# Halftone Characterization in the Frequency Domain

*Robert Ulichney*
*Digital Equipment Corporation*
*Maynard, Massachusetts*

## Abstract

For some classes of signals, characteristics that are complex in one domain, can be made simple in its transform domain. Such is the case with halftone patterns. Metrics to assess halftone features in the spatial frequency domain are presented. These include the radially averaged power spectrum, and the composite Fourier transform. These are used to analyze samples from various methods classified in order of decreasing entropy: white noise, blue noise, ordered dither with recursive tessellation arrays, and ordered dither with clustered-dot arrays. Emphasis is placed on the advantages of blue-noise halftoning and candidate processes that strive to achieve it.

## Introduction

The business of digital halftoning is one of perceptual illusions—the art of exploiting the perceptual system's lack of acuity at higher spatial frequencies to produce the sensation of many gray or color levels from few levels. Employing the frequency domain to assess the quality and nature of halftone patterns of fixed gray levels can provide important insight.

Figure 1 illustrates the result of halftoning with processes from four basic categories of techniques. The same continuous-tone input image was quantized to two levels, and printed at 100 by 100 pixels/inch to amplify the nature of the process used. The four images are ordered by decreasing correlation or entropy. Starting with highest entropy process, white-noise dithering was used in (a), and ending with the most correlated process, clustered-dot ordered dither used in (d). Even though these two cases are at the extremes of entropy or perceived randomness, they both suffer from textures due to long wavelengths, or low frequency energy.

The other two cases tend to reduce low frequencies. The image in (c) is the result of ordered dither with an array generated by the method of recursive tessellation[1]; the patterns described by Beyer[2] result from this method. The method of recursive tessellation builds a dither array by repeatedly subtiling the plane to find the centers of voids. This algorithm was generalized further to allow different starting points, and is called the void-and-cluster method of generating dither arrays.[3,4] When starting from all black or all white, this method will generate the same pattern used in Figure 1(c). While such dithering does have less low-frequency texture, it still suffers from regular rectangular artifacts

Blue noise dithering[6], with an example shown in Figure 1(b), is generally considered to be the most pleasing in that it does not display the regular structures seen in methods (c), nor does it suffer from the long wavelengths present in white noise. There are a number of techniques that fall into this category. Mitsa and Parker [5] were the first to produce ordered dither arrays with blue noise properties by directly synthesizing dither patterns to match a target energy distribution in the frequency domain. The much simpler void-and-cluster method mentioned earlier can also be used to generate blue noise dither arrays. While the results of either of nese methods, among others, will produce a result similar to that shown in Figure 1(b), the method used in this case was a form of modified error diffusion using a serpentine raster and perturbed weights[6].

For the purpose of examining the properties in the frequency domain, it is convenient to partition methods based on the span of the patterns generated.

