Image Sampling with the Bayer Color Filter Array

Russ Palum
Eastman Kodak Company
Rochester, New York

Abstract

Digital color imaging systems sometimes use three separate sensors to record red, green, and blue scene information, but single-sensor systems are more common. Single-sensor color imaging systems have a color filter array (CFA) applied to a monochrome imager. One very popular filter array, the Bayer pattern, was invented at Eastman Kodak Company by Bryce Bayer in 1976 and has been in Kodak products ever since. When the patent rights expired the pattern became an industry standard. This CFA pattern has twice as many green filtered pixels as red or blue filtered pixels. This is done to better preserve image sharpness. The spatial configuration of the Bayer pattern is also tailored to an orientation property of human vision. In this tutorial talk the image sampling properties of the Bayer pattern are described for each color record. The advantages of the pattern will be discussed and some interpolation techniques will be described.

Introduction

A small section of a typical (CFA) is shown in figure 1. The colors, red, green, and blue are represented by the letters R, G, and B. The color filter patches are aligned with each pixel in the array. This pattern was invented at Eastman Kodak Company about 25 years ago and is called a Bayer pattern after the inventor (US pat. No. 3971065)[1]. Digital color images usually consist of arrays of numbers in three color planes: red, green and blue. The number of pixels in each color plane equals the total number of pixels of all colors in the CFA covered array. When a color filter array is applied to a sensor, information is only measured for one color at each pixel location, but a value is required for all three colors at each pixel location. The values for the other colors are interpolated from the surrounding pixels and placed in the appropriate color plane. The interpolation algorithm can be very sophisticated and can lead to resolutions very close to the resolution of the monochrome sensor. Most of the resolution as seen by the eye is due to luminance, which is dominated by the green channel. To produce high resolution, the Bayer CFA has twice as many green pixels in the pattern compared to red or blue pixels.

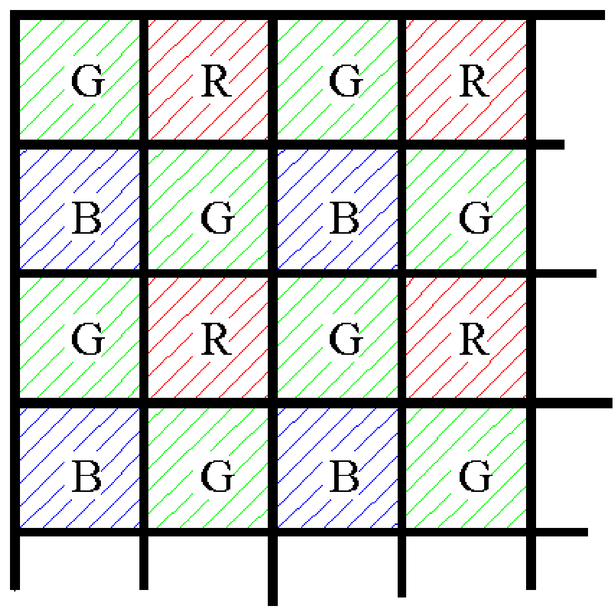

Figure 1. Bayer CFA pattern

The simplest way to interpolate the Bayer CFA is to divide the pattern into 2 by 2 pixel blocks containing a green pixel in the upper left and lower right, a red pixel in the upper right, and a blue pixel in the lower left. All four of the red plane pixels are set equal to the red value in the corresponding 2 by 2 block, all four of the green pixels are set equal to the average of the 2 green pixels in the block and all four of the blue pixels are set equal to the blue value in the 2 by 2 block. This technique does not have very good resolution and it produces colored edges on neutral objects. A representative image using this technique is shown in figure 2a and the same image interpolated using adaptive interpolation is shown in figure 2b. Two images of a section of a Macbeth chart are shown in figure 3. The top image uses the simple interpolation technique and the bottom image uses the adaptive technique. The simple technique produces yellow edges to the left and bottom of bright objects and cyan edges to the right and above bright objects.

