

# Color Filter Arrays Based on Mutually Exclusive Blue Noise Patterns

Wei Zhu\* and Kevin Parker

*Electrical Engineering Department, Hopeman Building, University of Rochester, Rochester, New York 14627*  
E-mail: [zhu@kodak.com](mailto:zhu@kodak.com)

and

Michael A. Kriss

*Center for Electronic Imaging Systems, University of Rochester, 240 East River Road,  
Rochester, New York 14623-1212*

Received July 21, 1997; accepted January 6, 1999

---

The ordered color filter arrays (CFA) used in single-sensor, color digital still cameras introduce distracting color artifacts. These artifacts are due to the phase-shifted, aliased signals introduced by the sparse sampling by the CFAs. This work reports the results of an investigation on the possibility of using random patterns as a CFA for single-sensor, digital still cameras. From a single blue noise mask pattern, three mutually exclusive, random CFAs are constructed representing the red, green, and blue color filters. An edge adaptive method, consisting of missing-pixel edge detection and boundary-sensitive interpolation is employed to reconstruct the entire image. Experiments have shown that the random CFA alleviates the problem of the low-frequency color banding associated with ordered arrays. This method also has the advantage of better preserving color-free, sharp neutral edges and results in less deviation-from-neutral on high-frequency, monochrome information. © 1999 Academic Press

---

## 1. INTRODUCTION

The emergence of affordable, digital still camera systems has focused attention on the quality of the hard copy images obtained from such cameras. Unlike conventional photographic camera systems, digital still cameras introduce artifacts associated with sampling. Due to cost considerations, most digital still cameras use a single CCD imaging sensor to record the scene [1–3]. The color is encoded by means of a color filter array (CFA). The resulting sparsely sampled images are used to form a complete image by employing sophisticated interpolation algorithms [4]. However, once aliasing is introduced, even

This article was originally scheduled to be part of the special issue on High Fidelity Media Processing.

\* Current address: Eastman Kodak Company, Rochester, NY 14650-1816. Corresponding author.

sophisticated interpolation algorithms cannot compensate for the aliasing introduced by the sparse color sampling. Figure 1 shows the resulting image of a test target (consisting of vertical bars, horizontal bars, and checkerboards) taken with a digital still camera using a Bayer CFA, shown in Fig. 2. Note the low-frequency color banding. This banding is due to the combination of aliasing (where high-frequency information is recorded as low-frequency images) and the spatial phase shift between the color images recorded by the CFA. One of the major challenges of digital camera systems is to find cost-effective ways to eliminate the artifacts introduced by these regular CFAs.

## 2. INTERPOLATION

The images captured by a color filter array contain spatially multiplexed, sampled color images. These multiplexed images must be separated into three sparse (red, green, and blue) images, each of which requires an interpolation algorithm to create a continuous image. An example of this is shown in Fig. 3, where the Bayer CFA is used to encode the color. In this example the following two kernels were used to implement the interpolation[4]: for the green image,

$$\begin{pmatrix} 0 & 0.25 & 0 \\ 0.25 & 1.0 & 0.25 \\ 0 & 0.25 & 0 \end{pmatrix};$$

for the blue and red images,

$$\begin{pmatrix} 0.25 & 0.5 & 0.25 \\ 0.5 & 1.0 & 0.5 \\ 0.25 & 0.5 & 0.25 \end{pmatrix}.$$

The simple black-and-white checkerboard pattern in Fig. 3 is separated into three sparsely sampled color images. The above interpolation kernels are used to form continuous images which are combined to form the final color image at the upper right. The edge patterns are typical of the Bayer CFA using these interpolation kernels. Other, more sophisticated interpolation algorithms can be used to improve the edge patterns, but they cannot eliminate the low-frequency color banding introduced by the sampling process [5, 6].

## 3. ALIASING

Aliasing takes place when an image (or signal) that contains high-frequency information is sampled under twice the Nyquist rate,  $f_N$  [7–9]. The Nyquist frequency is defined as one-half the sampling rate,  $f_S$ . Any information associated with spatial frequencies beyond the Nyquist frequency will be aliased to lower frequencies. To better understand this consider the Fourier transform of a sampled image, as shown in Fig. 4. In this example the input spectrum is assumed to be Gaussian, as shown on the left. The signal is captured by a linear array of square pixels of dimension  $d = 0.01$  mm as shown at the bottom of the figure. The sampling rate is  $f_S = 1/d = 100$  cycles per millimeter and the Nyquist frequency is  $f_N = 1/2d = 50$  cycles per millimeter. The aliased spectrum is shown on the right. The infinite number of replicas (which are a direct result of the sampling) overlap to create a spectrum that is highly aliased. The final image (signal) is represented by the spectrum

between  $-f_N$  to  $+f_N$ . A measure of the aliasing can be calculated by computing the power of nonaliased signal to that of the aliased signal within the frequency band defined by the lower and upper Nyquist frequencies [10, 11]. This will be covered in detail below.

To further understand the nature of aliasing consider Figs. 5 and 6. Figure 5 shows a sine wave of 10 cycles per millimeter and its spectrum. If this signal (image) is sampled at a rate of 16 times per millimeter  $f_S = 16$ , the resulting signal and spectrum is shown in Fig. 6. Note that the aliased frequency,  $f_A$ , is 6 cycles per millimeter. The general relationship between the observed aliased frequency and the original frequency,  $f_O$ , is given by

$$f_A = f_O \pm n f_S.$$

where  $n$  ( $n = 1, 2, 3, \dots$ ) defines which of the infinite number of spectrum replicas will cause the aliased frequency to fall within the frequency band  $-f_N$  to  $+f_N$ .

The aliasing defined by Figs. 5 and 6 can be seen in Fig. 1. The broad, low-frequency banding in both the vertical and horizontal directions are aliased images of the original, high-frequency bar patterns. The aliased images are due to the peaks in the spectrum of the second replica, falling within the spectral band  $-f_N$  to  $+f_N$ .

One method of reducing the effect of aliasing is to use an optical prefilter [12, 13]. The optical prefilter reduces the high-frequency content of the image before it is sampled by the CCD image sensor. While this does reduce the occurrence of color artifacts due to aliasing, it significantly reduces the sharpness of the final image.



