

## ***FULL WELLS AND BLUE SMOKE:***

### **DEMYSTIFYING HIGH-PERFORMANCE CCD CAMERA SPECS AND TERMS**

**By David Gardner**

**From Advanced Imaging, April 1999**

At one time, selecting a CCD camera was a fairly simple process—the primary technical requirement was simply to capture a visually-acceptable image. Most cameras had resolution and noise performance which were roughly equivalent to standard television, and the most significant factor in the buying process was *price*. Now, however, as imaging systems push their way into medical, scientific and high-end inspection applications, performance requirements go far beyond merely getting an image. In contrast to the earlier video cameras, modern imaging systems must be quantitatively specified in terms of dynamic range, linearity, sensitivity and full well capacity. These performance characteristics define the boundary between scientific camera characterization and a qualitative “feeling” for image quality.

<br>

To many potential users of these cameras, even for engineers, the language of camera characterization is new and, frequently, confusing. You’ve probably noticed that terms such as “full well capacity” and “shot-noise limited” are bandied about by camera manufacturers, but their true correlation to image quality is rarely understood well by the end-user.

<p>To make matters worse, a quick review of camera data sheets will reveal that there are currently as many ways of presenting technical specifications as there are camera vendors. The end result of this confusion is that imaging system integrators frequently resort to side-by-side shoot-offs between multiple camera types in hopes of finding one that will work in their individual applications. While in some cases these comparisons are needed, in most, *they are not*. In an effort to accurately compare camera performance, the imaging system integrator must schedule multiple camera vendors, re-configure cabling, re-design optics and modify system software, resulting in delays, added expense and a significant level of frustration.</p>

So, the purpose of this article is to remove the mystery involved in camera characterization so that camera users can intelligently decide which camera best fits their applications. In gaining a better understanding of some key camera terminology, the

camera user can decide which specifications are important to a given application and make purchasing decisions more confidently.

#### **“Dynamic range” and bits**

Dynamic range is perhaps the most abused technical parameter found in camera data sheets. At Salvador Imaging—and most other high-performance camera vendors such as Princeton Instruments and Photometrics (both new wings of Roper Scientific Imaging)—dynamic range is defined as the ratio of the largest signal that the CCD can handle (linearly) to the read-out noise (in the dark) of the CCD camera system.

From an intuitive point of view, dynamic range defines the brightest and darkest data in a given image which can reliably and faithfully be reproduced by the camera system. Dynamic range may be presented as a ratio, bit depth or equivalent dB rating. For example, a camera digitized to 12 bits with 2 counts of rms read noise would have a dynamic range of  $4096/2 = 2048:1$ , which is equivalent to  $20\text{Log}(4096/2)=66$  dB or 11 bits.

There are, of course, a number of situations where having a large dynamic range is advantageous. Radiology, motion-picture film scanning, and microscopy are all image applications where the image scene typically contains both very bright and very dim image information. In these applications, a large dynamic range (12 bits or more) allows the user to capture and reproduce extremely subtle gray-scale variations in both the bright and dark areas of the image. In contrast, an 8-bit camera used in these same applications would either lose data entirely by clipping gray-scale values at one end of the range, or provide insufficient quantization accuracy to reproduce subtle changes in gray level.

Having a large dynamic range is also very helpful in situations where the experimenter has no *a priori* knowledge of scene illumination and the experiment is difficult or costly to repeat. In this case, the large dynamic range allows the user to have a wider window of gray-scale values in which the scene may occur without loss of image data. Thus, the

experimenter has a higher probability of getting useful data on the first try. As an example, by using a 12-bit camera, the experimenter is allowed up to 16 times as much error in estimating the scene illumination than with an 8-bit camera.

Other imaging applications benefit from a large dynamic range in more subtle ways. For example, in semiconductor inspection, there is frequently a need to very accurately identify edge information. Typically, the image data to be captured is either black or white, which (at first) would seem to imply the need for only a very limited dynamic range (1 bit). In practice, however, system integrators use sub-pixel gray-scale variations to provide positional information far beyond the resolution of a single pixel. In this case, electronic noise, signal jitter and other camera parameters can dramatically effect sub-pixel measurements. In this and many other applications, the benefit of a high-dynamic-range camera really is not directly related to its ability to capture a wide range of gray-scale values. Instead, it has to do with the fact that such a camera must inherently be more stable, linear and noise-free over time and temperature.

A common *misconception* is that dynamic range equates to digitization level. For example, a camera with an 8-bit A/D converter is assumed to have 8 bits of dynamic range. Unfortunately, this misconception is being kept alive through the proliferation of technically incorrect data sheets. To understand why the two are not the same, consider two cameras which are digitized to 12 bits, or 4096 gray-scale levels. The first camera has an inherent rms noise of 8 A/D counts, and must have a minimum exposure large enough to overcome this basic noise floor to be detected. Although this camera is digitized to 12 bits, the effective dynamic range is actually  $4096/8=512:1=9$  bits.