Figure 2a. nearest neighbor interpolation

Figure 2b. adaptive interpolation

The example in figure 4 shows why this happens. The color of the edge is determined by the phase of the edge on the pattern. If the black half of the edge falls on the red/green column or row, the block of two columns of pixels corresponding to the edge will be cyan (more blue than green). This happens because the white half of the edge illuminates the green and blue pixel indicating the block of 4 pixels should be cyan. The black half of the edge covers the red and green pixel, indicating there is no red in the block of 4 pixels. If the black half of the edge falls on the blue/green column or row, the block of two columns of pixels corresponding to the edge will be yellow with more red than green (actually orange). For the case in figure 4, if the edge is moved one column to the right the edge will be properly black and white but if the black and white areas are reversed the edge will be yellow.

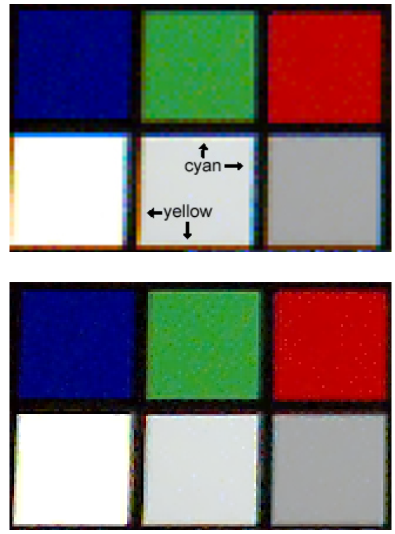

Figure 3. Macbeth chart

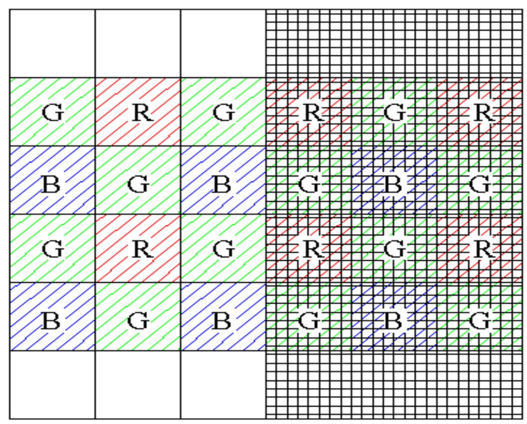

Figure 4a. black and white bar pattern

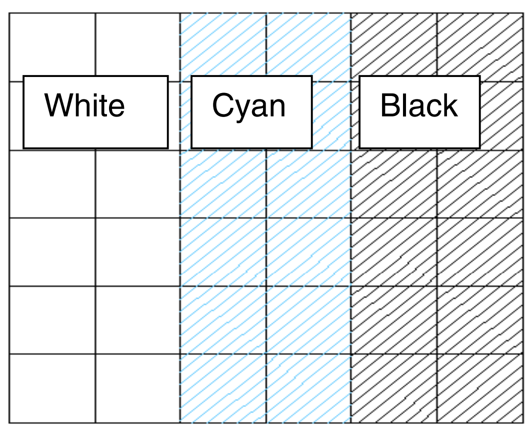

Figure 4b. image of black and white bar pattern

Adaptive Interpolation

To interpolate a Bayer image the green plane is adaptively interpolated first. The missing green pixels are interpolated in the direction the green values change least.[2] The red minus green plane and the blue minus green plane are interpolated next. The green plane is then added to the red minus green and the blue minus green to produce separate red, green, and blue planes. The resolution is increased with this technique because the resolution comes from the green plane, which has a higher sampling frequency than the red and blue plane. Colored edges are reduced because all three planes change when the green plane changes. It should be clear that for neutrals the R, G, and B values are equal so the R minus G plane and the B minus G plane will be filled with zeros for neutrals. In this case the level is clearly dependent only on the green plane.

Some of the Details

Figure 5 shows a section of Bayer pattern centered on a red pixel. The known green pixel values are copied to the green plane. The missing values (those corresponding to red and blue pixels) have to be interpolated from the known values. A green plane value corresponding to the center red Bayer pixel is missing and must be interpolated. The interpolation is done by first finding the smallest green gradient. The difference is taken between the east west greens and the north south greens. The green plane pixel value that corresponds to the red pixel location is the average of the two green pixels that have the smallest difference or gradient. The same procedure is performed for the green plane pixels that correspond to the blue pixel locations.

