Digital Halftoning using Green-Noise Masks*

Daniel L. Lau and Gonzalo R. Arce
Department of Electrical and Computer Engineering at the University of Delaware
Newark, DE 19720, USA
lau@ece.udel.edu, arce@ece.udel.edu

Abstract

In this paper, we introduce a novel technique for generating green-noise halftones–stochastic dither patterns composed of homogeneously distributed pixel clusters. Although techniques employing error diffusion have been proposed previously, the technique here employs a dither array referred to as a green-noise mask which greatly reduces the computational complexity formerly associated with green-noise. Compared to blue-noise masks, halftones generated with green-noise masks are less susceptible to printer distortions. Because green-noise constitutes patterns with widely varying cluster sizes and shapes, the technique introduced here for constructing these green-noise masks is tunable, that is it allows for specific printer traits, with small clusters reserved for printers with low distortion, and large clusters for printers with high. Being that blue-noise is a limiting case of green-noise, this new technique can even create blue-noise masks.

Introduction

Digital halftoning is the technique employed to convert images from continuous-tone to binary. Desktop printers such as laser and ink-jet, as well as commercial planographic and screen presses capable of only binary output, rely on digital halftoning in order to produce the illusion of continuous-tone. Typically, halftoning is done in one of two ways as either amplitude (AM) or frequency (FM) modulated halftoning. Both present their advantages and disadvantages.

AM halftoning is a point operation that forms a regular grid of clustered dots (Fig. 1 (top)). The resulting patterns are, therefore, very robust and computationally inexpensive but also of low quality. Visually, AM halftones are blurred versions of the continuous tone original.

FM halftoning, by isolating minority pixels in a random fashion, produces binary images with much higher spatial resolution (Fig. 1 (middle)), creating an image with much less blurring than their AM counterparts. Once a neighborhood operation, FM halftoning can be performed with minimal computational expense through the use of blue-noise dither arrays.[4] The problem with FM, though, is that by isolating minority pixels, FM halftones maximize the perimeter-to-area ratio of printed dots, thereby maximizing the susceptibility to inevitable process variation.

With recent advancements in halftoning technology, binary images can now be represented using the AM-FM hybrid, green-noise.[1] Green-noise halftones are stochastic dither patterns composed of homogeneously distributed minority pixel *clusters* (Fig. 1 (bottom)). By clustering pixels of like color, green-noise can sacrifice spatial resolution for pattern robustness. Green-noise can, therefore, create images with less blurring than AM and with more resistance to printer distortions than FM.

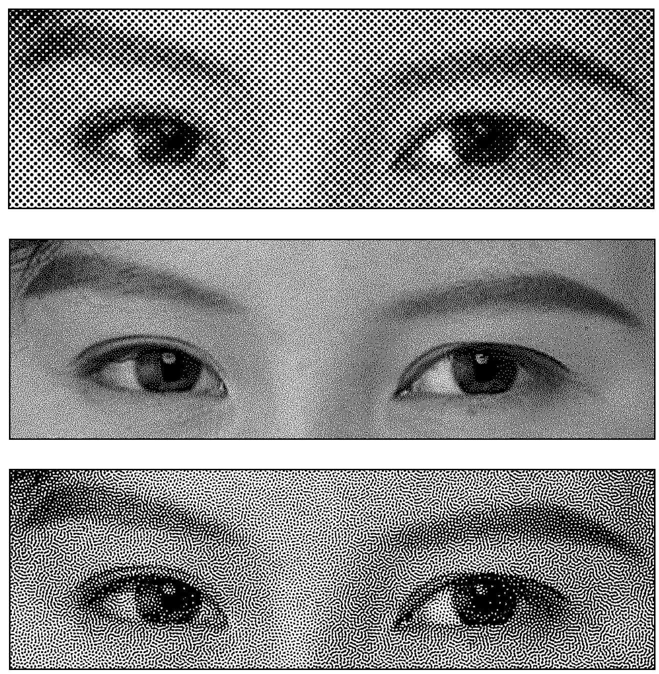

Figure 1. Eyes image using (top) AM halftoning using 8x8 cells, (middle) FM halftoning using blue-noise, and (bottom) AM-FM halftoning using a green-noise mask.

Until now, halftoning with green-noise has implied error-diffusion based halftoning.[3] Such schemes are neighborhood operations and, therefore, carry a high computational cost.[5] But in this paper, we introduce a novel technique for constructing green-noise masks. These masks make it possible to generate green-noise halftones with the same low complexity as AM and FM.[2] Being a tunable process, the construction algorithm can produce masks for a wide range of printing devices such that reliable devices have small clusters and unreliable devices have large.

