

On Filter Techniques for Generating Blue Noise Mask

Kevin J. Parker and Qing Yu

Dept. of Electrical Engineering, University of Rochester, New York

Meng Yao, Color Print and Image Division

Tektronix Inc., Wilsonville, Oregon

Abstract

We use the term "Blue Noise Mask" to mean a halftone threshold array that has been constructed to produce pleasing, unstructured (but not white noise) binary patterns. An efficient way to produce a Blue Noise Mask is to construct a set of binary patterns that are constrained by neighboring levels. These binary patterns are appropriate for different grey levels, and can be summed to form a single threshold array. The key to production of these binary patterns is filtering operations, which have received appropriate attention as an important element in the final quality of the halftone patterns. In this paper, we survey the peer reviewed, scientific literature concerning filter design, emphasizing both the frequency domain and the corresponding image domain operations. Some limitations of filtering are examined, along with some novel filter types. Finally, the role of post-filtering, or adjunct, criteria are examined for special effects.

Introduction

The state of halftoning by the mid 1980s was reviewed in an excellent paper by Ulichney.¹ He also demonstrated the efficacy and advantages of studying halftone patterns, not just in the image domain, but also in the transform domain, particularly examining the radial average power spectrum along with a measurement of anisotropy. He also introduced the term "blue noise" to describe certain binary patterns for which the radially averaged power spectrum (RAPS) contained only small amounts of energy in low frequencies. The definition of principal frequency was linked to the average spacing of minority pixels in a halftone pattern (and therefore to the average grey level). It may be possible to divide the commonly used techniques of that time into different categories [see Figure 1]. One category pertains to the nature of the halftoning process as either a point process (i.e., a halftone screen) or a neighborhood process (algorithm). Another category pertains to the rendered patterns as either locally periodic (repeating small kernels) or locally aperiodic (no local repetitions except over large "tiles"). Figure 1 gives some examples such as Bayer's dither² and clustered dot dither. Floyd and Steinberg's error diffusion³ is shown in two quadrants, since at many grey levels the algorithm produces locally unstructured halftone patterns, but at a few grey levels it actually produces locally repeating patterns. Of particular interest is the quadrant representing point processes (halftone screens) that produce locally unstructured halftone patterns. This contains the highly desirable combination of the speed of point processes with an unstructured pattern. However, white noise is generally unacceptable since it contains objectionable "clumps" or structures that correspond to low spatial frequency energy in the RAPS. Unfortunately, it is not a straightforward task to specify a function that has non-white properties when thresholded at any level, since in a mathematical sense, the thresholding of a band limited function generally produces a "white noise" result.⁴ Thus, constructing a halftone screen to produce desirable, fine, unstructured, blue noise binary patterns at many grey levels

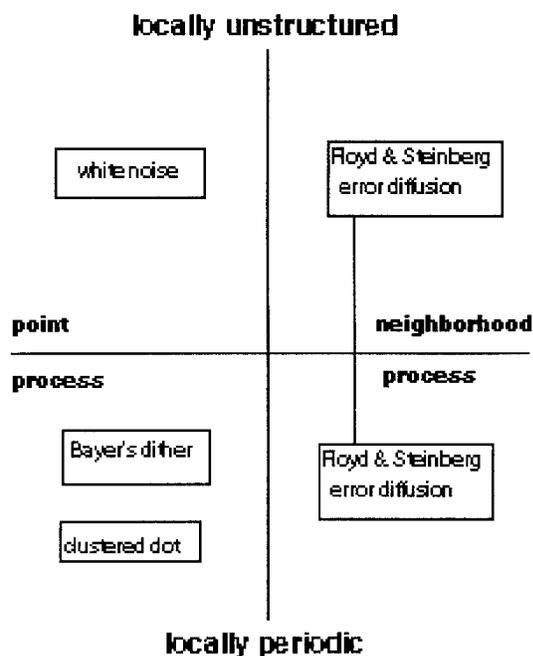


Figure 1. Some commonly used categories of halftoning circa 1980.

was an unrealized goal in the mid 1980's.