The potential for aliasing, PA, for a given digital camera can be calculated [10, 11]. The actual amount of aliasing depends on the scene content. For example, the image of a uniform area will show no aliasing, the image of a face will show some aliasing in the fine detail of the hair, and the image of a picket fence will show extreme aliasing. Consider the general case of a one-dimensional CCD imaging array, as shown in Fig. 7. The imaging element (pixel) is a square area of edge  $d$  and the pixel centers are separated by a distance  $p$  (the pitch). Let  $\alpha = d/p$  be the ratio of the pixel edge to the pitch. The potential for aliasing is given by the ratio of the power of the aliased signal to that of the power of the nonaliased signal in the frequency range of  $-f_N$  to  $+f_N$ :

$$PA = |\text{Aliased power}/\text{Nonaliased power}|_{-f_N}^{+f_N}.$$

If one assumes that the image spectrum is flat (no losses due to lens or optical prefilters), then the amount of aliasing is solely defined by the physical layout of the CCD imaging array. For the case shown in Fig. 7 the value for PA is given by


$$PA = \frac{\pi - \int_{-\pi\alpha/2}^{\pi\alpha/2} \text{Sin}^2(s)/s^2 ds}{\int_{-\pi\alpha/2}^{\pi\alpha/2} \text{Sin}^2(s)/s^2 ds}.$$

Figure 8 shows a plot of PA versus  $\alpha$ . For small values of  $\alpha$ , that is small pixels relative to the pixel pitch, the aliasing will be very large, while as  $\alpha$  approaches unity the potential for aliasing approaches a value of  $PA = 0.32$ . Digital still cameras on the market today use two types of CCD area arrays. The frame transfer cameras use all of the surface area and are best described by  $\alpha = 1$  while interline transfer CCD imaging arrays use only a fraction of the area for imaging and are best described by values for  $\alpha$  between 0.75 and 0.5. Thus, digital still cameras that use interline transfer CCD imaging arrays will introduce more color artifacts due to aliasing than those using frame transfer CCD imaging arrays. Figure 9







Input

Output



Bayer CFA






3

**FIG. 1.** Images from a digital still camera using a Bayer CFA. The image on the left was taken from 10 feet and the one on the right from 12 feet.

**FIG. 2.** Bayer (left) and SONY (right) color filter arrays.

**FIG. 3.** Interpolation scheme for a Bayer CFA.



**FIG. 16.** Three stochastic color filter arrays.

**FIG. 17.** Reconstructed images: (a) white noise CFA. (b) blue noise CFA.

**FIG. 18.** CFAs generated from blue noise mask: (a) BNM, (b) equal numbers of red, green, and blue sites, (c) twice as many green sites as red or blue sites.

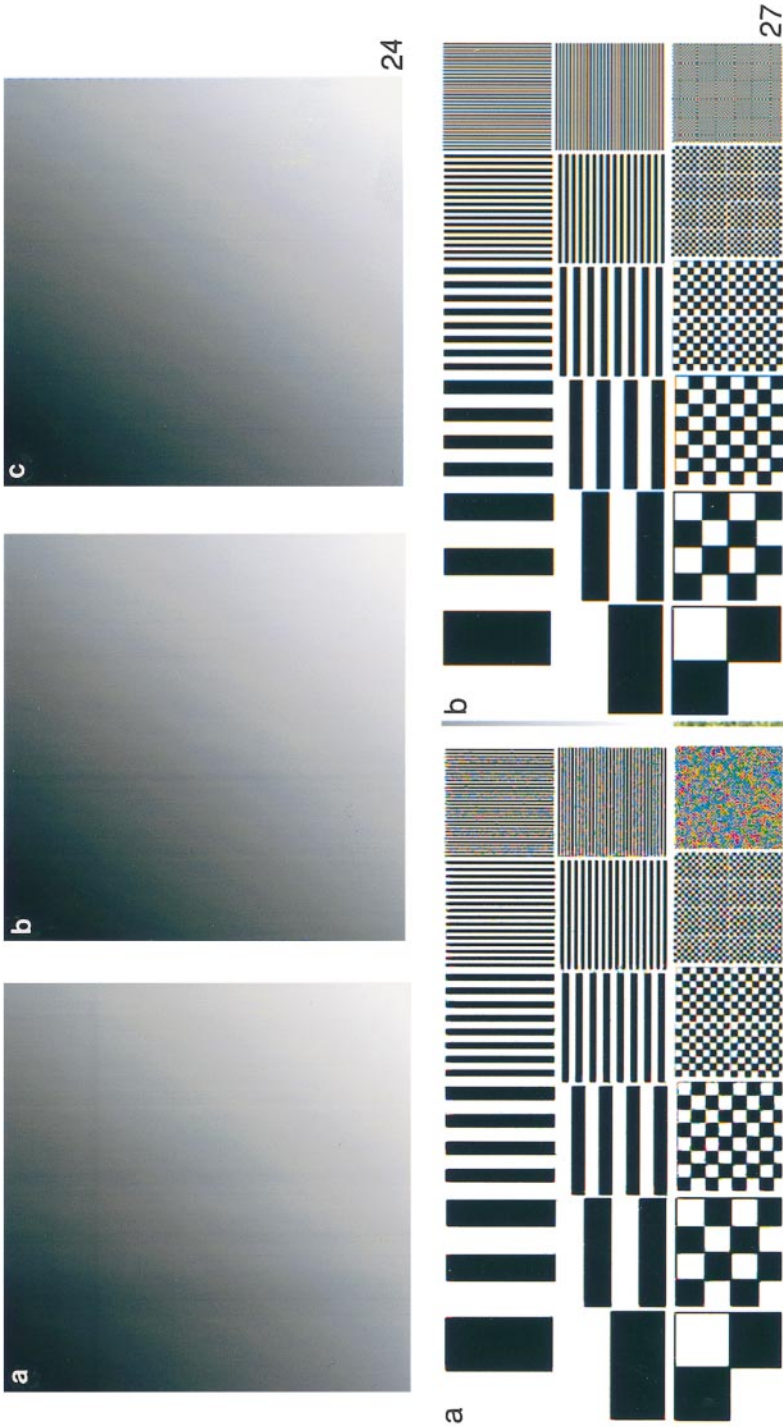



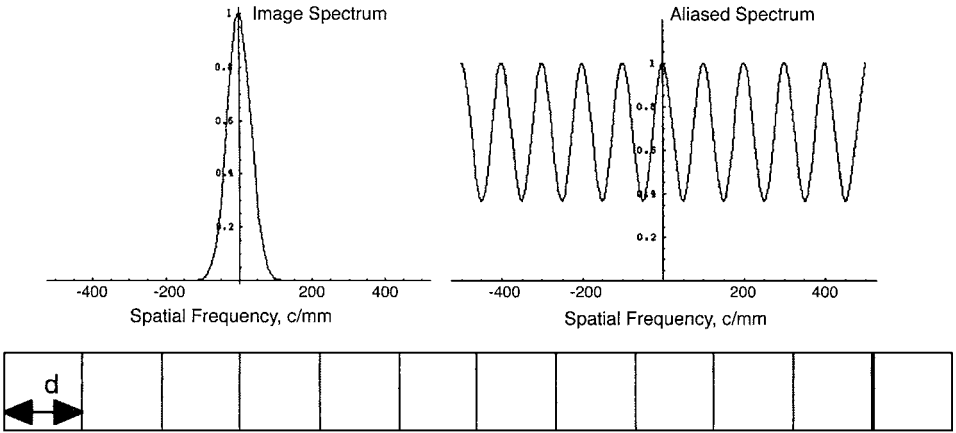
FIG. 24. Gray ramp test: (a) original, (b) reconstructed from BNM CFA, (c) reconstructed from Bayer CFA.