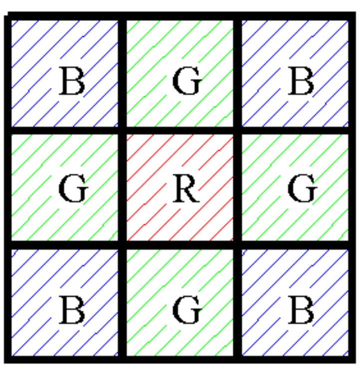

Figure 5. red pixel and surrounding pixels

Next, a sparsely populated red minus green and blue minus green plane are formed by subtracting the interpolated greens from the known red and blue pixels. The missing values are interpolated by simple averaging of 2 to 4 pixels depending on location in the plane. For example, 2 blue differences are averaged vertically for the blue plane value corresponding to the east and west greens in figure 5, and 4 blue difference values are averaged for the blue plane location corresponding to the red pixel. A blue minus green difference plane is shown in figure 6. Notice the low spatial frequency content. It is not necessary to adaptively interpolate this plane because the frequency content is so low. To display this plane the differences were divided by 2 and offset by 127 so they could be displayed as a standard 0 to 255 grayscale image. The actual values, for an 8 bit image, can be between –255 and 255 although the peak of the distribution of values is near zero. Adding the green plane to the difference planes produces the RGB image in figure 2b.

Figure 6. blue minus green plane

Spatial Frequency, Nyquist Frequency, and Sampling Frequency

With some reservation, the first row of black and white bars in figure 7 represents a spatial frequency. The black and white bars simplify the discussion but an actual spatial frequency varies sinusoidally from black to white. The grid or array under the bars represents a segment of an image sampling array like a CCD array. The bottom row of black and white bars represents the image of the top row of bars as reproduced by a digital imaging system.

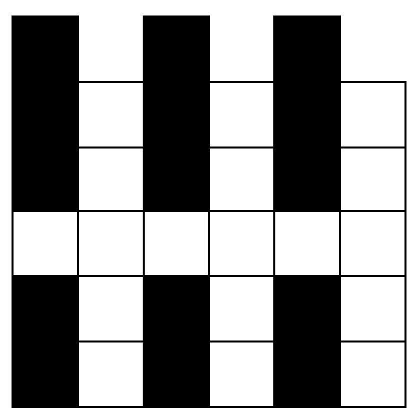

Figure 7. bar pattern at Nyquist frequency

Assuming the length of the array is 1 millimeter, the spatial frequency of the top row of bars is 3 cycles per millimeter. There are 6 pixels in the one millimeter length array, so the sampling frequency of the array is 6 samples per millimeter.

In order to reproduce a spatial frequency there must be a pair of pixels for each cycle. One pixel is required to respond to the black half cycle and one pixel is required to respond to the white half cycle. One pixel can never be used to represent more than a half cycle. Based on this requirement the highest frequency the array can reproduce is one half the array sampling frequency, this frequency is commonly called the Nyquist frequency. There is some contention over this name but in this document one half the sampling frequency will be called the Nyquist frequency. Determining the Nyquist frequency is not straightforward in a two dimensional array (area array).

An area array has a range of Nyquist frequencies even if the pixels are square and the space between the pixels is the same in both directions. When a CFA is applied to an area sensor the Nyquist frequency depends on color as well as direction.

A one dimensional array (linear array) has only one Nyquist frequency and one sampling frequency. The Nyquist frequency and the sampling frequency depend only on the pixel pitch, Xs. Two linear arrays are shown in figure 8, one has contiguous pixels (bottom) and the other is a sparse array (top). The pixels in a sparse array are not as wide as the pixel pitch. In both cases the space between pixel centers is the same. The sampling frequency is

$$\frac{1}{Xs}$$

in both cases and the Nyquist frequency is half of this or

$$\frac{1}{2 \bullet Xs}$$

in both cases.

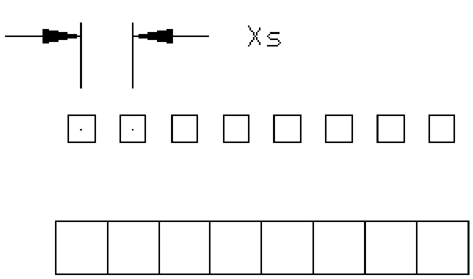

Figure 8. linear array

Consider figure 9, an area array. Starting with equal pixel pitch in X and Y, it is clear that the Nyquist and sampling frequency are the same on both the X and Y axis. The black bars are at the Nyquist frequency. Notice the sampling distance on the diagonal, Xs-diagonal, is much smaller so the sampling frequency is higher due to the inverse relationship between sampling frequency and sampling distance. Given a pixel spacing p, the sampling distance Xs is equal to p for the X and Y axis but Xs is

$$\frac{\sqrt{2}}{2} \bullet p \quad or \ (0.707 \bullet p)$$

on the diagonal axis. The sampling distance is reduced by this amount due to the geometry.