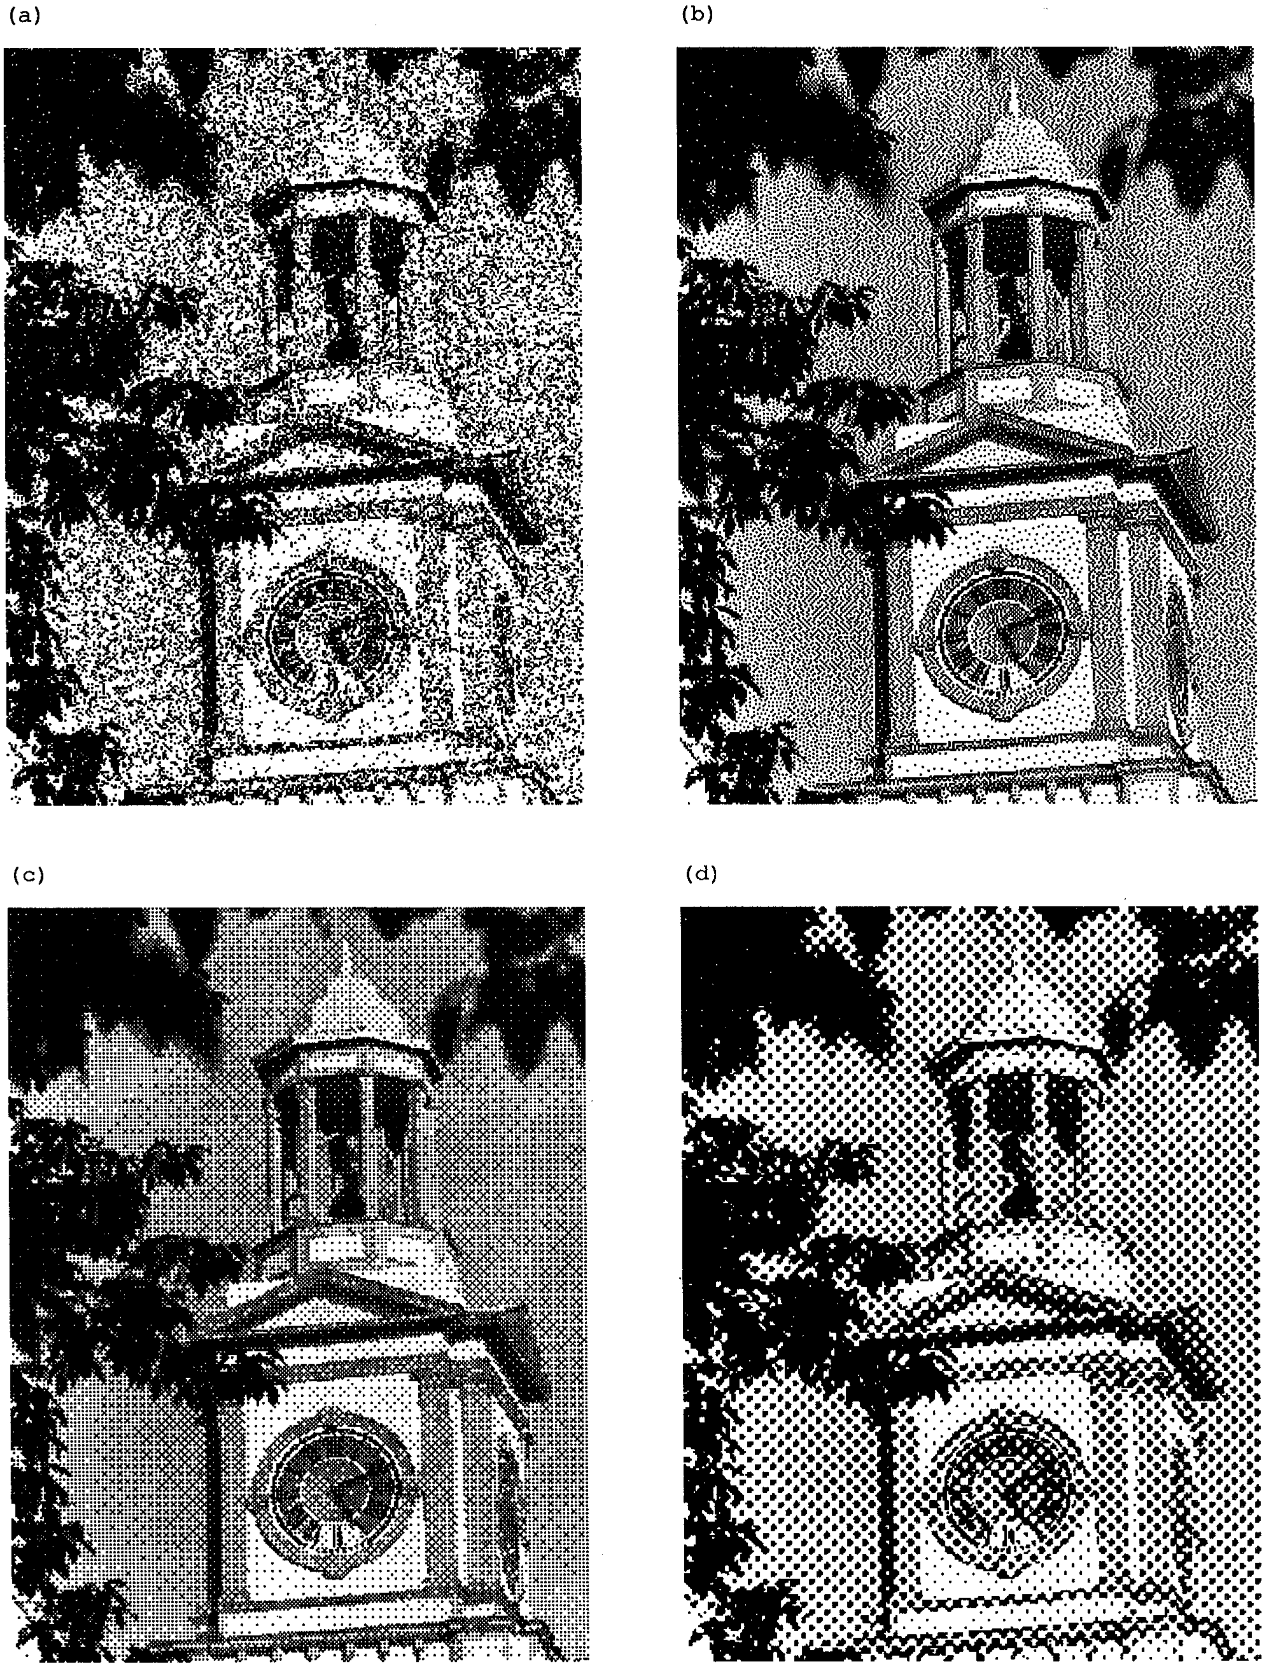

*Figure 1. Examples of halftoning with the four basic categories of techniques ordered by decreasing entropy. (a) White noise, (b) Blue noise, (c) Recursive tessellation arrays, and (d) clustered-dot. Images are shown at 100x100 pixels/inch.*