FIG. 27. Reconstructed test images: (a) BNM CFA, (b) Bayer CFA.




**FIG. 22.** Reconstructed images: (a) equal numbers of red, green, and blue sites, (b) twice as many green sites as red or blue sites.


**FIG. 28.** Reconstructed test images: (a), (c) BNM CFA; (b), (d) Bayer CFA.




**FIG. 4.** A linear array,  $d = 0.01$  mm samples a signal with a Gaussian spectrum. The resulting aliased spectrum shows the increased low-frequency signals due to the insufficient sampling rate.



**FIG. 5.** A continuous sine wave signal and its spectrum.



**FIG. 6.** The sine wave signal in Fig. 5 but sampled 16 times per millimeter and its spectrum. The aliased frequency is 6 cycles per millimeter compared to the actual value of 10 cycles per millimeter. The Nyquist frequency is 8 cycles per millimeter.



**FIG. 7.** A linear CCD imaging array.






FIG. 8. The potential for aliasing, PA, for linear CCD imaging arrays as a function of the ratio  $\alpha = d/p$ .




FIG. 9. The potential for aliasing is reduced when the input spectrum is not flat. Figure 10 shows the spectrum assumed for the plots above.

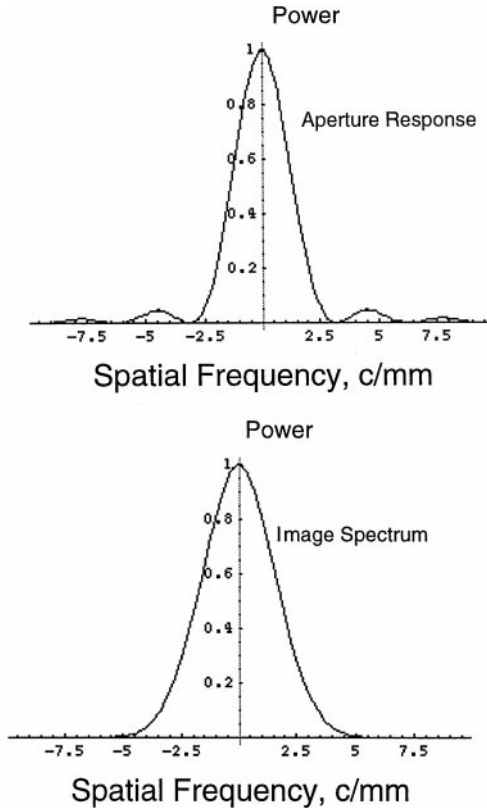


FIG. 10. The power of the aperture response and the image spectrum used to calculate the potential for aliasing in Fig. 9.


shows the potential for aliasing when a nonflat spectrum is assumed. Figure 10 shows the power of the Gaussian spectrum and the power associated with the square pixel aperture. If the frequency range of the image spectrum is further reduced, the potential for aliasing will get correspondingly smaller.

So far, it has been a one-dimensional potential for aliasing analysis. The equivalent two-dimensional analysis will predict even greater potential for aliasing as a function of  $\alpha$ . Furthermore, when a CFA is used to encode the color the values for  $\alpha$  become smaller, thus introducing a greater potential for aliasing.


The above calculations do not take into consideration the impact of system resolution and magnification. As the number of elements (pixels) in an imaging array increases, the aliased information is reproduced at higher frequencies (in the print). If the aliased information lies beyond the eye's effective low-pass filter range, no color artifacts will be seen [10, 11, 14]. If the image is enlarged, the aliased signals (and accompanying color artifacts) will be observed. The system response to aliased information is complex and is addressed in the literature [10, 11].

#### 4. STOCHASTIC PATTERNS

The work in stochastic digital halftones provides the key to a solution to the problem of aliasing outlined above [15–18]. Conventional halftone patterns have very regular arrays. When these patterns are scanned (or imaged with a CCD-based camera) the regular pattern gives rise to aliased images. The simplest way to see this effect is to consider an image of a uniform gray patch that is rendered by a conventional digital halftone. The patch on the left in Fig. 11 is such a rendering for a 53% gray level. If this patch is imaged by a CCD array (assume a frame transfer device) that has lower resolution than the basic structure of the halftone pattern, the resulting image will contain a pattern of low-frequency patterns due to the aliasing. The image on the right is a simulation of when the pattern on the left is imaged with a CCD sensor that has pixels three times as large in the linear dimension as the basic pixel structure of the halftone.



**FIG. 11.** On the left is a digital halftone rendering of a 53% gray level. The image on the right is a simulation of capturing the image on the left with a frame transfer type CCD camera with pixels that are three times the linear dimension of the basic halftone structure on the left.

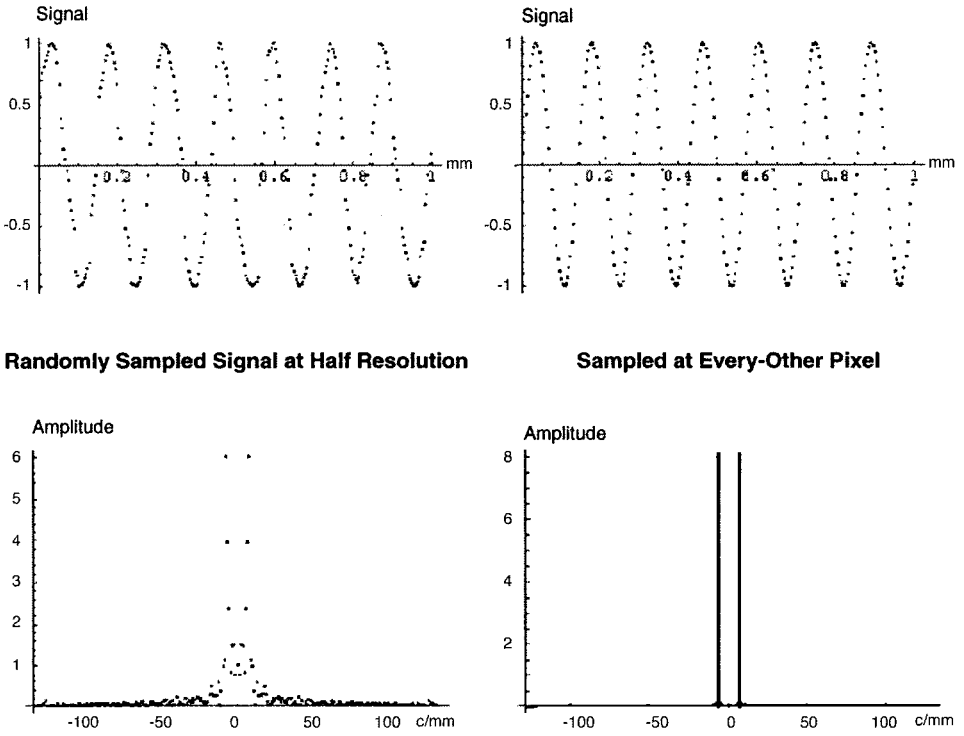


**FIG. 12.** The image on the left is a rendering of a 53% gray level with a stochastic digital halftone. The image on the right is the output of a CCD camera with pixels three times the linear dimension of the basic digital halftone structure. Compare these results with those in Fig. 11.

