On the Manipulation of Power Spectra of Halftone Patterns

T. Mitsa[1], K. J. Parker[1] and R. Ulichney[2]

[1]Electrical Engineering Department
University of Rochester, NY 14627
[2]Digital Equipment Corp.
Maynard, MA

Abstract

It has been shown that visually pleasing patterns of black and white pixels in halftone images are characterized by an isotropic, unstructured, dispersed dot pattern known as "blue noise". A relatively unexplored area of research is the examination of how different halftone patterns with different variations on the "blue noise" power spectra are perceived by human observers. A newly developed algorithm for power spectrum matching with binary patterns is used as a tool for gauging the sensitivity of human observers to different features of the power spectrum such as the location of the low frequency cut-off. The algorithm also makes possible the construction of novel halftone screens such as the "Blue Noise Mask".

1) Introduction

As has been previously discussed by Ulichney [1, 2] image power-spectrum characteristics play an important role in characterizing the appearance of halftone patterns. Blue-noise, or high frequency white noise characteristics in the Fourier domain correspond to unstructured, homogeneous patterns in the image domain, whereas white noise corresponds to grainy patterns with visually annoying low-frequency clumps. It is well established today that visually pleasing, homogeneous patterns, like the ones produced by some error diffusion techniques, have blue-noise power spectrum characteristics. In the halftoning literature today, the phrase "blue-noise pattern" implies a binary pattern that has the following characteristics:

- In the Fourier domain, most of the energy is concentrated in the high frequencies, while the low frequencies have negligible energy [1, 2]. An important note here is that the terms "low frequencies" and "high frequencies" are used to denote the relative position of the frequencies with respect to the Principal Frequency location; i.e., frequencies below the Principal Frequency are "low" and frequencies above the Principal Frequency are "high".

- In the Fourier domain, the Principal Frequency is located at the position that corresponds to the average distance of the minority pixels in the binary pattern.

- In the image domain, the binary pattern is isotropic; i.e. there are no preferred directions in the distribution of black and white dots.

However, this definition is somewhat vague, because it does not specify the exact energy amount of the Principal Frequency peak in order to have a pleasing blue-noise pattern. Another interesting question is how perturbations of the Principal Frequency position in the Fourier transform affect the appearance of the binary pattern in the image domain. In this section, we will address these issues and examine:

- The strictness of the definition of the Principal Frequency as the cutoff.

- The effect of different sizes of the Principal Frequency peak on the visual appearance of the binary pattern.

For that purpose, we design binary patterns that have the same first-order statistics (average gray level) but have spectra that differ in the location, height and width of the Principal Frequency peak.

2) The Binary Pattern Power Spectrum Matching Algorithm

The BIPPSMA algorithm generates desired binary patterns by enforcing different constraints in the image domain and Fourier domain. In the Fourier domain, we design the power spectrum of the binary pattern such that it possesses a particular Principal Frequency and has a specific energy distribution.

In the image domain, we apply the following constraints. The image must have a Binary state; i.e., only zeros and ones are permitted as pixel values. The image must also have homogeneous first-order statistics; i.e., the ratio of zeros to ones within a neighborhood is approximately equal to the gray level this binary pattern represents. BIPPSMA implements these constraints in an iterative fashion, by alternating between the two domains, until an error measure is satisfied. The result is a binary pattern that closely matches the desired power spectrum.

Specifically, the steps of BIPPSMA for the l^{th} iteration are:

1. Compute the Fourier transform $H_l(u, v)$ of the binary pattern $h_l(i, j)$ and obtain the power spectrum $P_l(u, v)$ by computing the square of the magnitude of the Fourier transform:

$$P_l(u,v) = |H_l(u,v)|^2.$$

The power spectrum is computed in this fashion and not by periodogram averaging, because we choose to work specifically with $h_l(i, j)$ as a signal realization of a random process, and thus treat h as a deterministic signal.

2. Compute the radially averaged power spectrum $P_{lr}(f_r)$ of $P_l(u, v)$ by averaging over concentric annuli of radius f_r centered around $u=v=0$.

