The Construction and Evaluation of Halftone Patterns with Manipulated Power Spectra

Theophano MITSA†, Robert ULICHNEY‡, and Kevin J. PARKER†

† *Dept. of Electrical Engineering*
University of Rochester
Rochester, NY 14627
email: {mitsa, parker}@ee.rochester.edu
‡ *Digital Equipment Corp.*
146 Main St., MLO5-2/U44, Maynard, MA 01754
email: ulichney@gauss.enet.dec.com

Abstract: We present a novel algorithm for manipulating the power spectrum of halftone dot patterns. A study is performed to determine the influence of quantitative parameters on subjective quality ratings. We focus on the role of the low-frequency cutoff in the manipulated spectrum in effecting pattern appearance. The new algorithm is useful for understanding the features of halftone patterns which are important in subjective evaluations, and can also be of use in the design of dispersed dot dither patterns.

Keywords: halftone, dot profiles, digital halftoning, power spectral analysis.

1 Introduction

We define the *dot profile* as the binary pattern that results from the halftoning of a constant gray level. The visual appearance of the dot profile depends on the halftoning process that was used. For example, dot profiles that were produced using the ordered dither method appear periodic, whereas dot profiles that were produced using error diffusion have a more *unstructured* visually-pleasing appearance. It is now accepted that visually pleasing halftone dot profiles such as those produced by the error diffusion method with perturbed weights, have certain properties that can be derived from the ensemble average power spectrum of the binary patterns. These properties include low anisotropy and a *blue noise*, or reduced low frequency, spectrum [Ulichney 87].

Anisotropy can be quantified as a measure of relative variance of power spectrum samples at fixed radial distances from the zero frequency (*dc*) term. Low anisotropy implies a homogeneous distribution of pixels and a lack of directional orientation in both the image domain and the Fourier transform domain.

Blue noise spectra imply the lack of large clumps or structures in the image domain, which corresponds to a lack of low frequency content in the Fourier transform domain. The spatial frequency that corresponds to the average spacing of *minority* pixels in the image domain for a fixed gray level is referred to as the Principal Frequency. For gray levels between white and middle gray, minority pixels will be black; likewise, for gray levels between middle gray and black, minority pixels will be white.

Good *blue noise* patterns have very little spectral energy below the Principal Frequency. The Principal Frequency defines a low-frequency cutoff. As spatial frequency is two-dimensional, this cutoff is the radius of a circular stop band centered on zero frequency.

It should be pointed out that when a white noise signal, which has a flat spectrum, is used as a threshold for halftoning, the objectionable long-wavelength textures that appear in the resulting mezzo-tint-like image are due to the passage of low frequency energy. The low-frequency cutoff for the power spectrum for all gray levels is effectively zero.

In this study, we design dot profiles that have the same first order statistics (average gray level) and similar isotropic spectra as achieved with variations on the error diffusion method, but differ in that each image has a different low-frequency cutoff. The results show that, when spectra are synthesized with the algorithm presented here, observers prefer a dot profile with a lower cutoff frequency than that predicted by using the average spacing criteria that defines the Principal Frequency.

2 Generation of dot profiles

For a fixed gray level, the Principal Frequency is given by [Ulichney 87]:

$$F = \sqrt{g}/S \qquad \text{for } g \leq .5$$

or

$$F = \sqrt{1-g}/S \qquad \text{for } g > .5$$

where F is the Principal Frequency expressed as cycles per unit length, g is the gray level expressed as a fraction of white to total pixels, and S is the horizontal or vertical distance between addressable points on the display. (Square pixels are assumed.)

We introduce the following modification to this definition for the case of gray level greater than 0.5. We will use a low-frequency cut-off:

$$F = K * \sqrt{1-g}/S$$

where K is a scale factor. For typical *blue noise* patterns generated with an error-diffusion based process, $K = 1$.

For this study, we have chosen to work with dot profiles for a fixed gray level of $g = 0.87$, or 87% of the pixels are printed as white, and 13% of the pixels are printed black.

