms just described, refer to Fig. 11. The left shows a halftone of a natural image generated with Ulichney's error diffusion with perturbed weights, the right is the same image halftoned with DBS. The characteristics observed previously in the gray-scale ramp reflect here as well. The image on the left presents more geometric artifacts and a noisy texture whereas the one on the right is smooth and presents very few artifacts.

III. BLUE-NOISE MODEL FOR MULTITONE DITHER PATTERNS

A. Spectral Statistics of Multitones

Much like binary dither patterns, multitone dither patterns representing a constant gray level g can be modeled as stochastic processes. Assume that a multitone dither pattern is created using N different inks of intensities (g_1, g_2, \dots, g_N) sorted according to intensity starting with the lightest. A white pixel (where nothing is printed) is said to have intensity $g_0 = 0$ while a black pixel is printed using intensity $g_N = 1$. The dither pattern, therefore, contains pixels of $N + 1$ different intensities.

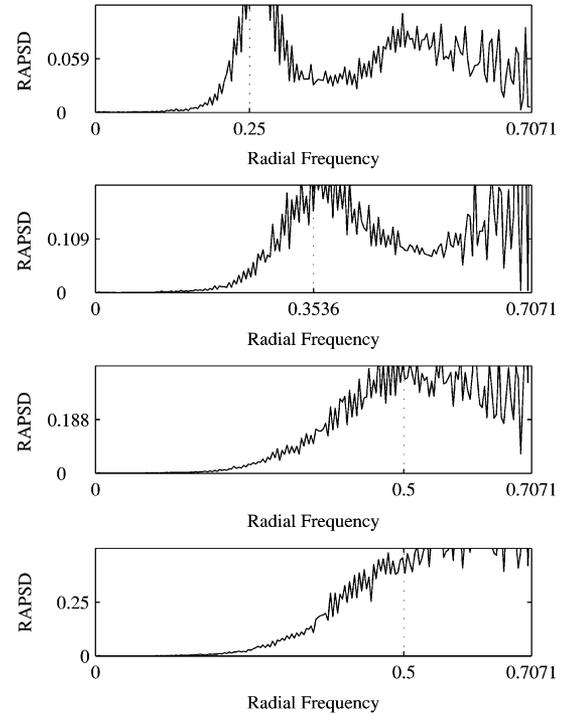


Fig. 10. RAPSD of halftones generated with DBS for gray levels $(1/16)$, $(1/8)$, $(1/4)$, and $(1/2)$.

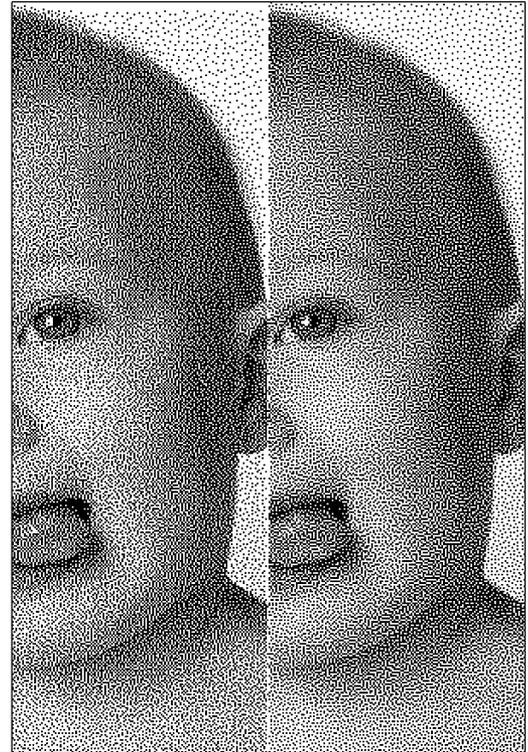


Fig. 11. Section of a blue-noise halftone of a 8-bit grayscale image generated with: Error diffusion with (left) perturbed weights and (right) DBS.

Each multitone pixel $M[\mathbf{n}]$ is thus considered a realization of a discrete random process obeying a probability density function

$$P(M[\mathbf{n}] = g_i) = p_i \quad |_{i=0}^N \quad (11)$$

where the probabilities $p_i|_{i=0}^N$ indicate the proportion of pixels of the corresponding inks included in the multitone, such that $\sum_{i=0}^N p_i = 1$. Furthermore, the probabilities are such that the mean or expected value of $M[\mathbf{n}]$, $E(M[\mathbf{n}]) = \sum_{i=0}^N p_i g_i = g$, while the variance is given by

$$\begin{aligned} \text{Var}(M[\mathbf{n}]) &= E(M^2[\mathbf{n}]) - (E(M[\mathbf{n}]))^2 \\ &= \left(\sum_{i=0}^N p_i g_i^2 \right) - g^2 = \sigma^2. \end{aligned} \quad (12)$$

The analysis and synthesis of multitones presents new challenges when compared to halftones. First, effects in the spatial domain should be evaluated since the average intensity or the textures of the dither pattern can be affected by the possible superposition of dots of different intensities or by clustering of different kinds of pixels. Second, the spectral domain analysis of multitones becomes more complex as the number of inks increases. Patterns formed with dots of the same ink will have their own spectral profile and their combination generates spectral cross terms.

A multitone can be thought of as the superposition of a series of halftone patterns printed on top of each other with different inks. In that sense, it can be related to color halftoning where three or more halftones, one for each one of the primaries used by the printing system, are superimposed in order to generate the appearance of a continuous tone color image. One of the phenomena observed when overlapping halftones is the appearance of *moiré*, a low-frequency interference pattern that introduces structural artifacts observed initially in AM color halftones. Lau *et al.* [28] showed that the *moiré* phenomenon appears in dispersed dot patterns as random fluctuations in texture referred to as *stochastic moiré*. The variations were shown to come from the lack of correlation between the dot locations in the different overlapped patterns. A similar observation was described by Wang and Parker [29]. They noticed that the superimposition of two blue-noise patterns did not necessarily result on a good quality pattern. In the spectral domain, they noticed that the spectrum of the combined pattern is a function of the spectrum of the individual patterns and the correlation between them. Their conclusion was that, in order to obtain a good quality combined pattern, the energy in the cross correlation must compensate for the energy present in the individual patterns that should not appear in their superposition.

The observations above provide a strong motivation to incorporate the correlation between different inks into the analysis and synthesis of multitones. To this end, a simple yet elegant method is proposed based on the threshold decomposition representation of signals [30], [39]. Threshold decomposition states that a discrete signal taken on one of k possible values can be represented as the weighted sum of $k - 1$ binary signals. For the case of multitoning, define M as the multitone dither pattern and the series of halftones $H_i|_{i=1}^N$ as

$$H_i[\mathbf{n}] = \begin{cases} 1, & \text{if } M[\mathbf{n}] \geq g_i \\ 0, & \text{else.} \end{cases} \quad (13)$$

The halftone H_i represents the level i threshold decomposition of the multitone M . According to this definition, a printed pixel

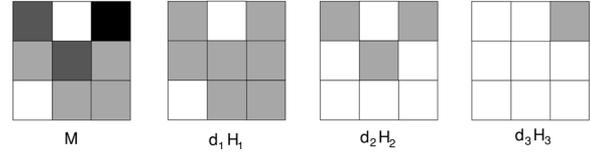


Fig. 12. Decomposition of a 3-ink multitone M in a series of halftones $H_i|_{i=1}^3$ satisfying the stacking constraint.

in H_i indicates that a printed pixel of intensity g_i or darker appears in the multitone in the same position. This also implies that there is a printed pixel in the same position in H_j for all $j \leq i$. That is, the halftones in the threshold decomposition of M stack. The multitone can be described in terms of its threshold decomposition representation as

$$M[\mathbf{n}] = \sum_{i=1}^N d_i H_i[\mathbf{n}] \quad (14)$$

where $d_i = g_i - g_{i-1}|_{i=1}^N$ are the relative differences between the intensities of the printable inks. An example of how this decomposition is performed is shown in Fig. 12. The multitone M is a 3×3 image printed with three inks with intensities $(g_1, g_2, g_3) = (1/3), (2/3), 1$.

