Journal of Electronic Imaging

SPIEDigitalLibrary.org/jei

Parametrically controlled stochastically seeded clustered halftones: design strategies and applications

Edgar A. Bernal
Robert P. Loce
Shen-Ge Wang
Zhigang Fan

Parametrically controlled stochastically seeded clustered halftones: design strategies and applications

Edgar A. Bernal,^{a,*} Robert P. Loce,^a Shen-Ge Wang,^b and Zhigang Fan^a

^aXerox Research Center Webster, Xerox Corp., 800 Phillips Road, MS128-27E, Webster, New York 14580

^bQualcomm MEMS Technologies Inc., 2581 Junction Avenue, San Jose, California 95134

Abstract. We present a stochastically seeded, clustered halftoning method with parametric control of dot shapes as well as seed placement adaption to local image structure. While the literature on other randomly structured screens (blue noise, green noise, and stochastic halftoning) is quite extensive, there is very little discussion on optimization of dot shape for robustness or preferred tile appearance. The parametric control presented here can enable an imaging engineer to optimize robustness for a given printing system. The halftoning method that we introduce defines dot centers as seeds that are placed typically in a random, high spatial frequency configuration. Spot functions are defined about these randomly placed seeds, where the spot function allows control of dot cluster growth, touch points, cluster angle, and eccentricity. The spot function can also be applied to regular and irregular polygonal halftone tiling. The adaptability of the seed placement provided by the halftoning method allows for better edge rendition than conventional isotropic methods and enables the implementation of data-embedding techniques. The resulting monochromatic screens can be seamlessly integrated with novel successive filling techniques in order to provide dot-off-dot vector halftoning capabilities. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JEI.23.1.013010](https://doi.org/10.1117/1.JEI.23.1.013010)]

Keywords: halftoning; stochastic screen; blue noise; green noise; spot function; digital watermarking; dot-off-dot; vector halftoning; successive fill.

Paper 13524P received Sep. 13, 2013; revised manuscript received Dec. 9, 2013; accepted for publication Dec. 16, 2013; published online Jan. 28, 2014.

1 Introduction

New options in halftone geometry have arisen as printed pixel resolution has increased. One area of development has been the so-called stochastically seeded clustered halftones, where nuclei (seeds) are placed in a pseudorandom manner [frequency modulation (FM)] up to a particular gray level and then the nuclei are grown in size [amplitude modulation (AM)] above that level. There are several patents and publications that describe methods for growing these dots, but the published methods do not allow parametric control of the shape of the dots. Hence, while dot growth for periodic halftones is a well-developed technology for optimizing tone robustness, noise resistance, and aesthetic appearance, dot growth for stochastically seeded screens has yet received little attention in the literature. This analysis is overdue given availability of 2400 dots per inch (dpi) laser printers (e.g., Xerox iGen 150 Press, Xerox Color 550/560, Xerox DocuColor 7000/8000, etc.), which have the addressability to control dot shape at high dot frequencies.

One class of published stochastically seeded halftones is designed with a focus on frequency-domain characteristics. These green noise-like methods^{1,2} may be well suited for optimizing spectral characteristics but do not address the classical printing considerations associated with dot shape. A second class of stochastically seeded clustered screens uses random seeds and then applies a fixed threshold array to control growth around the seeds.^{3,4} While these methods attempt to control growth in the spatial domain

where better control is possible, a fixed threshold array on random seeds does not consider neighboring structure and thus tends to produce high graininess and suboptimal touch points. A third class uses parameters to control dot growth within a Voronoi tessellation formalism.^{5–10} These methods do define a spatial tessellation and attempt to control growth and touching between the tiles for the purposes of print-to-print stability and uniformity. More recently, a clustered-dot halftone technique based on direct binary search was proposed.¹¹ Although the proposed technique does allow control of halftone shape and texture parameters, such as homogeneity and clusteredness, the technique only allows indirect control of these attributes via the adjustment of filter sizes.

Existing methods offer much less control than is available for growing periodic dots. The lack of control not only affects stability and uniformity, but also does not allow dot shaping for aesthetic purposes, such as using rounder dots for faces, squarer and sharper dots for graphics, extended highlight dots (which do not touch until the shadow range is reached), and extended shadow dots (which touch early in the gray-scale range and allow control of hole shape).¹²

The halftoning technique introduced in the present paper addresses the need for highly controllable dot growth in the stochastic setting. We extend methods employed for control of periodic halftones to irregular seed structures. A spot function formalism is used, where the function defines thresholds to be applied to an input image on a pixel-by-pixel basis to control the dot cluster growth, touch points, cluster angle, and eccentricity in the halftoned image. We show that this same spot function formalism can also be applied to

*Address all correspondence to Edgar Bernal, E-mail: edgar.bernal@xerox.com

regularly placed seeds to produce various unconventional halftone tilings. The halftoning technique presented herein is an irregular extension of the recently proposed hexagonal spot function.¹³ It may also be considered an evolutionary descendent of Ron Pellar's work on spot functions, which employed orthogonal cosine functions.^{14,15}

In addition, we introduce a seed placement method that adapts positions of dot clusters to the spatial frequency content and local dominant orientation of detail within an input image. As with the spot function advancements, our seed placement method can be considered an irregular tessellation extension of periodic halftoning algorithms that modulate seed frequency based on image content.¹⁶ We also show how judicious manipulation of seed placement frequency allows for efficient data-embedding, both in a perceptible and in an invisible manner. Last, we extend the application of the proposed spot function to achieve dot-off-dot vector halftoning via the application of a novel successive fill strategy.

