Anti-Aliasing Filter Analysis for Digital Cameras

Russ Palum; Eastman Kodak Company, Rochester, NY/USA

Abstract

Anti-aliasing filters have been used in video cameras since Pritchard patented the first filter for RCA in 1971. Anti-aliasing filters have been used in consumer digital still cameras at least since 1996, when Kodak introduced the Kodak Digital Science DC40 camera *(now called the* Kodak DC40 digital camera*). This paper will look at the reasons for having an anti-aliasing filter, some of the artifacts that remain with an anti-aliasing filter, and the system parameters that make an anti-aliasing filter unnecessary.*

Sampled Systems

Digital cameras are an example of a two-dimensional sampled imaging system. Sampled imaging systems can misrepresent high spatial frequencies as low spatial frequencies. This artifact is called "aliasing." To prevent aliasing there must be at least two pixels for each cycle at the highest frequency the camera is intended to capture. There must be a pixel to capture the dark part of the cycle and a pixel to capture the bright part of the cycle. Aliasing occurs when the spatial frequency content in the optical image exceeds half the sampling frequency. This is sometimes called the Nyquist frequency. A bar target at the Nyquist frequency is shown in Figure 1. If the target had more bars there

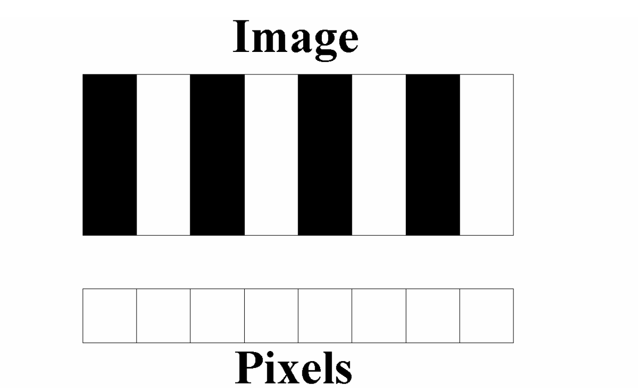

Figure 1*. Bar target at Nyquist frequency*

would not be enough pixels to represent the required number of black-to-white transitions.

To prevent aliasing the optics have to attenuate spatial frequencies above the Nyquist frequency. Lenses can be designed to attenuate spatial frequencies above Nyquist by increasing the size of the point spread function, but it is very hard to get field and f/number independent control of the point spread function. An anti-aliasing filter is usually added to the optical path to provide a fixed amount of change to the point spread function that limits the spatial frequency content of the image.

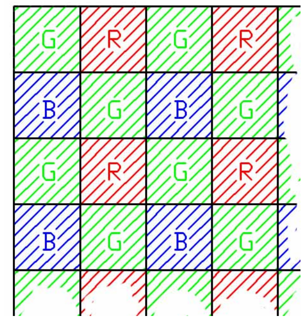

Figure 2*. Bayer pattern*

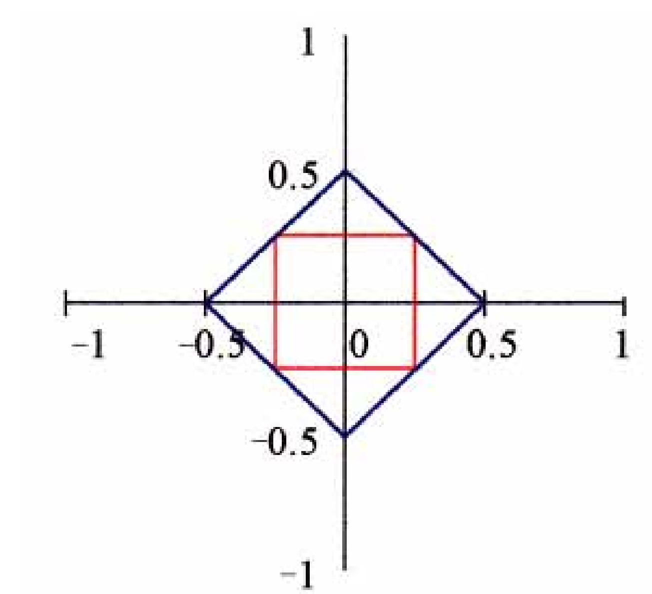

Figure 3*. Nyquist domain*

Nyquist Domain Graph

A Nyquist domain graph is the locus of the Nyquist frequency in a two-dimensional Fourier domain. The MTF of the optical system should go to zero outside the Nyquist domain. A small part of a Bayer imager is shown in Figure 2, and Figure 3 is the corresponding Nyquist domain graph. The units are in cycles per sample. A magnitude of one corresponds to one cycle per sample. One half cycle per sample is the Nyquist frequency. Computationally, each point on the graph corresponds to one divided by two times the color-dependent pixel pitch in each direction. The Nyquist domain graph for the Bayer pattern has an outer diamond that corresponds to the green pixels and an inner square that corresponds to the red and blue pixels. Ideally, the system MTF should go to zero outside the Nyquist domain. A more complete discussion of the Nyquist domain graph is in the first reference [1].