The set of halftones $H_i|_{i=1}^N$ can be described as a set of correlated stochastic processes whose marginal densities are

$$P(H_i[\mathbf{n}]) = \begin{cases} \sum_{j=i}^N p_j, & \text{for } H_i[\mathbf{n}] = 1 \\ \sum_{j=0}^{i-1} p_j, & \text{for } H_i[\mathbf{n}] = 0 \end{cases} \quad (15)$$

with means and variances given by

$$\mu_i = \sum_{j=i}^N p_j \text{ and } \sigma_i^2 = \mu_i(1 - \mu_i). \quad (16)$$

The mean of the multitone can be expressed as a function of the characteristics of the halftones H_i as

$$E[M] = E \left[\sum_{i=1}^N d_i H_i \right] = \sum_{i=1}^N d_i \mu_i = g \quad (17)$$

and since the random processes H_i are correlated, the variance of its linear combination is

$$\begin{aligned} \text{Var}(M) &= \text{Var} \left(\sum_{i=1}^N d_i H_i \right) \\ &= \sum_{i=1}^N d_i^2 \text{Var}(H_i) \\ &\quad + 2 \sum_{i=1}^N \sum_{j=i+1}^N d_i d_j \text{Cov}(H_i, H_j) \end{aligned} \quad (18)$$

where $\text{Cov}(H_i, H_j) = E(H_i H_j) - E(H_i)E(H_j)$ is the covariance of the random processes H_i and H_j . The product $H_i H_j$ with $j \geq i$ is equal to H_j , thus the covariance reduces to

$$\text{Cov}(H_i, H_j) = \mu_j(1 - \mu_i), \text{ for } j \geq i. \quad (19)$$

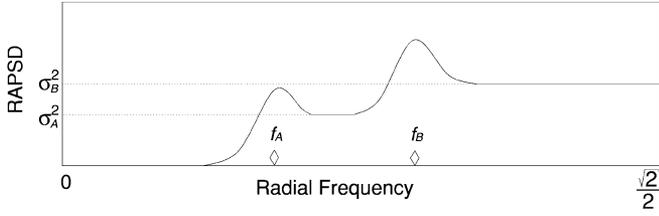


Fig. 13. Optimal RAPSD for a 2-ink multitone dither pattern. The frequencies f_A and f_B are the principal frequencies of the halftone patterns obtained by the threshold decomposition of the multitone, σ_B^2 is the variance of the multitone and σ_A^2 is the variance of the halftone pattern with the lowest principal frequency.

Replacing (19) in (18) yields

$$\begin{aligned} \text{Var}(M) &= \sum_{i=1}^N d_i^2 \sigma_i^2 + 2 \sum_{i=1}^N \sum_{j=i+1}^N d_i d_j \mu_j (1 - \mu_i) \\ &= \sum_{i=1}^N d_i^2 \sigma_i^2 + 2 \sum_{i=1}^N \sum_{j=i+1}^N d_i d_j \sqrt{\frac{\mu_j(1 - \mu_i)}{\mu_i(1 - \mu_j)}} \sigma_i \sigma_j. \end{aligned} \quad (20)$$

The variance of the multitone thus results as the weighted sum of the variances of each one of the halftones in the threshold decomposition plus a weighted sum of cross terms that indicate the interactions between dots of different intensities.

B. Multitone Blue-Noise Spectra

Assume a multitone M is represented as the superposition of a series of halftones as indicated in (14). Suppose the H_i s are blue-noise binary dither patterns; their spectra will have the shape indicated in Fig. 3 with amplitudes σ_i^2 as in (16) and principal frequencies f_i given by

$$f_i = \begin{cases} \sqrt{\mu_i}, & \text{for } \mu_i < \frac{1}{4} \\ \frac{1}{2}, & \text{for } \frac{1}{4} \leq \mu_i \leq \frac{3}{4} \\ \sqrt{1 - \mu_i}, & \text{for } \mu_i > \frac{3}{4} \end{cases}. \quad (21)$$

If the multitone dither pattern was generated as white noise, its spectrum would be flat and have an amplitude σ^2 as indicated in (12). Since the multitone is a linear combination of stacking binary blue-noise patterns, the spectrum of the aggregate should preserve some of the spectral characteristics of the individual patterns. For example, these patterns have a low-frequency response close to zero and a flat high-frequency response originated by the elimination of clustering and the preservation of the high-frequency components of white noise. These characteristics are also required for a multitone dither pattern. The mid-frequency range, however, should exhibit accumulation (peaks) of energy around the principal frequencies of each one of the halftones in the threshold decomposition representation. An example for the case of two inks is shown in Fig. 13.

It has been stated earlier that there must be correlation between the patterns $H_i|_{i=1}^N$ in order for the multitone to be visually pleasant. That correlation has yet to be characterized. The theory of statistical signal analysis provides a series of measures

that allow the analysis and quantification of the relationship between two or more signals in the spectral domain. Since the covariance allows us to study the similarities or differences between two signals in the time/spatial domain, it is natural for its Fourier transform, the *cross-spectral density function* (CSD), to be the first choice when analyzing the correlation of two signals in the spectral domain. This function is complex-valued and, in consequence, we should resort to analyze its magnitude and phase. The magnitude of the CSD is known as the cross-amplitude spectrum and it represents the average value of the product of the components of each signal for each frequency. The phase of the CSD, the phase spectrum, is the average phase-shift between the components of the two signals at each frequency [31]. Since the CSD is the Fourier transform of the cross correlation of the two signals, it can be calculated directly in the frequency domain by multiplying their PSDs, in consequence, spectral peaks corresponding to only one of the two signals can appear in the CSD even when there is no real relationship between the signals at that point. In order to avoid these kind of effects, another measure of the correlation of two signals in the frequency domain is required. The *magnitude squared coherence* function (MSC) is defined as the normalized modulus of the cross-power spectrum [32]

$$K_{xy}^2 = \frac{|P_{xy}(f)|^2}{P_x(f)P_y(f)} \quad (22)$$

where $P_{xy}(f)$ is the CSD of the patterns, $P_x(f)$ and $P_y(f)$ are the PSDs of the signals x and y respectively. Equation (22) is the frequency domain equivalent of the correlation coefficient and it can be interpreted as a measure of the correlation of two signals at each frequency. The correlation coefficient is defined as

$$r_{xy}^2 = \frac{(\text{Cov}(x, y))^2}{\sigma_x^2 \sigma_y^2}. \quad (23)$$

Applying this definition to a pair of sub-halftones $H_i, H_j (j \geq i)$ and replacing the variances and covariances with the values calculated previously leads to

$$r_{ij}^2 = \frac{(\text{Cov}(H_i, H_j))^2}{\sigma_i^2 \sigma_j^2} = \frac{\mu_j(1 - \mu_i)}{\mu_i(1 - \mu_j)} \quad (24)$$

so the terms in the last sum in (20) can be calculated as $d_i d_j r_{ij} \sigma_i \sigma_j$.

The MSC has several interesting properties, it is bounded between 0 and 1, being 0 for independent processes and 1 for signals that are linearly related (one is the result of filtering the other with a linear filter). In general, the MSC represents the portion of power of a signal at a given frequency that can be accounted for by its linear regression on the other. The coherence is invariant to linear filtration and is symmetric ($K_{xy}^2 = K_{yx}^2$). To apply the MSC to the analysis of multitones, one must calculate it for the different pairs of halftones $H_i|_{i=1}^N$ defined in (13). The method of averaged periodograms used to estimate the PSDs can also be used to estimate the CSDs. Since the patterns analyzed are isotropic, the radial average of the MSC is representative of the behavior of the function in the 2-D plane and can