We are able to modify a given *seed* dot profile such that it has a low-frequency cutoff with K less than 1, while the first order statistics and the isotropy are preserved after the modification. This is done by an algorithm which is described in the next section. Our proceedure was applied to a white noise pattern which was thresholded at $g = 0.87$ (Figure 1), to produce three binary patterns with $K = 0.54$

(Figure 2), $K = 0.707$ (Figure 3), and $K = 1$ (Figure 4). For $K = 0.54$ the cutoff frequency (in cycles per unit length) is equal to 0.1947, for $K = 0.707$ it is equal to 0.2549, and for $K = 1$ the cutoff is equal to 0.3606. Note that for the $K = 0.54$ case, shown in Figure 2, the cutoff is roughly one half of that obtained through error diffusion methods. The third case, $K = 1$, shown in Figure 4, represents an average spacing expected from error diffusion methods, whereas Figure 3, $K = 1/\sqrt{2}$, falls between the other two.

Figure 5 shows the radially averaged power spectra for the cases shown in Figures 2, 3 and 4. In these plots the power spectra are in units of the variance of the spatial pattern, where for gray level g, is simply given by $g(1-g)$. In our case ($g = .87$) the variance is 0.1131. Note that all three cases yield a fairly good *blue-noise* shape.

3 The Algorithm for Manipulation of Power Spectra

The general approach of this algorithm is that we wish to enforce different constraints in different domains. In the transform domain, we want the power spectrum of a halftone pattern to possess a circular stop band with a particular low-frequency cutoff. In the image domain, we apply the constraints of binary states (only values of 1 or 0 are permitted), and stationary first order statistics (the average number of 0 and 1 pixels must remain constant within all neighborhoods). These multiple constraints are implemented in an iterative operation which applies transform domain filtering to a binary pattern, then applies image domain operations to result in a modified binary pattern. The process is repeated until an error criterion is satisfied. The result is a binary pattern that closely matches the desired power spectrum.

Specifically, the sequence of steps is:

1. Take the 2-dimensional Fourier transform of the dot profile, $p(i,j)$, and obtain the power spectrum $H(u,v)$; where i,j are the pixel coordinates and u,v are the transform variables. The power spectrum is estimated by means of Bartlett's method [Bartlett 55] of averaging the magnitude squared of the Fourier transform, $P(u,v)$ of samples of $p(i,j)$.
2. Compute the radially averaged power spectrum $H(f)$, where the variable $f^2 = u^2+v^2$, as described in [Ulichney 87]. The spectral estimate, $H(u,v)$ is segmented into thin concentric annuli about the zero frequency term. Each annuli is associated with a radial frequency, f, the radial distance from zero frequency, and the energy within the annuli is averaged resulting in $H(f)$.
3. Design a function $D(f)$, that has the following relationship between the actual radial spectrum, $H(f)$, and the desired radial spectrum $H'(f)$:

$$H'(f) = D(f)H(f)$$

(Note that since our *seed* image, $p(i,j)$ was white noise, $D(f)$ is essentially $H'(f)$ in the first pass.)