3. Design a 1-D filter $D_{lr}(f_r)$ such that, when applied to $P_{lr}(f_r)$, it will produce the desired radially-averaged power spectrum, which will be denoted as $P'_{lr}(f_r)$. This can be expressed as follows:

$$P'_{lr}(f_r) = |D_{lr}(f_r)|^2 \times P_{lr}(f_r).$$

4. Produce a 2-D real and even, radially symmetric filter $D_l(u,v)$, by computing the square root of $|D_{lr}(f_r)|^2$ and replicating it for all angles in the Fourier domain, such that the entire 2-D Fourier space is filled. The radial symmetry of the filter in the Fourier domain corresponds to homogeneous and unstructured patterns in the image domain.

5. Apply the filter $D_l(u,v)$ to the Fourier transform $H_l(u,v)$ of the binary pattern

$$H'_l(u, v) = |H_l(u, v)| \times D_l(u, v).$$

6. Compute the inverse Fourier transform of $H'_l(u, v)$ to obtain $h'_l(i, j)$, a modified pattern that has the desired transform characteristics, but is no longer binary, because of the filtering operation:

$$h'_l(i, j) = IFT\{H'_l(u,v)\}$$

where *IFT* denotes the inverse Fourier transform.

7. Form a difference, or error array: $e_l(i, j) = h'_l(i, j) - h_l(i, j)$

8. Rank order all errors, considering separately those cases where $h_l(i, j) = 0$ and $h_l(i, j) = 1$. It is important to note here that , in the case of zeros, we rank order only the negative errors.

9. For the N_p pixel pairs (of 0 and 1 pixels) with the largest error magnitudes, where N_p is a given number, replace a zero with a one and a one with a zero. It is important to note here that before converting the value of a pixel, we check its neighborhood mean in order to ensure that no clumps will be caused by the conversion. This pairwise exchange achieves the following:

- By converting pixel pairs with the largest error compared to the desired $h'_l(i, j)$ pattern, we more closely match the desired filtered pattern.
- The binary state of the signal is preserved, since no new values are introduced.
- The mean is preserved, because the number of replaced zeros is equal to the number of replaced ones. Thus, BIPPSMA is a mean-preserving process.

After the exchange of zeros and ones, we obtain a modified pattern $h_{l+1}(i, j)$ which will be the starting pattern for the next iteration.

10. Examine the mean square error (MSE) between the desired pattern $h'_l(i, j)$ and the modified binary pattern $h_{l+1}=1(i, j)$.

If $MSE(l+1), \leq MSE(l)$, proceed to the next iteration and use $H_{l+1}(i, j)$ as the starting pattern in Step 1. Otherwise, exit the algorithm and choose $h_l(i, j)$ as the binary pattern that satisfies best the applied constraints. In a typical run on a 256 x 256 array at level g=0.87, the algorithm converged in approximately twenty iterations. The number of pixels that are replaced at each iteration, N_p, was chosen to be N_p= 256. Different values of N_p were tried, and for smaller values the number of iterations increased, while for bigger values clumping problems occurred.

3) BIPPSMA Results

BIPPSMA was applied to binary patterns, with dimensions 256 x 256, at two different constant gray levels g_1 = 0.87 and g_2 = 0.95. For both gray levels, the starting pattern was white noise thresholded at levels g_1 and g_2. The starting white noise pattern for level g_1 is shown in Figure 1. We experimented with different variations in the power spectrum of the desired signal. In order to examine the response to the Principal Frequency cutoff, we introduced a scaling factor K such that:

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

We experimented with three scaling factors, $K1 = 0.543$, $K2 = 1\sqrt{2}$ *and* $K3 = 1$. The third scaling factor corresponds to the standard definition of the Principal Frequency. In order to examine the effect of different amounts of Principal-Frequency energies on the appearance of the binary patterns in the image domain, we experimented with different heights and widths of the Principal Frequency peak. Specifically, we tried two different values for the width of the Principal Frequency peak, 4 and 11 discrete frequencies respectively. For the height, we tried the ratios of 1.5 and 2.5 between the maximum of the Principal-Frequency peak and the high-frequency baseline of 1.0. We also created binary patterns for both g_1 and g_2 that had no Principal Frequency peak.