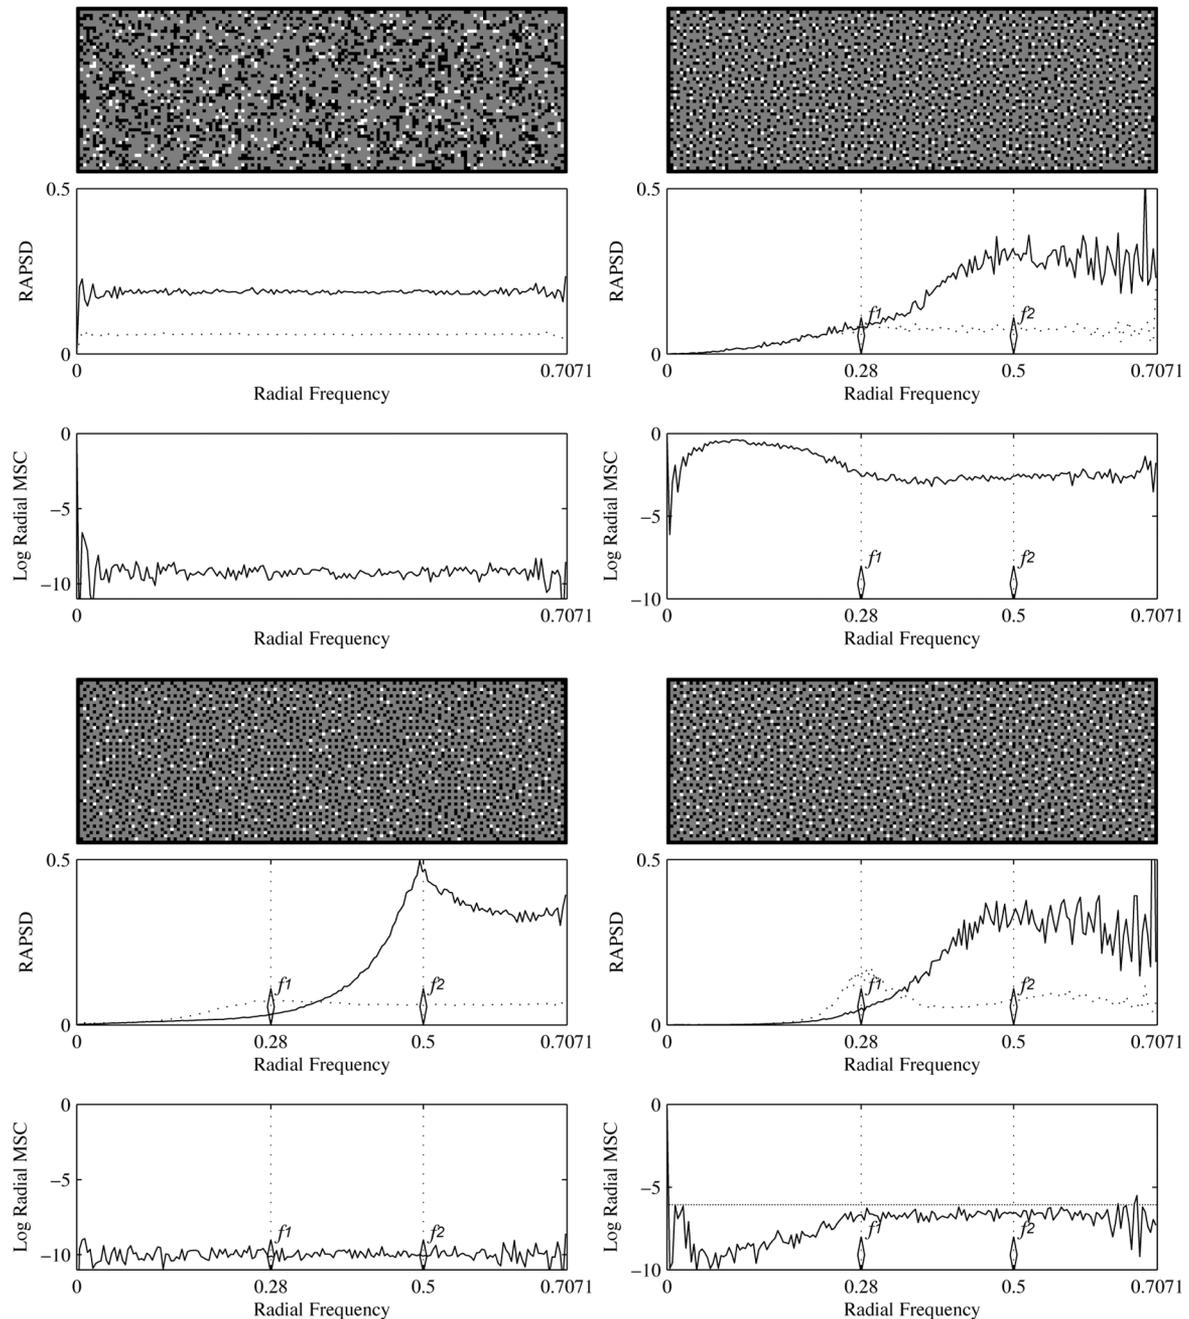


Fig. 14. Radial MSC of multitones of gray 150 generated as the superposition of (top-left) two independent white noise patterns, (bottom-left) two independent blue-noise patterns, (top-right) a suboptimal multitone generated with DBS, and (bottom-right) an optimal blue-noise multitone, with the RAPSD and the radial MSC of the patterns used for their generation.

be analyzed instead. Assuming that K periodograms are used to estimate the MSC, a value of $(1/K)$ or $-K$ dB is considered noise and indicates that the two patterns are independent.

To study the behavior of the MSC, a series of multitones were generated using different mechanisms. A representative example is shown in Fig. 14. The figure shows multitones of gray 150 generated with the same inks and ink concentrations, but with different methods. It also shows the RAPSD of the corresponding sub-halftones and the logarithm of their radial MSC. The top-left plot is a multitone generated using independent white-noise. The RAPSD of the sub-halftones are flat, as

expected and so is their MSC, whose low value (-10 dB corresponding to $K = 10$ periodograms) reflects no correlation between the patterns for any frequency. The bottom-left plot corresponds to the superposition of two blue-noise patterns generated independently. The pattern looks noisy and the lack of correlation between the patterns shows as an almost constant low level for the coherence, similar to the one observed for white noise. On the other hand, the RAPSD plots show that each of the patterns used to create this multitone are blue-noise, this indicates that some correlation needs to be introduced between the blue-noise patterns used to generate the multitone. The top-right

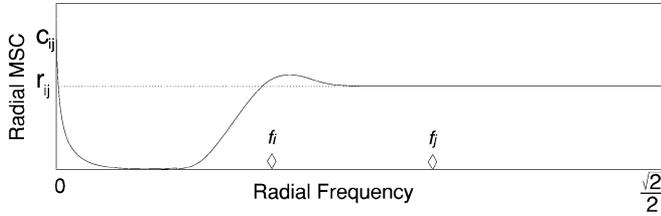


Fig. 15. Radial MSC of an ideal blue-noise multitone. f_i and f_j indicate the principal frequencies of the halftone patterns being evaluated, r_{ij}^2 is their correlation coefficient, and C_{ij} is calculated as in (25).

plots show a multitone generated with DBS for multitoning as in [33]. The pattern looks more uniform but presents clustering of minority pixels (white pixels for H_1 and black pixels for H_2). Again, the RAPSD of the sub-halftones shows they are blue-noise but the coherence shows an inadequate correlation between them (high values for the lower frequency band) that results in the artifacts mentioned before. The bottom-right pattern was generated following the blue-noise theory developed in this work. The pattern is the most visually pleasant of the four. Its radial MSC plot shows low coherence values for the lower frequencies and an almost constant value for all frequencies above the lowest cut-off frequency of the multitone. A large value of the MSC can also be observed in the very low-frequency band. This is related to the DC component of the patterns and the lack of energy for very low frequencies. The value in zero can be obtained by evaluating (22) in $f = 0$. The CSD can be replaced by the product of the PSDs of the subpatterns. The value of the PSD of each pattern in $f = 0$ is the square of their DC value, that is, the mean of the pattern. This results in

$$\begin{aligned} K_{ij}^2 |_{f=0} &= \frac{(P_i(0)P_j^*(0))(P_j(0)P_i^*(0))}{P_i(0)P_j(0)} \\ &= \mu_i^2 \mu_j^2 = C_{ij}. \end{aligned} \quad (25)$$

The conclusion obtained from these results is that the spectral coherence of the sub-halftones in a blue-noise multitone should be low for the low-frequency band and rise to the value of the correlation coefficient r_{ij} for all frequencies above the lowest principal frequency of the patterns being evaluated. An ideal plot of this function is shown in Fig. 15.

