Stochastic Moiré

Daniel L. Lau
Lau Imaging Technologies, Cambridge, Massachusetts
Arif M. Khan and Gonzalo R. Arce
Department of Electrical and Computer Engineering, University of Delaware
Newark, Delaware

Abstract

In digital halftoning, moiré refers to the periodic interference pattern created by superimposing the cyan, magenta, yellow, and black halftone patterns of periodic amplitude modulated screens. This paper introduces the concept of stochastic moiré—the aperiodic interference pattern created by superimposing the aperiodic halftone patterns of frequency modulated screens. Given two overlapping stochastic halftone patterns, this paper details how to extract a continuous-space surface that completely characterizes the spatial fluctuations in color/texture of stochastic moiré that exist in the halftone. To quantitatively measure the visibility of stochastic moiré, this paper introduces a statistic for characterizing the fluctuations in color/texture that result from stochastic moiré and shows how this metric can be used to identify the better of two arrangements of dots at minimizing the visual distortion stochastic moiré creates.

Introduction

Amplitude modulated (AM) halftoning refers to algorithms that create the illusion of continuous-tone in bi-level imaging devices by producing a regular pattern of round dots that vary in size according to tone such that dark shades of gray are represented by large printed dots and light shades of gray are represented by small dots. In color printers, continuous shades of color are produced by superimposing the halftone patterns of cyan, magenta, yellow, and black inks. While aligning each AM grid such that the round dots of each color are superimposed directly on top of one another (a process referred to as dot-on-dot) produces a full spectrum of colors,[1] even slight mis-registration (mis-alignment of the grids) can drastically degrade the visual quality of the printed image. Instead of overlapping dots, each AM screen is typically given its own orientation or screen angle. The problem now is not the distortion created by mis-registration, but rather that caused by the introduction of a visual interference pattern created by superimposing two or more regular patterns, a phenomenon known as *moiré*.

Periodic moiré created by the superposition of AM patterns has been studied in great detail by Amidror *et al*[2] who model AM screens as cosinusoidal gratings and their superposition in the spatial domain as convolution in the spectral domain, Fig. 1. Modeling binary dither patterns as continuous-tone cosinusoidal gratings leads to a Fourier domain representation that is composed exclusively of three purely real impulses. Because the superposition of black with any other color leads to black, Amidror *et al* adopt a multiplicative model where the pattern created by superposition is determined by a pixel-wise multiplication operation between the two patterns being superimposed. In the Fourier domain, this multiplicative relationship leads to a convolution of the impulses of each grating. Approximating the human visual system (HVS) in the Fourier domain as a circular rect-function, the optimal screen offset angle is the angle that keeps all but the DC impulse outside the cut-off frequency of the HVS, Fig. 2.

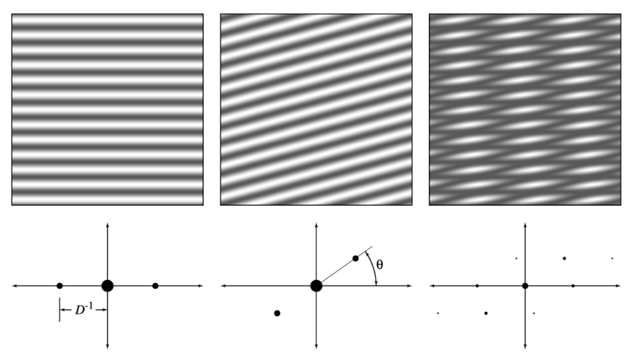

Figure 1. The periodic moiré approximated by the superposition of two cosinusoidal gratings such that the resulting interference pattern is modeled as the convolution of impulses in the Fourier domain.

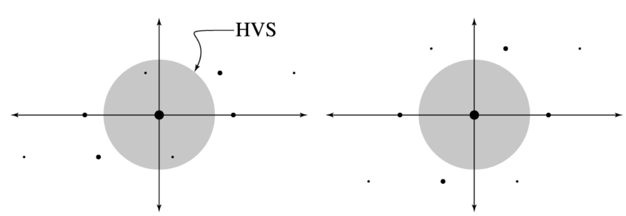

Figure 2. The periodic moiré approximated by the superposition of two cosinusoidal gratings and a circular step function modeling the HVS where (left) has visible moiré artifacts while (right) has a screen offset angle that virtually eliminates moiré from the halftone.

