A Novel Halftoning Technique; the Blue Noise Mask

Theophano Mitsa and Kevin J. Parker

Electrical Engineering Department

Abstract

A novel halftoning technique is presented where the halftoning is achieved by a pixelwise comparison of the gray scale image against a nonimage array, the blue noise mask. The blue noise mask is constructed such that it has unique properties in the image and transform domain. When the blue noise mask is thresholded at any level, the resulting binary pattern has the correct first order statistics, and also its power spectrum has blue noise (high frequency) characteristics which are visually pleasing. The construction of the blue noise mask is described and experimental results are shown. Also, results from a psychovisual study are provided where subjects rated halftoned images that have the same first order but different second order statistics.

1. Introduction

Many printing devices and displays, such as laser printers, FAX machines, desktop publishing systems, and liquid crystal displays, are not capable of reproducing gray scale images because they are bilevel. Gray scale images are converted to binary images using halftone techniques. Halftoning renders the illusion of various shades of gray by using only two levels, black and white, and it can be implemented either digitally or optically. Optical halftoning is used in such applications as lithography (newspaper printing).

In the case of digital halftoning, the halftoning processes are classified into point and neighborhood processes [Ulichney 1987] according to the number of pixels required to calculate one output pixel in the halftoned image. In point processes, the halftoning is accomplished by a simple pixelwise comparison of the gray scale image against a nonimage periodic array (mask). The most common point process is ordered dither. Halftoning using neighborhood processes is not done by a simple pixelwise comparison, but usually requires filtering operations that involve a number of input pixels in order to calculate one output pixel. The most common neighborhood process is error diffusion. In error diffusion every pixel of the gray scale image is compared against a fixed threshold and then the resulting binary value is subtracted from the original value of that pixel. The resulting difference (error), after being multiplied by appropriate weights, is diffused into a given size neighborhood.

Halftoning often degrades the perceived image quality and leads to fine-image-detail loss. The major image-quality trade-off in all halftoning processes is between achieving gray scale rendition and also maintaining edge and generally high spatial frequency information. Another important criterion for the evaluation of a halftoning process is its speed, which is directly related to its computational complexity. Point processes such as ordered dither are fast with good high frequency rendition, but the output image usually suffers from artifacts such as periodic structures and false contours. Neighborhood processes (error diffusion) give better quality images but are slow compared to point processes. The question we posed, and that naturally follows from the above, is: Can we develop a halftoning process that combines the speed of point processes with the quality of neighborhood processes? In this paper we propose a new halftoning process where the halftoning is achieved by a pixelwise comparison against a nonimage array, the "blue noise mask." We gave the name blue noise mask to our halftoning array because, when thresholded at any level, it results in an aperiodic and isotropic pattern with small low frequency components, which in the halftoning literature is known as a "blue noise pattern." Blue noise patterns have visually pleasing properties and, up to now, they could be produced only by using neighborhood processes such as error diffusion.

Although halftoning using a blue noise mask can also be implemented optically, in the following we will describe the digital case because of its wider applications. In Section 2 we describe the properties of the patterns that result after thresholding the blue noise mask at fixed levels (dot profiles) and present the blue noise principles. In Section 3 we describe the construction of the blue noise mask from the dot profiles. In Section 4 we describe a psychovisual study in which subjects rated images with the same first but different second order statistics. Experimental results are also presented.

2. Theory

2.1 Dot Profiles and the Blue Noise Mask

We denote as $f(i,j)$ the gray scale image, $m(i,j)$ the blue noise mask, and $h(i,j)$ the halftoned image. For an $M \times N$ B-bit gray scale image, the halftoned image can be of the same dimensions $M \times N$ (no interpolation is necessary) and a $J \times K$ B-bit blue noise mask can be employed, where $J \leq M$ and $K \leq N$. The blue noise mask is built with "wraparound" properties, such that the dimensions of the blue noise mask can be smaller than those of the gray scale image. The pixelwise comparison can proceed modulo J and modulo K in the respective directions with no apparent discontinuities or obvious periodicities. For example, for a 256 × 256 8-bit class of images a 128 × 128 8-bit blue noise mask can be used. It is important to note here that the blue noise mask is completely independent of the image to be halftoned, and once constructed the same B-bit blue noise mask can be used for any B-bit gray scale image.