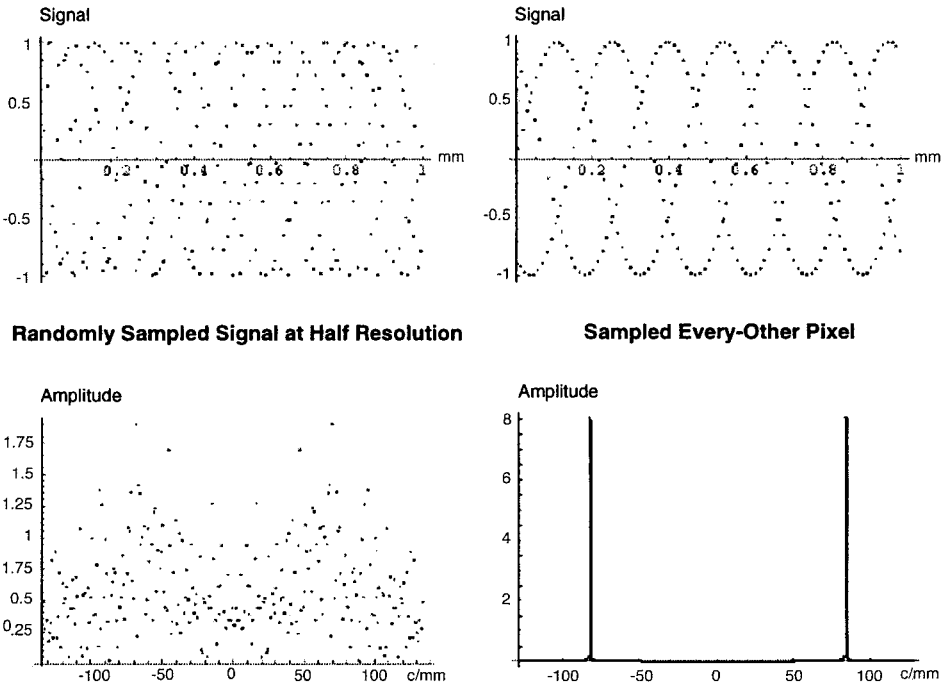
When one renders the 53% gray level by means of any one of a number of stochastic patterns and then images the results with the same CCD camera, the results are shown (as a simulation) in Fig. 12. While the CCD camera image is very different from the original, it does not have the periodic, low-frequency patterns that dominate the image in Fig. 11.

These results lead one to assume that the reverse should also be true. That is, if one were to capture an image with a random array, the resulting image would not show the aliasing of regular, high-frequency patterns in the original image. Consider the following one-dimensional example; refer to Figs. 13 and 14. A signal (image) array of 512 elements is constructed. The actual signal is a sine wave of some specified frequency. Two samplings of this array are made, one with a regular pattern of every other element (256 samples in all) and a random sampling along the array (with an average of 256 samples over many trials). The Fourier transform of the two sampled signals (without any interpolation) are then taken. In the case of Fig. 13 the signal frequency is low enough that no aliasing takes place (within the 512 possible samples). Note that the Fourier transform of both the regular and random samples produce a spectrum that clearly indicates a dominant frequency. The randomly sampled signal does show a buildup of some noise around the prime frequency and the peaks have less amplitude than that of the regularly sampled spectrum. Figure 14 shows the case for a highly aliased signal. The plot of the regularly sampled signal shows some sense of the aliased sine wave, while the randomly sampled signal does not. The Fourier transform of these two signals clearly indicates a strong aliased peak for the regularly sampled signal, while that of the randomly sampled signal shows a spectrum more typical of a noise pattern. Notice that the amplitude associated with the randomly sampled spectrum is considerably lower than that of the regularly sampled spectrum. Thus, the effect of the random sampling is to eliminate the aliased signal and replace it with a noise-like spectrum of lower amplitude.


The same concept can be extended to two-dimensional arrays. Figure 15 shows two such random patterns. The one on the left has 50% of the positions available for sampling while the one on the right has 65% of the sites available for sampling. Once the image has been formed through either of these stochastic sampling arrays, the missing points will have to be filled in by means of a sophisticated interpolation algorithm.



**FIG. 13.** Signal and spectrum of a sine wave sampled in a random and regular manner. The input signal is not aliased by the 512 element, linear imaging array.



**FIG. 14.** Signal and spectrum of a sine wave sampled in a random and regular manner. The input signal is highly aliased by the 512 element, linear imaging array.



**FIG. 15.** Two stochastic arrays to be used to capture an image. The one on the left has 50% of the imaging sites available while the one on the right has 65% of the imaging sites available.

In the above configurations the imaging sites that are “turned off” can be used to encode a different color. Thus, a spatial multiplexed, random color imaging array can be used to replace the more conventional CFA used in today’s digital still cameras. Figure 16 shows three examples of such arrays.

These examples were generated by using a simple random number generator. The one on the left has roughly equal red, green, and blue imaging elements, while the one in the center has about 50% green imaging elements and 25% each red and blue imaging elements. The one on the right is a subtractive random CFA with equal numbers of cyan, magenta, and yellow imaging sites. Subtractive CFAs have the advantage of passing more light than red, green, and blue CFAs, but the required matrixing back to red, green, and blue signals introduces increased noise and color artifacts [19].

The method used to determine which stochastic CFA is best, the choice of interpolation algorithm, and a method to correlate observed quality with a quality metric are the subjects of the rest of this paper.

## 5. METHODS

As can be seen from the above discussions, the use of randomized patterns as a color filter array for a single-sensor color digital still camera eliminates the low-frequency banding (aliasing) artifacts associated with ordered color filter arrays and seems to be a plausible approach in our efforts in minimizing aliasing. The next step is to determine what kind of random pattern is best. An important factor in the decision making is that the filters generated from the random pattern for each color plane should sample the entire image space as regularly as possible in order to avoid losing too much information in any given area of the image.

### 5.1. Color Filter Array Generation

One obvious solution to the above question is to generate a white noise pattern and threshold it into three different color filter arrays. (See Fig. 16.) For example, if the random pattern has values ranging from 0 to 1, then all the pixels with value 0 to 0.33 can be the red sensors, all the pixels with value 0.33 to 0.66 can be the green sensors, and the rest can be

blue sensors. One drawback of using white noise patterns is that it has uneven clumps which result in very sparse sampling of the image in certain places of the image, and very dense sampling in others. This greatly degrades the quality of the final reconstructed image since in certain regions of the image not enough information is present for a reasonable estimate of the original image value. The details of the reconstruction process will be discussed shortly. See Fig. 17 for a comparison of a sample reconstructed image using a white noise pattern as color sensors and using blue noise patterns as color sensors. The blue noise CFA resulted in a superior reconstructed image, and the clumping is obvious in the reconstructed image using white noise as the basis for a CFA.