The remainder of this paper is organized as follows. In Sec. 2, the spot function design procedure is introduced. The application of the presented spot function design strategy to arbitrary tessellations of the plane is illustrated in Sec. 3. Section 4 introduces a modification to the baseline spot function design strategy that enables improved edge and detail rendition. Section 5 introduces a data-embedding methodology enabled by the flexible spot function manipulation that is intrinsic to the proposed halftoning technique. We develop an extension of the scalar halftone spot function design to a dot-off-dot vector halftone strategy in Sec. 6. We summarize our findings in Sec. 7.

2 Spot Function Design

The spot function that we have developed is parametrically controlled and based on triangle tessellations. It is well suited for controlled growth of halftone spots arrayed on a nonregular grid as encountered in stochastically seeded clustered halftone screens. The key steps in generating the spot function, as illustrated by the schematic diagram in Fig. 1, are as follows:

1. Determining seed locations: set of points from a source such as a stochastic screen.
2. Defining a triangle tessellation based on seed locations.
3. For each pixel, determine the triangle in which it resides.
4. For each pixel, determine its distances from its triangle bases.
5. Use the distances to determine the value of a function that generates halftone thresholds. The function should have parameters that can be adjusted to control its sharpness and slope.

Fig. 1 Schematic illustration of the steps involved in the creation of the halftone spot function values as a function of distances to bases of encompassing triangles.

2.1 Determining Seed Locations

Seed locations define the nuclei for growth of the spot function and the principal frequency of the screen. The seeds could be locations on a regular grid or irregular locations obtained from random perturbations of regular grids or from the application of an FM halftoning method to a uniform gray-level patch. Without loss of generality, we demonstrate the performance of the technique with seed locations obtained from a stochastic screen halftone of a uniform gray patch with area coverage T . Figure 2 shows a sample seed array.

2.2 Determining Triangular Tessellation and Parameters

A triangular tessellation of a plane is a juxtaposition of triangular shapes in a pattern with no overlaps and no gaps. In order to tessellate the pixel area containing pixel locations, seeds are used as vertices of the triangles, and every seed is used as a vertex. If the seeds were arranged in a regular pattern, such as a hexagonal grid, then the tessellation would be regular and formed by congruent triangles. Figure 3 shows the triangular tessellation of a portion of the seed pattern from Fig. 2.

2.3 Determining Triangle Membership

By definition, a triangular tessellation completely fills the space with no overlapping; therefore, every pixel in the image must belong to a triangle in the tessellation. For every pixel P in the image, a determination is made as to which triangle encompasses it. Figure 4 shows an example of a pixel P encompassed by triangle $P_1P_2P_3$, where P_1 , P_2 ,

Fig. 2 Example of seeds defined as activated pixels, obtained by applying a stochastic halftone threshold array to a patch with 6% area coverage.

Fig. 3 Resulting triangular tessellation of a portion of the seed pattern from Fig. 2.

and P_3 are the triangle vertices. There is no need to determine the encompassing triangle of a pixel if the pixel is a vertex. Once triangles are defined, relatively straightforward geometric relationships can be used to determine useful triangle parameters, such as triangle heights H_1 , H_2 , and H_3 .

2.4 Determining Distances Between Pixel and Triangle Bases

Distances between pixels and triangle bases are determined from projections from the pixel location onto the triangle altitudes, or similarly, distances to the respective bases. An altitude of a triangle is a straight line through a vertex and perpendicular to the opposite side. The length of the altitude is the height with respect to the given vertex. The opposite side is called the base of the altitude. The distances h_1 , h_2 , and h_3 are the distances from P to their respective bases, which are opposites P_1 , P_2 , and P_3 , respectively, as illustrated in Fig. 4.

2.5 Determining Spot Function Values

The distances calculated in Sec. 2.4 are used as inputs to a function that generates halftone thresholds. The function has

Fig. 4 Triangle $P_1P_2P_3$ encompassing pixel P . The triangle has heights H_1 , H_2 , and H_3 shown in dashed blue lines, and projected pixel heights h_1 , h_2 , and h_3 shown in dashed red lines.

parameters that can be adjusted to control its sharpness and slope. In our algorithm, the function is a weighted sum of three cosines that are functions of the three respective distances. Algebraic powers of the distances control sharpness of each dot touch point with its neighboring dots. Cosine weights control the dot touch sequencing, such that contact with neighboring dots can occur at different gray levels, thereby avoiding instability that may occur due to simultaneous touching. The spot function itself can be used to halftone an image, or a sampled version of it can be applied as a threshold array for efficient implementation in a printer. Given the pixel and triangle heights, we compute the spot function value Q of that pixel according to

$$Q_1 = a_1 \cos[2\pi(h_1/H_1)] + a_2 \cos[2\pi(h_2/H_2)] + a_3 \cos[2\pi(h_3/H_3)], \quad (1)$$

where a_i are parameters that control the elongation of the dot along the direction defined by H_i .