Spatial Statistics of Green-Noise

Before constructing masks, this section summarizes the statistical framework, introduced by Lau *et al*,[1] for studying stochastic halftone patterns and, using spatial statistics, characterizes the green-noise halftoning model. In the Lau *et al* framework, an aperiodic halftoning process is defined as a stochastic model, Φ, governing the location of minority pixels, or points x_i, within a binary dither array. A sample of Φ, ϕ is written as $\phi=\{x_i : i=1,\dots,N\}$ with $\phi(B)$ a scalar quantity defined as the number of minority pixels, x_is, in a subset B of the dither array. Using this notation, $\phi(x)=1$ indicates that the pixel at sample x is a minority pixel (pixel x is 1 for $0\le g<0.5$ and 0 for $0.5\le g<1$) while $\phi(x)=0$ indicates that the pixel at sample x is not.

As a scalar quantity, the first order moment or the expected value of $\phi(x)$ is the *intensity*, $I(x)$, which is the unconditional probability that sample x is a minority pixel. For a binary pattern representing gray level g, $I(x)=g$ for $0\le g<0.5$ and $1-g$ for $0.5\le g<1$. A second statistic for characterizing Φ is the quantity $K(x;y)$ defined as:

$$\boldsymbol{K}(x;y)=\frac{\mathbf{E}\{\phi(x)\,|\,y\in\phi\}}{\mathbf{E}\{\phi(x)\}}, \tag{1}$$

the ratio of the conditional expectation that sample x is a minority pixel given that sample y is a minority pixel to the unconditional expectation that sample x is a minority pixel. Referred to as the *reduced second moment measure*, $K(x;y)$ is a measure of the influence that a minority pixel at sample y has on pixel x. If $K(x;y)>1$ then sample x is more likely to be a minority pixel given y while if $K(x;y)<1$ then sample x is less likely to be a minority pixel given y.

For a stationary point process Φ, $K(x;y)=K(r,\theta)$ where r is the distance between samples x and y and θ is the direction from y to x. For stationary point processes which are also isotropic[†], $K(r,\theta)$ is independent of θ and is commonly referred to as the *pair correlation*, $R(r)$, which is defined explicitly as:

$$\boldsymbol{R}(x)=\frac{\mathbf{E}\{\phi(\Omega_y(r))\,|\,y\in\phi\}}{\mathbf{E}\{\phi(\Omega_y(r))\}}, \tag{2}$$

the ratio of the expected number of minority pixels located in the ring $\Omega_y(r)=\{x : r<|x-y|\le r+dr\}$, under the condition that a minority pixel exists at sample y to the unconditional expected number of minority pixels located in the ring $\Omega_y(r)$.

For studying inter-point relationships between the pixels of a binary dither pattern, the pair correlation is a very powerful metric. This can be seen by looking at the well studied blue-noise model where given that the goal of blue-noise halftoning is to space minority pixels as homogeneously as possible,[5] the pair correlation of the ideal blue-noise dither pattern is of the form of Fig. 2 (top). This pair correlation shows three distinct features. The first feature, labeled (a), is a pair correlation $R(r)=0$ near $r=0$ indicating a strong inhibition of minority pixels close to one another. The second, labeled (b), is a pair correlation approaching 1 for increasing r indicating a decreasing correlation between minority pixels as r gets larger and larger. The final feature, labeled (c), is a series of peaks at integer multiples of λ_b, the principle wavelength of blue-noise. These peaks indicate a frequent occurrence of the inter-point distance λ_b.

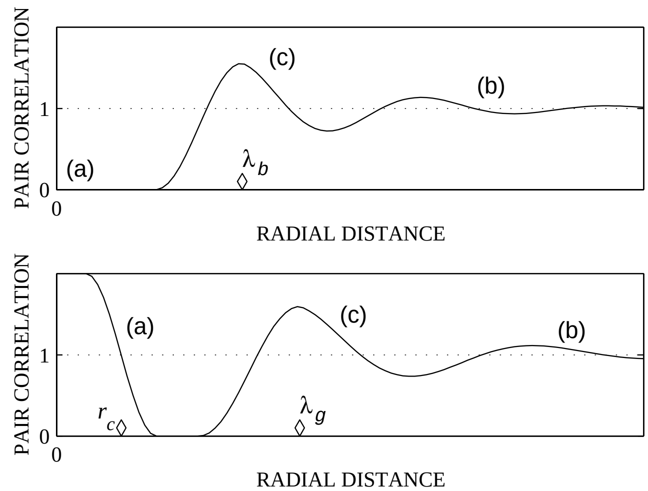

Figure 2. The ideal pair correlation for (top) blue-noise and (bottom) green-noise.