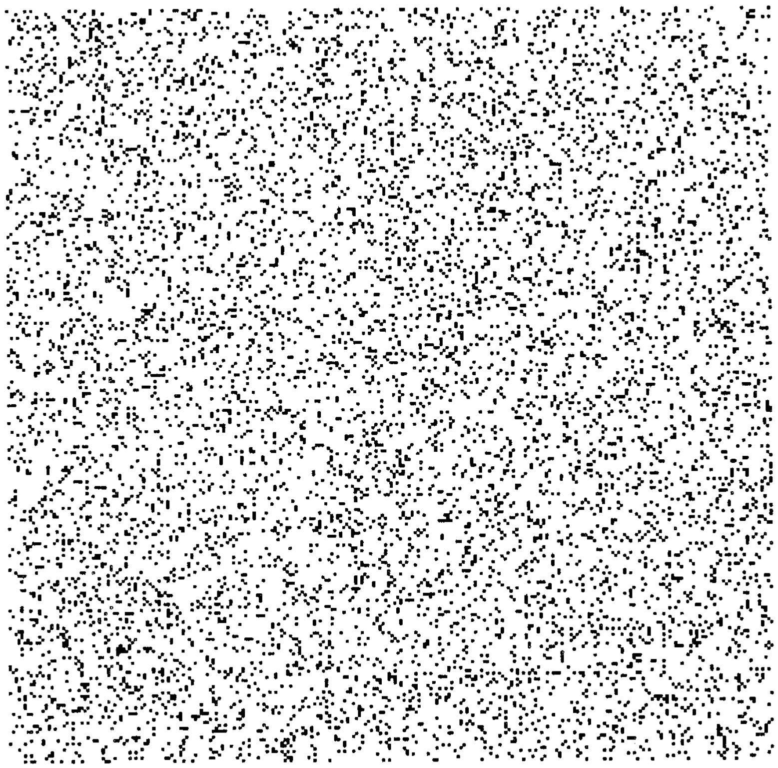

Figure 1. White noise *seed pattern*

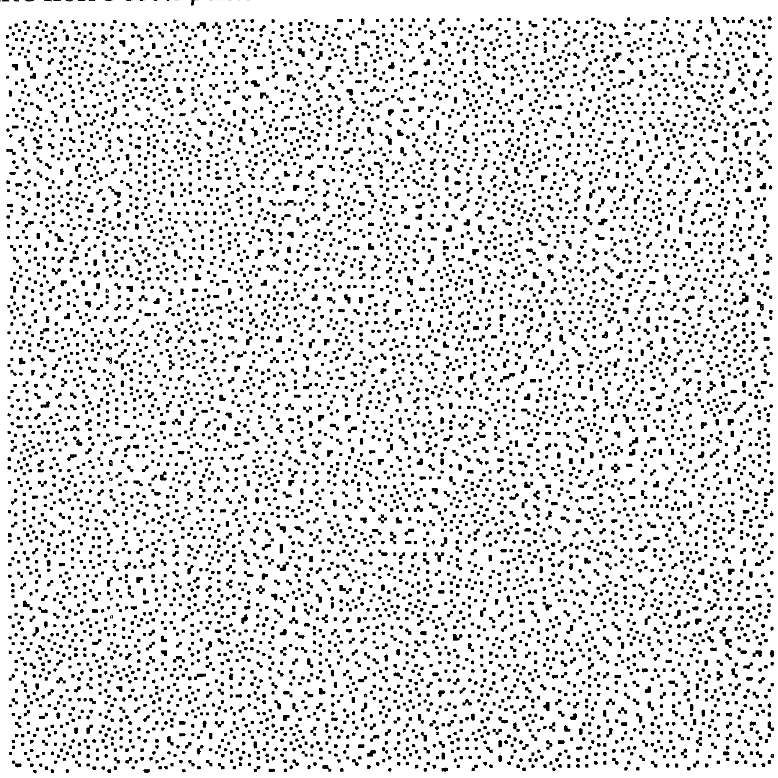

Figure 2. Synthesized pattern with $g = .87$ and $K = 0.540$.

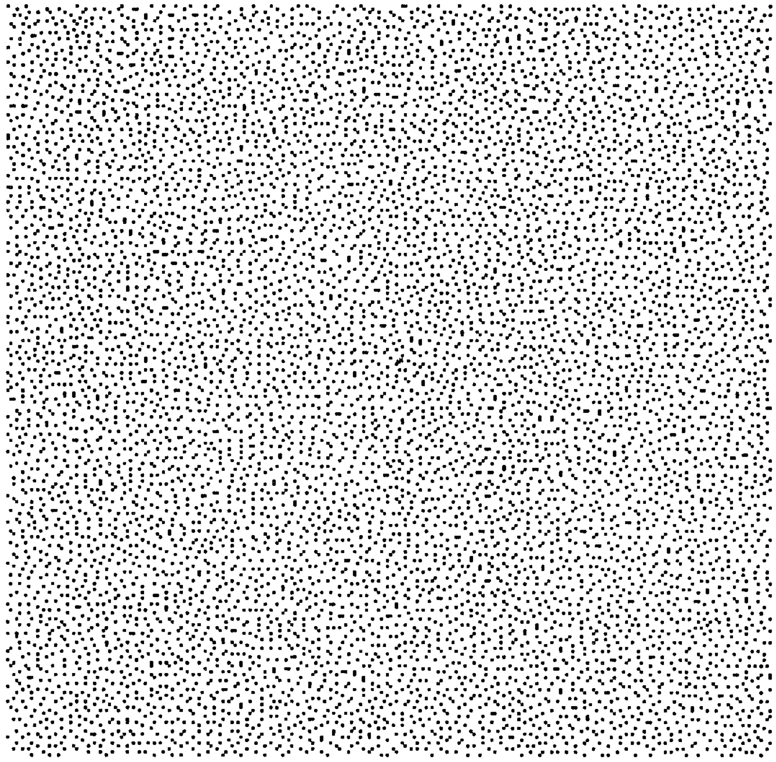

Figure 3. Synthesized pattern with $g = .87$ and $K = 0.707$.