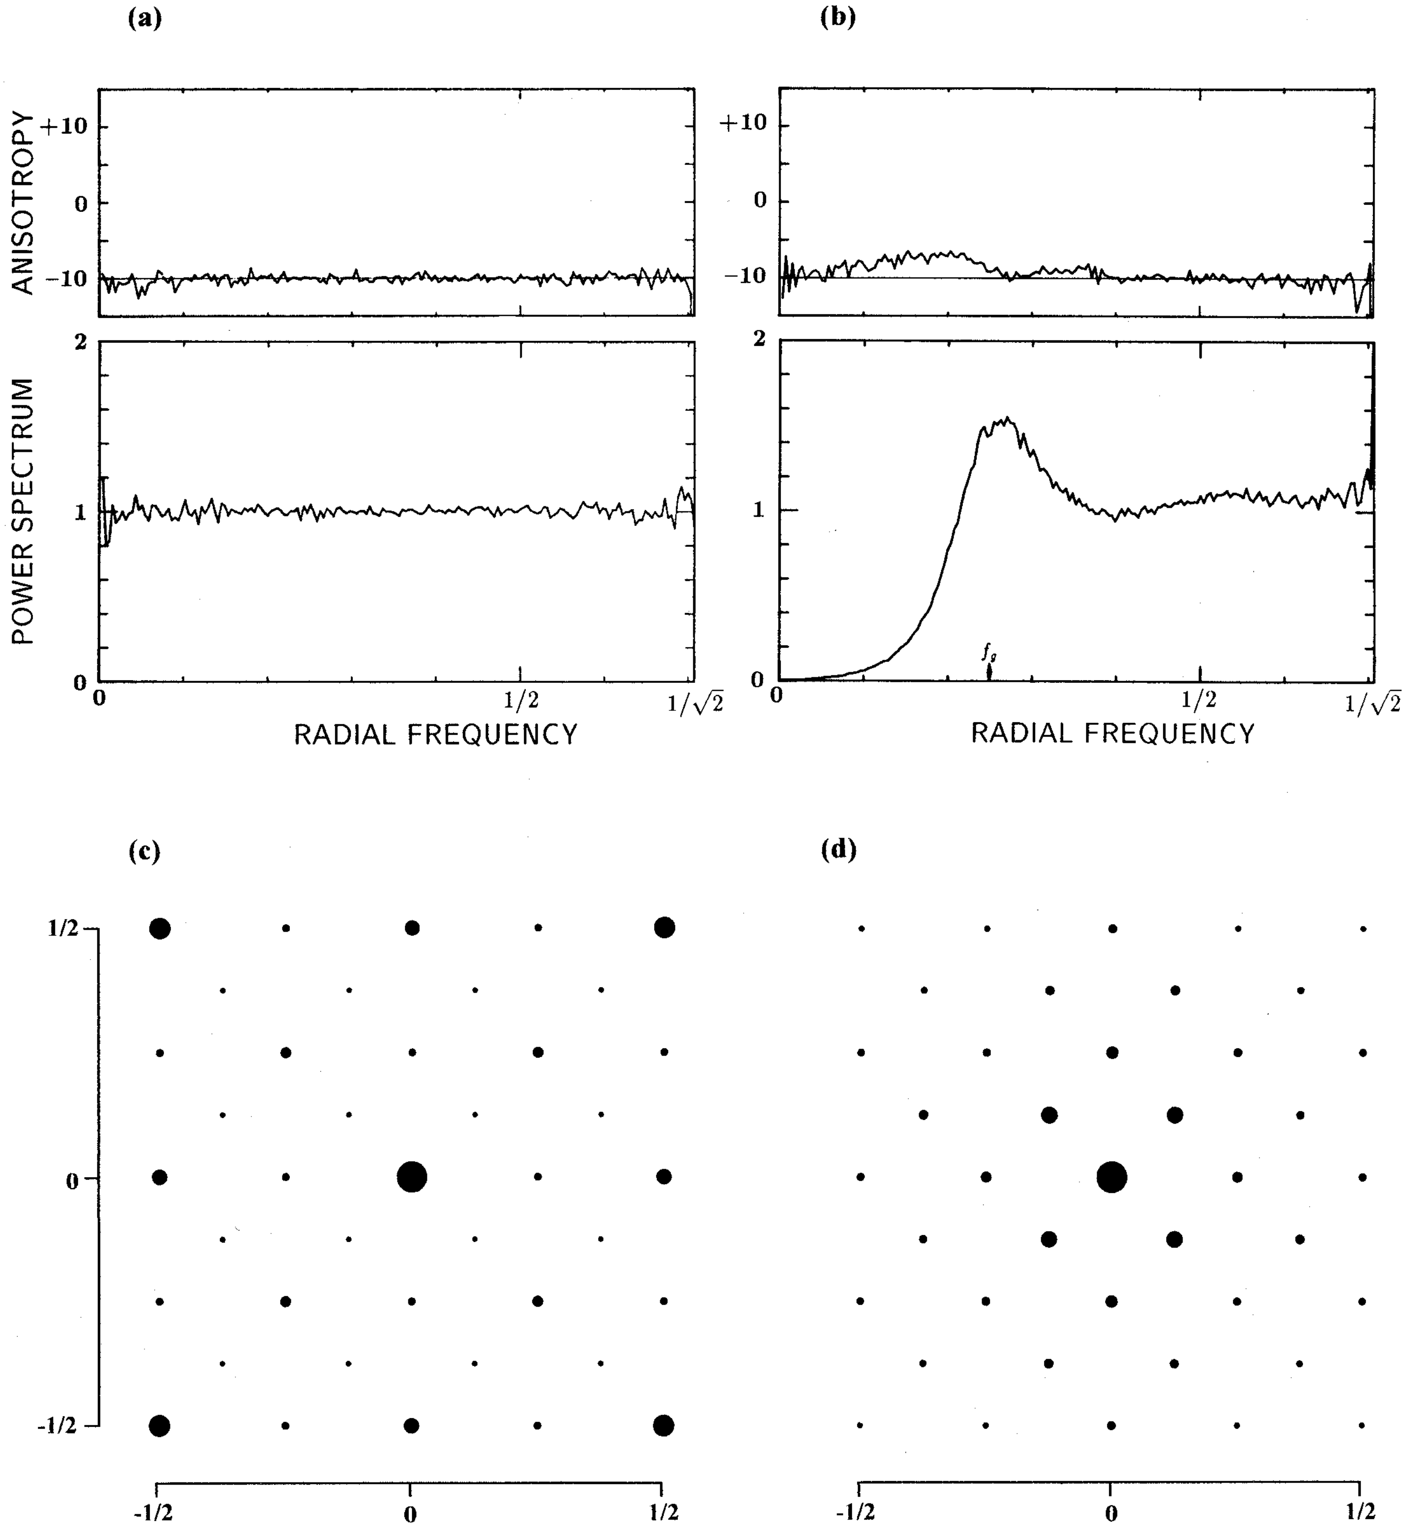

*Figure 2. Frequency domain characterization of the methods used in Figure 1. Radially averaged power spectra are shown in (a) for white noise, and (b) blue noise. Composite Fourier transforms are shown in (c) for recursive tessellation dither arrays, and (d) for clustered dot dither. In all cases spatial frequency is in units of* $5^{-1}$ *cycles/unit-length, where S is the pixel period.*

## Small-Period Patterns

Frequency domain characterizations of patterns resulting from the halftoning methods of used in Figure 1 are shown in Figure 2. The patterns examined are those that result from dithering fixed gray levels. For methods such as (c) recursive tessellation arrays, and (d) clustered-dot, fixed gray levels yield small (typically less than 16 by 16) periodic patterns equal to the size of the dither array. For such small-period patterns, it is convenient to display the average of the magnitude of the Fourier transforms for all fixed gray levels in a composite Fourier transform[6], as shown in Figure 2(c) and (d). In these plots, the areas of dots represent the average magnitude of the frequency coefficient in that location. The DC term is in the center.