In frequency modulated (FM) halftoning, the illusion of continuous-tone is produced by a random arrangement of same sized dots where the average distance between two dots varies with tone such that dark shades of gray are represented by a tight packing of dots and light shades of gray are represented by a loose packing. The optimal FM patterns are the ones that spread the printed dots as homogeneously as possible—creating a dither pattern composed exclusively of high frequency blue-noise.[3] In theory, superimposing two or more stochastic dither patterns does not lead to the appearance of moiré, but superimposing two FM screens does lead to a halftone with low frequency artifacts—creating a noisy (uncorrelated) appearance.

Amidror *et al*'s spectral analysis is not valid for analyzing stochastic halftones due to its reliance on the simple characterization of cosinusoidal gratings in the spectral domain. The spectral representation of stochastic halftones are very complicated, containing a full spectrum of real and imaginary components. Only through the computational convenience of digital computers can the spectrums of stochastic halftones be managed.

In this paper, we study the phenomenon of low-frequency graininess created by superimposing stochastic dither patterns, referring to the phenomenon as *stochastic moiré*. Like moiré in AM halftoning, stochastic moiré cannot be eliminated but only minimized. In this paper, we show that stochastic moiré is defined by a relationship between in- and out-of-phase interactions between two or more dither patterns, and the optimal halftoning algorithms are the ones that space the in- and out-of-phase points either as close together or as far apart as possible.

Spatial Analysis of Moiré

The basic premise of Amidror *et al*'s work on periodic moiré is that the superposition of the two AM screens is a multiplicative relationship in the luminance image. In the Fourier domain, this multiplicative relationship leads to a convolution of the Fourier spectrums. Amidror *et al* go further to simplify the problem by replacing binary screens with cosinusoidal gratings, a surface whose Fourier spectrum is characterized by a handful of impulses orderly arranged on the frequency plane. This model has effectively simplified periodic moiré to a framework that can be implemented by hand, not computation. The optimal alignment of screens is now a simple problem of moving impulses as far away as possible from the origin of the spectral plane.

Periodic Moiré

Noting Ref. 1, moiré is a periodic pattern of over-lapping and non-overlapping dots. Making Amidror *et al*'s substitution of cosinusoidal gratings in place of binary screens, our basic premise for characterizing moiré (both periodic and stochastic) is that moiré is a pattern of in-phase and out-phase alignments between two gratings. Noting Fig. 3, the locations where the two gratings are in-phase are the points where the two peaks (solid lines) intersect or where two valleys (dashed lines) intersect. Out-of-phase points are the locations where a peak of one pattern intersects a valley of the other.

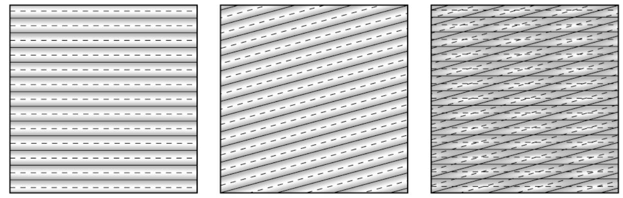

Figure 3. The periodic moiré created by superimposing two cosinusoidal gratings where the peaks and valleys of each grating is indicated as a solid or dashed line respectively.

Shown in Fig. 4 are the in-phase (Xs) and out-of-phase (Os) points between two cosinusoidal grating. The parameter r represents what we refer to as the *principle wavelength* of moiré and it is the average distance from an in-phase point to its nearest out-of-phase point. For this pair of periodic gratings, the distance r has a deterministic solution:

$$r = \cos(\theta), \qquad (1)$$

where θ represents the relative screen angle between the two screens. From our understanding of the human visual system, minimizing the visibility of moiré becomes a problem of either maximizing r ($\theta = 0°$) or minimizing r ($\theta = 90°$) such that the moiré pattern fluctuates too slowly per unit length to be noticed by the HVS or too rapidly. Using two-dimensional gratings, Fig. 5 shows the locations of the in and out-of-phase points. Here minimizing r means setting $\theta = 45°$.

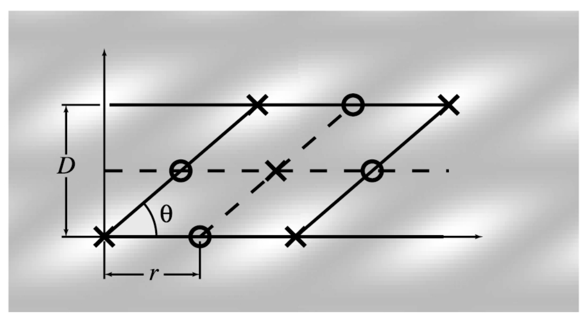

Figure 4. Diagram showing the location of in-phase (Xs) and out-of-phase (Os) components where the minimum distance between Xs and Os is r.