In addressing this clumping problem, we propose the use of blue noise masks (BNM) as the random pattern to be used to generate the color filters [17, 18, 20, 21]. Blue noise masks are dither patterns generally used in image halftoning and are designed to disperse similar values evenly throughout the pattern, therefore greatly reducing the clumping phenomenon associated with white noise patterns. This results in superior reconstructed images [15, Chap. 8]. To generate the three color filter arrays, one blue noise mask is multiply thresholded (with two thresholds) to obtain three mutually exclusive patterns. These patterns can then be used as the three-color filter arrays. A blue noise mask typically has values between 0 and 255; thus, all the pixel locations with BNM values between  $0-T_1$  will be red sensors, all the locations with values between  $T_1-T_2$  will be green sensors, and all the locations with values between  $T_2-255$  will be blue sensors. The thresholds can be varied to adjust the number of sensors for each color filter. See Fig. 18.  $T_1$  and  $T_2$  in Fig. 18b are selected to assign the same number of red, green, and blue color filters, while  $T_1$  and  $T_2$  in Fig. 18b are selected to favor green color filters. After passing the input image through the three color filter arrays, three separate images representing the red, green, and blue color planes are obtained. In each of the color planes, there are locations with known pixel values, corresponding to the locations where a sensor for that particular color is present. The other pixel locations are initially considered to have unknown values. It is the task of the reconstruction stage to fill in these empty locations and reproduce the input image.

## 5.2. Reconstruction from CFA Sampled Images

The advantage to using an ordered filter array is that the reconstruction phase is relatively simple. Prior knowledge about the locations of each color sensor and the periodicity in these locations facilitates the interpolation process [3, 4]. However, in the case of blue noise mask filter arrays, the random nature of the color filters makes accurate interpolation a more challenging task. A simple interpolation between the nearest available pixels tends to blur the sharp edges in the image and makes the reconstructed image unacceptably “soft.” Therefore, an adaptive interpolation scheme must be used that takes into consideration the edge information in the input image. We propose a reconstruction scheme which first performs a missing-pixel edge detection to obtain edge information about the input image, and then uses an adaptive interpolation process to reconstruct the original image. (See Fig. 19 for the flowchart of the algorithm.)

### 5.2.1. Missing-pixel edge detection.

Traditionally, the Sobel operators are one method of edge detection on images where all pixel values are known. See Fig. 20a. However, in this case, where there are missing pixels at random locations, a  $3 \times 3$  Sobel operator might not be able to capture any edge information due to the small filter size. Generally speaking, if the



FIG. 19. Flowchart of reconstruction process.

above-mentioned thresholds  $T1$  and  $T2$  are selected to give the same number of sensors for each color filter, then approximately  $1/3$  of the pixels in each color plane must have known pixel values. With a  $3 \times 3$  Sobel filter, the probability that none of the three pixels on one side of the filter has a known pixel value is about 0.3 for the uniform case and about 0.125 for the case with 50% green sites. This will result in very inaccurate edge detection due to large amounts of missing data. To accommodate such problems, we extend the Sobel operator in an effort to reach more pixels with known pixel values. For vertical edge detection we use a  $3 \times 5$  or  $3 \times 7$  filter, and for horizontal edge detection we use a  $5 \times 3$  or  $7 \times 3$  filter, as shown in Figs. 20b, c. By having longer sides, we greatly increase the probability (to as high as 0.99) that at least one of the pixels on the long side will have a known pixel value. Since the known pixel locations are sparse, most likely, different numbers of pixels will contribute on either side of the filter. In order to compensate for such asymmetries, the sums from the positive and the negative sides are normalized by the number of contributing pixels on either side. Spurious edges due to uneven distribution of known pixels on both sides are thus avoided by this normalization process. The missing-pixel edge detection step again indicates the clear advantage of using blue noise patterns versus white noise patterns as color filters. If there is considerable clumping in the random color CFA, resulting in an extended region where there are no known pixel values, the edge detection will not yield




FIG. 20. Edge detection filters: (a) traditional Sobel operator, (b)  $3 \times 5$  and  $5 \times 3$  edge detection filter, (c)  $3 \times 7$  and  $7 \times 3$  edge detection filter.

satisfactory results in that region. Since the BNM does not contain such clumps, it generally assures the proper performance of our edge detection algorithm.

Using the above-mentioned edge detection, we obtain separate horizontal and vertical edge maps of the input image for all three color planes. Since we are interested in just knowing whether or not there is an edge at a certain pixel, the edge map can be reduced to binary form without losing its purpose. Thus, we threshold the edge map obtained from the previous step to obtain our binary edge map. The threshold level should be selected so as to avoid as many false edges as possible, while at the same time including all real edges.

The thresholding process eliminates many of the false edges. However, since the edge detection is performed on only a fraction of the total number of pixels in the input image, there still will be some spurious false edges retained from thresholding. These spurious edges tend to be small isolated structures. Because our final reconstruction step depends heavily on the accuracy of the edge detection, we have devised a method to remove these spurious edges. The method that we employ in this case is morphological opening. Morphological opening will remove all structures smaller than a specified size. A larger size will cause more real edges to be removed, while a smaller size will not eliminate all false edges.

*5.2.2. Edge-adaptive reconstruction.* After obtaining the finalized edge maps for both horizontal and vertical directions for each of the three color planes, the final reconstruction step can be performed. The reconstruction process is identical for all three color channels, and it can be summarized as follows:

For all pixel locations with missing data:

If the pixel location is on both a horizontal and a vertical edge,

Find the nearest known pixel in a crosshair manner and replicate that pixel's value.

Else, if the pixel location is on a horizontal edge only,

Find the nearest known pixel in the horizontal direction and replicate that pixel's value.


Else, if the pixel location is on a vertical edge only,

Find the nearest known pixel in the vertical direction and replicate that pixel's value.

Else, if the pixel location is not on any edges

Find all the known pixels in a neighborhood making sure that no edge pixels are included, and their average value is taken as the value for the missing pixel location.

See Fig. 21. The darkened pixel represents the unknown pixel to be reconstructed, and the numbers indicate the order in which the algorithm searches for a pixel to replicate to the missing pixel. As can be seen in the above algorithm, the reconstruction step depends heavily on the accuracy of the edge detection.



**FIG. 21.** Pixel replication scheme for different edge situations: (a) missing pixel is on both a horizontal and a vertical edge, (b) missing pixel is on vertical edge only, (c) missing pixel is on horizontal edge only.



## 6. EXPERIMENTS AND RESULTS

The above-mentioned techniques were tested on a variety of test images.