An additional set of parameters $\gamma_i > 0$ can be introduced for control over the roundness of the dot sides and the sharpness of corner touch points.

$$Q_2 = a_1 \cos[2\pi(h_1/H_1)^{\gamma_1}] + a_2 \cos[2\pi(h_2/H_2)^{\gamma_2}] + a_3 \cos[2\pi(h_3/H_3)^{\gamma_3}]. \quad (2)$$

Another level of control can be enabled by inverting the spot function. The inversion allows for well-controlled holes, which can be desirable for dark image subject matter.

$$Q_3 = -\{a_1 \cos[2\pi(h_1/H_1)^{\gamma_1}] + a_2 \cos[2\pi(h_2/H_2)^{\gamma_2}] + a_3 \cos[2\pi(h_3/H_3)^{\gamma_3}]\}. \quad (3)$$

Data normalization is an additional step that typically must be performed with spot functions. Once all pixels in the matrix are processed to obtain Q values, the resulting thresholds are gray-level shifted and scaled to fit the data range, such as $[0,255]$ for an 8-bit system or $[0,1023]$ for a 10-bit image path. The effect of the growth control parameters is illustrated in Figs. 5 to 10, which show gray ramps halftoned with the aforementioned spot function having different parameter values. The images in the figures when printed at 2400 spots per inch have a principal frequency somewhat >200 cells per inch (cpi); from Fourier analysis of images, we determined it to be ~ 230 cpi. We note that although constructing the spot function is a computationally expensive process, it only has to be performed once for a given set of spot function parameters. The resulting fill order matrix can then be stored for future use and used within the framework of traditional screening methods to halftone arbitrary input images. We also note that, while halftone dot

Fig. 5 Gray ramp halftoned with the spot function from Eq. (2) with $a_i = 1$ and $\gamma_i = 1$.

Fig. 6 Gray ramp halftoned with the spot function from Eq. (2) with $a_i = 1$ and $\gamma_i = 0.6$. Observe that $\gamma_i < 1$ produces a pincushion appearance. This could improve stability in some marking processes.

Fig. 7 Gray ramp halftoned with the spot function from Eq. (2) with $a_i = 1$ and $\gamma_i = 1.5$. Observe that $\gamma_i > 1$ tend to make dots rounder and delay touching at corners. This could be preferred for some marking processes or for subject matter consisting primarily of highlights.

Fig. 8 Gray ramp halftoned with the spot function from Eq. (2) when a_i is varied with H_i , and $\gamma_i = 1$. There are numerous ways in which the a_i coefficients can be controlled. This image demonstrates increasing a_i proportionally with H_i .

Fig. 9 Gray ramp halftoned with the spot function from Eq. (2) when a_i is varied with H_i and $\gamma_i = 1$. This image demonstrates increasing a_i inversely proportionally with H_i .

and hole symmetry is sometimes desired (e.g., a screen that transitions from dots with a predetermined shape to holes with the same shape at the midtones) in order to avoid undesirable textures or coarse patterns in the print,¹⁷ the stochastic nature of the proposed technique prevents it from resulting in perfectly symmetric patterns. For a technique that allows control of periodic and, if so desired, symmetrical halftone structures, we refer the reader to Sec. 3.

3 Polygonal Boundary Halftones

The spot functions of Sec. 2.5 can also be applied to regular or irregular polygonal tessellations of the plane, resulting in enhanced dot shape control that enables applications such as watermarking, security printing, and special effects. The polygonal tiling can be regular (repetitive) or irregular and can be varied in a manner adapted to the image content or to data that is being embedded. The extension of our triangular-tessellation-based technique to arbitrary polygonal tessellations of the plane is possible because polygons,

Fig. 10 Gray ramp halftoned with the spot function from Eq. (3) with $a_i = 1$ and $\gamma_i = 1$. Observe the well-controlled holes. This parameterization could be preferred for dark subject matter.

Fig. 11 (a) Pentagonal tiling of the plane. (b) Gray ramp halftoned with the spot function associated with the tiling from (a).

Fig. 12 (a) Plane tiling with hexagons, squares, and triangles. (b) Gray ramp halftoned with the spot function associated with the tiling from (a).

regular and irregular, can in turn be tessellated by triangular shapes. The spot function is based on triangulation of the polygons and center-to-boundary distances rather than on distances from a center alone, which allows for greater control over shape and overall appearance of the printed tiled halftone. Figures 11 and 12 show two sample plane tilings as well as the binary image that results from halftoning a gray ramp with the associated spot function.

Figures 13 and 14 show binary images resulting from halftoning a gray ramp with spot functions associated with irregular pentagonal and hexagonal plane tilings, respectively. Both halftones share the same seed locations.