In Figures 2, 3 and 4, we show binary patterns that correspond to gray level $g = 0.87$, and that are all designed to have no Principal Frequency peak but different Principal Frequency locations. For Figure 2, the Principal Frequency is $f_g = 0.194$ ($K1 = 0.543$), for Figure 3 $f_g = 0.254$ ($K2 = 0.707$), and for Figure 4 $f_g = 0.360$ ($K3 = 1$). The corresponding normalized radially-averaged power spectra $P_r(f_r)/\sigma_g^2$ are shown in Figure 5. The DC term is not shown. The anisotropies of these binary patterns were calculated and were all relatively low.

It is interesting to compare the spectra of these patterns with the desired spectra for each pattern, shown in Figure 6. As can be seen from this figure, BIPPSMA successfully creates binary patterns with power spectra that match the desired ones, with the exception of some low-frequency leakage and some additional energy at the Principal Frequency location for $K = 1$. Low-frequency leakage is inevitable using this particular algorithm for the following reason: BIPPSMA replaces continuous, filtered values of $h_l'(i, j)$, with binary values in $h_{l+1}(i, j)$, according to the error array. This substitution of binary for continuous value can be viewed as a quantization procedure, which classically is treated as a white-noise process. However, as shown in Figure 5, the energy in the low frequencies is relatively small for $K1 = 0.543$ and for $K2 = 0.707$, compared to the energy in the high frequencies. A degree of low-frequency leakage is actually a stabilizing feature in BIPPSMA, because in this way we avoid the problem of dividing with zero in Step 3.

Other examples were systematically studied. In one case we created a binary pattern that corresponds to $g = 0.87$ and is designed to have a Principal Frequency peak with an eleven-pixels width, and a ratio of 1.5 between the highest value of the Principal Frequency peak and the high-frequency baseline. Also created was a binary pattern that corresponds to $g = 0.87$, and is designed to have a Principal Frequency peak with a four-pixels width and a ratio of 2.5 between the highest value of the Principal Frequency peak and the high-frequency baseline. The desired radially-averaged power spectra for both of these patterns are shown in Figure 7, while their actual radially-averaged power spectra are shown in Figure 8. As can be seen from these figures, the actual power spectra match well the desired power spectra, and both were relatively isotropic.

4) Psychovisual Results

Fifteen (256 x 256) binary patterns with different location, height, and width of the Principal-Frequency peak were created for level $g_1 = 0.87$. Also nine images with a different location and height of the Principal-Frequency peak were created for $g_2 = 0.95$. The Principal-Frequency energy affects the general high-frequency to low-frequency energy ratio in the spectrum of the pattern, and for that reason, images with different Principal-Frequency energies will be referred to, in the following, as images with different HLR (High-to-Low Ratio). We did not vary the width for $g_2 = 0.95$, because of the very little effect of different widths on the visual appearance of a binary pattern. The binary patterns were printed at 70 dots/in on an Apple Laserwriter and presented in a random order to ten subjects under indentical conditions. The rating scale we used, is a frequently used scale [3] from 5(best) to 1(worst). The viewing distance was standardized at approximately 10 inches.

We found that different amounts of HLR had a very small effect on the rating of the patterns, typically yielding less than 0.5 difference in ratings using g=0.87 and g=0.95 gray level halftone images with HLR between 1 and 2.5. On the other hand, a different Principal-Frequency location significantly affected the rating of the pattern. Patterns, for which $K = 0.707$, were rated the best (4.5 average). Patterns with $K = 0.543$ and $K = 1$ were rated significantly lower (2.5 average), because of the visually annoying clumps that dominate them.

5) Conclusion

We have introduced a novel algorithm, BIPPSMA, for the power-spectrum manipulation of binary patterns. Using this algorithm, we created binary patterns with the same first-order statistics, but different second-order statistics. A psychovisual test was conducted in order to establish a connection between the visual appearance of a binary pattern and its second-order statistics. Our conclusions based on the psychovisual test results are:

- Using BIPPSMA, best patterns were rated the ones with $K = 0.707$. It is important to note here, that the conclusion that patterns with $K = 0.707$ are rated better than those with $K = 1$ is not a general conclusion, but is connected with the implementation of BIPPSMA.

- The size of the Principal-Frequency peak does not necessarily play a leading role in the visual appearance of a binary pattern.

The ability to manipulate binary patterns towards a desired power spectrum is also useful in construction of novel halftone screens, such as the "Blue Noise Mask" (4, 5).