Thus for a B-bit image $f(i,j)$, the blue noise mask array $m(i,j)$ is a B-bit array such that, when thresholded against $f(i,j)$, up to 2^B levels of varying distribution of black and white dots can be represented on a rectangular grid. The binary pattern that results after thresholding the blue noise mask at a constant level g is called the dot profile for that level [Allebach and Liu 1977]. The dot profiles are arrays that have the same dimensions as the mask array and consist of ones and zeros. The ratio of ones to zeros is different for every dot profile and depends on the gray level that dot profile represents. In the following, the gray level g is normalized with respect to 2^B and thus $0 \le g < 1$. In our notation, 0 represents black and 1 represents white. According to this notation, the higher the gray level, the more ones and less zeros will be contained in the dot profile. We denote as $p[i,j,g]$ the value of the dot profile [Allebach and Liu 1977] at pixel location (i,j) and for the gray level g. For any gray scale image the corresponding binary image $h(i,j)$ can be constructed as follows in terms of the dot profiles: For every pixel in the gray scale image array $f(i,j)$ that is at the (i,j) location and has a value $f_{ij} = g$, the corresponding pixel h_{ij} in the binary image array $h(i,j)$ has a value that is given by the value of the g-level dot profile at the (i,j) location. Thus:

$$h_{ij} = p[i,j,g] = \begin{cases} 1 & g > m_{ij} \\ 0 & g \le m_{ij} \end{cases} \tag{1}$$

The dot profiles for every level are designed and combined in such a way as to build a single valued function, the blue noise mask. More specifically, these dot profiles are not independent of each other, but the dot profile for level $g_1+\Delta g$ is constructed from the dot profile for level g_1 by replacing some selected zeros with ones. For example, for an $N \times N$ B-bit mask with maximum pixel value given by 2^B-1, and with $\Delta g = 1/2^B$, the number of zeros that will change to ones, in order to go from level g_1 to level $g_1+\Delta g$, is $N^2/2^B$. The dependence of the dot profiles can be expressed as follows [Allebach and Liu 1977]:

$$\text{if } g_2 > g_1 \cap p[i,j,g_1] = 1 \Rightarrow p[i,j,g_2] = 1 \tag{2}$$

As we create the dot profiles sequentially, another array called the cumulative array is incremented in such a way as to keep track of the changes in the dot profiles from gray level to gray level. We denote as $c[i,j,g]$ the value of the cumulative array at location (i,j) and for gray level g. When thresholded at any level $g_1 \le g$ the cumulative array $c[i,j,g]$ gives us the dot profile for level g_1. After we have created the dot profiles for all levels the cumulative array becomes the blue noise mask. It is important to note here that both the dot profile and the cumulative array are 3-D functions and depend on location as well as gray level. On the other hand, the blue noise mask is a 2-D function and depends only on location.

2.2 Blue Noise

The dot profiles for every level are designed to be locally aperiodic and isotropic binary patterns with small low frequency components, which in the halftoning literature are known as blue noise patterns [Ulichney 1987]. The absence of low frequency components in the frequency domain corresponds to the absence of disturbing artifacts in the spatial domain. The radially averaged power spectrum of a blue noise pattern is shown in Figure 1 (adapted from [Ulichney 1987]). By radially averaged power spectrum we mean the sample average of the power spectrum $P(f)$ in the annulus $\left|\,|f| - f_r\right| < \Delta/2$:

$$P_r(f_r) = \frac{1}{N_r(f_r)} \sum_{i=1}^{N_r(f_r)} P(f) \tag{3}$$

where f_r is the radius of the annuli and $N_r(f_r)$ is the number of samples that fall within the annuli. As shown in Figure 1, the cutoff frequency f_g is known as the Principal Frequency and its inverse $\Delta_g = 1/f_g$ is known as the Principal Wavelength [Ulichney 1987]. The Principal Wavelength is proportional to the distance between the minority pixels at every level g, where by minority pixels we mean the black dots for levels $g > 1/2$ and the white dots for levels $g < 1/2$. The Principal Frequency assumes its highest value for the level $g = 1/2$, since at that level the populations of the black and white dots are equal and thus very high frequency components appear in the binary image. The Principal Frequency depends on the gray level g, as follows:

$$f_g = \begin{cases} K\sqrt{g}/R & \text{for } g \le \frac{1}{2} \\ K\sqrt{1-g}/R & \text{for } g > \frac{1}{2} \end{cases} \tag{4}$$

where K is a scaling factor and R is the distance between addressable points on the display, and the gray level g is normalized between 0 and 1. Error diffusion produces binary patterns whose Principal Frequencies fit the above formula for $K = 1$. The dot profiles for the blue noise mask