It is possible to make anti-aliasing filters that are color-dependent, but these filters are expensive [2]. In practice, the filter is usually matched to the green channel, and the MTF goes to zero at the green channel Nyquist. The red and blue channels can still alias. For most systems, aliasing is a small periodic change in brightness, but if a colored edge falls on a pixel boundary, a color interpolation error may occur. Single-pixel highlights or very high-contrast transitions tend to produce color interpolation errors. These errors are serious because they are not merely a small change in hue but quite often the name of the color changes.

Anti-Aliasing Filters

There are a number of types of optical anti-aliasing filters; arrays of cones or pyramids can go in front of the lens or at the stop, and regular or random diffraction gratings can also be used anywhere in the optical path. The most common filter is the birefringent quartz filter. Quartz filters can result in different patterns depending on the number of layers and the orientation. Three-layer quartz filters can produce four, seven, or eight spots. The most common quartz filter is the four spot filter. The four spot filter produces the best cutoff with the least blur. The diffractive filters tend to be a uniform circle or square that must be as wide as 2 pixels to cut off at the Nyquist frequency. The four spot filter only has to be one pixel wide in each direction to cut off at the Nyquist frequency. The birefringent filter patterns with more than four spots basically fill in the four spot pattern so they act somewhat like the uniform circle or square.

An intuitive approach to analyzing anti-aliasing filters in one dimension starts with projecting an image point into the scene. A four spot filter is reduced to two spots in one dimension. Figure 4 shows an image point projected backwards through the lens and

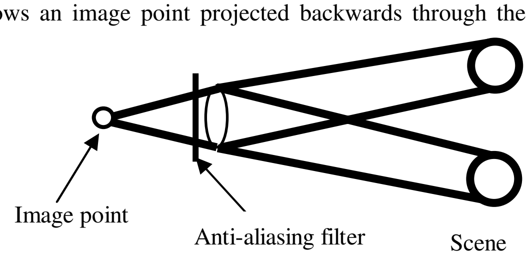

***Figure 4**. image point projected into object plane*

the anti-aliasing filter into the object plane. Two object points contribute to each image point. The modulation should be zero if these object points are superimposed on an object that produces an image at the Nyquist frequency. The cosine image in Figure 5 is shown at the Nyquist frequency. Each pair of matching dots (black, gray, and white) is separated by one pixel pitch. The signal due to each pair of dots always adds to one; this will suppress any modulation as a result of the cosine image. The dots are shown at a number of discrete positions but the modulation is zero everywhere as the dots slide along the cosine. The uniform spot must be two pixel pitches wide to always add to zero. Notice the gray rectangle is one pixel pitch wide and it will sum to a value that varies as it slides over the cosine. The black rectangle will always sum to the same number because it is one period wide at the Nyquist frequency.