ACKNOWLEDGEMENTS

This work was supported by the University of Rochester Department of Electrical Engineering.

REFERENCES

1) R. A. Ulichney. *Digital Halftoning*. MIT Press, Cambridge, Massachusetts, 1987.
2) R. A. Ulichney. Dithering with Blue Noise. *Proceedings of the IEEE*, 76(1):56-79, January 1988.
3) J. O. Limb. Distortion Criteria of the Human Viewer. *IEEE Trans. on Systems, Man, Cybernetics*, SMC-9(12):778-793, December, 1979.
4) T. Mitsa and K. J. Parker. Digital Halftoning Using a Blue Noise Mask *SPIE* Vol. 1452 Image Processing, pp. 47-56, 1991.
5) T. Mitsa and K. J. Parker, Digital Halftoning using a Blue Noise Mask, IEEE ICASSP '91, CH2977-7, pp. 2809-2812, 1991.

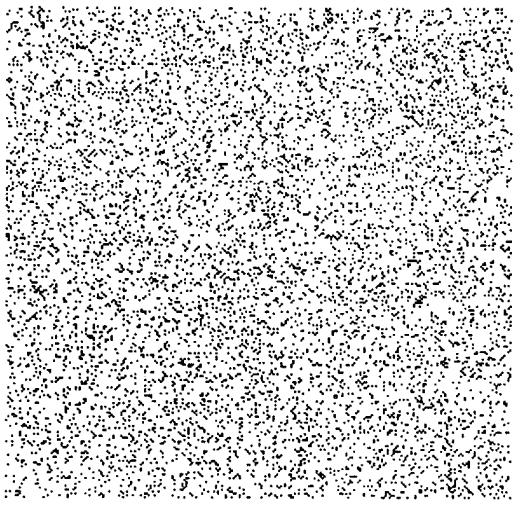

Fig. 1 "Seed" binary pattern used in BIPPSMA

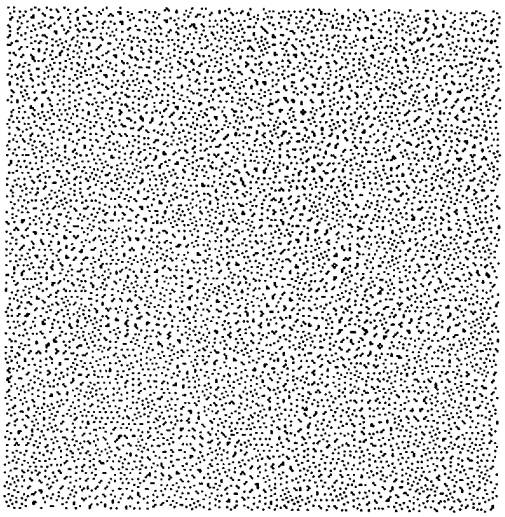

Fig. 2 BIPPSMA output for low Principal Frequency (K1)

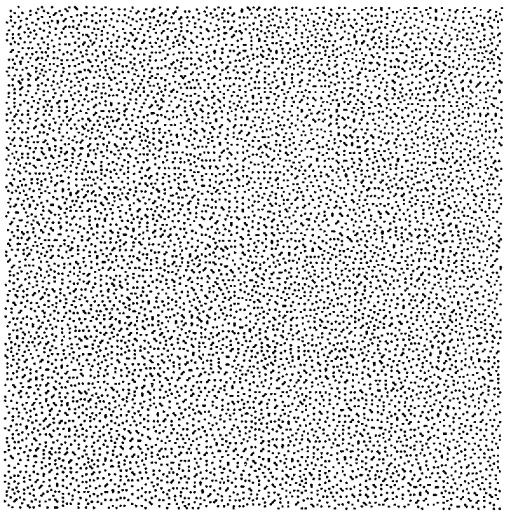

Fig. 3 BIPPSMA output for medium Principal Frequency (K2)