### 6.1. Experiments to Determine Thresholds

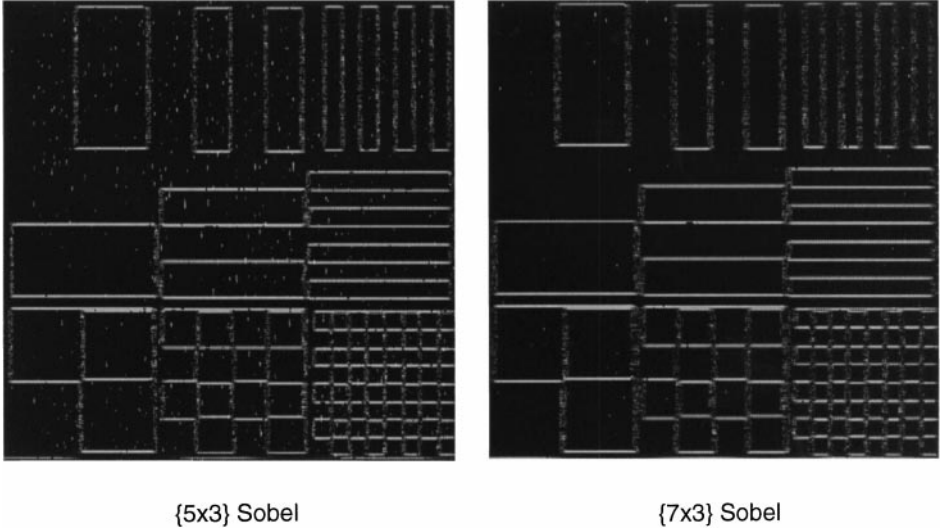
The first few experiments that were performed determined the number of color filter elements to be assigned to each color plane, the size of the edge detection window, the threshold level for the edge map, and the size of the morphological opening structure.

It was mentioned in the development of the Bayer CFA pattern that the eye is most sensitive to details in the green channel. That was the reason behind assigning 50% of the available color filters to the green channel in the Bayer pattern. We also took this approach and assigned 50% of the color filters to the green channel and 25% for each of the other two channels. It is obvious that for our missing-pixel edge detection, the more known pixels there are, the better the performance. Therefore, one great side effect to this assignment was that we obtained a highly reliable edge map from the green channel. If we assume that the edge map from the green channel is representative of all three color channels, then we can use only the edge detection results from the green channel in the reconstruction of all three color channels. Not only does this approach decrease the computation load, but it also offered more reliable edge maps which lead to better results in the reconstruction process. Figure 22 shows a comparison between using 50% green color filter, 25% red and 25% blue color filters, and using 33% for all three-color filters. An enlarged section of the final reconstructed image from both methods are shown. As can be seen, the reconstructed image using the 50%, 25%, 25% combination gave significantly better results.

As mentioned before, the missing-pixel edge detection needs larger windows in order to find accurate edges. Some of the window sizes that we tried are  $3 \times 5$  and  $3 \times 7$  for vertical edges and  $5 \times 3$  and  $7 \times 3$  for horizontal edges. Results show that even though the larger window gives slightly more accurate edge maps, the difference is visually minimal. This again demonstrates that the blue noise random color filters are evenly distributed so that, even in the smaller window size case, enough known pixels are found to obtain reasonable edge map results. See Fig. 23 for edge maps found by the  $5 \times 3$  window and the  $7 \times 3$  window.

Next, the edge map is converted from an 8-bit image to a binary one by thresholding the edge map in order to eliminate a majority of false edges. The threshold level tested ranged from 30 to 90 in 20 increments. A lower threshold will allow more random structures to be edges, while a higher threshold might eliminate some true edges. We found that a threshold level of 50 reaches a good compromise between false and missed real edges. Even though higher threshold levels, such as 70, also produced good results, we chose the level to be 50 because the lower threshold will preserve more real edges, and the additional morphological opening step can remove some more false edges not eliminated by the lower thresholding level. However, if we further lower the threshold to 30, enough false edges are retained such that even the opening process will not be able to remove them, causing inferior reconstruction results.

The morphological opening structure size also needs to be determined. Again there is a trade-off in the size selection process. A larger structure will remove more false edges, but at the cost of removing true, fine edges, and vice versa for smaller opening structures. Since edges are usually only one pixel wide, the opening structure only needs to be one-dimensional, but the length needs to be decided. Several different lengths were tested, but we found that a structuring element of length 10 gives the best performance.



**FIG. 23.** Edge maps obtained from missing-pixel edge detection using two different Sobel operators, each designed to find horizontal edges.


## 6.2. Experiments Using Metrics

Once the above-mentioned parameters were experimentally determined, we compared the performance of the blue noise color array filter pattern to the Bayer CFA pattern using two different metrics. In the following experiments, unless it is specifically mentioned, the edge detection windows are  $7 \times 3$  and  $3 \times 7$  for horizontal and vertical edge detection, respectively; the threshold used is 50, and the opening structure is  $1 \times 10$  and  $10 \times 1$  for horizontal and vertical edges, respectively.

The first metric we used is the deviation from neutral [4]. It was found that in black and white (neutral) regions, spurious color artifacts are more objectionable to the human eye than artifacts that are also neutral in color. Therefore, the deviation-from-neutral metric can be used in monochrome regions to measure the performance of the color filter array. Deviation from neutral can be determined by transforming the RGB values for each pixel into a trilinear plot using the formulas:

$$\begin{aligned} x &= B \cos(30^\circ) - G \cos(30^\circ) \\ y &= R - B \sin(30^\circ) - G \sin(30^\circ) \\ r &= \sqrt{x^2 + y^2}. \end{aligned}$$

The first experimented image is a gray ramp. See Fig. 24. The color distribution map in the form of a trilinear plot for both the Bayer CFA pattern and blue noise random CFA pattern is shown in Fig. 25. In these plots, all neutral pixel values should fall on the origin of the plot. Therefore, the further the point wanders from the origin, the more colored the pixel values are. Furthermore, the average value for  $r$  over all pixels can be calculated. A large value for the average  $r$  signifies that more pixels did not maintain their neutrality during the filtering and reconstruction process. The gray ramp can be considered as a “litmus” test for color filters; if the reconstructed image can maintain the overall neutrality of the gray ramp, then the color filter pattern can be considered for further testing. As can be seen



**FIG. 25.** Histograms of the trilinear plots of the data obtained from the gray ramps in Fig. 24. The neutral peak in the center of each plot represents the vast majority of the data while the nonneutral points represent (for the most part) the pixels on the edges of the gray ramp. The deviations from neutral for the blue noise CFA is more uniformly (randomly) distributed about the neutral point, while that of the Bayer CFA represents the more constrained colored edges produced by the Bayer CFA and subsequent interpolation.

from Fig. 24, both the Bayer CFA pattern and the blue noise random CFA pattern preserved the overall neutral appearance of the gray ramp. But in terms of our metric, the blue noise pattern performed better, mainly due to the fact that at the borders of the gray ramp the blue noise random pattern preserved the sharpness of the edges, while the Bayer pattern caused color banding.