Uniform Spot Filter and Four Spot Filter

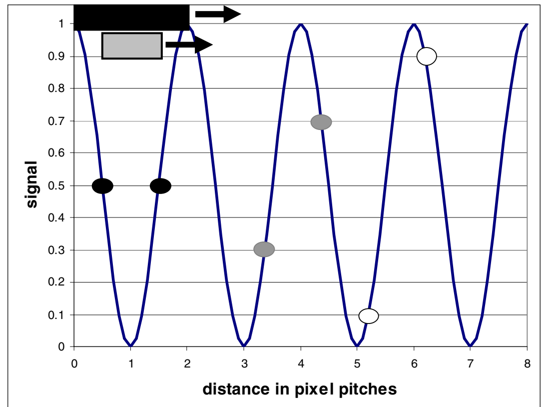

***Figure 5**. Two-spot and uniform anti-aliasing filter*

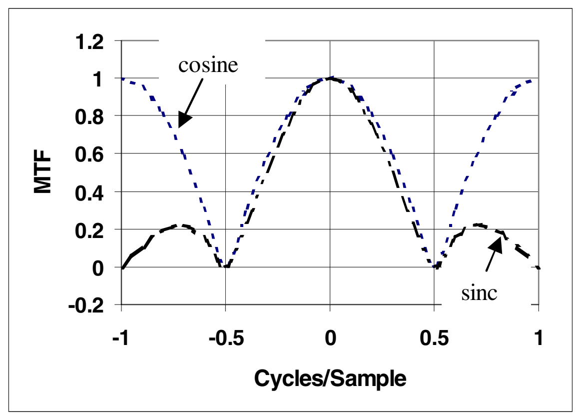

***Figure 6**. Two-spot and uniform spot MTF*

The uniform spot does have an advantage. The MTF of a uniform spot is the modulus of a sinc function that oscillates between zero and lower and lower peaks past Nyquist. The four spot pattern MTF is the modulus of a cosine that goes back to one at the sampling frequency and multiples of the sampling frequency. In general, the lens MTF attenuates past Nyquist to prevent aliasing when the four spot filter opens up. The uniform spot MTF and the four spot pattern are shown in Figure 6. Notice the cosine has a little better MTF below the Nyquist frequency.

The Lens Revisited

The point spread function for a diffraction-limited lens is an Airy disk. Figure 7 is an image of a modified Airy disk; the outer rings have been brightened so they are visible. Eighty-four percent of the power in the Airy disk is in the center bright spot, therefore, the diameter is usually taken to be the diameter of the center spot. The diameter of the center spot is:

$$D = 2.44 * \lambda * N \quad (1)$$

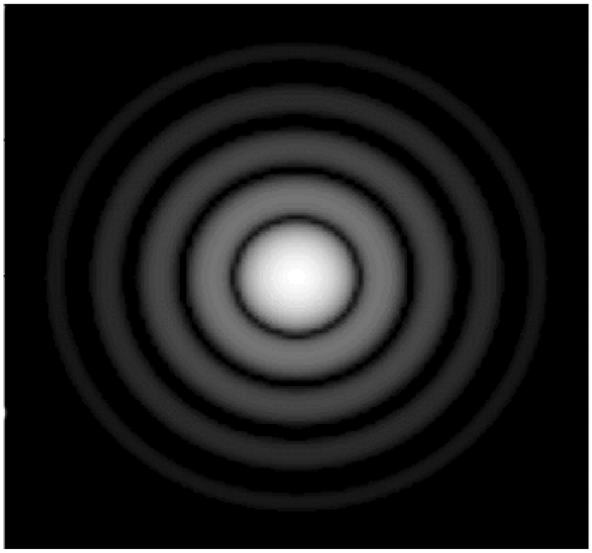

***Figure 7**. Airy disk*

where D is the diameter, λ is the wavelength, and N is the f/number. D will be in the same units as the wavelength.

The lens point spread function limits aliasing in some cameras. The pixels used in consumer cameras have become smaller. Some cameras have pixels around 2 μm. At pixel sizes below 2 μm, the diffraction-limited (read theoretically perfect) lens MTF is sufficient to prevent aliasing and color interpolation error. Imager lenslets limit the lens f/number to about f/3 because larger cone angles overfill the pixel active area. At f/3 the lens has a 4 μm point spread function. The system MTF drops to about 20% when there are 3 pixels per point spread function so at f/3 the lens will suppress aliasing for pixels a little larger than 1 μm. At f/8 the lens will have an 11 μm point spread function and will suppress aliasing almost completely for a 2 μm pixel pitch imager.

Interpolation Error

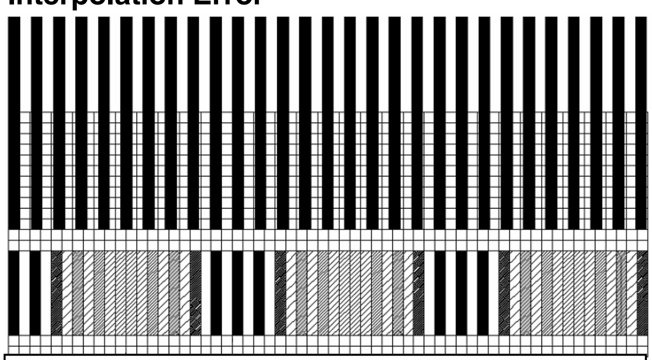

***Figure 8**. interpolation error*

Interpolation error looks like aliasing but it is not aliasing. The bar pattern in Figure 8 results in an image pattern with low frequency modulation even though it is sampled at more than 2 samples per cycle. This happens because the phase of the samples changes from one cycle to the next. Another example is shown in Figure 9. The sampled sine wave in the upper graph of the figure