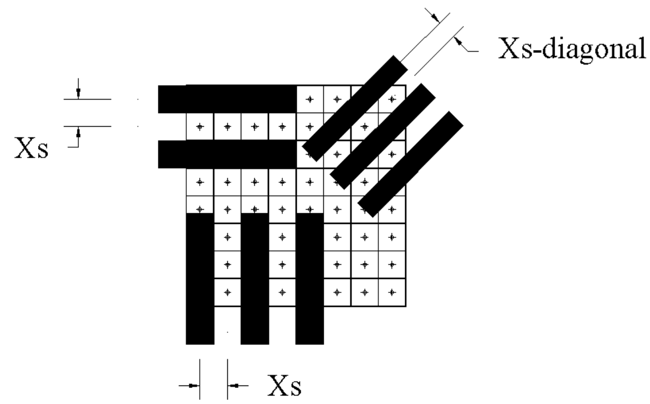

Figure 9. area array Nyquist frequency

Nyquist Domain Graph

A Nyquist domain graph can be drawn to show the Nyquist frequency in any direction. The length of a line drawn to a point on the graph from 0,0 indicates the Nyquist frequency at the angle of the line. Square pixel spacing produces a square Nyquist domain graph. A Nyquist domain graph for the array in figure 9 is shown in figure 10. The vertices of the square are on the diagonals and the sides fall on the X and Y axis. The graph is normalized to the sampling frequency

$$\left(\frac{1}{p}\right)$$

The X and Y axis Nyquist fall at

$$0.5 \bullet \frac{1}{p}$$

and the radial distance to the diagonal is greater by a factor of

$$\sqrt{2}$$

because the sampling pitch is smaller by this factor, remember the sampling pitch and the sampling frequency have an inverse relationship. The longer distance to the corner indicates a higher Nyquist frequency on the diagonal. A monochrome square array has better resolution on the diagonals than on the X and Y axis.

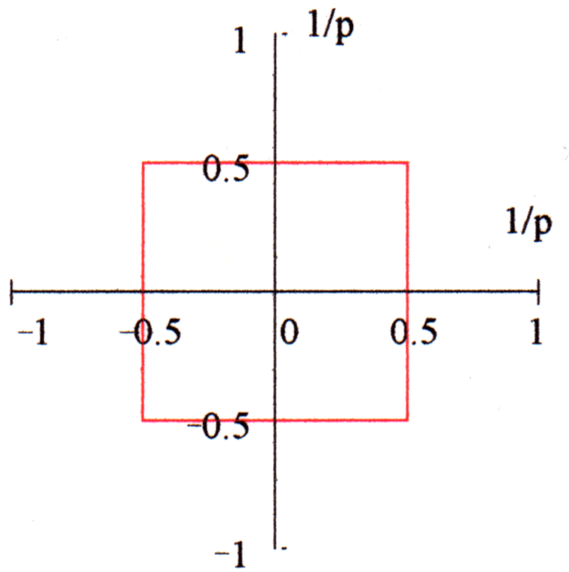

Figure 10. Nyquist domain graph, units of cycles/unit length. Normalized to the sampling frequency

Bayer Pattern

What happens when the Bayer pattern is applied to the array? First there are three different patterns, red, green and blue. The red and blue patterns are effectively the same because they have the same sampling frequency and Nyquist frequency. The green pattern is different because of its orientation and pixel pitch.

The following discussion will be based on the pixel pitch (p) of the monochrome sensor. The red and blue pixel patterns are shown in figure 11, notice that the pixel pitch and orientation are the same for the red and blue pattern, the only difference is a 1 column and 1 row offset. The pixel pitch in the red and blue is twice the monochrome pixel pitch on both the X and Y axes.