In contrast to blue-noise, green-noise has a pair correlation of the form of Fig. 2 (bottom). This pair correlation also displays three distinct features. The first, labeled (a), is a pair correlation $R(r)>1$ for $r\le r_c$ indicating a clustering of minority pixels where r_c is the *cluster radius*. This cluster radius is related to the average number of minority pixels per cluster, $\underline{M}$, as:

$$\pi r_c^2=\underline{M}. \tag{3}$$

Eqn. (3) states that the area enclosed by a circle of radius r_c is equivalent to the area covered by a cluster of size $\underline{M}$ pixels. The second feature of the ideal green-noise pair correlation, labeled (b), is a pair correlation which approaches 1 with increasing r indicating, like the blue-noise pair correlation, a decreasing correlation of minority pixels with larger and larger separation between them. The final feature, labeled (c), is a series of peaks at λ_g, the principle wavelength of green-noise, indicating a frequent occurrence of the inter-cluster distance λ_g.

BIPPCCA

In this paper, we introduce the iterative procedure BIPPCCA (BInary-Pattern-Pair-Correlation-Construction-Algorithm) to construct binary dither patterns with arbitrary pair correlations and a given intensity by randomly converting pixels of an arbitrarily sized array from a majority (0) to a minority (1) value. Progressively building upon the previous iteration, BIPPCCA begins with an $M\times N$ all zero array, ϕ, with one pixel selected at random and converted to a minority pixel. Given the dither pattern ϕ

with minority pixels $\{x_i : i=1,2,\ldots\}$, BIPCCA assigns a probability of becoming a minority pixel to each majority pixel in ϕ. BIPPCCA then replaces the maximum likely majority pixel (the majority pixel with the highest corresponding probability) with a minority pixel. The process is then repeated until the dither pattern ϕ of size $M\times N$ has $I\times M\times N$ minority pixels where $I=g$ for $0\leq g<0.5$ or $1-g$ for $0.5\leq g<1$.

BIPPCCA is able to construct ϕ such that the resulting dither pattern has a desired reduced second moment measure by adjusting the probabilities of majority pixels being converted to minority pixels at each iteration according to the current set of minority pixels in ϕ and $\underline{R}(r)$, the pair correlation *shaping function*. $\underline{R}(r)$ is a user specified function derived from the desired pair correlation with values of $\underline{R}(r)>1$ increasing the likelihood of minority pixels being placed a distance r apart and values of $\underline{R}(r)<1$ decreasing the likelihood. Recall from the previous section that $R(r)>1$ indicates that given a minority pixel at location y all samples x, for which $|x-y|=r$, are more likely to be a minority pixel than any point z for which $R(|z-y|)<R(r)$. So given that a minority pixel is placed at y, BIPPCCA increases the likelihood of a pixel becoming 1 for all pixels x for which $R(r)>1$. It also decreases this likelihood for all z for which $R(|z-y|)<1$.

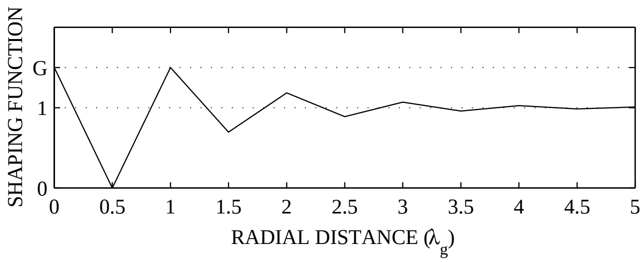

Figure 3. The pair correlation shaping function used to construct green-noise dither patterns.

The function $\underline{R}(r)$ is referred to as the spatial *shaping function* due to its influence in "shaping" the pair correlation of the resulting pattern. Used in this paper for $\underline{R}(r)$, Fig. 3 shows a very simple approximation of the ideal pair correlation for isotropic green-noise with peaks at integer multiples of λ_g, the principle wavelength of green-noise, and valleys mid-way between peaks. We note that although more elaborate $\underline{R}(r)$ could be proposed, this model was selected because of its very simple structure. In simulations, 1.01 proved a good value as the maximum amplitude (labeled G in Fig. 3) for $\underline{R}(r)$.

In addition to $\underline{R}(r)$, the probabilities of majority pixels being converted is also influenced, at each iteration, by the current concentration of minority pixels in ϕ through C, the *concentration array*. C ensures homogeneity in ϕ by decreasing the probabilities of becoming a minority pixel for majority pixels in areas of dense minority pixel concentration and increasing the probabilities for majority pixels in areas of sparse minority pixel concentration.