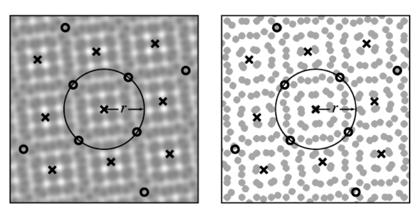

Figure 5. The in-phase and out-of-phase points for periodic moiré created by superimposing (left) two 2D cosinusoidal gratings and (right) two binary AM halftone patterns.

Now considering the periodic AM patterns for a given color composed of a given proportion of **A** and **B**, if all the dots of **A** are in-phase with the dots of **B**, the halftone creates a certain color and texture for the viewer. If dots of **A** are all out-of-phase with **B**, the halftone creates a different color and/or texture for the viewer. If we choose θ such that *r* is minimized (θ = 45°), the screens of **A** and **B** create the moiré that minimizes the visibility of artificial textures created by superimposing **A** and **B**. Given the low-pass nature of the human visual system (HVS),[4] minimizing *r* creates a pattern that maximizes the frequency of spatial fluctuations in color/texture. Maximizing these fluctuations moves the spectral content of the halftone pattern to where the HVS is least sensitive, creating an apparent image composed exclusively of a DC color/texture.

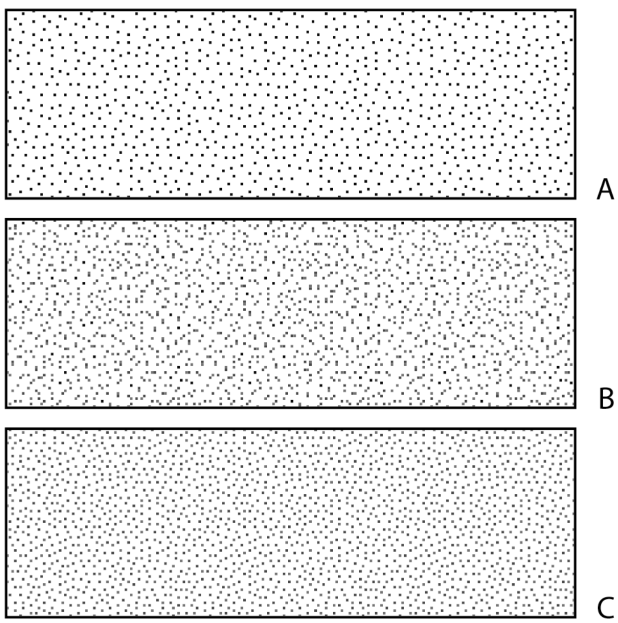

Figure 6. Three two-color halftone patterns illustrating (A) perfectly overlapping or in-phase pixels, (B) uncorrelated pixels, and (C) perfectly non-overlapping or out-of-phase pixels.

Aperiodic Moiré

Now while this framework of in-phase and out-of-phase points offers no advantage over Amidror *et al*'s with regard to periodic moiré, it can be extended to stochastic halftoning. Before doing so, we need a concise definition of what stochastic moiré is and what it is not. We start with Fig. 7 where several stochastic halftone patterns are shown each composed of 6.25% (1/16) cyan coverage and 6.25% magenta coverage. Pattern **A** shows the instance where minority pixels of cyan and magenta are homogeneously distributed such that magenta pixels are placed midway between cyan pixels. Shown in pattern **C** is the case where magenta pixels are placed directly on top of all cyan pixels. While it maybe argued by a given observer that pattern **A** creates a hue different from that of pattern **C**, we ignore those differences focusing on the fact that pattern **A** certainly portrays a different texture from that of **C**. Any differences in hue are assumed to be corrected by manipulating the percentage of coverage of the various inks.

Now while **A** and **C** have different textures, we make no claims as to whether **A** or **C** is a better texture. The real study of stochastic moiré occurs in pattern **B**. Here the magenta and cyan dither patterns are completely uncorrelated with some minority pixels (of both patterns) overlapping and some falling directly in between those of the other pattern. All other minority pixels fall some where in between these two extremes. From pattern **B**, we define stochastic moiré as the change in texture that occurs from one point to another within a given halftone. Like periodic moiré, the optimal stochastic halftone is the one that either minimizes or maximizes the amount of fluctuation in this texture per unit area. In the case of pattern **A** and **C**, the amount of fluctuation is minimized with patterns **A** and **C** offering equivalent optimality regarding moiré.