Important research combining the power spectrum and human visual system (HVS) concepts with binary patterns was reported by Sullivan, Ray & Miller.⁵ Their approach generated locally unstructured binary patterns, one 32×32 square, "tileable" pattern for each grey level, and they used a cost function with HVS weighting to guide a Monte-Carlo approach with simulated annealing in the creation of individual binary patterns. Although this work did not focus on the production of a halftone (threshold) array, and did not employ the direct filtering approach that is the major subject of this paper, the work serves as an important benchmark whereby individual binary patterns were manipulated using image domain and transform domain concepts in order to produce a desired result.

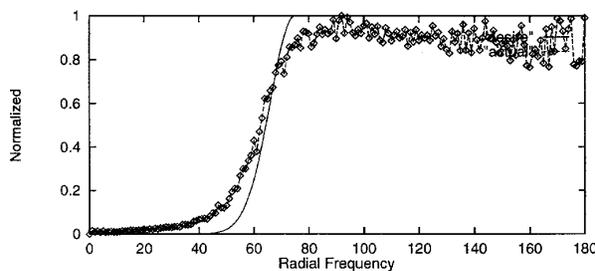


Figure 2. Filter $D(p)$ and actual blue noise RAPS $B_g(p)$ for a blue noise pattern at level $g = 210$

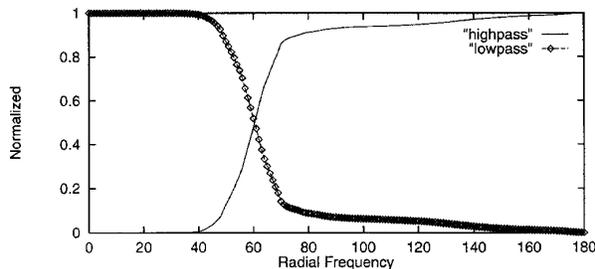


Figure 3. Lowpass filter generated from Fig. 2 and Eq. 2 and 3.

Direct Filtering of Binary Patterns and the Generation of a BNM

A novel approach to the generation of a halftone array was reported by Mitsa & Parker⁶ They produced a 256×256 halftone array of 8 bit numbers which, when thresholded at any level, would produce an unstructured dispersed dot binary pattern. The novel algorithm for generating this BNM recognized that in a halftone array, the binary pattern at any grey level $g + 1$ can be thought of as being "built up" from the binary pattern at level g . And furthermore, they introduced the concept of filtering a binary pattern from level g to select the location of pixels that would be the "best candidates" for addition of majority pixels required for level g

+ 1. Once binary patterns for all grey levels had been sequentially produced in this way from some "seed" level, the binary patterns could be summed to produce a threshold array or Blue Noise Mask. The role of filters in this procedure warrants close attention, and a summary of some important results is given in this section. Mitsa & Parker^{6,7} selected the location of pixels that should be changed from minority to majority values by finding the extremes of an error function $e[i,j]$. This error function was generated by directly filtering the binary pattern of level g and then subtracting the binary pattern from the filtered pattern. Specifically, in the image domain:

$$e(i, j) = [h_{hp_g}(i, j) * b_g(i, j)] - b_g(i, j) \quad (1)$$

where $h_{hp_g}(i, j)$ is a high pass filter selected for level g , $b_g(i, j)$ is the binary pattern, and $*$ denotes convolution with circular "wraparound" properties. In the transform (or spatial frequency) domain:

$$E(k, l) = B(k, l)[H_{hp_g}(k, l) - 1] \quad (2)$$

Since $H_{hp_g}(k, l)$ was described as a blue noise (high pass) filter, then the overall filter $[H_{hp_g}(k, l) - 1]$ is a low pass filter. Thus, an essential feature of this filtering is that binary pattern "clumps" in the image domain (corresponding to low frequency energy in the transform domain) can be located by directly filtering the binary pattern. Also, the highpass region (or lowpass region) can be related to the principal frequency f_g which is a function of the grey level g . Mitsa and Parker also suggested that the filter H_{hp_g} could be made adaptive to directly shape the RAPS of the binary pattern for level g . That is:

$$H_{hp_g}(p) = \sqrt{\frac{D(p)}{B_g(p)}} \quad (3)$$

where $D(p)$ is the desired blue noise RAPS for level $g + 1$ and $B_g(p)$ is the known RAPS for level g . As with any filtering approach, care should be taken to avoid unwanted discontinuities in the transform domain, or "ringing" in the image domain. This approach is depicted in Figures 2 and 3. Note that this specific filter is computationally more involved than simple low pass filters, but the approach demonstrates the central requirement of providing a low pass filter with a cutoff frequency linked to grey level g by the principal frequency¹:

$$f_g = \begin{cases} \sqrt{g/R}, & \text{for } g \leq 1/2 \\ \sqrt{1-g/R}, & \text{for } g > 1/2 \end{cases} \quad (4)$$

where R is the distance of the addressable points on the display. This relationship is plotted in Figure 4. The dashed lines in Figure 4 correspond to the factor of $1/\sqrt{2}$ (in principal frequency, or $\sqrt{2}$ in average separation) that was discussed by Mitsa and Parker⁷ as an empirical choice of transition cutoff frequency for some filters. In a later paper, Parker, Mitsa and Ulichney⁸ demonstrated how, at a single grey level, changes in the filter cut-off frequency could

produce different types of final halftone patterns. Specifically, for f_c representing the cut-off frequency of a filter, let $f_c = Kf_g$; where K is an adjustable scaling factor. For $K = 0.5$ to 1.0 , different "textures" of binary patterns were produced, with $K = 1/\sqrt{2}$ employed for general use. Ulichney⁹ further explored the filter issue, choosing a 2-D Gaussian filter, implemented in the image domain for direct operation on the binary pattern. He demonstrated that, for smaller arrays less than 32×32 for example, even a single Gaussian filter (not adjusted for grey levels as in the previous work) could produce a useful BNM for some applications. Of course, the seed pattern and Gaussian width require careful choice in order to produce a desirable result. In general, it is beneficial to vary the cut-off filter (or Gaussian width parameter) with grey level.^{7,10} Yao and Parker¹⁰ also demonstrated that a variety of low pass filter shapes could produce desirable halftone patterns, so long as the filter parameters were adjusted for appropriate cut-off with respect to the principal frequency. Mitsa and Brothwarte¹¹ further developed the concept of a filter bank (for different grey levels) as a wavelet concept. Low pass filters were each adjusted for a specific cut-off below the principal frequency ($K = 2$), and this was related to a wavelet-type filter bank. Dalton¹² described the use of band-pass filters to produce textured binary patterns. Thus, the general lessons from this work are that the direct filtering of binary patterns can be useful in selecting pixels that can be changed to produce a desired result in the image domain, and correspondingly approximate a desired power spectrum in the transform domain. Some other remarks on filtering are contained in recent patent documents, but these have not yet been published in the peer-reviewed scientific literature, and therefore are outside the scope of this discussion. To illustrate the varieties of filters that have been used, Figure 5 and 6 depict the frequency domain and corresponding spatial domain filter from Mitsa and Parker,^{6,7} Ulichney,⁹ and Yao and Parker,¹⁰ in each case the filter represented is the one that would be used for level 210 out of 256. Only isotropic filters are shown, but anisotropic filters have also been used.¹⁰

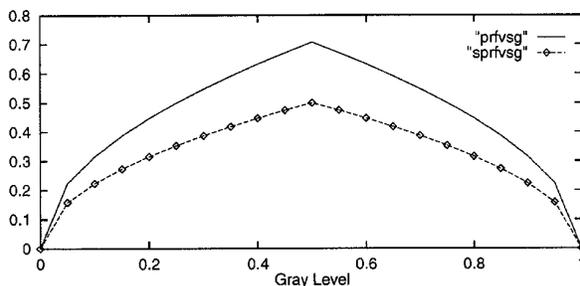


Figure 4. Principal frequency for each gray level

We note that the specification of a filter as a low pass filter in the image domain provides perhaps the easiest way to explain the algorithm to new designers. However, the benefit of specifying the filter in the transform domain, and as

an initially high pass operation, is that the final RAPS of the binary pattern will generally approximate the shape of the high pass filter that is specified in Eqs. 2 or 3.^{7,8,10} Thus, a halftone designer considering the final RAPS of the binary pattern can envision changes resulting from different filter shapes by considering the filter as an (initially) high pass frequency domain filter as taught in the early reference.^{6,7}