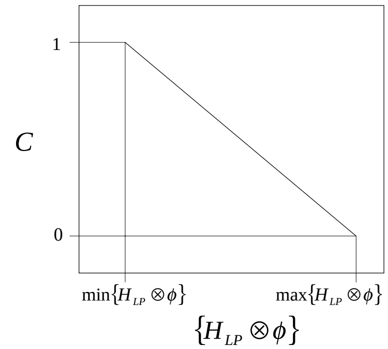

Figure 4. Mapping function used to construct the concentration matrix C.

In BIPPCCA, the concentration of minority pixels is measured as the output after applying the low-pass filter, H_{LP}, to ϕ. In selecting a low-pass filter, an obvious choice for H_{LP}, as suggested Ulichney,[6] is the Gaussian filter such that:

$$H_{LP}(r)=exp(-r^2/2\sigma^2), \tag{4}$$

for some constant σ. Note that in order for a minority pixel to have an influence on neighboring clusters, the filter H_{LP} should have a higher σ for small I where clusters are farther apart than for large I where clusters are closer together. In this paper, such a relationship between H_{LP} and I is ensured by setting $2\sigma^2=\lambda_g$. The concentration matrix is then constructed from $\text{conv}_{\text{circ}}\{H_{LP},\phi\}$, the output after filtering ϕ with the low-pass filter H_{LP} using *circular convolution*, according to the mapping of Fig.4.

The steps of BIPPCCA for generating an $M\times N$ binary dither pattern representing intensity level I are described as follows:

1. Initialize all pixels of an $M\times N$ array, ϕ, to zero.
2. Randomly select one pixel of ϕ, and convert that pixel, $\phi[m,n]$, to one.
3. Create an $M\times N$ array, U, of uniformly distributed random numbers such that $U[i,j]\in(0,1]$ is the probability that $\phi[i,j]$ will become a minority pixel.
4. Given the most recently converted pixel, $\phi[m,n]$, scale the value $U[i,j]$ for all pixels $\phi[i,j]=0$ by $\underline{R}(r)$ such that $(U[i,j])_{new}=(U[i,j])_{old}\cdot\underline{R}(r)$ where r is the minimum wrap-around distance between the two pixels $\phi[m,n]$ and $\phi[i,j]$.
5. Construct the concentration matrix C using the mapping of Fig.4 from $\text{conv}_{\text{circ}}\{H_{LP},\phi\}$, the output after filtering ϕ with the low-pass filter H_{LP} using circular convolution.
6. Locate the majority pixel in ϕ with the highest probability (the pixel $\phi[m,n]=0$ such that $(U[m,n]\cdot C[m,n])>(U[i,j]\cdot C[i,j])$ for all $1\leq i\leq M$ and $1\leq j\leq N$ and that $\phi[i,j]=0$), and set that pixel, $\phi[m,n]$, to one.
7. If the number of minority pixels in ϕ is equal to $I\times M\times N$, then the algorithm quits with the output pattern given by ϕ; otherwise, continue at step 4.

Constructing Green-Noise Masks

The basic premise to constructing green-noise masks is to generate a set, $\{\phi_g : 0 \le g \le 1\}$, of dither patterns, constructed in any order using BIPPCCA, from a set, $\{\underline{R}(r;g) : 0 \le g \le 1\}$, of shaping functions with one pattern and one shaping function for each possible discrete gray-level g (256 levels for 8-bit gray-scale images). The dither array, *DA*, is then constructed by assigning to each pixel a threshold according to the spatial arrangement of binary pixels within $\{\phi_g : 0 \le g \le 1\}$. Note that because BIPPCCA generates patterns such that minority pixels are represented by pixels equal to one, the dither pattern ϕ_g, for $0.5 < g \le 1$, are generated by inverting the pixels of a dither pattern created by BIPPCCA for gray-level $1-g$.