Restricting our analysis to the spatial domain, one way to characterize fluctuations in texture is to identify the in-phase and out-of-phase points. In-phase points are the locations where minority pixels of both patterns overlap, while out-of-phase points are the locations where a minority pixel of one pattern falls directly between minority pixels of the other. A minority pixel is directly between minority pixels of the other pattern when it coincides with a vertex point of the voronoi mesh defined by the minority pixels of the other pattern, Fig. 6. We choose these vertices as the location of valleys because they are mid-way between minority pixels, with any displacement moving the valley closer to one minority pixel and farther from the others. The principle wavelength of stochastic moiré will again be the average distance between an in-phase point and its nearest out-of-phase point.

The problem with such a metric, as just described, is that it relies solely on the extreme phase points for which there may be very few. The minority pixels that fall between extremes tell us something about how well aligned the two dither patterns are, and we need to make use of this information. If we return to our definition of moiré as being a pattern of overlapping and non-overlapping dots, in-phase minority pixels are instances of overlapping pixels while out-of-phase pixels are clearly non-overlapping. We can even say that these out-of-phase pixels are as non-overlapping as minority pixels can possibly be for a given halftone. Minority pixels that are neither of these extremes are non-overlapping, but we will say that they are not as non-overlapping as out-of-phase pixels. How non-overlapping a minority pixel is can be measured as the distance from that minority pixel to its nearest neighboring minority pixel in the other pattern.

In formal terms, the dither patterns $\phi_A = \{a_i : i = 1,2,\ldots\}$ and $\phi_B = \{b_j : j = 1,2,\ldots\}$ define a discrete-space, 2-D function $D[n]$ such that:

$$D[n] = \sum_{i=1}^{N} d_i \delta[n - a_i], \text{ where} \qquad (2)$$

$$d_i = \min_j \left\| a_i - b_j \right\|_2 . \qquad (3)$$

In its eqn. (2) form, $D[n]$ represents a continuous-space signal $\mathcal{D}(x)$, an image we will refer to as the *stochastic moiré surface*, that has been sampled on a stochastic sampling grid defined by ϕ_A. If the sampling grid has a Poisson distribution, as would be the case of a blue-noise dither pattern with principle wavelength λ_a, we can obtain $\mathcal{D}(x)$ from $D[n]$ using an ideal low-pass filter with cutoff frequency $0.5\ \lambda_a^{-1}$.

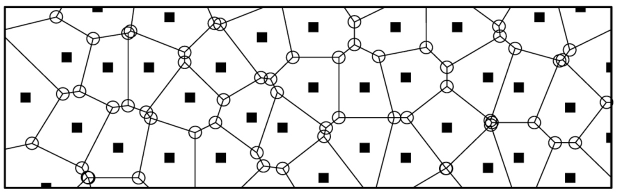

Figure 7. Diagram showing the location of out-of-phase (circles) pixels corresponding to a stochastic dither pattern with minority pixels marked as black squares.

Given our stochastic moiré surface, characterizing stochastic moiré is now reduced to the problem of characterizing the spatial fluctuations in the continuous-space, 2-D surface $\mathcal{D}(x)$. How we measure spatial fluctuations in this new monochrome image is arbitrary, but it is important to take into account our understanding of the human visual system to put a higher cost on mid-frequency fluctuations as opposed to very low or very high. Perhaps the easiest way to measure the visibility of stochastic moiré is to measure the visual cost[5] of the stochastic moiré surface as:

$$VC(\phi_A, \phi_B) = \text{Average Power}(\mathcal{D}^f(u) \times HVS(u)) \qquad (4)$$

where $\mathcal{D}^f(u)$ if the Fourier transform of $\mathcal{D}(x)$ and $HVS(u)$ is the spectral, low-pass filter model of the human visual system.

For an illustration of eqn. (4) as our new stochastic moiré metric, Fig. 8 (top) shows a halftone pattern representing the cyan-magenta color (0.11,0.11) along with its stochastic moiré surface. This particular surface has a visual cost of 0.0854, assuming a 300 dpi printer and a viewing distance of 10 inches. The pattern shown in Fig. 8 (bottom) has, on the other hand, a visual cost of 0.0432. Comparison of the two visual costs suggests an improvement in color rendition from the (top) pattern of Fig. 8 to the bottom. Visual inspection certainly concurs.