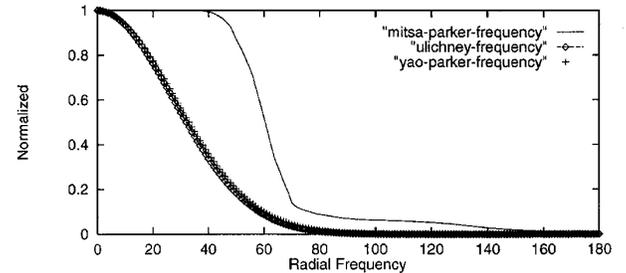


Figure 5. Different filters in frequency domain

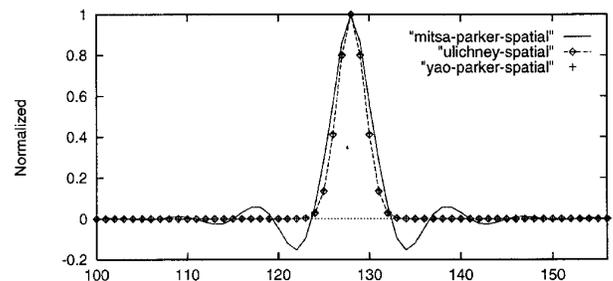


Figure 6. Different filters in spatial domain

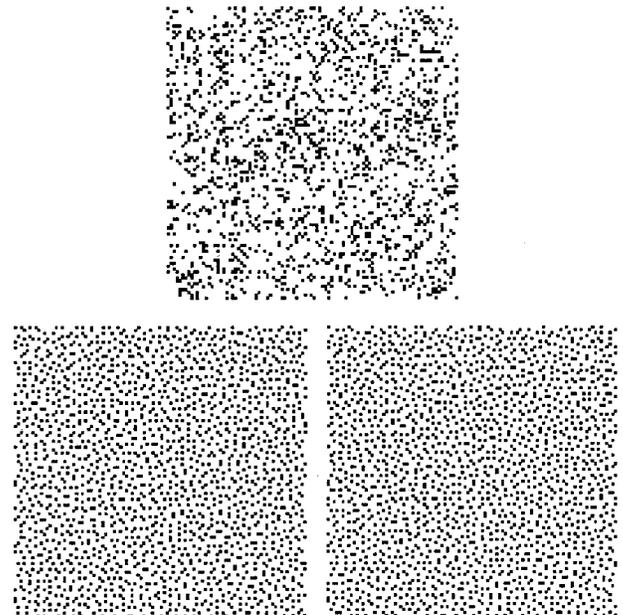


Figure 7. Upper: white noise "seed pattern"; Left : from Gaussian filter; Right: from Gaussian filter with band-reject component

Other Filtering and Post-Filtering Techniques

In this section, we explore further modifications to the filters and to the selection of pixels using additional criteria or process, the latter simply termed "post-filtering techniques." One interesting filter application is the manipulation of the spectral peak at the principal frequency. In filtering the binary pattern with a simple low-pass filter, such as a Gaussian, the low frequencies are minimized by changing certain identified pixels. By careful choice of a band-pass filter, one can enhance selective frequencies in the transform domain, through the proper identification of certain pixels in the image domain. For example, let us take a binary pattern at level 210 shown in Figure 7. Now let us use two different filters with the algorithm for identifying and swapping pixel values:

a) Gaussian:

$$F(u,v) = e^{-\frac{u^2 + v^2}{2\sigma^2}} \quad (5)$$

b) Gaussian with a band-reject component:

$$F(u,v) = e^{-\frac{u^2 + v^2}{2\sigma^2}} - ae^{-\frac{(\sqrt{u^2 + v^2} - f_g)^2}{2\sigma^2}} \quad (6)$$

Figure 8 and 9 show the frequency domain and corresponding spatial domain filters for each, where $a=0.02$, $f_g \approx 110$. Each of these filters was repeatedly applied to the "seed" pattern with pixel swapping until the MSE stopped decreasing. The resulting patterns are given in bottom part of Figure 7 and the corresponding RAPS in Figure 10.