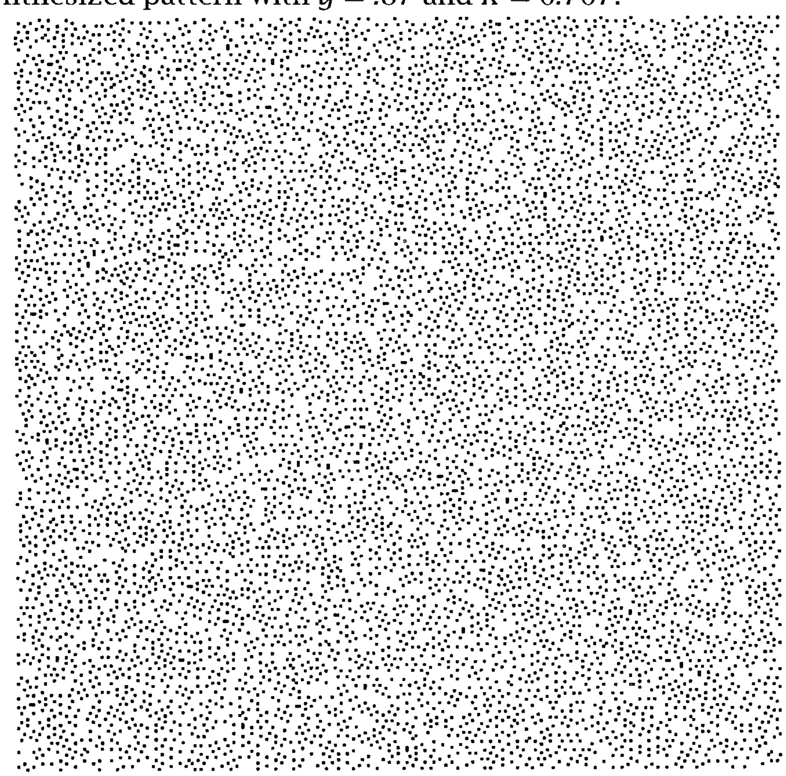

Figure 4. Synthesized pattern with $g = .87$ and $K = 1.000$.

4. Produce a real and even, radially symmetric filter

$$D(u, v) = \sqrt{D(f)}$$

5. Apply the filter to the transform of the dot profile to obtain a new, filtered transform $P'(u,v) = D(u,v)P(u,v)$. While applying the unwindowed filter $D(u,v)$ in this way is not the most accurate filter method, the iterative nature of our approach ultimately mitigates all errors. We note here that by careful choice of the shape and cutoff frequency of $D(u,v)$, the new spectrum $H'(u,v)$ approaches the desired characteristics.
6. Take the inverse Fourier transform of $P'(u,v)$ to obtain an new dot profile $p'(i,j)$. This profile has the desired transform characteristics but is no longer binary because of the filtering operation.
7. Form a difference, or error array $e(i,j) = p'(i,j) - p(i,j)$.
8. Rank order all errors , considering separately those cases where $p(i,j) = 0$ and $p(i,j) = 1$. Form pairs of these ranked pixels, one from the $p = 0$ list and one from the $p = 1$ list, from largest to smallest errors.
9. For some N pixel pairs, defined by the largest errors, replace a 0 with a 1 and a 1 with a 0. In this way, those pixel pairs with the largest error, compared to the desired $p'(i,j)$ pattern, are converted to more closely match the desired, filtered pattern. At the same time, we preserve the binary state of the pattern and the mean values.
10. The process can be repeated until some measure of error is satisfied.

In our experience, using 256×256 arrays, the transform and image domain steps can be completed in reasonable time using $N = 50$ or greater pairs and with reasonable error convergence occurring in 5 to 20 iterations. We note that other criterion, such as the calculation of local neighborhood means, can also be included to influence the choice of pixels to be changed.