C. Blue-Noise Multitoning

Several algorithms for multitoning have been proposed in the literature, mostly as extensions of previously developed halftoning algorithms. For example, the error diffusion algorithm was modified by replacing the binary thresholding by a multilevel quantizer (Gentile *et al.* [34]), correlated error diffusion was applied to channels that represented the available inks (Faheem *et al.* [35]), screening was extended to multitoning using Bayer dither arrays [34] and clustered-dot dither [36]. Iterative algorithms for multitoning have been proposed based on neural networks [37] or as an extension of DBS [33]. The latter is also applied to the design of a multitoning dither array. Some of these algorithms introduce the concept of a gray level schedule/distribution. Its objective is to define and control the amount of each of the printable inks used to generate a certain gray level. This concept gives such algorithms some

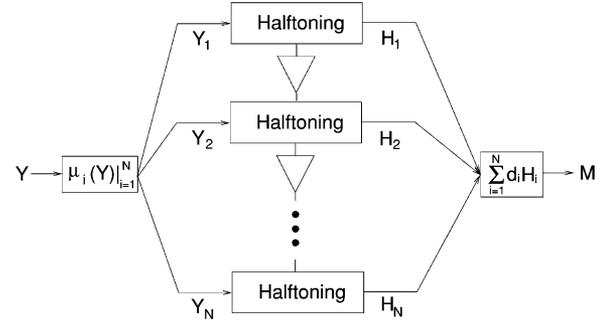


Fig. 16. Blue-noise multitoning. A continuous tone image Y is divided in N components that can be halftoned with any algorithm in a correlated fashion to generate a set of halftones, the threshold decomposition representation of the final multitone. The set of halftones $H_i |_{i=1}^N$ is recombined to generate the multitone using (14).

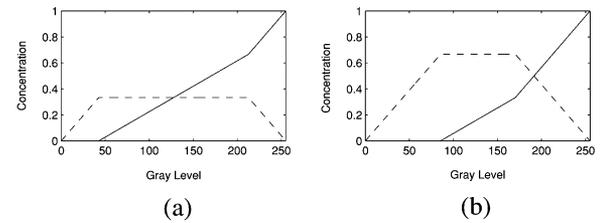


Fig. 17. Two different concentrations of (solid) black and (dashed) gray inks to use with blue-noise multitoning.

extra versatility since these schedules can be defined in several different ways. On the other hand, the results obtained with different schedules are compared and evaluated subjectively, no criteria for optimality is proposed, and the spectral analysis of the results is absent or very limited.

In order for a multitone to be optimal according to the theory just shown, the dots of different inks should be located in a correlated fashion in order to attain the spectral profile required. In previous sections, threshold decomposition was used to break down multitones into halftones in order to facilitate its analysis. A similar scheme can be applied to the synthesis of multitones of continuous tone pictures to ensure optimality. Assume a constant patch of intensity g is to be reproduced using the inks $g_i |_{i=1}^N$, in proportions $p_i |_{i=1}^N$. The intensity of the patch can be represented as

$$g = \sum_{i=1}^N g_i p_i(g) = \sum_{i=1}^N d_i \mu_i(g) \quad (26)$$

where $\mu_i(g) = \sum_{j=i}^N p_j(g)$ and $d_i = g_i - g_{i-1}$. If a patch of intensity $\mu_1(g)$ is halftoned using blue noise, the resulting dither pattern will have the same statistics as H_1 , the level one threshold decomposition representation of an ideal multitone as defined in (13). The process is repeated for a patch of intensity $\mu_2(g)$ with the constraint that the resulting halftone should stack on the first one. The resulting dither pattern holds the same properties required by H_2 . If the procedure is repeated for the remaining $\mu_i(g) |_{i=3}^N$ ensuring that the i th halftone stacks on the $i - 1$ st, the result is a series of N halftones that meet all the requirements indicated in (13) to (17). In consequence, a linear combination of these halftones results in an optimal blue-noise

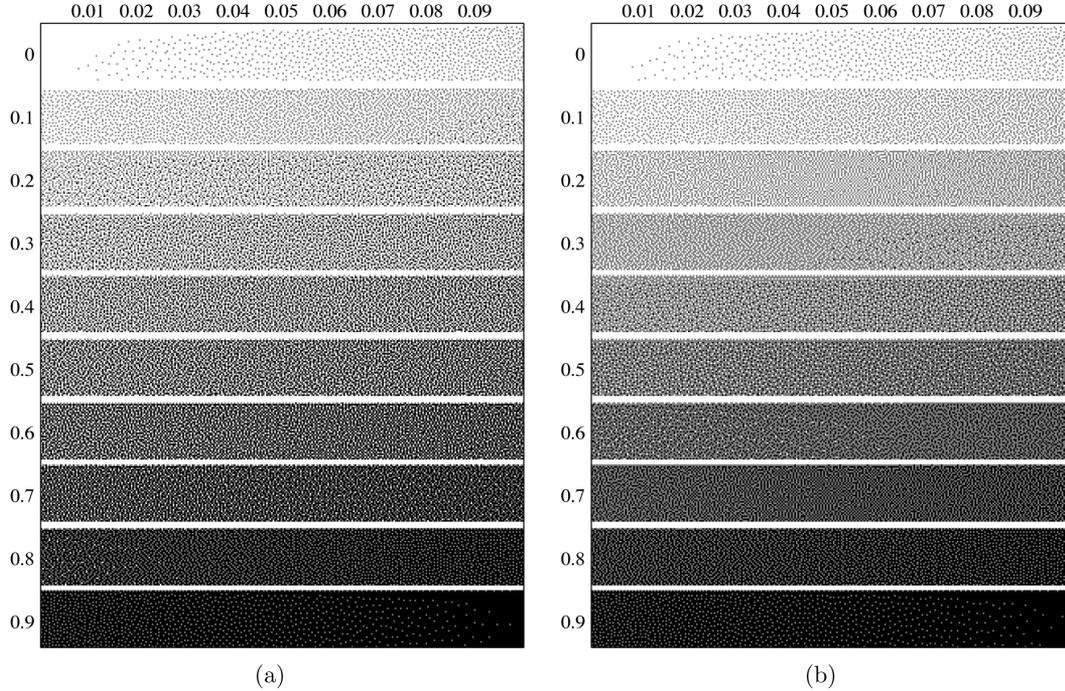


Fig. 18. Multitones of a gray-scale ramp generated with blue-noise multitoning error diffusion using the gray level concentrations in Figs. 17(a) and (b), respectively.

multitone. The procedure to multitone a continuous tone image Y , shown in Fig. 16, will be as follows.

- 1) The ink intensities g_i are obtained from the printing mechanism to be used and their concentrations $p_i(g)$ should be determined *a priori* by the user.
- 2) Define the set of sub-images $Y_i|_{i=1}^N$, where $Y_i[\mathbf{n}] = \mu_i(Y[\mathbf{n}])$, and $\mu_i(g)$ is defined as in (16).
- 3) Using the selected halftoning algorithm, halftone Y_1 to obtain H_1 .
- 4) Halftone Y_2 to obtain H_2 using the same algorithm and taking into account the stacking constraint. Follow the same procedure with the remaining $Y_i|_{i=3}^N$ (a more detailed description of this process for error diffusion and DBS is presented later on.)
- 5) Use (14) to obtain the final multitone.

The mechanism proposed is fairly simple since its implementation is based on the repeated execution of well known halftoning algorithms. The division of the original image into subimages can be implemented with a look-up table and the synthesis of the final multitone from the subhalftones is just a linear combination. Examples of this structure applied to well-known halftoning algorithms follow.²

1) *Blue-Noise Multitoning With Error Diffusion*: In order to generate blue-noise multitones by means of error diffusion, the stacking constraint should be involved in the quantization of the pixels such that

$$H_i[\mathbf{n}] = \begin{cases} 1, & \text{if } Y_i[\mathbf{n}] + H_i^e[\mathbf{n}] \geq \frac{1}{2} \text{ and } H_{i-1}[\mathbf{n}] = 1 \\ 0, & \text{else} \end{cases} \quad (27)$$

²These algorithms were introduced by the authors in [38], they are reintroduced here since the blue-noise multitone theory justifies the results shown in the previous publication.

where $H_i^e[\mathbf{n}]$ is the error diffused to the pixel $H_i[\mathbf{n}]$ and $i = 1, \dots, N$. If $i = 1$, it is assumed that $H_0[\mathbf{n}] = 1 \forall \mathbf{n}$.