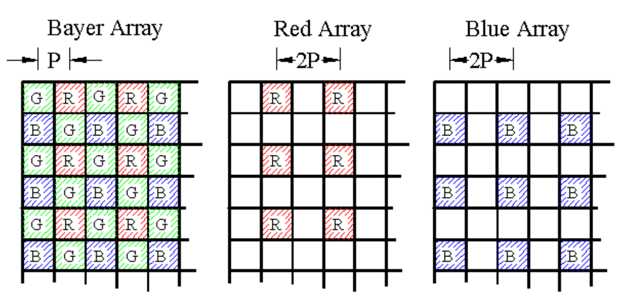

Figure 11. red and blue array

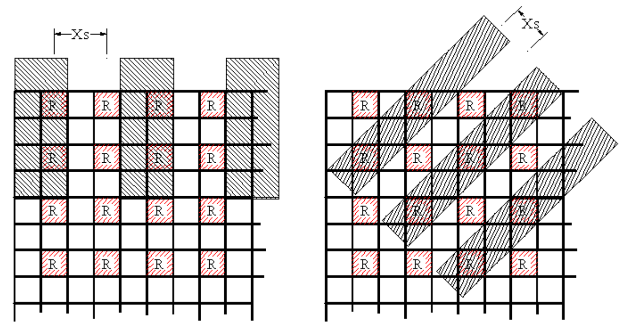

Figure 12. red array line pairs at the Nyquist frequency

Line patterns at the Nyquist frequency are shown in figure 12 for the red pixel array. Xs is equal to 2 • p in this case so the sampling frequency is

$$\frac{1}{2 \bullet p}$$

and the Nyquist frequency is one half of the sampling frequency or

$$\frac{1}{4 \bullet p}.$$

The diagonal spacing is

$$\frac{\sqrt{2}}{2}$$

times the X-Y pixel spacing so the diagonal spacing is narrower than the X-Y spacing by

$$\frac{\sqrt{2}}{2}$$

which makes the Nyquist frequency higher by

$$\frac{1}{\frac{\sqrt{2}}{2}}$$

which equals

$$\sqrt{2}.$$

The Nyquist domain for the red and blue arrays are nearly the same as the monochrome array Nyquist domain except the square is half as large on both axes as shown in figure 13.

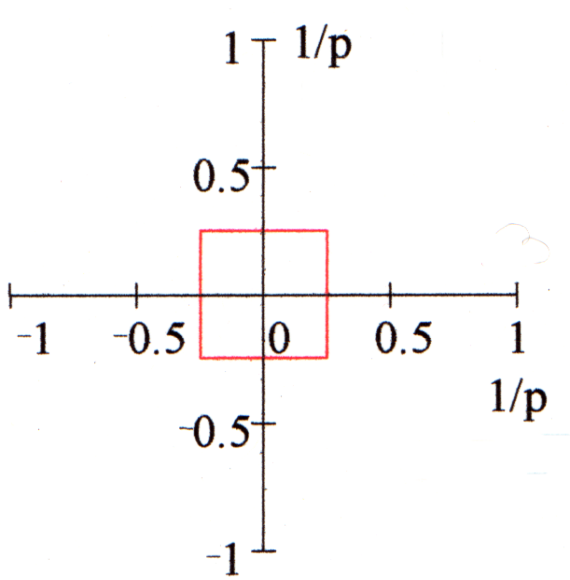

Figure 13. red and blue array Nyquist domain, units in cycles per unit length, normalized to the monochrome sampling frequency

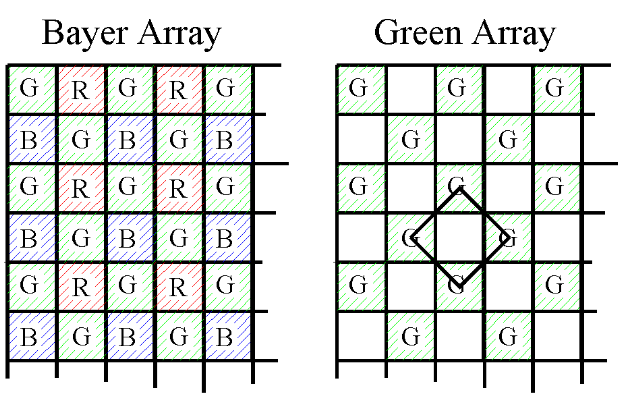

Figure 14. bayer green array

Green Array

The green array is rather interesting. It is a checker board pattern, but a checker board pattern is really a square pattern turned 45 degrees as shown in figure 14. The green pattern has a higher sampling frequency on the X and Y axes than on the diagonal axes because it is a square pattern turned 45 degrees. Line pairs at the sampling frequency are shown in figure 15. The sampling pitch (Xs) in X and Y equals the monochrome pixel pitch (p). Xs on the diagonals equals