4 Seed Placement for Enhanced Edge and Detail Rendition

The traditional stochastically seeded clustered halftone method from Sec. 2 uses dispersed-dot dithering techniques,

Fig. 13 Gray ramp halftoned with the spot function associated with an irregular pentagonal tiling.

Fig. 14 Gray ramp halftoned with the spot function associated with an irregular hexagonal tiling.

such as error diffusion or stochastic screening to determine seed locations, which typically results in a limited range of seed frequencies. Improvement in rendering image detail can be achieved by introducing texture- and orientation-adaptive seeding: the stochastic screen that best correlates with the local image characteristics is chosen among a set of pre-designed screen candidates. The resulting halftone image will possess a larger number of seeds in areas with busier spatial activity, effectively increasing the principal frequency of the halftone where improved detail rendition is required. Also, the location of the seeds can be made to follow the general orientation of the image so that rendition of edges and lines is improved regardless of their orientation. Figure 15 shows a block diagram illustrating the steps involved in the improved algorithm. In general, the key steps in defining the halftone structure are (1) for each pixel in the image, determine the local spatial activity and its dominant orientation, (2) determine seed locations according to the orientation and activity, and (3) calculate the halftone spot function based on the tessellation determined by the seed locations and binarize the image using the spot function.

4.1 Determining Dominant Orientation and Texture of Local Image Area

Image orientation and activity can be estimated in several ways. For example, the magnitude of the local gradient can serve as an estimate of the amount of spatial activity and its angle can be used to estimate the dominant orientation. Transform-based methods, such as principal component analysis, Gabor filtering, and Fourier, cosine, or wavelet decomposition, can also be used. In this paper, we propose the use of nonzero frequency (AC) coefficients in the discrete cosine transform (DCT) domain coefficients (which are readily available from JPEG-compressed data) to make

Fig. 15 Schematic representation of the presented improved halftoning algorithm with texture- and orientation-adaptive seeding.

that determination. Figure 16 shows a graphical representation of the first-level DCT basis functions used in JPEG, whose coefficients are used by the algorithm to determine dominant orientation and texture. Intuitively, a larger coefficient (0,1) indicates variation in the vertical direction of the local image block (typically 8×8 pixels). Similarly, the absolute values of coefficients (1,0) and (1,1) are proportional to the horizontal and diagonal edge activity of the local image block, respectively. Given these relationships, the magnitude of the AC coefficients can be used to estimate the principal orientation of activity within the image block. For simplicity, the current implementation of the algorithm approximates the local orientation to one of four possible angle values, from -45 to 90 deg in steps of 45 deg by using only the first-level JPEG decomposition coefficients, but a more accurate estimate of orientation could be achieved by using a larger number of coefficients. The estimated angle of orientation corresponds to that of the JPEG coefficient with the largest absolute value. If the largest coefficient is that in position (1,1), then the sign is also taken into account: a positive coefficient is indicative of a -45 -deg angle, whereas a negative coefficient is indicative of a 45 -deg orientation. More accurate estimates of local dominant directions can be achieved by using lower-level decomposition coefficients.

JPEG coefficients also carry information on the amount of local spatial variation. The current implementation of the algorithm estimates spatial activity by summing the first-level AC coefficients of the DCT decomposition of each image block (coefficients at coordinates (0,1), (1,0), and (1,1) in Fig. 16). Specifically, the first-level DCT coefficients are indicators of the main spatial component in each image block.

4.2 Determining Seed Locations According to Orientation and Spatial Activity

A halftone screen is designed for each of a selected set of dominant image spatial activity orientations. Once the dominant direction of the spatial activity in the local image block is determined, its corresponding halftone screen is used to halftone an 8×8 pixel block. The effects of this operation are twofold: first, the location of the seeds will be more aligned along the direction of the overall orientation of the image block; second, image blocks with larger spatial variation will have a larger number of seeds, effectively

Fig. 16 Graphical representation of first-level DCT basis functions.

Fig. 17 Anisotropic stochastic screens used to determine seed locations (principal frequencies) for image blocks at (a) 0 deg, (b) 90 deg, (c) -45 deg, and (d) 45 deg.

increasing the local principal frequency of the halftone in an orientation consistent with high image activity, thus improving detail rendition in those areas. For demonstration purposes, four anisotropic stochastic screens (one for each angle estimated in step 1) were designed, as illustrated in Fig. 17.

The four anisotropic screens have a common subset of isotropically distributed dots to facilitate transitions between regions of different orientation; Fig. 18 shows this common subset of isotropic dots, which is also the identical output of the above four stochastic screens when halftoning a uniform input image at the low common gray level.

Figure 19(b) shows the image resulting from halftoning the image of Fig. 19(a) with the algorithm described above.

Figure 20 illustrates the enhanced edge and detail-rendering capabilities of the texture/orientation-adaptive seeding algorithm introduced in Sec. 4. Note how the number of seeds increases on edge regions, independent of the gray values of the image, effectively increasing the principal frequency of the halftone around areas of increased spatial activity. Also note the adaptive anisotropy, where the seed arrangement follows the edge orientation, even for cases where the edge has continuously changing orientation.