2) *Blue-Noise Multitoning With DBS*: In order to incorporate DBS as the halftoning algorithm to use in the multitoning structure in Fig. 16 a few considerations need to be made. Assume DBS is applied to the sub-half-tone H_i . When a toggle or a swap is performed, it is mandatory to ensure that the stacking constraint is maintained. In consequence, a change of the pixel $H_i[\mathbf{n}]$ from a “1” to a “0” will require that all the pixels $H_j[\mathbf{n}]|_{j=i+1}^N$ are changed to zero. If the change is the opposite (from a “0” to a “1”), all the pixels $H_j[\mathbf{n}]|_{j=1}^{i-1}$ have to be changed to “1.” Since a change in a pixel in one of the sub-half-tones implies a change in several of them, the quality metric used to determine if a change is accepted needs to include all sub-half-tones. Such a metric could be defined as

$$E = \sum_{i=1}^N E_i$$

$$\text{where } E_i = \int |\hat{H}_i(\mathbf{x}) - \hat{Y}_i(\mathbf{x})|^2 dx. \quad (28)$$

The efficient implementation of the algorithm described by Analoui and Allebach should be applied to each sub-half-tone independently, taking into account the previous considerations.

IV. SIMULATIONS

To test the effectiveness of the algorithms described in the previous section a series of examples is shown as follows: A grayscale ramp is multitoned using both, error diffusion and DBS with the gray level concentrations indicated in Fig. 17 and inks $g_1 = 127$ (dashed) and $g_2 = 255$ (solid).

Fig. 18 shows the results obtained with blue-noise multitoning with error diffusion and Fig. 19 the ones from blue-noise

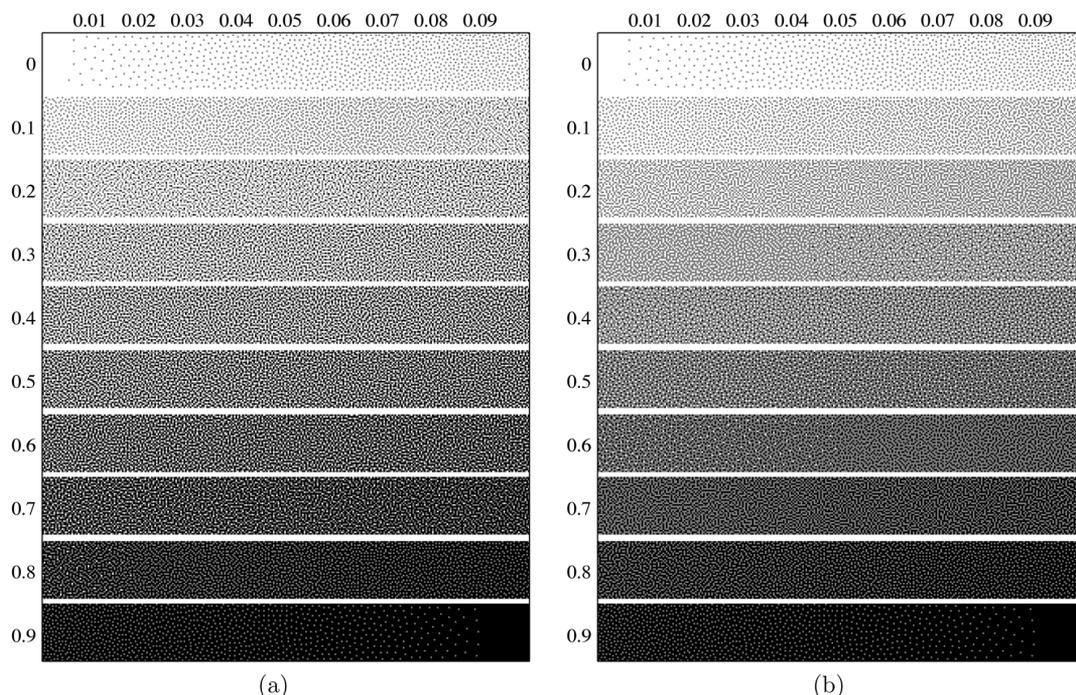


Fig. 19. Multitones of a gray-scale ramp generated with blue-noise multitoning DBS using the gray level concentrations in Fig. 17(a) and (b), respectively.

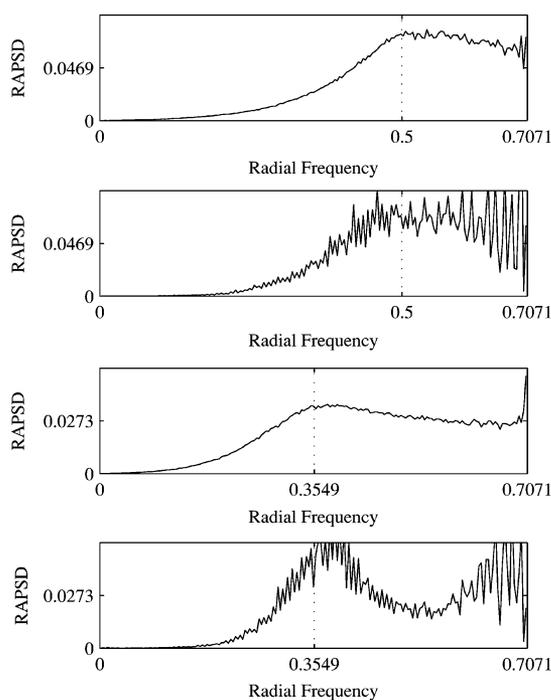


Fig. 20. RAPSD of multitones generated for both gray level distributions in Fig. 17. Top to bottom: ED for gray level (1/16), DBS for gray level (1/16), ED for gray level (1/8), DBS for gray level (1/8).

multitoning with DBS. The improvement in the quality obtained from the introduction of the gray ink is fairly evident when these plots are compared with Figs. 7 and 9. The RAPSD of these multitones for gray levels (1/16) and (1/8) are shown in Fig. 20. Only one plot is shown for each intensity/method since the ink levels for both gray level distributions are the same for these intensities. These tones are reproduced as patterns of

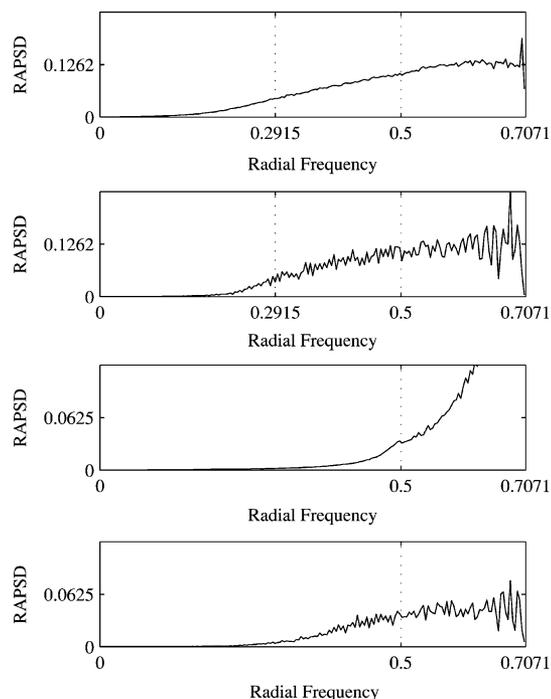


Fig. 21. RAPSD of multitones for gray level (1/4). Top to bottom: ED with the gray level distribution in Fig. 17(a), DBS with the same gray level distribution, ED with the gray level distribution in Fig. 17(b), DBS with the same gray level distribution.

gray dots over a white background and, in consequence, their RAPSD correspond to the ones of blue-noise halftones with the same principal frequencies, as indicated in the figure. Since this patterns are created using one ink, no analysis of the spectral coherence is necessary.

Fig. 21 shows the RAPSD of the multitones obtained with both methods and both gray level concentrations for gray level

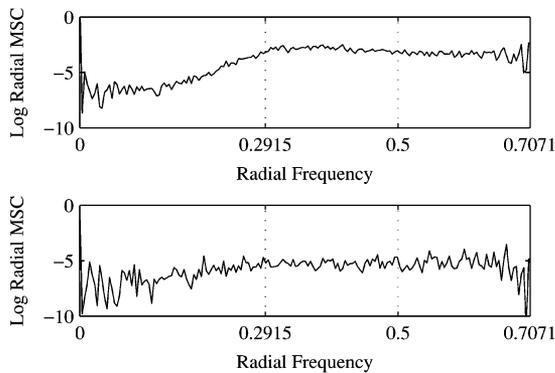


Fig. 22. MSC of the multitones generated with (top) blue-noise error diffusion and (bottom) DBS for gray level (1/4) and the gray level distribution in Fig. 17(a).