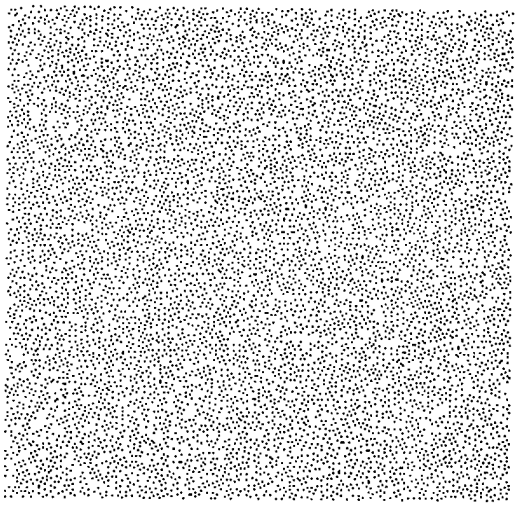

Fig. 4 BIPPSMA output for high Principal Frequency (K3)

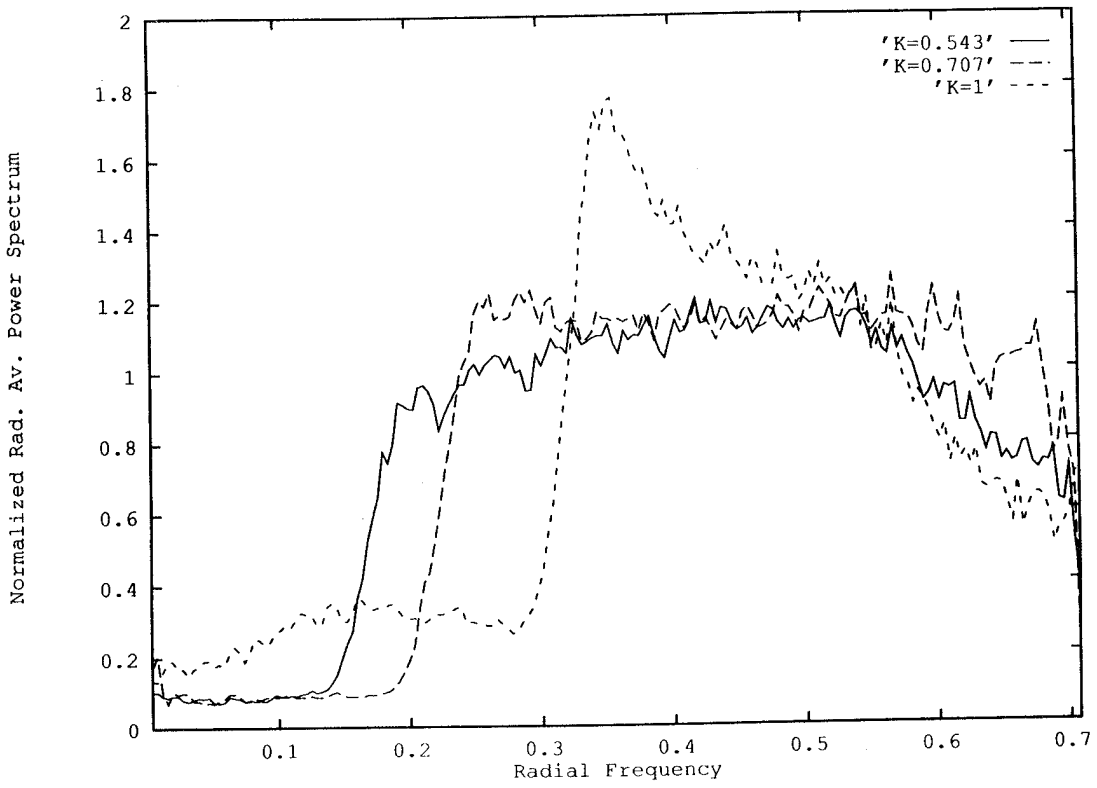

Fig. 5 Radial average power spectra magnitudes (vertical) vs radial frequency (horizontal) corresponding to the binary images shown in Figs. 2, 3, and 4.