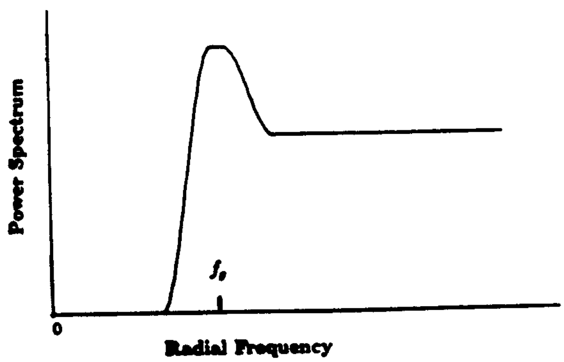

Figure 1 Blue noise power spectrum.

were designed with $K = 1/\sqrt{2}$. The choice for this value of K is based on practical filter design considerations, as will be explained in the next section, and also on the psychovisual results given in Section 4.

3. Construction of the Blue Noise Mask

For an $M \times N$ B-bit image with 2^B-1 as the maximum pixel value, the blue noise mask is constructed by building the dot profiles "on top of" each other. We start from the dot profile for level $g = 1/2$, which is a suitable blue noise pattern, and then we create the dot profiles for levels $g > 1/2$. The cumulative array is first created at level $g = 1/2$ and at that level it contains a value of 2^{B-1} at every pixel that corresponds to a zero in the dot profile $p[i,j,1/2]$ and a value of $2^{B-1}-1$ otherwise:

$$\forall p[i,j,1/2] = 0 \Rightarrow c[i,j,1/2] = 2^{B-1} \tag{5}$$

$$\forall p[i,j,1/2] = 1 \Rightarrow c[i,j,1/2] = 2^{B-1} - 1 \tag{6}$$

In that way, when the cumulative array is thresholded at $g = 1/2$ (2^{B-1} unnormalized), the resulting binary pattern is the dot profile $p[i,j,1/2]$. In general, the dot profile for level $g_1+\Delta g$ is built from the dot profile for level g_1 by changing $X = (M \times N)\,\Delta g$ zeros to ones, provided that X is an integer. The choice of zeros that will be changed to ones is done as follows:

Step 1. Take the 2-D Fourier transform of $p[i,j,g_1]$ and obtain $P[u,v,g_1]$ where u,v are the transform coordinates.

Step 2. Apply a blue noise filter $D[u,v,g_1]$ to $P[u,v,g_1]$ and in this way obtain:

$$P'[u,v,g_1] = P[u,v,g_1] \times D[u,v,g_1] \tag{7}$$

The blue noise filter is designed such that the radial average of $P'[u,v,g_1]$ is of the form shown in Figure 1. The Principal Frequency is given by (4) with $K = 1/\sqrt{2}$ and $g = g_1+\Delta g$. This choice for K is based on practical filter design considerations because $D[u,v,g_1]$ is applied in a radial average sense and this value of K is the highest value for which all the annuli are subscribed inside the Fourier transform domain (no truncation effects).

Step 3. Take the inverse Fourier transform of $P'[u,v,g_1]$ and obtain $p'[i,j,g_1]$, which is no longer binary but has much better blue noise properties.

Step 4. Form the difference $e[i,j,g_1] = p'[i,j,g_1]-p[i,j,g_1]$. We will refer to this difference as the error array.

Step 5. Rank order all the zeros (pixels for which $p[i,j,g_1] = 0$) according to the value of $e[i,j,g_1]$ for each pixel.

Step 6. Set a limit, $l_\epsilon = \epsilon$, for the magnitude of the highest acceptable error. This limit is usually set equal to the average magnitude error. Change to ones $(M \times N)\,\Delta g$ pixels that contain a zero and have an error higher than the defined limit:

$$p[i,j,g_1] = 0 \cap e[i,j,g_1] > l_\epsilon \Rightarrow p[i,j,g_1+\Delta g] = 1 \tag{8}$$

Before changing the value of a pixel from zero to one, additional spatial criteria, such as neighborhood mean and runlengths, are also checked in order to ensure that no visually annoying structures will be created after the value of this pixel has changed.

