Characterizing Digital Cameras with the Photon Transfer Curve

By: David Gardner
Summit Imaging

Introduction

Purchasing a camera for high performance imaging applications is frequently a confusing and irritating process. Two cameras with seemingly identical printed specifications can perform completely differently during a side by side "Shoot-off". One of the sources for this confusion is that a camera specified as 12 bits from one vendor may mean simply that a 12 bit A/D converter is used somewhere in the camera – while to another (more rigorous) camera vendor, 12 bits means a dynamic range of 72dB.

The inability to perform a quantitative comparison based upon technical specifications alone can be solved in one of two ways. The first (and most common) approach would be to do side by side comparisons of all cameras which "look" like they might work. This brute force approach will work, but it is very expensive and time consuming for the camera system integrator. In addition to coordinating the delivery of the camera, frame grabber, cables and power supply, the system integrator must familiarize thereself with new software and determine a suitable method for measuring camera performance. While this approach might be intriguing as a doctoral thesis topic, it is very inefficient for the camera consumer. Moreover, it places all of the burden of proving technical capability on the consumer rather than the camera manufacturer.

An alternative approach is to place the burden of technical proof on the camera manufacturer. With this plan, a standardized test procedure is used during camera manufacturing to provide consistent, quantitative and verifiable performance data such as read noise, dark current, full well capacity, sensitivity, dynamic range, gain, and linearity. Fortunately, a test method of this type has existed for well over a decade and is known as the Photon Transfer Curve (PTC). The PTC characterization method is used by NASA's Goddard Space Flight Center, NASA Jet Propulsion Labs and leading camera manufacturers around the world to allow for an apples-to-apples comparison of key performance parameters. The fundamentals of PTC calibration are derived from simple system theory where knowledge of a system's input and output signals are used to derive the characteristics of the system itself.

Measuring Noise with Noise

From a very basic measurement point of view, the PTC essentially works as follows: 1) the camera itself is a system block with light as an input, and digital data as an output, 2) we know that the only noise introduced at the input is shot noise due to the nature of photons themselves – and we can predict exactly what that noise will be at a given illumination level, and 3) any difference between the noise at the input and the noise at the output, must have been caused by the camera (or sensor) electronics.

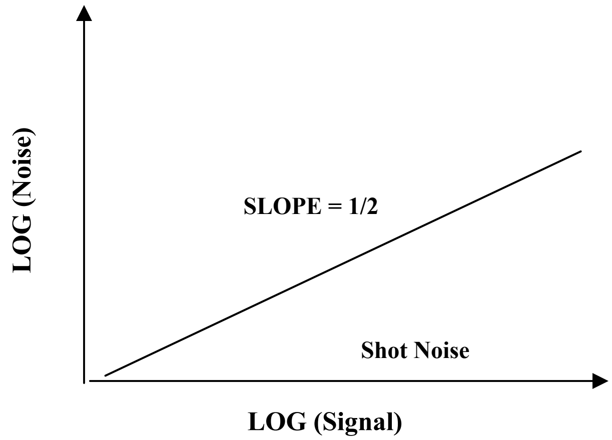

Figure 1. Rms shot noise of input light plotted on a Log-Log curve.

The use of noise as a test stimulus to the camera is very convenient because the natural input signal for an imager is light, and the noise characteristics of light are very well known. In Figure 1, the shot noise characteristics of light are plotted as a function of illumination level on a Log-Log graph. The rms value of shot noise is equal to the square root of the mean number of photons incident on a given pixel. Thus, the shot noise profile becomes a straight line with a slope of ½ on the Log-Log curve (since Log $X^{1/2}$ = 1/2Log X). Keep in mind that this noise is inherent the nature of light itself and has nothing to do with the camera design.

In contrast to Figure 1 which shows only the noise associated with the input signal (light), Figure 2 shows the PTC which contains the typical noise profile seen at the output of a digital camera. In this figure you can see three distinct noise regions of the CCD camera system: read noise, shot noise, and fixed pattern noise. As discussed above, the PTC compares the differences in Figure 1 and Figure 2 to determine the operational characteristics of the camera itself. In the paragraphs below, the unique characteristics of each of the regions will be discussed:

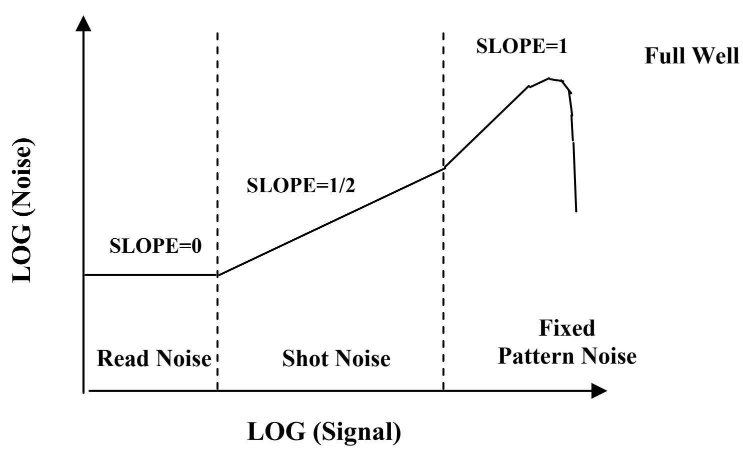

Figure 2. Photon Transfer Curve.

<u>Read noise:</u> Read noise is represented by the first (flat) region of the graph shown as Figure 2, and is the random noise associated with the CCD output amplifier and it's (external) signal

processing electronics. This is also referred to as the noise floor of the camera, and represents the baseline noise in total darkness.

Shot noise: The second region of the graph in Figure 2 shows shot noise which is inherent in the light itself – it does not originate in the camera. As the input light level increases in amplitude, the noise at the camera output rises out of the "read noise" region and becomes dominated by shot noise. Shot noise is directly related to the input illumination, and is proportional to the square root of that signal. In this region of operation, the camera is operating as a shot-noise limited system where S = the input illumination (signal)

Fixed Pattern noise: The right most region of the PTC in Figure 2 shows the fixed pattern noise, which becomes dominant at relatively high levels of illumination. This noise results from differences in sensitivity between pixels, or Photoresponse Nonuniformity (PRNU.) This noise is directly proportional to input signal strength, so the slope in this region is 1. That is:

$$N_{FP} \propto \left(S \times PRNU\right)$$

Full-Well: As illumination levels are increased, the individual CCD pixels are unable to hold any additional charge without "spilling-over" into adjacent wells. At this point on the noise curve, output noise abruptly drops because charge sharing between adjacent pixels averages the signal and suppresses random noise. At the point where the photon transfer curve peaks, as shown in Figure 2, the CCD is said to have reached full-well. Normally, a camera system is calibrated such that the maximum A/D output is achieved at or near full-well.

The full photon transfer curve shown in Figure 2 is shown to illustrate the various noise regions. However, in practice, the fixed pattern noise portion of the curve is eliminated during measurement, and the shot noise region is extended to the point of full well operation. This is done to provide linearity and sensitivity data at higher illumination levels and is accomplished by subtracting two illuminated fields. Subtracting the two images eliminates fixed pattern variations since they are present in both images. With fixed pattern noise eliminated the PTC in Figure 3 continues along with a slope of ½ until full well occurs.

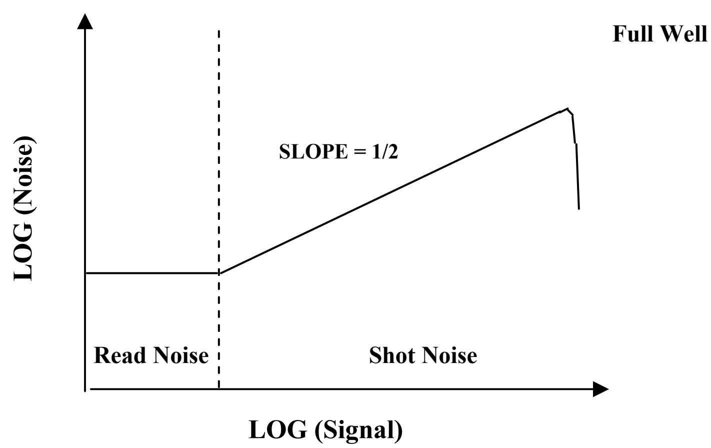

Figure 3 – PTC with Fixed Pattern Noise Eliminated through Frame Subtraction

Data collection

During the Photon Transfer Curve measurement, the CCD is exposed to a precisely controlled source of light which provides a flat (uniform) illumination field. This is done by using an integrating sphere with a monochromic light source such as LED's or filtered white light. The integrating sphere guarantees better than 99% uniformity at the output port so that data collected across the array is consistent. The use of monochromatic light is important to remove the effects of color temperature and quantum efficiency variations.