(1/4). The plots on top correspond to the gray level concentrations in Fig. 17(a). For this case, the patterns are generated using 33% gray pixels and 8.5% black pixels approx. The RAPSD for this case presents the continuous ascending pattern predicted, with the steps being more obvious in the picture corresponding to DBS. The corresponding MSC for these patterns is shown in Fig. 22. A behavior similar to the one predicted can also be noticed. The patterns present lower values of the MSC for values under the first principal frequency and higher values from there on. The plots on the bottom of Fig. 21 correspond to the gray level concentration in Fig. 17(b). In this case, the pattern is still made up of only gray and white pixels; hence, the RAPSD resembles the one of a blue-noise halftone and the analysis of the MSC is not applicable.

Figs. 23 and 24 show the RAPSD and MSC of the patterns obtained when applying both methods with both gray level concentrations to a pattern of intensity (1/2). For the concentrations in Fig. 17(a), the multitone is obtained as the superposition of two halftones with principal frequency 0.5. The RAPSD plots show an increase in the power around this value of frequency. This is the case for the plots of the MSC as well. For the concentrations in Fig. 17(b), the principal frequencies of the patterns are almost the same so the behavior of the RAPSD and the radial MSC are very similar in both cases.

Finally, Fig. 25 shows the results obtained applying blue-noise multitoning error diffusion and DBS to a natural image with different gray level concentrations. The improvement of the textures and the general appearance of the pictures is fairly evident when compared with the halftones in Fig. 11. One thing that can be noticed in this figure is the remarkably different results that can be obtained when multitoning a picture using different gray level concentrations. This leads to a critical question: is it possible to find an optimal gray level distribution to reproduce a tone for a given set of inks? This issue has been researched by the authors and it will be reported elsewhere.

V. CONCLUSION

To date, multitone dither patterns are designed through the extensions of well-known halftoning algorithms. While binary halftone patterns are well understood and optimally designed, the multitone patterns attained through the simple extensions of

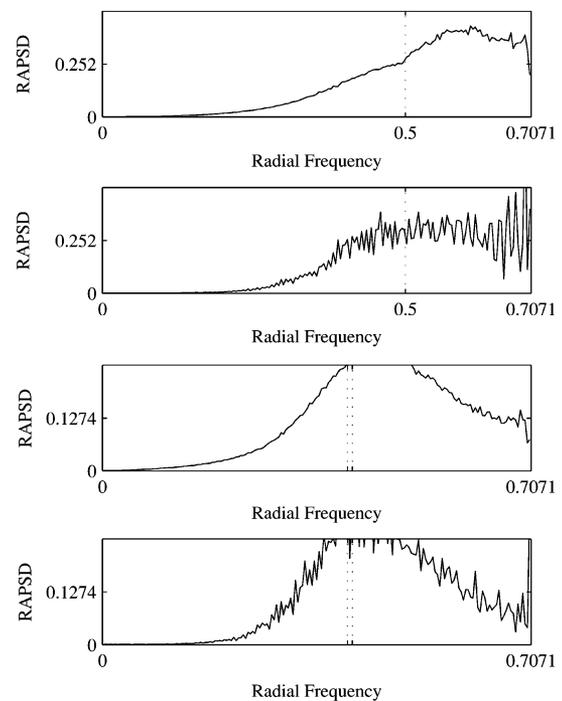


Fig. 23. RAPSD of multitones generated with blue-noise error diffusion and DBS for different gray level concentrations for gray level (1/2). Top to bottom: ED for the concentrations in Fig. 17(a), DBS for the same concentrations, ED for the concentrations in Fig. 17(b), DBS for the same concentrations.

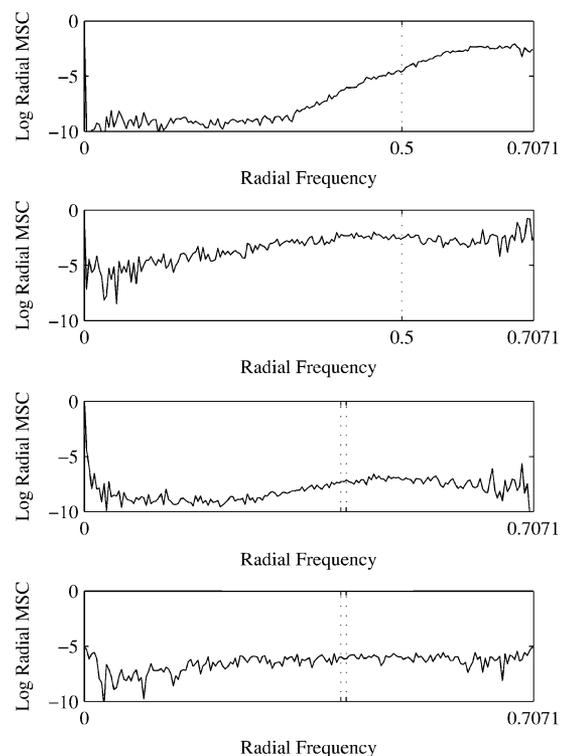


Fig. 24. MSC of multitones generated with blue-noise error diffusion and DBS for different gray level concentrations for gray level (1/2). Top to bottom: ED for the concentrations in Fig. 17(a), DBS for the same concentrations, ED for the concentrations in Fig. 17(b), DBS for the same concentrations.

binary halftone methods lack a theoretical spectral analysis like the one developed by Ulichney. Multitoning methods to date