For comparison purposes, Fig. 21 shows the same halftone image area of Fig. 20(a) but obtained with the uniform isotropic principal frequency algorithm from Sec. 2. Figure 21(a) shows the halftoned image obtained with the common subset of seeds (Fig. 18), where the low-frequency

rendition of the uniform sky would be printed in a stable low-noise manner, but the high-frequency content of the image is poorly represented by the low principal frequency of the screen. Figure 21(b) shows the opposite effect, where a union of the screens in Fig. 17, representing an isotropic high principal frequency screen, was used to better resolve fine image detail. The opposite problem occurs where the

Fig. 18 Isotropic set of dots common to the four anisotropic screens depicted in Figs. 17(a) to 17(d).

Fig. 19 (a) Continuous tone image and (b) binary image resulting from halftoning the image in (a) using the algorithm of Sec. 4.

uniform sky is likely to print with more noise and less stability. These two screens represent the principal frequency range covered by the adaptive algorithm, which will produce the positive effects of these screens without incurring the negative [compare Fig. 20(a) to Figs. 21(a) and 21(b)].

5 Visible and Invisible Watermarking via Seed Frequency Modulation

Halftone-based watermarking allows data embedding in images as part of the printing process. The embedded data can be used for a variety of applications, such as ownership protection, authentication, covert communication, and any other usage that can benefit from additional information. Most existing halftone-based watermarking methods are designed for traditional stochastic screens (FM),^{18,19} periodic clustered-dot screens (AM),^{20–24} or error diffusion.²⁵ Seamless visible and invisible data-embedding can be achieved by judicious manipulation of seed frequencies within the context of the algorithm introduced in Sec. 2. Specifically,

two different FM to AM transition thresholds can be assigned to areas of the image according to the presence or absence of a watermark mask; the difference between the thresholds determines whether the watermark is visible or not.

5.1 Watermark Encoding

Encoding of a watermark is achieved by halftoning image content with two screens with different principal frequencies, the screens being selected by the watermarks. Figure 22 illustrates the steps involved in the creation of a watermarked halftone image, from a continuous-tone input image and a binary watermark mask. Recall that in the method introduced in Sec. 2, seeds were placed according to a binary patch corresponding to the stochastic screen halftoning of a uniform gray patch with area coverage T . In the watermark-embedding application, two different patches determine seed placement, the patches having been obtained from stochastic screen halftoning of two gray patches with area coverage

Fig. 20 Illustration of the enhanced edge and detail rendition capabilities of the algorithm from Sec. 4.

Fig. 21 Performance of the traditional stochastically seeded clustered stochastic screen from Sec. 2 at (a) low-end and (b) high-end of the principal frequency range.

Fig. 22 Schematic illustration of the steps involved in the creation of a digital watermark.

T_W and T_{NW} . At every pixel of the input image, seed location, and consequently, seed placement frequency (principal frequency), is determined based on the gray value of the pixel as well as on the presence or absence of watermark at that location. For locations where watermark (no watermark) is present, seeds are placed according to the binary patch obtained from the halftoning of a gray patch of value $T_W(T_{NW})$.

Figures 23 and 24 provide an illustrative example of an encoded halftone image. Figure 23(a) shows the continuous-tone input image, and Fig. 23(b) shows a sample binary watermark mask. Figures 24(a) and 24(b) show pixel-size versions of the halftoned version of the image region highlighted in Fig. 23(a) for two different transition threshold gaps. In the case of Fig. 24(a), $T_{NW} = 6\%$ and $T_W = 15\%$. In contrast, Fig. 24(b) has transition thresholds $T_{NW} = 12\%$ and $T_W = 15\%$. It can be seen that, as the absolute difference between the transition thresholds $|T_{NW} - T_W|$ decreases, the perceptibility of the watermark decreases as well. The visibility of the watermark is thus easily controllable by adjusting the relative seed frequency between the watermark and nonwatermark areas.

5.2 Watermark Decoding

Decoding of a watermark is achieved by identifying the principal frequency of the halftone screen used to halftone local regions in the halftoned image. In the context of this section, let us define the term dot as a set of 8-connected ON pixels and the term hole as a set of 8-connected OFF pixels. The decoding approach we propose is based on local counts of

Fig. 23 (a) Continuous tone image. (b) Binary watermarking mask.

Fig. 24 Binary images resulting from the halftoning of the image in Fig. 23(a) with the watermark in Fig. 23(b) using the introduced watermarking technique with (a) $T_{NW} = 6\%$, $T_W = 15\%$ and (b) $T_{NW} = 12\%$, $T_W = 15\%$.

dots and holes. Consider the curves from Fig. 25(a), which illustrate the number of dots and holes for a traditional spot function designed according to the procedure from Sec. 2. For a spot function with transition threshold T , and as the gray level increases from 0, the number of dots increases steadily until gray level T is reached, at which point it stabilizes. At a gray level where the dots start to touch, the number of dots decreases until all pixels in the screen are turned ON, at which point a single dot remains. In contrast, the number of holes remains steady at unity from a gray level of zero until a gray level where dots start to touch, at which point it increases steadily up to a level where dots grow so much that they start absorbing holes. From that point on, the number of holes decreases steadily until it reaches zero. Note that the shapes of the dot and hole count curves for a specific screen vary according to the choice of T and the spot function utilized.