In contrast, the second camera has a read noise floor of 1 A/D count. Since the noise in this camera is one A/D unit, the dynamic range becomes essentially the range of the A/D, or 12 bits. Thus, while both cameras are digitized to the same level, there is an 8-to-1 difference in dynamic range!

#### **Signal-to-noise ratio and full well capacity**

Consider an application in which you must measure the light falling on a particular pixel to a precision of 1% (SNR = 100). Even with a perfectly noiseless camera and noiseless detector, the random nature of photons imposes an effective noise level (shot noise) which *must* be captured in order to make the measurement. The nature of photon generation demands that for the most precise measurements, the

CCD must be able to handle large numbers of photons per pixel, and the number of electrons which can be contained in a pixel is referred to as the “full well capacity”.

Because the shot noise in a camera grows with the square root of the number of photo-electrons, the maximum signal-to-noise ratio of a properly-designed camera will be “shot-noise limited”—the maximum signal-to-noise ratio of the camera is limited by the inherent statistical nature of light rather than the read noise floor of the camera electronics. This point is important because it means that the signal-to-noise ratio (SNR) of a camera (the ratio of signal to noise at some illumination level) may never be as high as the dynamic range (the ratio of the maximum possible signal to the minimum camera noise). So, a camera with a 12-bit dynamic range does not necessarily have 12-bit SNR at any point in its operating curve. In fact, to have a 12-bit SNR would mean that the CCD’s full well capacity would have to be at least  $4096 \times 4096 = 16,777,216$  electrons. Only a small number of highly-specialized sensors can meet that need.

This is one area where scientific-grade CCDs clearly differ from CCDs designed for other purposes. The amount of charge that a CCD can store in each pixel depends largely on the physical size of the pixel. For this reason, scientific CCDs usually have relatively large pixels, as large as 12 to 27 microns on a side. Since the cost of producing integrated circuits is strongly area-dependent, non-scientific CCDs are usually made as *small* as possible—pixel sizes of 8 microns or less—and a correspondingly smaller full well capacity.

For example, a consumer-grade CCD with a pixel size of 7 microns may have a full well capacity of around 40,000 electrons. In this case, the highest precision possible would be about 0.5%. In contrast, a scientific-grade CCD with a full well of 400,000 electrons would allow for over three times more accuracy.

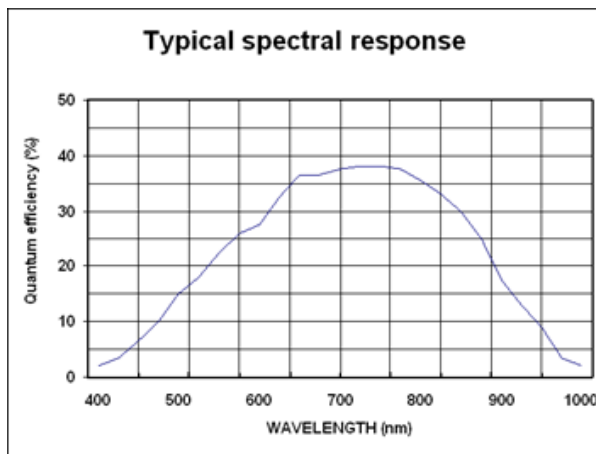
#### **Quantum efficiency and optical sensitivity**

Optical sensitivity quantifies the minimum light level which can be reliably detected above the camera’s read noise. For a given exposure length, the two factors which effect sensitivity are:

- 1) how many signal photoelectrons are generated within a pixel for a given illumination level, and

- 2) the minimum number of photoelectrons needed to be reliably detected over the camera's inherent noise floor.

The percentage of incident light which is converted into usable signal electrons is referred to as the Quantum Efficiency (QE). Figure 1 shows a typical CCD QE curve. A key aspect to note in Figure 1 is that the QE is a function of illumination *wavelength*. This dependence on wavelength means that a camera with a high QE in the green (550nm) wavelengths may have a substantially lower QE in the far red (800nm) or vice-versa. A common point of confusion results when a manufacturer simply states that they have a QE of, say 40%. This statement is, by itself, not only meaningless, it's potentially misleading.



**Figure 1. Typical CCD Quantum Efficiency Curve. Unless the wavelength is specified, QE is not properly defined.**

The second parameter which directly impacts camera sensitivity is the "read noise" of the camera. Lower read noise means that the camera needs fewer photoelectrons to overcome the read noise level. Thus, a camera with 25 electrons of effective read noise will require half as many photoelectrons to overcome the noise floor than a camera with a 50-electron read noise floor does. This is an important point because many camera users make the erroneous assumption that a camera with high QE will automatically be more sensitive than a camera with low QE.