Before introducing case studies of the stochastic moiré surface, its important to understand why ϕ_A must be a blue-noise process, and the reason lies in stochastic sampling theory whereby a Poisson point process is composed exclusively of a DC frequency component combined with spatial frequencies above $1/\lambda_a$, where λ_a is the minimum distance between minority pixels in ϕ_A. Perfect reconstruction of the continuous-space signal is guaranteed if the signal being sampled only contains spatial frequencies at or below $0.5/\lambda_a$, Fig. X.

Based on our definition of $D[n]$, the stochastic moiré surface is such that the maximum rate at which the surface can change is to go from a maximum value or peak at one sample to a minimum value at a nearest neighboring sample. Half the wavelength of this maximum rate component is, therefore, equal to λ_a just as the distance from a peak to valley in a cosine is equal to half the cosine's wavelength. The frequency of this maximum rate component is, therefore, equal to $0.5\ \lambda_a^{-1}$. We conclude that since the maximum spatial frequency contained within the stochastic moiré surface is $0.5\ \lambda_a^{-1}$ with the sampling rate being sufficiently high as to avoid aliasing. If, for instance, the reference pattern defining the sampling grid were from a white-noise pattern, then the stochastic moiré surface would be impossible to extract from the spectral components of the sampling grid.

Now suppose that both dither patterns are blue-noise patterns but with different minority pixel intensities. In such an instance, the sampling grid is defined by the pattern with a higher intensity. This follows from our earlier comment that about how fast the surface can change from a peak to a valley. Observing Fig. X, while it is not possible for the surface to move from a peak to a valley between any two neighboring minority pixels in ϕ_B, it is possible to go from a peak to a valley in less distance than between any two neighboring samples in ϕ_A. We note that this is an important observation since the lower intensity dither pattern will have a lower cutoff frequency. The casual observer may think the lower cutoff frequency would make for a more conservative estimate on the moiré surface.

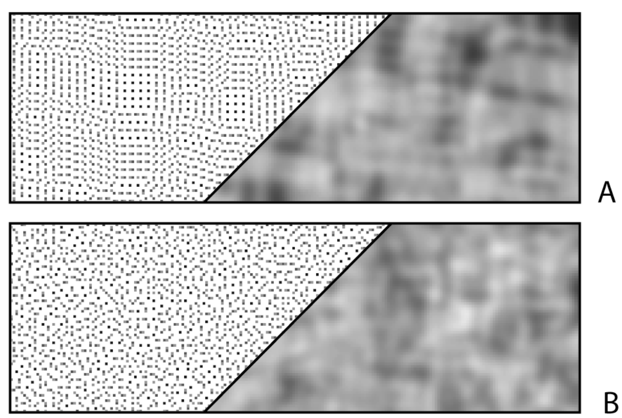

Figure 8. Two dither patterns and their corresponding stochastic moiré surfaces.

Case Studies

As a demonstration of eqn. (4)'s feasibility at quantitatively measuring stochastic moiré, this section compares the resulting visibility measures for 2-color halftoning using Floyd's and Steinberg's,[6] Jarvis *et al*'s,[7] and Ulichney's[3] error-diffusion. It is well documented that Floyd's and Steinberg's error-diffusion creates strong periodic textures at gray levels 1/4, 1/3, and 1/2 while Jarvis *et al*'s does so at gray level 1/3. In our 2-color halftones, these periodic textures lead to large patches of fixed phase between pattern **A** and **B**—leading to large near-DC components in the stochastic moiré surface and, therefore, large visibility measures.

Shown in Fig. 9 (top) and (middle) are the 1-D plots of visual cost for varying minority pixel intensities in **A** equal to **B**$\in(0,1/2]$ with large spikes corresponding to strong periodic textures. Through a perturbation of error filter coefficients, Ulichney's error-diffusion breaks up periodic textures in the dither pattern—creating a far more pleasant blue-noise halftone. In our 2-color prints, this perturbation eliminates strong spikes (bottom) and creates the optimal 2-color halftoning scheme among the three being compared here.