The second set of experimented images are horizontal stripes, vertical stripes, and checkerboards of increasing frequency. The results in terms of average  $r$  values versus frequency are shown in the Tables 1, 2, and 3 and also in Fig. 26. Three reconstructed images are compared: a reconstructed image using the Bayer CFA pattern, a reconstructed image using the BNM CFA and a  $5 \times 3$ ,  $3 \times 5$  edge detection window, and a reconstructed image using the BNM CFA and a  $7 \times 3$ ,  $3 \times 7$  edge detection window. As can be seen, in both the horizontal and vertical stripes case, the blue noise pattern performed much better. This is because the Bayer pattern will always cause low-frequency color banding in periodic structures, while the random nature of the blue noise pattern, in combination with edge-sensitive reconstruction, alleviates this problem. In the checkerboard patterns, the Bayer pattern fairs much better at high frequencies because the missing-pixel edge detection collapses at such high frequencies so that all pixels are determined to be edge pixels, and the reconstruction becomes highly inaccurate.

It is interesting to examine the reconstructed images of the highest frequency patch for all three patterns. See Fig. 27. The reconstructed images from using both the Bayer CFA pattern and the blue noise random CFA pattern show considerable artifacts; however, the nature of the artifacts are very different. In the Bayer case, the blue and orange color banding becomes

**TABLE 1**  
**Average  $r$  for Vertical Stripe Test Target**

	1	2	4	8	16	32
Bayer	4.977	6.607	11.671	21.461	42.721	84.453
BNM5	0.768	0.915	2.012	3.638	6.657	28.413
BNM7	0.679	0.809	1.657	3.116	5.532	25.916

*Note.* The heading for each column represents the number of line pairs per unit length in the test target.

**TABLE 2**  
Average  $r$  for Horizontal Stripe Test Target

	1	2	4	8	16	32
Bayer	4.581	7.449	12.042	32.516	41.485	83.952
BNM5	0.416	1.036	2.265	4.033	6.022	30.407
BNM7	0.331	0.920	2.222	3.637	5.388	29.162

highly visible in the horizontal and vertical stripe patterns, while in the blue noise case, the artifacts are random color bursts. It is difficult to argue which one is more preferable. In the checkerboard pattern, the blue noise pattern caused a totally random mosaic of colors in the resulting image, while the Bayer pattern caused a mostly neutral image, but with a totally different pattern from the checkerboard. Without systematic psychophysical testing, it is not possible to say which sampling and interpolation artifacts are more distracting.

The second metric we used to compare CFA performance is the human visual system (HVS) weighted mean squared error (MSE). This metric is used because it relates the quality of the appearance of test images to the human eye. HVS-MSE is calculated by first creating the eye function given by the formula shown below [14]. Then, the original input image, the image processed by Bayer CFA, and the image processed by the blue noise CFA are individually convolved with the eye function. Finally, standard MSEs are taken between the original and the two processed images,

$$\text{Eye}(f) = A \left[ \alpha + \left( \frac{f}{f_0} \right) \right] \exp \left[ - \left( \frac{f}{f_0} \right)^\beta \right],$$

where  $A$ ,  $\alpha$ ,  $\beta$ , and  $f_0$  are constants. For this experiment,  $A = 2.2$ ,  $\alpha = 0.192$ ,  $\beta = 1.1$ ,  $f_0 = (0.114)^{-1}$ , and  $f$  is scaled to represent a normal viewing distance of 10 inches.

See Fig. 28 for reconstructed images using Bayer CFA and blue noise CFA. In terms of the metric, in one case the Bayer pattern performed better, and in the other, the blue noise pattern performed better. See Tables 4 and 5. However, from these two test images, it is difficult to say which CFA sampling and reconstruction is preferred. For example, in the church test image, the sharpness in the Bayer pattern is compromised by the underlying color banding in the bricks.

## 7. DISCUSSION AND CONCLUSION

In order to eliminate the artifacts due to the phase-shifted, aliased signals introduced by the sparse sampling associated with traditional ordered color filter arrays, we investigated the possibility of using random patterns instead of traditional ordered patterns as a CFA for single-sensor, digital still cameras. It has been shown that by using color filter arrays

**TABLE 3**  
Average  $r$  for Checkerboard Test Target

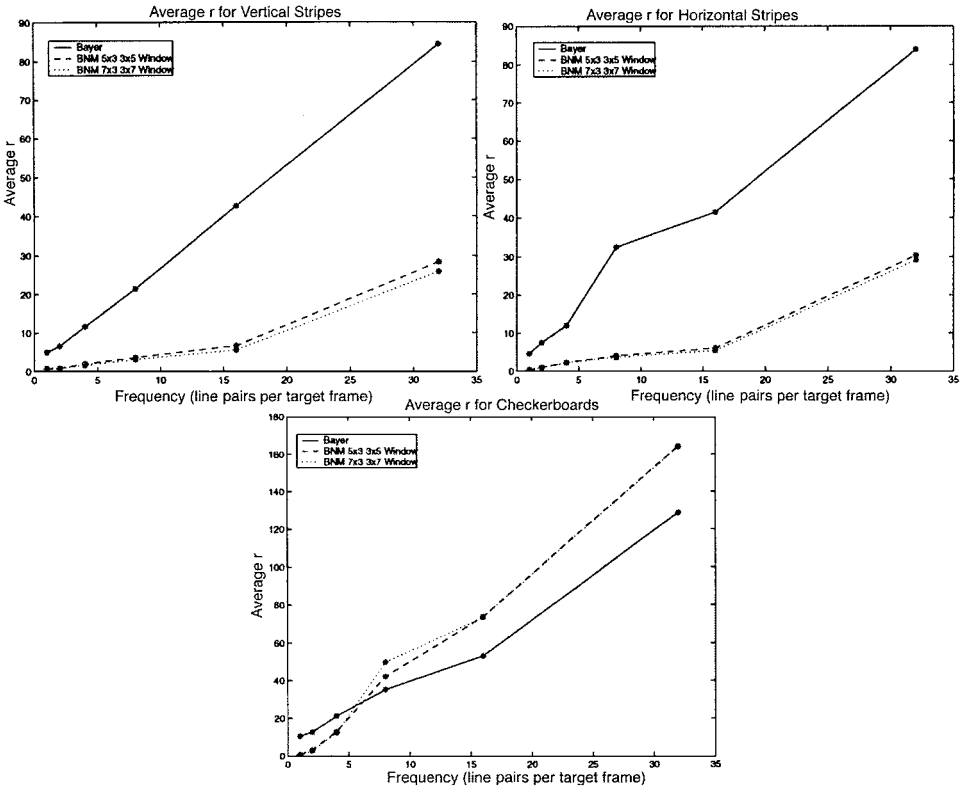
	1	2	4	8	16	32
Bayer	10.435	12.589	21.166	35.227	53.116	128.754
BNM5	0.693	3.184	12.874	42.165	73.721	164.216
BNM7	0.450	2.759	12.289	49.782	73.489	163.855