4 Visual Evaluation

The binary patterns of Figures 2, 3, and 4 were printed at approximately 70 dots per inch on a laser printer, and the images were shown under identical conditions to 10 volunteers. Image quality with respect to the homogeneity of the pattern was rated on a scale of 1 (worst) to 5 (best). The binary patterns were (256×256) pixels and the viewing distance was approximately 6 times the image size. The average score for Figure 3 was approximately 4.5, whereas the scores for Figures 2 and 4 were approximately 2.5 and 3, respectively. This indicates that the K factor of less than 1, $K = 0.707$, was associated with the most pleasing pattern, presumably because of the lack of majority or minority clumps (white voids or black clumps).

5 Discussion and Conclusion

We have developed a new algorithm for designing halftone dot profiles with modified power spectra. The results show that using the binary-pattern design procedure discussed above, it is possible to generate a binary pattern with a Principal Frequency scale factor of $K = .707$,, that is subjectively better than others, including a pattern with the same first order statistics (gray level) and the expected Principal Frequency scale factor of $K = 1.0$.. As can be seen from Figure 4 increasing the low-frequency cutoff beyond $K = .707$ results in larger majority pixel clumps. This may be due to the fact that, when using this particular algorithm, increasing the cutoff frequency effectively eliminates the minority clumps, however since the algorithm preserves the first order image statistics this may result in majority pixel clumps.

It should be pointed out that while the pattern generated with this iterative algorithm with $K = 1$ did have considerable low frequency *leakage*, as evidenced in the spectral plot in Figure 5. This could account for the clumping. One could argue that many of the $g = .87$ patterns based on modified error-diffusion methods found in [Ulichney 87, ch. 8], with effectively $K = 1$, appear more homogeneous and isotropic than our manipulated spectra at $K = .707$. The point of this work was not to generate the most pleasing pattern, but to explore the consequences of direct manipulation of the power spectrum.

This ability to manipulate halftone pattern power spectra can be useful in the construction of novel halftone screens. For example, in the *Blue Noise Mask* [Mitsa & Parker 91], a single valued, 2 dimensional function is constructed by adding dot profiles for consecutively increasing gray levels. The dot profile for each gray level uses the previous dot profile as a *seed image* and changes are made in the mean value and low frequency cut-off using the algorithm steps 1 through 8, followed by the addition or elimination of 1's as needed to change the mean value.

The ability to directly manipulate power spectra enables further studies to be performed on other seed images, such as error diffusion based blue noise patterns, and on a wide range of gray levels.

References

[Ulichney 87] R. Ulichney, *Digital Halftoning*, MIT Press, 1987.

[Bartlett 55] M.S. Bartlett, *An Introduction to Stochastic Processes with Special Reference to Methods and Applications*, New York: Cambridge University Press, pp. 274–284, 1955.

[Mitsa & Parjer 91] T.Mitsa and K.J.Parker, Digital halftoning using a Blue Noise Mask, *SPIE Conf.*, San Jose, CA,1991.