Step 7. Finally, after all $(M \times N)\,\Delta g$ zeros have been replaced with ones, the cumulative array is updated. The cumulative array is updated by adding one only to those pixels that still correspond to a zero in the dot profile $p[i,j,g_1+\Delta g]$:

$$c[i,j,g_1+\Delta g] = c[i,j,g_1] + \overline{p[i,j,g_1+\Delta g]}\,;\ g_1 \geq 1/2 \tag{9}$$

where the bar indicates a logical "not" operation changing zeros to ones and vice versa. In this fashion, when the blue noise mask is thresholded at constant level $g_1+\Delta g$, the resulting binary pattern is the dot profile $p[i,j,g_1+\Delta g]$. The above process is also shown in Figure 2. We repeat this procedure until we create the dot profiles for all the gray levels from $1/2+\Delta g$ up to 1. The levels from $1/2-\Delta g$ to 0 are created in the same way with the only difference being that now ones are changed to zeros and the cumulative array is updated as follows:

$$c[i,j,g_1-\Delta g] = c[i,j,g_1] - p[i,j,g_1-\Delta g]\,;\ g_1 \leq 1/2 \tag{10}$$

When the process has been implemented for all gray levels, the cumulative array contains the selected dot profile for each level, and is therefore the desired blue noise mask.

4. Psychovisual Testing

In this section we describe a psychovisual test that was conducted in order to establish a connection between the spectra of a binary pattern and its visual appearance. For that purpose, binary patterns with the same first order but different second order statistics were rated in a psychovisual test. These patterns were created with a procedure similar to the one described in Section 3. The difference in the second order statistics was in regard to the location, and/or height, and/or width of the Principal Frequency peak, and the criterion according to which the images were rated was homogeneity. The rating scale was from 1 (worst) to 5 (best) and the images were rated by ten expert subjects. The binary patterns were 512×512 and the viewing distance was approximately six times the image size.

Binary patterns for which the Principal Frequency had the same location but different height and width appeared similar, while binary patterns for which the Principal Frequency had the same height and width but different location had a substantially different visual appearance. According to the psychovisual test results the best patterns were the ones with $K = 1/\sqrt{2}$.

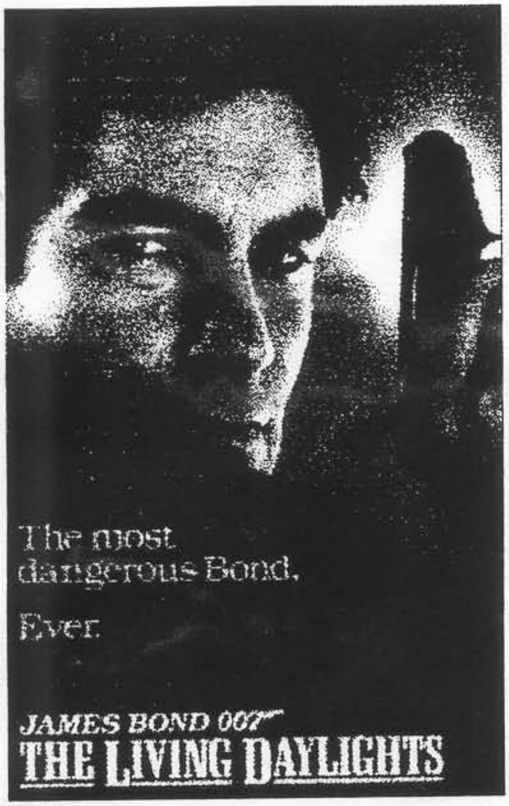

Figure 2 Gray scale image halftoned using the blue noise mask method.

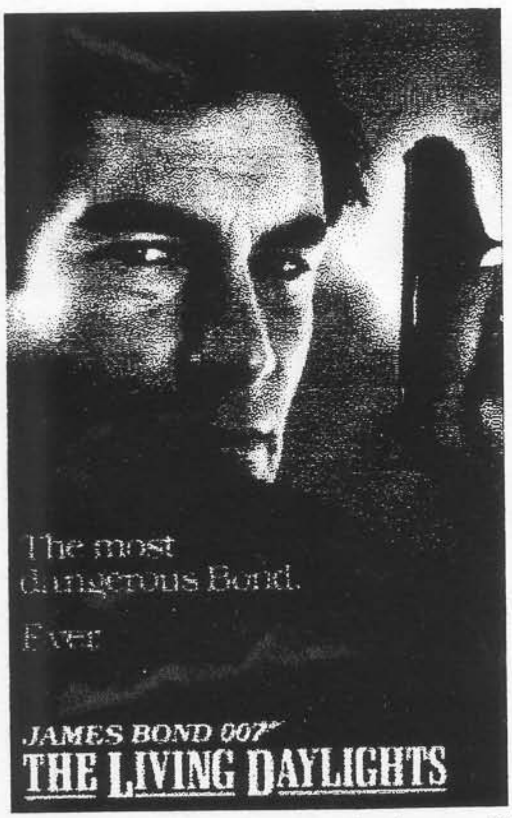

Figure 3 Gray scale image halftoned using the error diffusion method.