In each of these particular cases, the dither array, and thus resulting pattern size, consist of 32 elements. There are 32 unique elements in the transform domain. The high frequency edges are repeated for symmetry. The corners correspond to the highest possible frequency, that due to a checkerboard pattern. (Note that in all cases in Figure 2, frequency is expressed in term of the pixel, or sample, period S, in units of $S^{-1}$ cycles/unit length.) The clustered-dot, or classical screen pattern in (d) shows a considerable amount of frequency energy next to the DC term, and gradually tapers off. These coefficients correspond to the period of the pattern. Contrast this with the plot in Figure 2(c), where the highest concentration of non-DC energy is at the high frequency corners

## Large-Period Patterns

For methods that produce large or even aperiodic patterns, it is convenient to reduce the 3D frequency plot to two dimensions by averaging the power spectrum over concentric annuli. This results in a radially average power spectrum[6] as shown in Figure 2(a) and (b). In this particular case, the spectral plots are shown for a 12.5% fixed gray level. It is important to note that the total energy in the power spectra varies greatly with gray level, peaking at 50% gray and becoming zero at the black and white endpoints. The power spectrum is plotted in units of gray scale variance

$$\sigma_g^2 = g(1-g),$$

so that all plots are normalized.

While the normalized average within the concentric annuli is an important metric, the sample variance is also. It can be used to measure the degree to which the patterns are not radially symmetric or anisotropic as a function of radial frequency. Anisotropy is defined[6] as this sample variance divided by the square of the Power Spectrum. It is a noise-to-signal ratio, expressed in dB in Figure 2(a) and (b). As both the white noise and blue noise methods in this example are aperiodic and stochastic, a form of spectrum estimation was used by averaging 10 periodograms. This results in a theoretical lower bound of -10 dB anisotropy[6], shown as a reference line in the figure.

White noise, so named because it possesses equal energy at all frequencies shows its expected shape in Figure 2(a). The uncorrelated nature of white noise dither is a desirable property, but the low frequency energy is the culprit responsible for the objectionable textures. So what is wanted is high frequency white noise; I call this "blue noise" to contrast with the term "pink noise,"[7] used to describe low frequency white noise. At a fixed gray level, one can think of a principal wavelength that is the average separation of minority pixels is an isotropic distribution. In the frequency domain, this becomes the principal frequency. It is this gray-level dependent frequency that serves as the cut off frequency for blue noise, and is indicated by $f_g$ in Figure 2(b).

A good blue noise shape, and low anisotropy are evident in Figure 2(b). Of course, to fully analyze a candidate blue-noise method, such plots are needed for all, or at least a good representative sample of, gray levels. Directional artifacts are easily monitored with the anisotropy metric.

Frequency domain metrics such as the composite Fourier transform and radially average power spectra have proven to be a useful gage of halftone quality and characteristics.

## References

1. R. Ulichney, "Frequency Analysis of Ordered Dither," *Hard Copy Output, Proc. SPIE*, vol. **1079**, pp. 361-373, 1989.
2. B.E. Bayer, "An optimum method for two level rendition of continuous-tone pictures," *Proc. IEEE Int. Conf. Commun., Conference Record,* pp. (26-11)-(2615), 1973.
3. R. Ulichney, "The void-and-cluster method for dither array generation," *Proc. SPIE, Human Vision, Visual Processing, and Digital Display IV*, vol. **1913**, pp. 332-343, Feb. 1-4, 1993.
4. R. Ulichney, "Filter Design for void and cluster dither arrays", *Proc. SID Int. Symposium*, June, 1994.
5. T. Mitsa and K.J. Parker, "Digital halftoning technique using a blue-noise-mask," *J. Opt. Soc. Am. A*, vol. **9**, pp. 1920-1929, 1992.
6. R. Ulichney, Digital Halftoning, *The MIT Press*, Cambridge, 1987.
7. R. Ulichney, "Dithering with blue noise," *Proc. IEEE*, vol. **76**, no. 1, pp. 56-79, 1988.