Fig. 25 Dot and hole count curves for (a) a single spot function with transition threshold T ; and for two spot functions with thresholds T_1 and T_2 (b) when no erosion or dilation is applied and (c) after application of erosion and dilation operations.

Consider the case where two spot functions with two different transition threshold values T_1 and T_2 are used; the dot and hole count curves in this scenario are plotted in Fig. 25(b). Inspection of Fig. 25(b) reveals that the number of dots and holes at a portion of the midtones and shadows, respectively, can be used to discriminate between the screens, while the light shadows (levels below the smallest transition threshold, in this case T_1) and very dark shadows are not useable for decoding purposes. The useable gray-scale range can be increased by breaking the touch points between dots in the light shadows via morphological erosion and by filling in holes in the shadows via morphological dilation. The decision of whether to use erosion and dilation or not is made based on estimated area coverage information or on a comparison between local dot and hole count and predetermined thresholds. Figure 25(c) shows the new dot and hole count curves resulting from the application of the erosion and dilation operations.

Figure 26 shows the image resulting from the decoding of the invisible watermark shown in the image from Fig. 24(b). White pixels denote data identified as watermark. The retrieval of the watermark is not perfect because some portions of it fall in the unusable gray-scale range.

6 Dot-Off-Dot Vector Halftoning

In this section, we introduce a dot-off-dot vector halftoning method that is an extension of the monochromatic spot function introduced in Sec. 2. While the method proposed in Ref. 26 introduces the use of multiple uncorrelated green noise screens to achieve color halftoning, it does not perform vector halftoning *per se* and, consequently, is unable to completely avoid dot overlap whenever possible (e.g., when two colorants are used and average area coverage per colorant is $<50\%$ or when three colorants are used and the average area coverage per colorant is $<33\%$). The algorithm presented herein implements a successive fill strategy on the single stochastically seeded clustered halftone screens from Sec. 2 for up to three colorants. As will become clear, the strategy is equivalent to a first colorant dot being grown from the triangle vertices, a second colorant dot being grown from the triangle center, and a third colorant dot being grown from the midpoints of the triangle sides. Growth in this manner ensures separation of colorants up to a certain density level—whenever allowed by the intended area coverage—which is an inherent characteristic of dot-off-dot screens.

Fig. 26 Binary image resulting from the decoding of the embedded watermark from Fig. 24(b).

t_1 $(h_1, h_2, h_3)_1$	t_2 $(h_1, h_2, h_3)_2$	...	t_{N-1} $(h_1, h_2, h_3)_{N-1}$	t_N $(h_1, h_2, h_3)_N$
------------------------------	------------------------------	-----	--------------------------------------	------------------------------

Fig. 27 Vector of thresholds (in ascending order) associated with orthogonal distances between the location of the point and the sides of spatially varying triangles.

When the use of a fourth colorant is required at a given pixel, it can be accommodated via methods typically used for clustered screens, such as stochastic screening the lightest colorant.

6.1 Intuitive Description of the Successive Fill Strategy

The fill order for a scalar halftone spot function determined according to the procedure outlined in Sec. 2 establishes a threshold t associated with each location within a triangle, which can be uniquely described by the distances from the specific location onto each of the triangle sides, namely h_1 , h_2 , and h_3 . In traditional successive fill algorithms, the thresholds in the vector are determined by a pair of rectilinear coordinates within a rectangular cell based on simple modulo indexing into the cell; in contrast, the present vector elements are associated with the aforementioned orthogonal distances to sides of spatially varying triangles that are determined from two-dimensional point coordinates in space, as illustrated in Fig. 27.

For a given input pixel, the successive fill algorithm is implemented as follows:

1. For an input continuous-tone pixel having nonzero value for only one colorant, use thresholds defined from the left of the vector.
2. For an input pixel having nonzero value for two colorants, use thresholds defined from the left of the vector for a first colorant and from the right for the second colorant.
3. For an input pixel having nonzero value for at least three colorants, use thresholds defined from the left of the vector for a first colorant, from the right for

the second colorant, and the middle for the third colorant.

A filling strategy implemented in this manner results in a first colorant dot being grown from the seed locations, or triangle vertices, a second colorant dot being grown from the triangle center location, and a third coloration dot being grown from the midpoints of each of the triangle sides, as illustrated in Fig. 28.