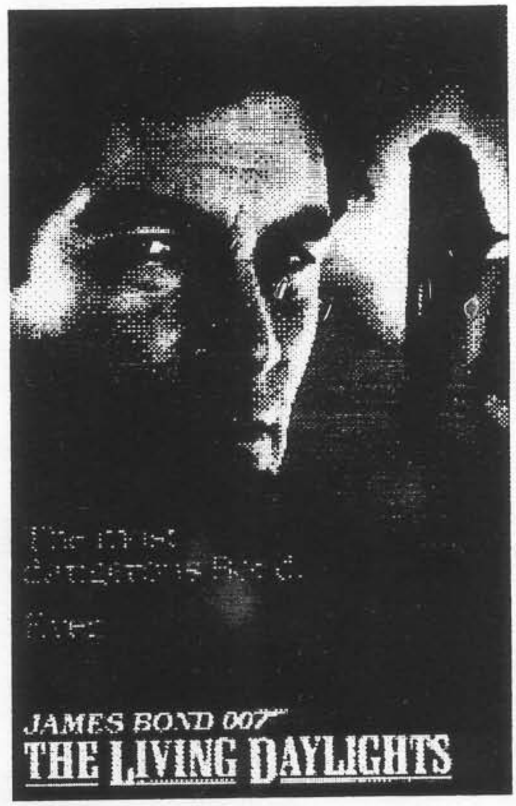

Figure 4 Gray scale image halftoned using the ordered dither method.

5. Results

In Figures 2, 3, and 4 we show a gray scale image (640×400) that was halftoned using the blue noise mask method, an error diffusion method (Ulichney's perturbed weights [1987]), and an ordered dither method [Bayer 1973], respectively. Ordered dither and the blue noise mask methods were approximately 50 times faster than the error diffusion method. In Figure 5 we show the results of the psychovisual test for level $g = 0.87$ and for three different Principal Frequency locations ($K1 = 1/\sqrt{2}$, $K2 = 1$, and $K3 = 0.543$). For each of these locations five different images were created corresponding to different peak heights and widths. As can be seen from the graph, the images that correspond to $K = 1/\sqrt{2}$ were rated best, regardless of the height and width of the Principal Frequency peak. This was one of the reasons for the choice of $K = 1/\sqrt{2}$ in the design of the blue noise mask.

6. Conclusion

We have presented a new halftoning method where the halftoning is achieved by a simple pixelwise comparison against a nonimage array, the blue noise mask. This method combines the speed of point processes such as ordered dither with the aperiodic, high quality results of neighborhood algorithms such as error diffusion.

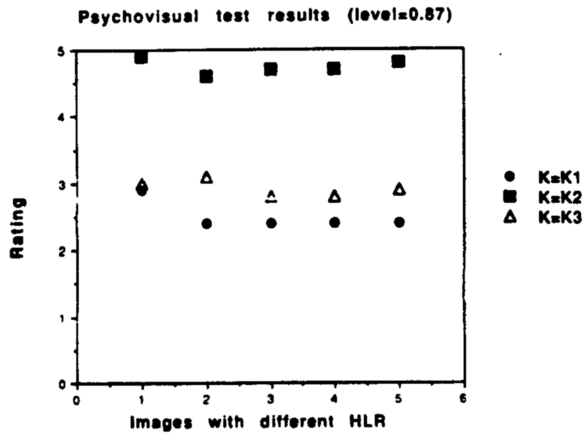

Figure 5 Psychovisual test results.

References

Ulichney, R. *Digital Halftoning*. MIT Press, 1987.

Allebach, J.P. and B. Liu, "Analysis of halftone dot profile and aliasing in the discrete binary representation of images," *J. Opt. Soc. Am. 67*, 1147-1154, 1977.

Bayer, B.E., "An optimum method for two-level rendition of continuous-tone pictures," *IEEE Int'l. Conf. Comm.*, 26-11–26-15, 1973.