**TABLE 4**  
**Table of HVS-MSE for Different Color**  
**Filter Arrays on a Lady Test Image**

	HVS-MSE
Bayer	40.293
BNM7	35.342

**TABLE 5**  
**Table of HVS-MSE for Different Color**  
**Filter Arrays on a Church Test Image**

	HVS-MSE
Bayer	28.907
BNM7	40.799



**FIG. 26.** Plot of average  $r$  versus frequency for different test patterns: (a) vertical stripes, (b) horizontal stripes, (c) checkerboards.

constructed from mutually exclusive blue noise patterns, we were able to eliminate the color banding artifacts associated with ordered filter arrays on images containing high-frequency periodic information. However, the reconstruction process is heavily dependent on the accuracy of the edge detection. Since the edge detection is performed on incomplete pixel information, and the parameters associated with the thresholding and morphological opening are decided heuristically, most of the artifacts occur at the boundaries inside the image. If there is an image with highly irregular edge characteristics, then these parameters might have to be adjusted to obtain satisfactory results. One possible future research area is to devise a method which automatically adjusts these parameters, based on the statistics of the input image, and chooses the most suitable ones for each input image. Additionally, the edge detection and interpolation algorithms could be improved by examining data from the color planes jointly, as opposed to separately. However, this would also add to the complexity of the reconstruction.


### ACKNOWLEDGMENTS

The authors thank Dr. Kevin Spaulding and Mr. Rodney Miller of the Eastman Kodak Company for their helpful insight into the development of efficient edge detectors.

### REFERENCES

1. M. A. Kriss, Electronic imaging: The challenge, the promise, *J. Soc. Photogr. Sci. Technol.* **50** (5), 1987, 357–378.
2. M. A. Kriss, K. Parulski, and D. Lewis, Critical technologies for electronic still imaging systems, in *Proceedings of SPIE, Vol. 1082, Applications of Electronic Imaging, 1989*, pp. 157–184.
3. M. A. Kriss, Image quality requirements for a multimedia environment, *J. Soc. Photogr. Sci. Technol.* **59** (1), 1996, 186–211.
4. M. A. Kriss, Color filter arrays for digital electronic still cameras, in *Proceedings of the IS&T 49th Annual Conference, May 19–24, 1996, Minneapolis*, pp. 272–278.
5. E. Shimizu *et al.*, The digital camera using a new compression and interpolation algorithm, in *Proceedings of the IS&T 49th Annual Conference, May 19–24, 1996, Minneapolis*, pp. 268–272.
6. T. Tanaka *et al.*, A new interpolation method of single-chip CCD color camera image, in *Proceedings, 5th International Conference on High Technology, WORLD TECHNO FAIR IN CHIBA'96, Sept. 11–14, 1996*, pp. 370–376.
7. A. V. Oppenheim and R. W. Schaffer, *Digital Signal Processing*, Prentice–Hall, Englewood Cliffs, NJ, 1975.
8. J. W. Goodman, *Introduction To Fourier Optics*, McGraw–Hill, New York, 1968.
9. R. C. Gonzales and R. E. Woods, *Digital Image Processing*, Addison–Wesley, New York, 1992.
10. M. A. Kriss, Image analysis of discrete and continuous systems: Film and CCD sensors, in *Proceedings of SPIE, Vol. 1398, CAN-AM Eastern '90, 1990*, pp. 4–14.
11. M. A. Kriss, Digital electronic still imaging: A system analysis and simulation, in *Proceedings of the IS&T 49th Annual Conference, May 19–24, 1996, Minneapolis*, pp. 317–321.
12. D. R. Cok, Reconstruction of CCD images using template matching, in *Proceedings, IS&T 47th Annual Conference/ICPS, 1994*, pp. 380–385.
13. J. E. Greivenkamp, Color dependent optical prefilters for the suppression of aliasing artifacts, *Appl. Opt.* **29**, 1990, 676–684.
14. A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice–Hall, Englewood Cliffs, NJ, 1989.
15. R. A. Ulichney, *Digital Halftoning*, MIT Press, Cambridge, MA, 1987.
16. R. W. Floyd and L. Steinberg, An adaptive algorithm for spatial grayscale, *Proc. SID* **17** (2), 1976, 75–77.
17. T. Mitsa and K. J. Parker, Digital halftoning technique using a blue-noise mask, *J. Opt. Soc. Am. A* **9**, 1992, 1920–1929.
18. R. Eschbach, (Ed.), *Recent Progress in Digital Halftoning, IS&T*, Springfield, VA, 1994.

19. K. A. Parulski, Color filters and processing alternatives for one-chip cameras, *IEEE Trans. Electron Devices* **ED-32**, 1985, 1381–1389.
  20. T. Mitsa and K. J. Parker, Digital halftoning using a blue noise mask, in *ICASSP 91: 1991 International Conference on Acoustics, Speech, and Signal Processing*, Vol. 2, pp. 2809–2812, IEEE, New York, May 1991.
  21. Q. Yu, K. J. Parker, and M. Yao, On filter techniques for generating blue noise mask, in *Proceedings, IS&T's 50th Annual Conference*, pp. 391–395, IS&T, Springfield, VA, May 1997.
- 



WEI (AMY) ZHU received both a B.S. and a M.S. in Electrical Engineering from the University of Rochester, Rochester, NY in 1996 and 1997, respectively. Currently, she is a research scientist at the Eastman Kodak Company. Her research interests are halftoning, digital cameras, and digital image and video libraries. She is also a member of Tau Beta Phi and Phi Beta Kappa.



KEVIN J. PARKER holds the B.S. degree in engineering science from SUNY at Buffalo. He received M.S. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology. He is professor and chair of the Department of Electrical Engineering at the University of Rochester, professor of radiology, and director of the Center for Biomedical Ultrasound at the University of Rochester. Dr. Parker is a Fellow of the Institute of Electrical and Electronics Engineers and a member of the Acoustical Society of America and the American Institute of Ultrasound in Medicine. Dr. Parker's research interests are in digital halftones, medical imaging, 3D visualization, and ultrasound imaging.



MICHAEL A. KRISS received his B.A., M.S. and Ph.D. in physics from the University of California at Los Angeles in 1962, 1964, and 1969, respectively. Dr. Kriss joined the Eastman Kodak Research Laboratories in 1969. During his 23 years at Kodak Dr. Kriss focused on all aspects of color imaging systems. He was the head of the Image Processing and Algorithm Development Laboratories prior to his retirement in 1992. Dr. Kriss joined the Center for Electronic Imaging Systems (CEIS) at the University of Rochester in 1994 where he is currently a senior scientist and associate director of the CEIS. Dr. Kriss is a member of the IEEE, IS&T, Sigma Xi, Sigma Pi Sigma, and the New York Academy of Science. Dr Kriss' research interests are in all aspects digital imaging systems.