6.2 Algorithmic Implementation of the Successive Fill Strategy

The algorithmic implementation of the algorithm outlined above is described next. For a given input pixel having colorant values $\mathbf{f} = (f_1, f_2, f_3)$, we determine halftoned pixel values $\mathbf{b} = (b_1, b_2, b_3)$ by applying the following procedure:

1. For a first nonzero colorant, apply threshold t_i corresponding to the distances from the pixel to the triangle sides, $(h_1, h_2, h_3)_i$. If more than one colorant is nonzero, the threshold is applied to the darkest or second darkest colorant. Denote the index of this first colorant as k , its input value as f_k , and its halftoned value as b_k . Then

$$b_k = 1 \quad \text{if } f_k \geq t_i, \quad \text{else } b_k = 0. \quad (4)$$

2. If a second colorant has a nonzero value, apply threshold $t_N - t_i$, where t_i is the threshold corresponding to the distances from the pixel to the triangle sides, $(h_1, h_2, h_3)_i$. The second darkest colorant is thresholded using this rule. Denote the index of this second colorant as l , its input value as f_l , and its halftoned value as b_l . Then

$$b_l = 1 \quad \text{if } f_l \geq t_N - t_i, \quad \text{else } b_l = 0. \quad (5)$$

3. If a third colorant has a nonzero value, apply threshold $|t_{N/2} - t_i|$ corresponding to its distances to the triangle sides, $(h_1, h_2, h_3)_i$. The lightest colorant is thresholded using this rule. Denote the index of this third colorant as m , its input value as f_m , and its halftoned value as b_m . Then

$$b_m = 1 \quad \text{if } f_m \geq |t_{N/2} - t_i|, \quad \text{else } b_m = 0. \quad (6)$$

Application of rule 3 to the lightest colorant is optional because the particular spot functions that we provided above will have the least compact form for this rule and, therefore, could print less reliably in this scenario. Figure 29 shows a schematic of this successive fill algorithm, which is significantly different from successive fill methods used for traditional stochastic screens.²⁷

Figures 30 and 31 show two- and three-colorant sweeps, respectively, halftoned with the algorithm described above.

Fig. 28 Origin of growth regions for image areas with color content in at least three different channels. The particular colorants in the figure are only used as an example.

Fig. 29 Schematic showing the threshold order used in the presented successive fill method.

Fig. 30 Two-colorant sweeps halftoned with the proposed algorithm. The dot-off-dot structures are maintained up to 50% area coverage per colorant.

Fig. 31 Three-colorant sweeps halftoned with the proposed algorithm. The dot-off-dot structures are maintained up to 33% area coverage per colorant.

The dot-off-dot structures are maintained up to 50 and 33% area coverage, respectively. This dot-off-dot structure produces smoother textures and a larger gamut while using less ink compared to scalar halftoning techniques applied independently to each color channel. At the same time, and given the clustered nature of the screens, it offers improved stability. One potential problem associated with dot-off-dot screens is their sensitivity to color shift due to misregistration between the color planes in the print engine. Some degree of level blending can be built into the spot functions to alleviate this sensitivity. The specific implications of misregistration between color planes on the quality of the presented dot-off-dot halftoning technique are beyond the scope of the present paper. For an in-depth study on the effects of misregistration relative to induced color shifts in clustered color halftones, we refer the reader to Ref. 28.

7 Discussion

We believe we have filled a gap in the literature on clustered stochastic screens by introducing parametric control of dot shape as well as seed placement adaption to local image structure. The halftoning method that we introduced defines dot centers as seeds that are placed typically in a random high spatial frequency configuration. Spot functions are defined about these randomly placed seeds, where the spot function allows control of dot cluster growth, touch points, cluster angle, and eccentricity. The spot function can also be applied to regular and irregular polygonal halftone tiling. The seed adaption aspect of the halftoning method allows for better edge rendition than conventional isotropic methods and enables seamless data-embedding by judicious manipulation of seed frequency. The seed placement strategy associated with the spot function can be seamlessly extended to accommodate data-embedding applications, both in visible and invisible modes. Last, the introduced scalar spot functions were paired with novel successive fill strategies to achieve a dot-off-dot vector halftoning strategy, which showcases the flexibility of the proposed methodology.

References

1. R. Levien, "Output dependant feedback in error diffusion halftoning," in *Proc. IS&T's 8th Int. Congress on Advances in Non-Impact Printing Technologies*, pp. 280–282, The Society for Imaging Science and Technology, Williamsburg, Virginia (1992).
2. D. L. Lau, G. R. Arce, and N. C. Gallagher, "Green-noise digital halftoning," *Proc. IEEE* **86**(12), 2424–2444 (1998).
3. S. Wang, "Stochastic (stochastic clustered) halftone screen design," in *Int. Conf. on Digital Printing Technologies*, pp. 516–521, The Society for Imaging Science and Technology, Seattle, Washington (1997).
4. Y. Ng et al., "Advances in NexPress digital printing technology," in *Proc. 23rd Int. Conf. on Digital Printing Technologies*, pp. 489–