Thus, selective enhancement (or suppression) of regions of the power spectrum can be designed by proper selection of filters. Note also that the difference in lowpass filters (Figs. 8 and 9) in either domain is very subtle but the resulting binary patterns have a recognizably different RAPS. The examples given here demonstrate the wide utility in designing binary patterns with particular characteristics. In post-processing techniques, one can introduce additional steps or criteria beyond the output of the filters. For example, in printing the concept of dot gain is important. Consider that dot gain is related to the area-to-perimeter ratio of printed spots. Generally speaking, halftone patterns with smaller area-to-perimeter ratios (i.e. finely dispersed patterns) will exhibit greater dot gain. To reduce dot gain, we can increase the area-to-perimeter ratio. A nonsymmetric BNM is made for that purpose using a post-processing technique after filtering the darker grey level binary patterns where dot gain is significant. Suppose we use the mid-level (128 if we assume that the total number of levels is 256) as the starting level in making a BNM. The starting pattern, which is obtained by repeatedly filtering a white noise pattern and removing the clumps, will have connected white dots and black dots as well as isolated dots. As we build the BNM

pattern from level 128 to level 127, we need to replace certain white dots by black dots. Normally, these white dots are chosen by examination of the extreme of a low pass filter operation, and connected white dots are more likely to be picked than isolated white dots. This will result in substantial dot gain, because as the number of white dots is decreasing, the number of isolated white dots is increasing, resulting in a decreased area-to-perimeter ratio. However, if we select the isolated white dots first as we move to lower levels and keep the connected white dots intact until there are no more isolated white dots, dot gain can be reduced to a certain degree. Figure 11 shows different results for a uniform BNM and a nonsymmetric BNM. In addition, other filtering and perturbation steps can be applied to enforce other criteria or characteristics.

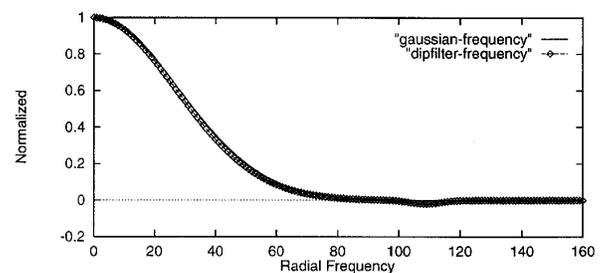


Figure 8. Two filters in frequency domain

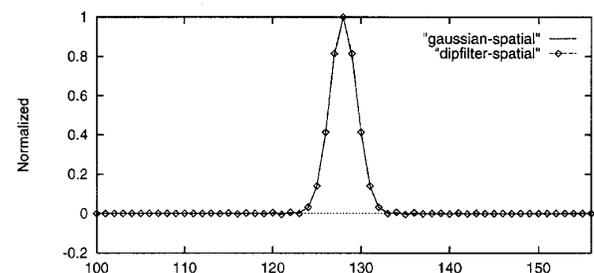


Figure 9. Two filters in spatial domain

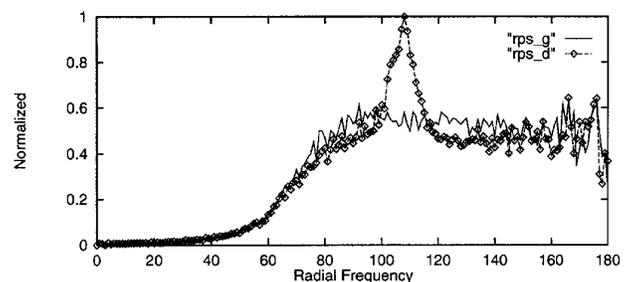


Figure 10. RAPS for the two blue noise pattern

Conclusion

A well accepted method of producing a BNM for halftoning

involves the creation of binary patterns from a "seed" pattern. An important concept is the direct filtering of binary patterns to select pixel locations for changes in the binary pattern. The filters may be adaptable to the grey level, and may be adaptable to the individual binary patterns that are created. Generally speaking, the filters described in the literature since 1991 have had smooth low pass characteristics, and almost all have been adjustable to grey levels. As is common in filtering theory, either the image domain or the corresponding transform domain filters may be utilized, depending on implementation preferences. This paper emphasizes that filtering and post-filtering operations can be designed to produce highly specific characteristics in binary patterns, and therefore a variety of specialized Blue Noise Masks.