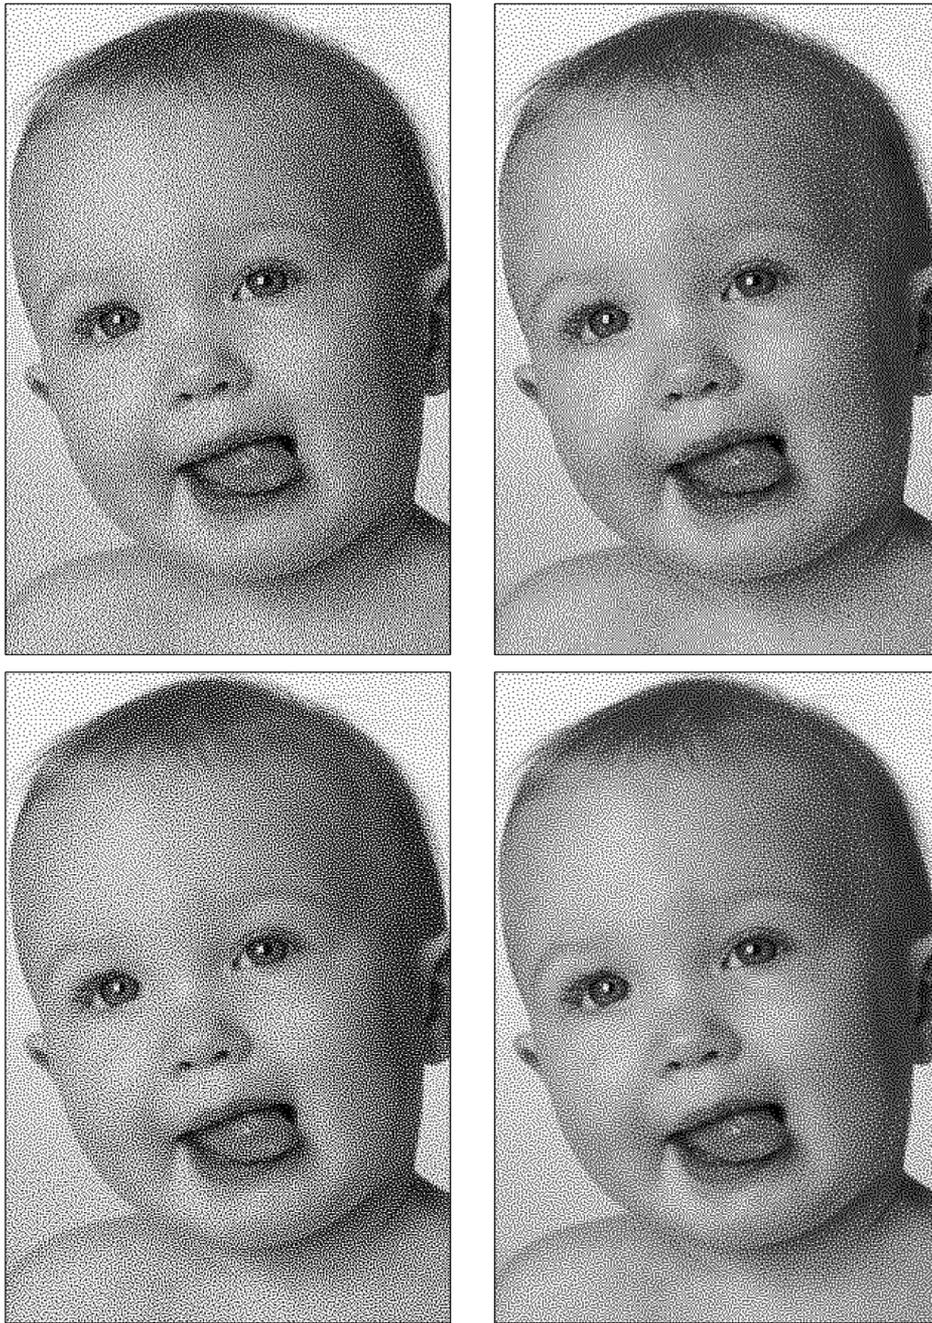


Fig. 25. Multitones of a natural image generated with: (top) blue-noise multitoning error diffusion and (bottom) DBS using the gray level concentrations in Fig. 17(a) (right) and Fig. 17(b) (left).

thus generate suboptimal patterns. This paper proposes a multitoning blue-noise model to serve as a standard by which multitoning algorithms are optimized, qualified, and categorized, in the same way that the classic blue-noise model for halftones does for binary patterns, where the better of two algorithms is the one whose output resembles the most the characteristics described in the model. To do so, we have introduced statistics for characterizing the spectral properties of multilevel dither patterns by treating the subject pattern as a stack of binary halftones, one for each of the available inks. Each one of these halftones could be characterized using blue-noise halftone dithering. The statistics of the multitone can be found as func-

tions of the characteristics of these patterns. To characterize the correlation between them we introduced a new spectral measure, the radial magnitude squared coherence function.

Using our newly introduced metrics, this paper's principal focus has been the introduction of a model characterizing the ideal spectral statistics of aperiodic, dispersed-dot, multilevel dither patterns that, like their binary counterparts, minimize low-frequency graininess.

To illustrate these results, a few examples are presented in Fig. 26. Here, patches of intensity 236, 19, and 127 are multitoned with DBS as presented in Section III-C2 and with DBS as initially introduced in [33]. The algorithms use the same HVS

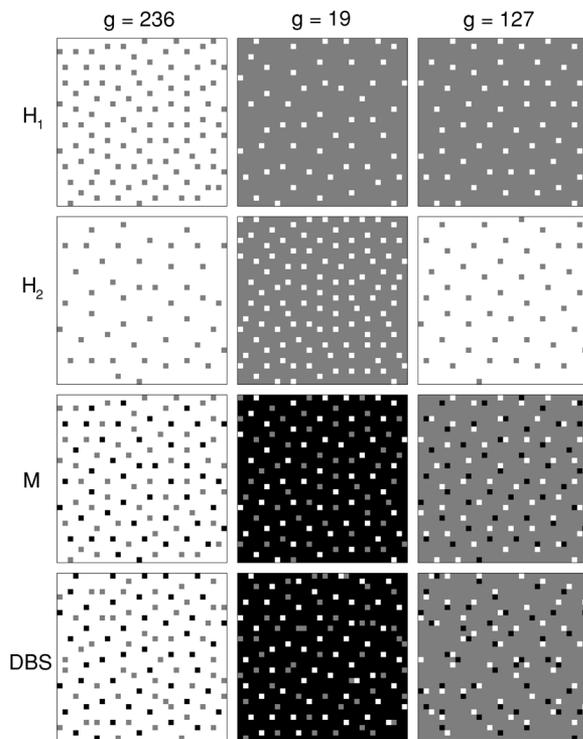


Fig. 26. Combination of stacking blue-noise halftone patterns to generate blue-noise multitones and final result compared with multitone DBS as in [33].

model and the same ink concentrations: 90% for the majority pixels and 5% for each one of the other two. The gray ink has an intensity $g_1 = 127$. The figure shows the sub-halftones H_1 and H_2 , followed by the blue-noise multitone (M) and the DBS multitone. The pictures shown are the 32×32 central sections of 256×256 original patches. The first column corresponds to the case of $g = 236$. In this case, the minority pixels in both sub-halftones will be printed pixels. H_1 is created by means of an unconstrained blue-noise halftoning algorithm to contain 10% printed pixels. H_2 is created as a subset of H_1 containing half of its pixels; such subset is chosen by the halftoning algorithm. What a good blue-noise generator will do is to select the pixels that are further apart from the ones available.

The second column of Fig. 26 corresponds to $g = 19$. This time H_1 is a blue-noise halftone pattern with 5% unprinted pixels. H_2 should contain 10% unprinted pixels but, due to the stacking constraint, half of them have to be in the same location than the unprinted pixels in H_1 . It is the responsibility of the halftoning algorithm to add the remaining unprinted pixels such that they do not form clusters with the original unprinted pixels in H_1 . This is again inherent to a blue-noise generator. The location of some of the minority pixels is forced in a pattern that is already blue-noise. The halftoning algorithm is free to locate the remaining pixels where necessary to avoid clustering and to create the required spectra.

Finally, the third column corresponds to $g = 127$. This case is different from the ones above in that the minority pixels in each sub-halftone form nonintersecting sets. H_1 is composed of 95% printed and 5% unprinted pixels while H_2 is just the opposite: 95% unprinted, 5% printed. Once H_1 is generated,

the halftoning algorithm must choose from the locations of the printed dots in H_1 , where to locate the printed pixels in H_2 . A good blue-noise generator should be able to choose those locations in such a way that the resulting pattern has its energy concentrated on the higher end of the spectrum, that is, avoiding clustering as much as possible. An algorithm like error-diffusion, that takes into account just a few pixels in the neighborhood of the one being evaluated, may locate minority pixels on H_2 (printed) close to minority pixels on H_1 (unprinted), but an algorithm that takes into account a larger vicinity of the current pixel (like DBS) should do a better job avoiding such clusters.

Now, compare the final results obtained with blue-noise DBS (third row of Fig. 26) with multitone DBS as in [33], shown in the last row of Fig. 26. Looking at the multitones in the first column, one notices how pixels in the blue-noise multitone are spread more evenly in the pattern. The minimal distance between two printed pixels in such pattern is one pixel and it appears only once (bottom right corner of the pattern). In the DBS pattern there are several pairs of pixels with such distance and there is even a pair of pixels located next to each other diagonally, that is, there is more clustering. The same happens with the patterns on the second column. In the third column, even though the quality of the blue-noise multitone is slightly lower than in the other two cases, it can be seen how the inclusion of the structure in Fig. 16 helps to break the big clusters that appear on a DBS pattern. These examples illustrate how the inclusion of the structure in Fig. 16 improves the output of well know halftoning algorithms.

REFERENCES

- [1] D. L. Lau and G. R. Arce, *Modern Digital Halftoning*. New York, New York: Marcel Dekker, 2002.
- [2] R. A. Ulichney, "Dithering with blue noise," *Proc. IEEE*, vol. 77, no. 1, pp. 56–79, Jan. 1988.
- [3] D. L. Lau and R. Ulichney, "Blue-noise halftoning for hexagonal grids," *IEEE Trans. Image Process.*, vol. 15, no. 5, pp. 1270–1284, May 2006.
- [4] G. R. Arce, *Nonlinear Signal Processing: A Statistical Approach*. Hoboken, NJ: Wiley, 2005.
- [5] R. Ulichney, "A review of halftoning techniques," in *Proc. SPIE-IS & T. Color Imaging V: Device-Independent Color, Color Hard-Copy, and Graphic Arts*, Jan. 2000, pp. 378–391.
- [6] B. E. Bayer, "An optimum method for two level rendition of continuous-tone pictures," in *Proc. IEEE Int. Conf. Rec. Communications*, Jun. 1973, pp. 11–15.
- [7] J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Trans. Syst., Man, Cybern.*, vol. 21, no. 1, pp. 33–38, Jan. 1991.
- [8] C. H. Chu, "A visual model-based halftone pattern design algorithm," presented at the IEEE Signal Processing Soc. 7th Workshop on Multi-dimensional Signal Processing, Sep. 1991.
- [9] T. Mitsa and K. J. Parker, "Digital halftoning technique using a blue-noise mask," *J. Opt. Soc. Amer. A*, vol. 9, no. 11, pp. 1920–1929, Nov. 1992.
- [10] R. A. Ulichney, "The void-and-cluster method for dither array generation," in *Proc. SPIE, Human Vision, Visual Processing, Digital Displays IV*, 1993, vol. 1913, pp. 332–343.
- [11] J. P. Allebach and Q. Lin, "FM screen design using DBS algorithm," presented at the Int. Conf. Image Processing, Lausanne, Switzerland, 1996.
- [12] K. E. Spaulding, R. L. Miller, and J. Schildkraut, "Methods for generating blue noise dither matrices for digital halftoning," *J. Electron. Imag.*, vol. 6, no. 2, pp. 208–230, Apr. 1997.
- [13] R. W. Floyd and I. Steinberg, "An adaptive algorithm for spatial grayscale," *Proc. SID*, vol. 17, no. 2, pp. 75–78, 1976.