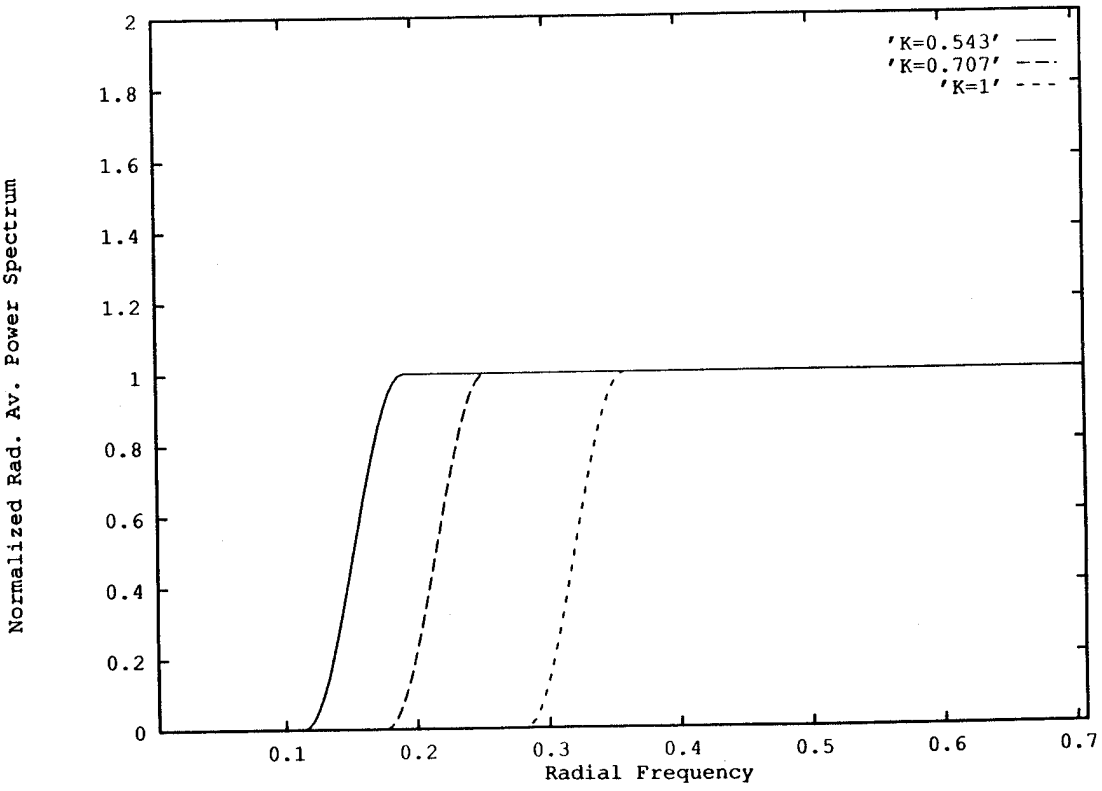

Fig. 6 Desired radial average power spectra for the binary patterns shown in Figs. 2, 3, and 4.