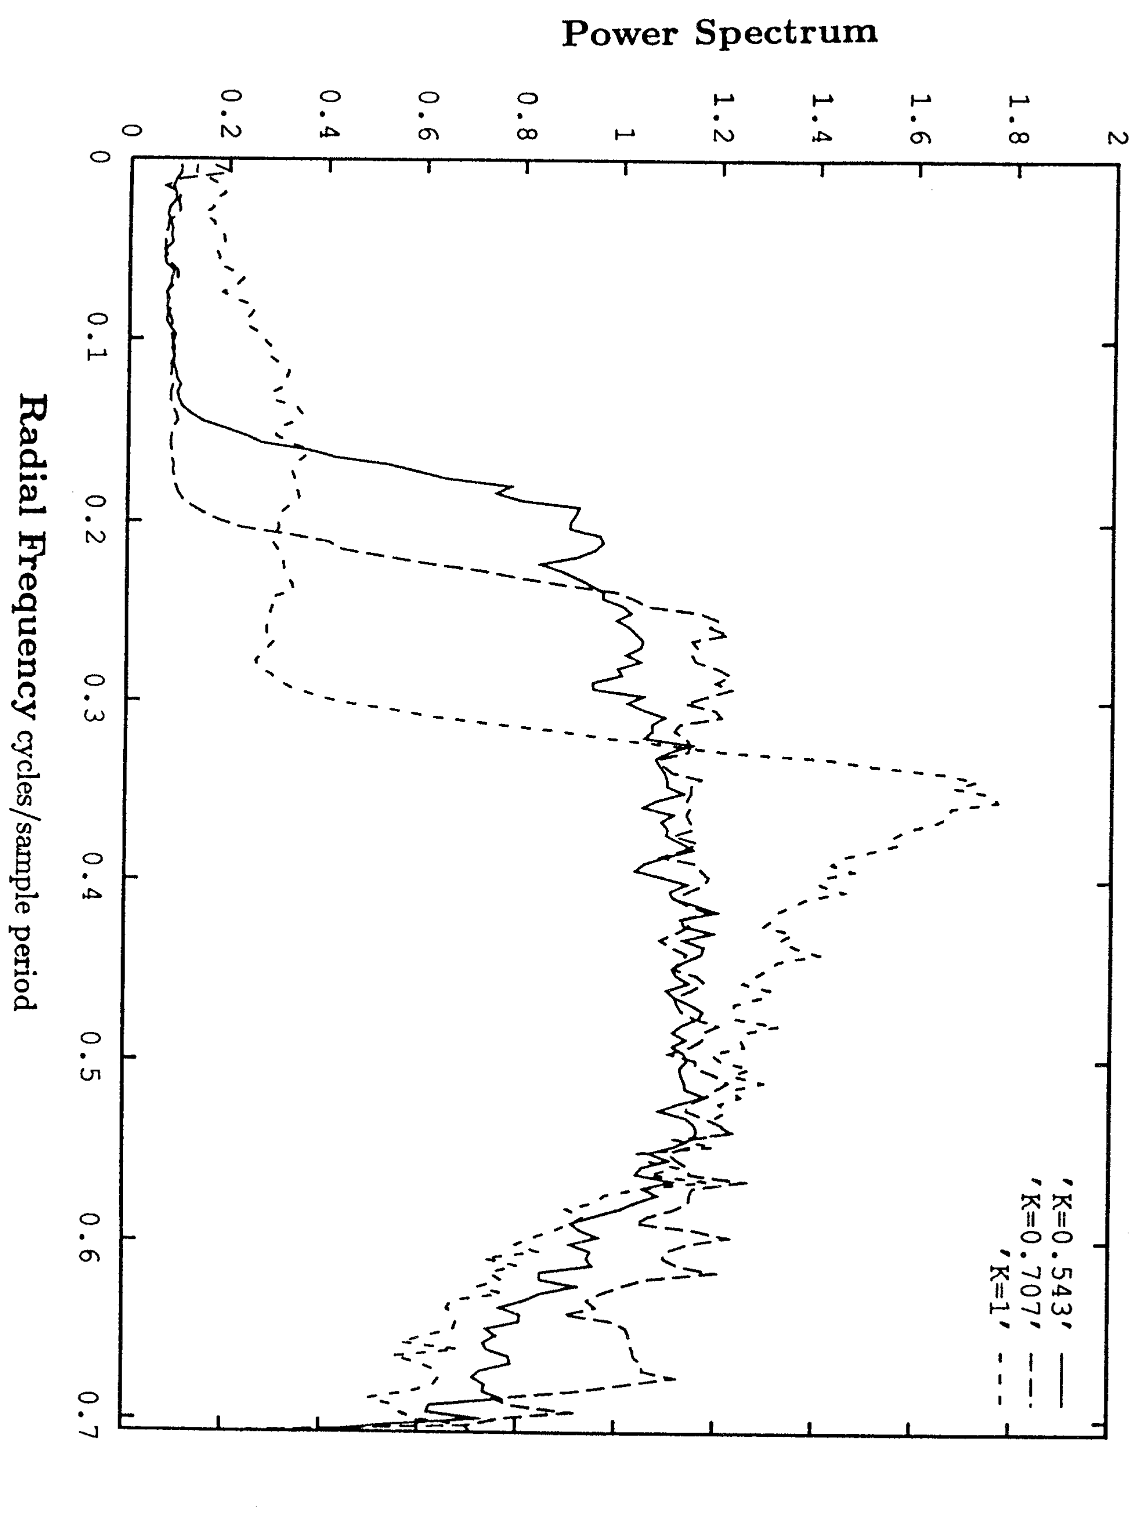

Figure 5. Radially averaged power spectra for the 3 manipulated cases.