- [14] 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. Graph. Image Process.*, vol. 5, pp. 13–40, 1976.
- [15] P. Stucki, "Mecca—A multiple-error correcting computation algorithm for bi-level image hardcopy reproduction," IBM Res. Lab. Zurich, Switzerland, Tech. Rep. RZ1060, 1981.
- [16] J. N. Shiau and Z. Fan, "A set of easily implementable coefficients in error-diffusion with reduced worm artifacts," in *Proc. Color Imaging: Device-Independent Color, Color Hard Copy, and Graphics Arts*, Mar. 1996, vol. 2658, pp. 222–225.
- [17] I. H. Witten and R. M. Neal, "Using peano curves for bilevel display of continuous-tone images," *IEEE Comput. Graph. Appl.*, vol. 2, pp. 47–52, May 1982.
- [18] L. Velho and J. M. Gomes, "Digital halftoning with space filing curves," *Comput. Graph.*, vol. 25, no. 4, pp. 81–90, Jul. 1991.
- [19] D. E. Knuth, "Digital halftones by dot diffusion," *ACM Trans. Graph.*, vol. 6, no. 4, pp. 245–273, Oct. 1987.
- [20] J. Sullivan, R. Miller, and G. Pios, "Image halftoning using a visual model in error diffusion," *J. Opt. Soc. Amer. A*, vol. 10, no. 8, pp. 1714–1724, Aug. 1993.
- [21] R. Eschbach and K. T. Knox, "Error-diffusion algorithm with edge enhancement," *J. Opt. Soc. Amer. A*, vol. 8, no. 8, pp. 1844–1850, Dec. 1991.
- [22] R. Eschbach, "Reduction of artifacts in error diffusion by means of input-dependent weights," *J. Electron. Imag.*, vol. 2, no. 4, pp. 352–358, Oct. 1993.
- [23] V. Ostromoukhov, "A simple and efficient error-diffusion algorithm," in *Proc. SIGGRAPH*, 2001, pp. 567–572.
- [24] P. Li and J. P. Allebach, "Tone-dependent error diffusion," *IEEE Trans. Image Process.*, vol. 13, no. 2, pp. 201–215, Feb. 2004.
- [25] D. L. Lau, G. R. Arce, and N. C. Gallagher, "Digital color halftoning with generalized error diffusion and multichannel green-noise masks," *IEEE Trans. Image Process.*, vol. 9, no. 5, pp. 923–935, May 2000.
- [26] N. Damera-Venkata and Q. Lin, "AM-FM screen design using donut filters," in *Proc. SPIE-IS&T. Color Imaging IX: Processing, Hardcopy, and Applications*, Jan. 2004, vol. 5293, pp. 469–480.
- [27] M. Analoui and J. P. Allebach, "Model based halftoning using direct binary search," presented at the SPIE: Human Vision, Visual Processing, and Digital Display 111, San Jose, CA, 1992.
- [28] D. L. Lau, A. M. Khan, and G. R. Arce, "Minimizing stochastic moire in frequency-modulated halftones by means of green-noise masks," *J. Opt. Soc. Amer. A*, vol. 19, no. 11, pp. 2203–2217, Nov. 2002.
- [29] M. Wang and K. J. Parker, "Properties of jointly-blue noise masks and applications of color halftoning," *J. Imag. Sci. Technol.*, vol. 44, pp. 360–370, Aug. 2000.
- [30] J. P. Fitch, E. J. Coyle, and N. C. Gallagher, Jr., "Median filtering by threshold decomposition," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 32, no. 6, pp. 1183–1188, Dec. 1984.
- [31] M. B. Priestley, *Spectral Analysis and Time Series*. London, U.K.: Academic, 1992.
- [32] R. Shiavi, *Introduction to Applied Statistical Signal Analysis*, 2nd ed. San Diego, CA: Academic, 1999.
- [33] G. Y. Lin and J. Allebach, "Multilevel screen design using direct binary search," *J. Opt. Soc. Amer. A*, vol. 19, no. 10, pp. 1969–1982, Oct. 2002.
- [34] R. S. Gentile, E. Walowitz, and J. P. Allebach, "Quantization and multilevel halftoning of color images for near-original image quality," *J. Opt. Soc. Amer. A*, vol. 7, no. 6, pp. 1019–1026, Jun. 1990.
- [35] F. Faheem, G. R. Arce, and D. L. Lau, "Digital multitoning using gray level separation," *J. Imag. Sci. Technol.*, vol. 46, no. 5, pp. 385–397, Sep/Oct. 2002.
- [36] R. Miller and C. Smith, J. P. Allebach and B. E. Rogowitz, Eds., "Mean-preserving multilevel halftoning algorithm," in *Proc. SPIE: Human Vision. Visual Processing and Digital Display TV*, 1993, vol. 1913, pp. 367–377.
- [37] P. R. Bakic and D. P. Brzakovic, "Extension of CNN based multilevel halftoning to color reproduction," presented at the Int. Conf. Image Processing, Oct. 1997.
- [38] J. Bacca Rodriguez, G. R. Arce, and D. L. Lau, "A new method for digital multitoning using gray level separation," presented at the Int. Conf. Image Processing, Oct. 2006.
- [39] G. R. Arce, "A General weighted median filter structure admitting negative weights," *IEEE Trans. Signal Process.*, vol. 46, no. 12, pp. 3195–3205, Dec. 1998.



J. Bacca Rodríguez (M'03) received the B.S. degree in electronics engineering from the Pontificia Universidad Javeriana, Bogotá, Colombia, in 2001, and the M.E.E. and Ph.D. degrees from the Department of Electrical and Computer Engineering, University of Delaware, Newark, in 2003 and 2007, respectively.

He is currently a postdoctoral researcher at the University of Delaware under Dr. G. R. Arce. His research interests include nonlinear signal processing, image processing, halftoning, human visual system modeling, and digital audio.



G. R. Arce (F'00) received the Ph.D. degree from Purdue University, West Lafayette, IN, in 1982.

Since 1982, he has been with the faculty of the Department of Electrical and Computer Engineering, University of Delaware, Newark, where he is the Charles Black Evans Distinguished Professor and Department Chairman. His research interests include statistical and nonlinear signal processing and their applications. He is a coauthor of the textbooks *Digital Halftoning* (Marcel Dekker, 2001), *Nonlinear Signal Processing and Applications* (CRC Press, 2003), and *Nonlinear Signal Processing: A Statistical Approach* (Wiley, 2004). He is a frequent consultant to industry and holds ten U.S. patents.

Dr. Arce served as an Associate Editor for several IEEE and OSA journals.



D. L. Lau (M'01) received the B.S. degree in electrical engineering (with highest distinction) from Purdue University, West Lafayette, IN, in 1995, and the Ph.D. degree from the University of Delaware, Newark, in 1999.

Currently, he is an Associate Professor with the University of Kentucky, Lexington, but has also worked as a DSP engineer at Aware, Inc., and as an image and signal processing engineer at Lawrence Livermore National Laboratory. His research interests include 3-D imaging sensors for fingerprint identification and multispectral color acquisition and display. His published works in halftoning include the introduction of the green noise halftoning model, as well as stochastic moiré.