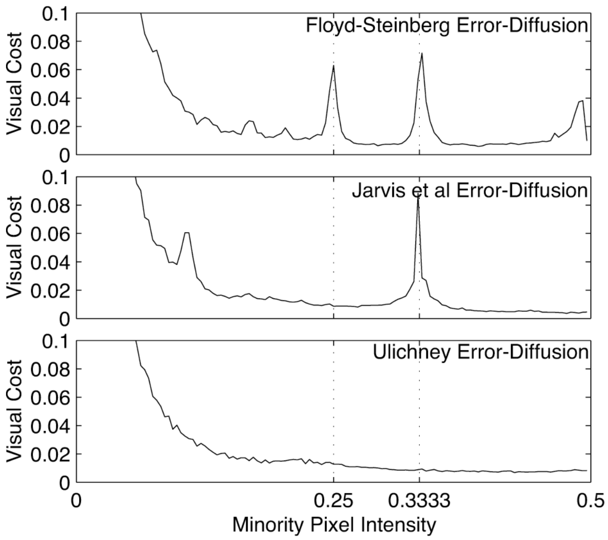

Figure 9. Plots of the visual cost versus minority pixel intensity level for (top) Floyd's and Steinberg's, (middle) Jarvis et al*'s, and (bottom) Ulichney's error-diffusion.*

Conclusions

Stochastic moiré is the visible interference pattern created by superimposing two aperiodic dither patterns. Using stochastic sampling theory, we showed how to extract a continuous-space stochastic moiré surface from two overlapping dither patterns. This surface represents a powerful and especially elegant tool for optimizing color halftoning algorithms. A very simple demonstration of the stochastic moiré surface was presented when we showed that, in general, Ulichney's perturbed error diffusion algorithm produced better two color prints than Floyd's and Steinberg's or Jarvis *et al*'s error diffusion.

With regards to optimality, there are two approaches to minimizing the visible artifacts associated with stochastic moiré. The first is to minimize spatial fluctuations per unit area such that the corresponding texture changes too slowly to be noticed by the human visual system. This approach is equivalent to dot-on-dot or dot-off-dot approaches to AM halftoning, and it, therefore, carries the same constraints on screen registration. The second approach to minimizing the visibility of artifacts is to maximize spatial fluctuations per unit area such that the texture changes too rapidly to be noticed by the viewer.

In a closing note, it's important to understand that this paper's only goal is to introduce stochastic moiré and a framework for characterizing it. This paper does not represent a technique for optimizing the halftones of two color prints, while such techniques may be obvious. We leave it to future works and to other authors to describe any such details. Ultimately, this paper alleviates the need for hunt-and-peck optimization of halftoning algorithms.

Acknowledgment

The contributions of Arif Khan and Gonzalo Arce were sponsored by Lexmark International, Inc.

References

1. H. R. Kang, Digital Color Halftoning, SPIE Optical Engineering Press, Bellingham, Washington, USA and IEEE Press, New York, USA, 1999.
2. I. Amidror, R. D. Hersch, and V. Ostromoukhov, "Spectral analysis and minimization of moiré patterns in color separation," Journal of Electronic Imaging, vol. 3, pp. 295-317, July 1994.
3. R. A. Ulichney, "Dithering with blue noise," Proceedings of the IEEE, vol. 76, pp. 56-79, Jan. 1988.
4. F. W. Campbell, R. H. Carpenter, and J. Levinson, "Visibility of aperiodic patterns compared with that of sinusoidal gratings," Journal of Physiology, vol. 190, 283-298, 1969.
5. J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," IEEE Transactions on Systems, Man, and Cybernetics, vol. 21, 33-38, Jan. 1991.
6. R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial gray-scale," Proceedings Society Information Display, vol. 17, no. 2, pp. 75-78, 1976.
7. J. F. Jarvis, C. N. Judice, and W. H. Ninke, "A survey of techniques for the display of continuous-tone pictures on bilevel displays," Computer Graphics and Image Processing, vol. 5, pp. 13-40, 1976.

Biography

Dr. Daniel L. Lau received his B.S. in Electrical Engineering from Purdue University, West Lafayette with highest distinction in 1995 and then his Ph. D. degree in Electrical Engineering at the University of Delaware in 1999. Daniel spent a year signal and image processing at the Lawrence Livermore National Laboratory in California and is currently a DSP engineer with Aware Inc, Bedford MA. He is also a private consultant on digital printing through his own firm, Lau Imaging Technologies. His published works in halftoning include an article in the December 1998 issue of the *Proceedings of the IEEE* and his own book *Modern Digital Halftoning* published by Marcel Dekker, January 2001.