$$\sqrt{2} \bullet p.$$

The sampling frequency in X and Y is 1/p. The sampling distance on the diagonals is larger by

$$\sqrt{2}$$

which makes the sampling frequency smaller by

$$\frac{1}{\sqrt{2}}$$

so the diagonal sampling frequency is

$$\frac{1}{\sqrt{2} * p}.$$

The Nyquist frequencies are half these values. A Nyquist domain graph for the red, green and blue array is shown in figure 16. The inner square is the Nyquist domain for the red and blue array and the outer diamond is the Nyquist domain for the green array. The axes on this graph are normalized to the sampling frequency for the monochrome sensor. The resolution on the X and Y axes is greater than the diagonal resolution for the green array.

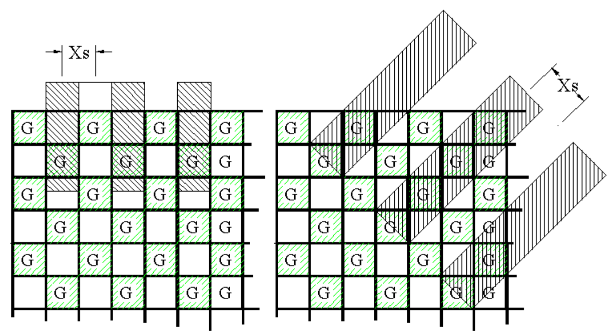

Figure 15.green array line pairs at the Nyquist frequency

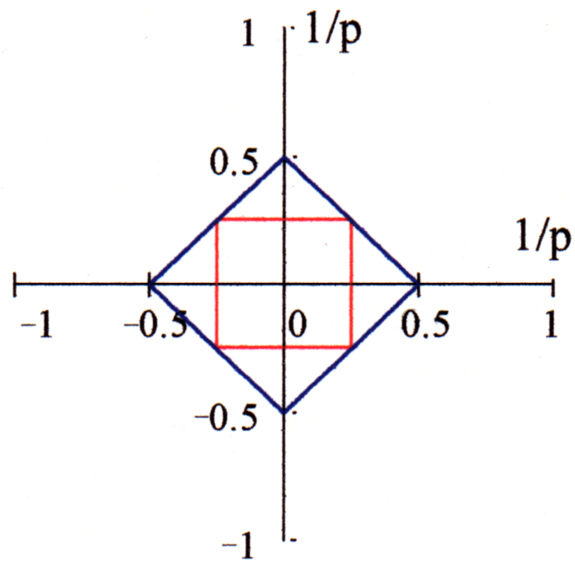

Figure 16. red, green, and blue Nyquist domain

Conclusions

For a Bayer imager the final system resolution is dominated by the green array. A standard Bayer CFA has higher resolution on the X and Y axes compared to the diagonal axes due to higher green sampling frequency on the X and Y axes. This matches the characteristics of the eye which is more sensitive to spatial frequencies on the horizontal and vertical axis.[3] Finally, adaptive interpolation can lead to X and Y axis resolution as high as a monochrome array with the same pixel pitch.

References

1. Bryce E. Bayer, Color Imaging Array, US Patent 3,971,065, July 20, 1976.
2. J. Adams, K. Parulski, K. Spaulding, Color Processing in Digital Cameras, J. *IEEE Micro*, vol. 18, no. 6, pp. 20-30.(1998).
3. A. Watanabe, T. Mori, S. Nagata, and K. Hiwatashi, Spatial sine wave responses of the human visual system, Vision Research, vol 8, pp. 1245-1263. (1968)

Biography

Russ Palum received a BS and MS in photographic science in 1979, and an MS in electrical engineering in 1988, both from the Rochester Institute of Technology. He joined the Eastman Kodak Company in 1977 and has worked on process development for molded glass optics, asphere metrology, scanner light source design, lens design, anti-aliasing filter design and most recently image data path software for small CMOS image arrays.