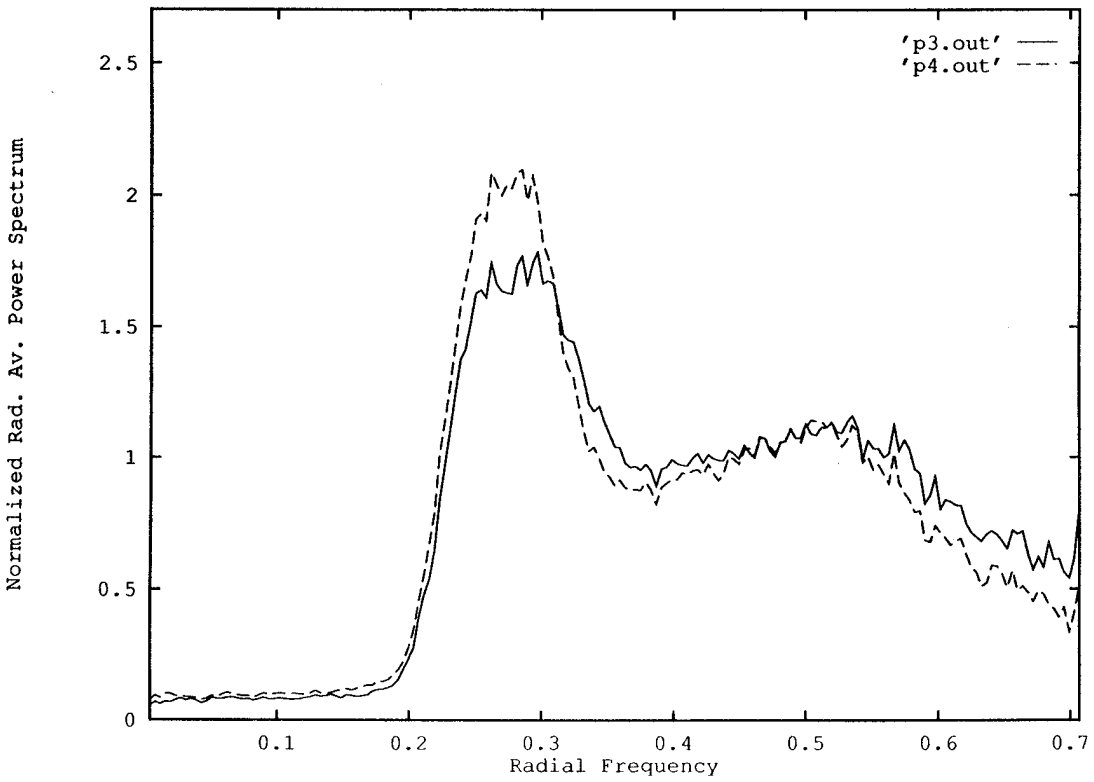

Fig. 7 Radial average power spectra for two different binary patterns with the same Principal Frequency and different energy at the Principal Frequency peak.

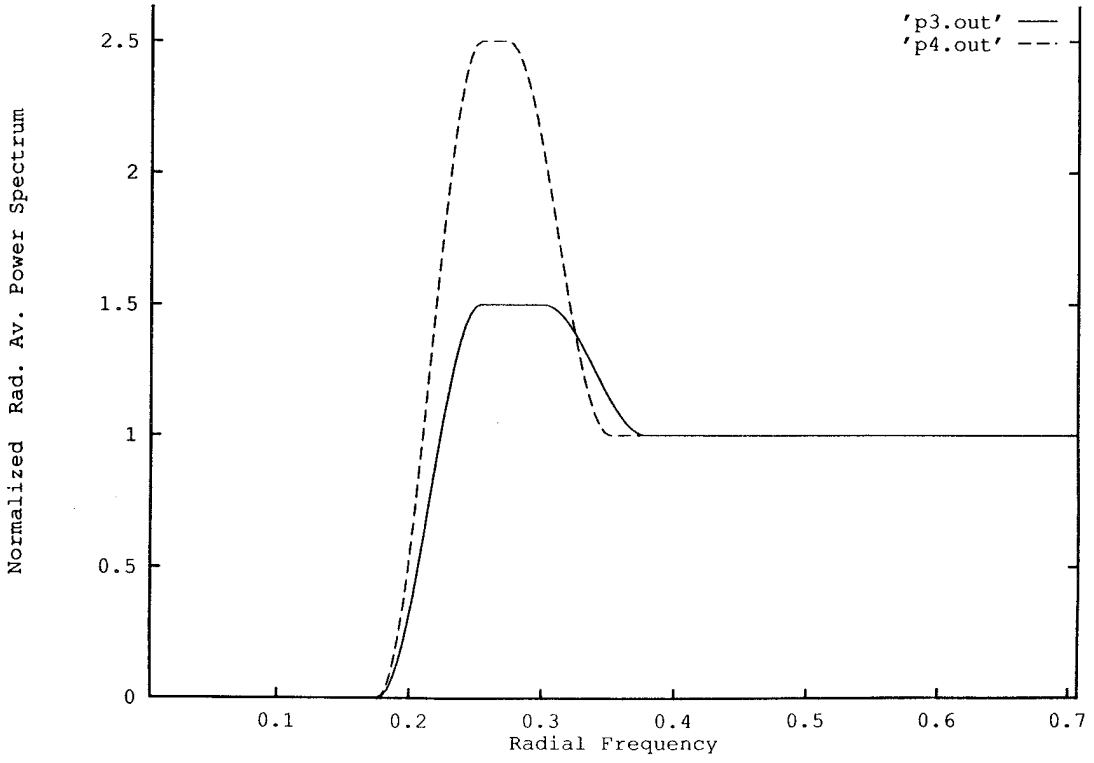

Fig. 8 Desired power spectra for the two binary patterns discribed in Fig. 7.