- 493, The Society for Imaging Science and Technology, Anchorage, Alaska (2007).
5. H. Asai, "Image recording method and image recording apparatus," U.S. Patent No. 7,821,672 (2010).
6. H. Asai, "Non-periodic clustered color halftone screen," *J. Inst. Image Electron. Eng. Jpn.* **35**(1), 566–575 (2006).
7. H. Asai, "Halftone dots, halftone dot forming method and apparatus therefor," U.S. Patent No. 6,989,913 (2006).
8. H. Lee et al., "Method of generating halftone screen and system thereof," U.S. Patent No. 7,619,777 (2009).
9. J. Huang and A. Bhattacharjya, "Method and apparatus for generating dispersed cluster screen," U.S. Patent No. 7,239,429 (2007).
10. V. Ostromoukhov, "Pseudo-random halftone screening for color and black and white printing," in *The 9th Int. Congress in Non-Impact Printing Technologies*, pp. 579–582, The Society for Imaging Science and Technology, Yokohama, Japan (1993).
11. P. Goyal et al., "Clustered-dot halftoning with direct binary search," *IEEE Trans. Image Process.* **22**(2), 473–487 (2013).
12. P. Fink, *PostScript Screening: Adobe Accurate Screens*, Adobe Press, Mountain View, California (1992).
13. S. Wang and R. Loce, "Moiré-Free color halftoning using hexagonal geometry and spot functions," *J. Electron. Imaging* **21**(1), 013017 (2012).
14. R. Pellar, "Electronic halftone generator," U.S. Patent No. 4,196,451 (1980).
15. R. Pellar, "Electronic halftone generator," U.S. Patent No. 4,149,183 (1979).
16. H. Z. Hel-Or, X. M. Zhang, and B. A. Wandell, "Adaptive cluster dot dithering," *J. Electron. Imaging* **8**(2), 133–144 (1999).
17. H. Kang, *Digital Color Halftoning*, SPIE Press, Bellingham, Washington (1999).
18. S. Wang and K. Knox, "Embedding digital watermarks in halftone screens," *Proc. SPIE* **3971**, 218–227 (2000).
19. K. Knox and S. Wang, "Digital watermarks using stochastic screens," *Proc. SPIE* **3018**, 316–322 (1997).
20. S. Wang, "Digital watermarking using phase-shifted stoclustic screens," U.S. Patent No. 6,252,971 (2001).
21. G. Sharma and S. Wang, "Show-through watermarking of duplex printed documents," *Proc. SPIE* **5306**, 670–684 (2004).
22. B. Oztan and G. Sharma, "Continuous phase modulated halftones and their application to halftone data embedding," in *Proc. Int. Conf. on Acoustics, Speech and Signal Processing*, Vol. 2, pp. 333–336, The Institute of Electrical and Electronics Engineers Signal Processing Society, Toulouse, France (2006).
23. S. Wang, "System and method for the generation of correlation-based digital watermarks," U.S. Patent No. 8,077,907 (2011).
24. O. Bulan et al., "Data embedding in hardcopy images via halftone-dot orientation modulation," *Proc. SPIE* **6819**, 68190C (2008).
25. D. Shaked and Z. Baharav, "Watermark rate and the distortion of error diffused halftones," HP Labs Technical Report, HPL-98-89, The Hewlett-Packard Company, Technion City, Haifa, Israel (2001).
26. 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.* **9**(5), 923–935 (2000).
27. G. Sharma and S. Wang, "Method and system for designing spatially-partitioned and correlated stochastic screens for color halftoning," U.S. Patent No. 6,501,567 (2002).
28. B. Oztan, G. Sharma, and R. Loce, "Misregistration sensitivity in clustered-dot color halftones," *J. Electron. Imaging* **17**(2), 023004 (2008).

Edgar A. Bernal is a senior research scientist at the Xerox Research Center in Webster, New York. He joined Xerox in 2006 with MSc and PhD degrees in electrical engineering from Purdue University in West Lafayette, Indiana. His early career focused on the areas of image processing, halftoning, image perception, watermarking, and color theory. His current research interests include computer vision, video compression, video-based object tracking, and the application of novel sensing technologies to healthcare and transportation.

Robert P. Loce is a research fellow and technical manager in the Xerox Research Center, Webster. He joined Xerox in 1981 with an associate degree in optical engineering technology from Monroe Community College. While working in optical and imaging technology and research departments at Xerox, he received a BS in photographic science (RIT 1985), an MS in optical engineering (UR 1987), and a PhD in imaging science (RIT 1993), and passed the US patent bar in 2002.

Shen-Ge Wang received his BS in instrumental mechanics from Changchun Institute of Optics in China and his PhD in optics from the University of Rochester. Formerly, he was a principal scientist working at the Xerox Research Center in Webster, New York, where his research was focused on digital halftoning, printer modeling, and digital watermark technologies. He has many publications and over 100 patents in those fields.

Zhigang Fan received his MS and PhD degrees in electrical engineering from the University of Rhode Island, Kingston, in 1986 and 1988, respectively. He joined Xerox Corporation in 1988, where he is currently a principal scientist in Xerox Corporate Research and Technology. His research interests include various aspects of image processing and recognition. He has authored and coauthored more than 90 technical papers, as well as over 200 patents and pending applications. He is a fellow of SPIE and a fellow of IS&T.