In order to avoid ambiguities in the assignment of thresholds to *DA*, the dither patterns, ϕ_g, are constructed under the *stacking constraint* that $\phi_k \subset \phi_g$ for all $k<g$ or that if given $\phi_k[m,n]=1$, then $\phi_g[m,n]=1$ for all $g<k$. As a consequence, the threshold assigned to each pixel $DA[m,n]$ is equal to the minimum g for which $\phi_g[m,n]=1$. In BIPPCCA, a dither pattern ϕ_k can be constructed given ϕ_g such that $\phi_k \subset \phi_g$ by constraining step 6 of BIPPCCA to only consider those pixels of ϕ_k for which $\phi_g[m,n]=1$; furthermore, a dither pattern ϕ_g can be constructed given ϕ_k such that $\phi_k \subset \phi_g$, if, in step 1, ϕ_g is initialized to ϕ_k and each value, $U[m,n]$, is scaled according to the location of minority pixels in ϕ_g as defined by $\underline{R}(r;g)$. In addition to the above modifications to BIPPCCA, the shaping function, $\underline{R}(r;g)$, must also observe the stacking constraint as minority pixel clusters in pattern ϕ_g cannot decrease in size from those in ϕ_k.

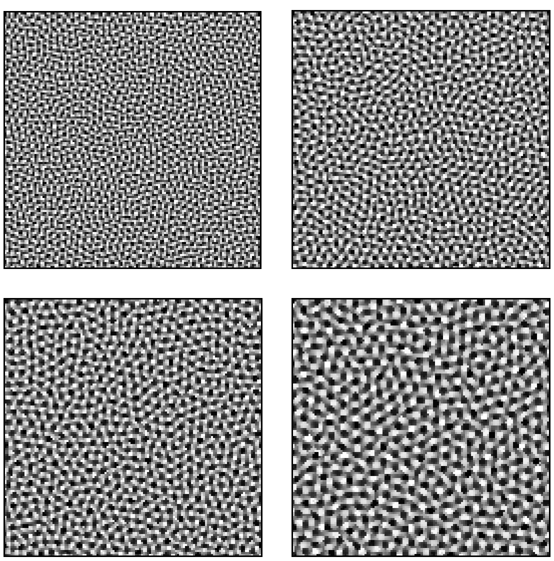

Figure 5. Green-noise masks constructed to produce halftoning patterns with increasing coarseness.

Conclusions

Unlike blue-noise halftoning, which distributes the minority pixels of a binary dither pattern as homogeneously as possible,[5] green-noise halftoning forms clusters of minority pixels that are distributed as homogeneously as possible.[1] Typically, green-noise patterns are generated via error-diffusion based techniques.[2] In this paper, we have described the algorithm BIPPCCA that constructs binary patterns according to a desired pair correlation. Although intended here for green-noise, BIPPCCA can be used to construct blue-noise patterns as well.

Under a stacking constraint, can use BIPPCCA to build green-noise masks, dither arrays designed to produce green-noise halftones by thresholding, pixel-by-pixel, a continuous tone original. Far less computationally complex than error-diffusion based algorithms, these green-noise masks can also be tuned to specific printer characteristics by adjusting pattern coarseness. Fig. 5 shows four instances of the green-noise mask, each with increasing coarseness over the previous.

References

• Lexmark International provided partial funding for this work.

† A point process Φ is isotropic if the statistical properties of Φ are invariant to rotation.

1. D. L. Lau, G. R. Arce, and N. C. Gallagher, “Green-Noise Digital Halftoning,” Proceedings of the IEEE, vol. 86, no. 12, December 1998, pp. 2424-2444.
2. D. L. Lau, G. R. Arce, and Neal C. Gallagher, “Digital Halftoning by means of Green-Noise Masks,” *Journal of the Optical Society of America A*, vol. 16, no. 7, July 1999, pp. 1575-1586.
3. R. Levien, “Output dependant feedback in error diffusion halftoning,” is *IS&T’s Eight International Congress on Advances in Non-Impact Printing Technology*, (Williamsburg, Virginia, USA), October 25-30, 1992, pp. 280-282.
4. K. Spaulding, R. Miller and J. Schildkraut, “Methods for generating blue-noise dither matrices for digital halftoning,” *Journal of Electronic Imaging*, vol. 6, no. 2, April 1997.
5. R. A. Ulichney, “Dithering with blue noise,” *Proceedings of the IEEE*, vol. 76, no. 1, January 1998, pp. 2424-2444.
6. R. A. Ulichney, “The void-and-cluster method for dither array generation,” *Human Vision, Visual Processing, Digital Displays IV, Proceedings SPIE*, vol. 1913, 1993.

Biography

Daniel L. Lau received his B.S.E.E. degree from Purdue University, West Lafayette, IN, with highest distinction in 1995. He then completed his Ph.D. in electrical engineering at the University of Delaware in May 1999. Daniel Lau has worked in image processing at the Lawrence Livermore National Laboratory, and his research interests include image processing, digital halftoning, multimedia, nonlinear filters and art conservation. His paper, “Green-Noise Digital Halftoning,” appeared in the December 1998 issue of the *Proceeding of the IEEE.*