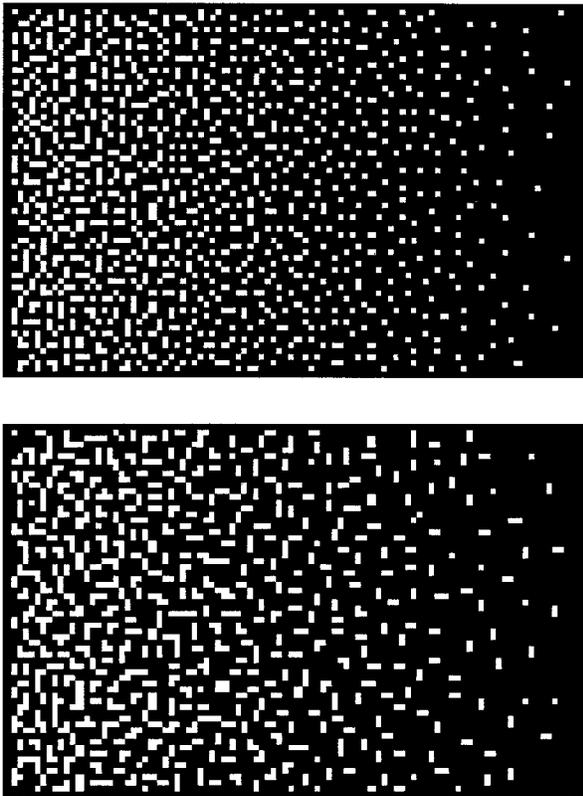


Figure 11. Gray ramps for symmetric (upper) and nonsymmetric (lower) BNMs for levels 95 - 0 but beginning with the same seed pattern at $g=128$. (not shown)

References

1. R. A. Ulichney, "Dithering with blue noise," *Proc. IEEE*, vol. **76**, pp. 56-79, 1988.
2. B. E. Bayer, "An optimum method for two-level rendition of continuous-tone pictures," *Proc. IEEE Int. Conf. Commun.*, pp. 2611-2615, 1973.
3. R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial grayscale," *Proc. Soc. Inf. Displ.*, vol. **17**, pp. 75-77, 1976.
4. E. Steinberg, R. J. Rolleston, and R. Easton, "Analysis of random dithering patterns using second-order statistics," *J. Elec. Imag.*, vol. **1**, no. 4, pp. 396-404, 1992.
5. J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Trans. Syst. Man Cybern.*, vol. **21**, no. 1, pp. 33-38, 1991.
6. T. Mitsa and K. J. Parker, "Digital halftoning using a blue-noise mask," in *Proc. Soc. Photo-Opt. Instrum. Eng.* **1452** (M. Civanlar, S. K. Mitra, and B. P. Mathur, eds.), pp. 47-56, 1991.
7. T. Mitsa and K. J. Parker, "Digital halftoning technique using a blue-noise mask," *J. Opt. Soc. Am.*, vol. **A**, no. 9, pp. 1920-1929, 1992.
8. K. Parker, T. Mitsa, and R. Ulichney, "A new algorithm for manipulating the power spectrum of halftone patterns," *SPSE's 7th Int. Congress on Non-Impact Printing*, pp. 471-475, 1991.
9. R. A. Ulichney, "The void-and-cluster method for dither array generation," *ISST/SPIE Symposium on Electronic Imaging Science & Tech.*, 1993.
10. M. Yao and K. J. Parker, "Modified approach to the construction of a blue noise mask," *J. Elec. Imag.*, vol. **3**, no. 1, pp. 92-97, 1994.
11. R. Mitsa and P. Brathwaite, "Wavelets as a tool for the construction of a halftone screen," *ISST/SPIE Symposium on Electronic Imaging Science & Tech.*, 1995.
12. J. C. Dalton, "Perception of binary texture and the generation of stochastic halftone screens," *ISST/SPIE Symposium on Electronic Imaging Science & Tech.*, 1995.