For example, assume you are evaluating two cameras, each with a 200,000 electron full well capacity. Camera A has a QE of 80% at 550 nm and a dynamic range of 200:1. Camera B has a QE of

only 20% at 550 nm and a dynamic range of 4000:1. It would seem at first glance that the camera with 80% QE would be the most sensitive because it is 4 times more efficient at converting the selected wavelength of light to signal electrons. However, this neglects the fact that the inherent read noise floor of Camera A is 1000 electrons (200,000 electrons/200) whereas Camera B has a read noise floor of 50 electrons (200,000 electrons/4000). Thus, although Camera A is four times more efficient in capturing the light, Camera B requires 20 times fewer (1000/50) photoelectrons to generate the same video signal. The end result is that the camera with the four-times lower QE is actually 5 times more sensitive!

### Linearity?

In an ideal linear camera, doubling the light intensity should result in a doubling of the video signal. More generally, if the light impinging on the detector surface is increased by some multiple, the corresponding video signal should increase by precisely the same factor. Specifying the linearity of a camera essentially describes how well it follows this ideal performance.

A linear camera means that its ability to resolve subtle variations in gray is the same whether the image detail is, on average, very bright or very dim. In other words, the sensitivity to a change in illumination is the same regardless of illumination level. Linearity is typically most important in applications where the image data will be processed by something other than the human vision system. Interferometry, X-ray, solder-joint inspection, bio-chip readers and spectrometry are all applications where linearity is crucial.

The linearity of the scientific-grade CCD itself when operated with proper signal condition circuitry is almost always better than 1%, except at extremely high signal levels. In fact, it is so difficult to measure the non-linearity of the CCDs that 1% linearity specification usually means that the CCD is more linear than the test fixture.

There are, however, a couple of things which the camera electronics can do which will result in very non-linear operation. First, if the camera has an anti-bloom function which is incorrectly calibrated, the system response will be very non-linear, resulting in compressed gray-scale data at higher illumination levels. Secondly, high-speed operation of CCDs sometimes reduces the achieved linearity because of slew rate limitations in either the signal conditioning electronics or in the CCD output stage itself.

Depending on the specific cause of the slew rate limitation, this effect may cause reduced gray scale resolution at one end or the other of the gray scale range. In the case where the low light detail is compressed, this effect will result in reduced sensitivity at low light levels and a dynamic range which appears to be higher than it really is.

**Sure measurement techniques**

Beyond keeping camera terminology straight, the next question for many system integrators needs to be, “How do I know that the printed specifications are accurate?”

Fortunately, there *is* a very accurate measurement technique—the Photon Transfer Curve—which quantifies each of the parameters we’ve discussed. Simplistically, the Photon Transfer Curve uses the knowledge that the rms shot noise associated with light itself will exactly equal the square root of mean signal level. During calibration, the camera is exposed to a controlled sequence of light intensities, and any noise measurement which deviates from the ideal shot-noise characteristic must be due to the camera itself.

The Photon Transfer Curve has been adopted by NASA’s Jet Propulsion Lab (as well as most scientific camera manufacturers) as the defining technique for camera calibration. The Photon Transfer measurement provides verifiable information on read noise, full well capacity, linearity, dynamic range and camera sensitivity, and is critical in verifying the accuracy of printed technical specifications. Proper use of the Photon Transfer Curve removes the guess work from camera selection, saving both time and money for the imaging system integrator.

Figure 2 shows a sample Photon Transfer Curve. From this curve, we can tell that the camera has a read noise floor of about 1.2 A/D counts rms. The camera gain is about 85 electrons per A/D count, full well is around 280,000 electrons and the dynamic range is  $4096/1.2 = 3413:1$  (71 dB). Further processing of the data in this curve will yield linearity and sensitivity information as well.




Figure 2. Photon Transfer Curve: a cure for blue smoke blues.

**Beyond ulcers**

Selecting the appropriate camera is key to the success of any imaging application. However, as critical as it is, the selection process does not need to go hand-in-hand with ulcer medication. Understanding what camera terminology means and how it can impact the application is the key to successful selection.

As a first step in evaluating any high-performance camera, it is strongly recommended that the system integrator request a full set of camera performance specifications, including a Photon Transfer Curve. If the camera manufacturer is unwilling or unable to provide the requested data, proceed with caution, as it is very difficult to characterize any high-end camera without these tools.

Finally, it should be noted that while this discussion has centered around CCD cameras, the same parameters and measurement techniques can be used to extract the illusive performance characteristics of CMOS or Active Pixel Sensors.

For further detail on the Photon Transfer Curve, as well as some excellent reference material on camera performance, please visit the Salvador Imaging Website at

<http://www.salvadorimaging.com/>