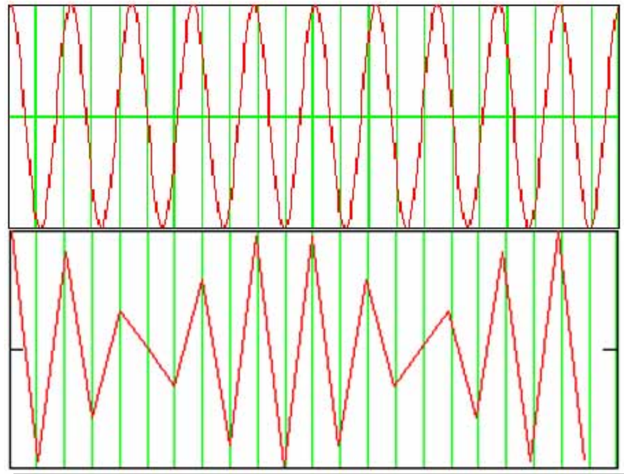

***Figure9**. Sampled image near Nyquist*

is sampled at greater than two samples per cycle, but the reconstruction below has low frequency modulation that looks like aliasing. The sampling meets the criteria of the Whitaker Shannon law as found in Gaskill [3]: "Any band-limited function can be specified exactly by its sampled values, taken at regular intervals, provided that these intervals do not exceed some critical sampling interval." To stay below the critical sampling interval there must be more than two sampling intervals for each cycle at the band limit. If this condition is met, the function can be reconstructed from the sampled data; if this criteria is exceeded, aliasing will result. Reconstruction [4] is often ignored, however, reconstruction is an important part of displaying a sampled image. The display should have more pixels than the camera, and sinc interpolation, or any other appropriate interpolation function, should be used to up-sample the captured image to the display size. The sinc convolution in the spatial domain is equivalent to multiplying by a rect function in the frequency domain. This suppresses higher order copies of the image frequency spectrum.

The display must have about four times as many pixels as the capture device in each direction to reduce reconstruction errors to about 10% modulation. Ten times as many pixels will reduce the modulation to a few percent. The worst case low frequency modulation can be determined by splitting two samples across the peak of a sine wave at the Nyquist frequency, as shown in Figure 10. It is rare for a scene to contain a large area at a spatial

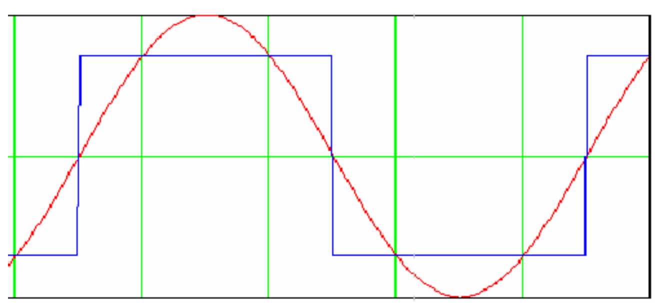

***Figure 10**. Worst case modulation*

frequency exactly at the Nyquist, so the phase of periodic objects tends to vary from sampling exactly at the peak to splitting the samples across the peak over a large number of pixels. This produces the low frequency modulation in Figures 8 and 9.

Conclusions

Aliasing can be avoided in sampled images by band-limiting the optical image to the Nyquist frequency, although most consumer cameras do not band-limit all of the channels to the Nyquist frequency. The four spot birefringent filter does a good job of preventing color interpolation error in digital cameras, however, the filter is usually chosen to prevent aliasing in the green channel and allows aliasing in the red and blue channel. An image might exhibit an artifact that looks like aliasing even if the optical image is band-limited to the Nyquist frequency; this artifact can be corrected if proper reconstruction is applied to the image.

References

[1] R. Palum, Image Sampling with the Bayer Color Filter Array, IS&T 2001 PICS Conference Proceedings, pp. 239–245. (2001).

[2] J.E. Greivenkamp, Optical spatial filter, United States Patent, 4,575,193, (March 11, 1986).

[3] J.D. Gaskill, Linear Systems, Fourier Transforms, and Optics (New York:John Wiley and Sons, 1978), pp.266–267.

[4] Ibid., p.271

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.