Prior to starting the test run, the operator adjusts the light source so that the camera is at its previously determined full-well illumination. At this illumination, the camera gain and offset are adjusted such that the A/D converter full-scale corresponds to the CCD full-well condition, within about 100 ADU's. With this initial calibration complete, the light source is then stepped from complete darkness to full well illumination in precisely measured increments. At each illumination level, two images are captured from these two images and then subtracted on a pixel by pixel basis and the variance is computed as discussed above. The exact illumination level for each measurement is recorded with a calibrated photodiode.

After data collection, the PTC is calculated by determining the rms random noise at the camera output. The noise is calculated at a specific (measured) light level by subtracting two frames (to remove PRNU and frame to frame offsets), and then taking the sum of the squares of all pixels in the subtracted image and dividing by two times the number of total pixels. The result of this calculation is the statistical variance of the time varying random noise in the image, where the factor of two in the divisor comes from the fact that the noise is initially doubled due to the subtraction process.

where:

$X1_i$ = the individual pixel values of the first frame, in ADU
$X2_i$ = the individual pixel values of the second frame, in ADU
M1 = the mean of all pixel values in the first frame
M2 = the mean of all pixel values in the second frame
Np = the number of pixels in the sample

$$Variance = \sigma^2 = \frac{\sum_{i=1}^{Np} \left([X1_i - M1] - [X2_i - M2] \right)^2}{2 \times Np}$$

$$and \ \ rmsnoise = \sigma = \sqrt{\sigma^2}$$

So you've measured a lot of noise – where do you get the other camera parameters?

Once the noise at each illumination level is calculated, a PTC is generated by plotting the camera's output rms noise vs. average signal level on a Log-Log curve. Each parameter is discussed below along with it's relevance to the PTC plot.

Read Noise

Read noise is the inherent electronic noise floor of the camera. It tells you the minimum signal you will ever see, and it is used to compute dynamic range. Read noise is directly available from PTC by recording the noise level at zero illumination. On the PTC, the read noise is shown in terms of the number of ADU's of rms noise in darkness –which can be multiplied by the gain of the system to yield the noise floor in e^-.

Gain

The gain of camera systems is typically expressed in ADU/e^-, that is, the number of A/D units per signal electron. To obtain gain, we note that for an increase in illumination of X, the camera's average signal level will change by GX. In contrast, the same increase in illumination level will cause the noise variance, to change by G^2X. The ratio of these two is the slope of the line measured when plotting noise variance vs. average signal on a linear graph. Thus, the camera gain, G, is obtained by fitting a line to this variance curve in the shot noise limited region and measuring the slope of that line.

Dynamic Range

The dynamic range is calculated as full-scale of the A/D range divided by the smallest detectable signal. Since the read noise sets this lower limit, the dynamic range is given by:

$$DR = \frac{ADU_{FS}}{\sigma_R}$$

The DR, expressed in dB, becomes:

$$DR = 20\log\frac{ADU_{FS}}{\sigma_R}$$

Full Well

This parameter is derived from the maximum A/D output (counts) and the gain. The formula is as follows:

$$FW = S_{MAX} / G$$

where:

S_{MAX} = maximum A/D signal before the PTC begins to slope downward
G = gain, as previously defined
FW = full well (in electrons)

Note that this calculation depends on the camera gain being accurately calibrated such that the maximum signal level occurred at full-well and resulted in the maximum A/D output.

Nonlinearity

This parameter is based on the error between the best-fit straight-line to the original data of the camera input vs. the camera mean response at each illumination level, in ADU's. The formula is as follows:

$$INL = \frac{\mathrm{E}_{MAX} - \mathrm{E}_{MIN}}{AD_{FS}} \times 100$$

where:

INL	= Integral Nonlinearity
E_{MAX}	= maximum (most positive) error from best-fit straight line
E_{MIN}	= minimum (most negative) error from best-fit straight line
AD_{FS}	= A/D converter full-scale value (counts)

Note that (E_{MAX} - E_{MIN}) represents the worst-case peak-to-peak amplitude of the error. The A/D converter full-scale value is the value of the highest measurement taken which is still within the linear performance range. Normally this is set to be as close to the actual full-scale capability of the A/D converter as is practical during camera calibration.

Effective Number of Bits (ENOB)

The effective number of bits (ENOB) is a standard term which indicates the useful number of digital bits which a system can deliver based on its SNR. The calculation basically converts the SNR, normally expressed as a $\log_{10}$ number, to $\log_2$, and is as follows:

$$ENOB = \frac{\log\left(\frac{ADU_{FS}}{\sigma_R}\right)}{\log(2)}$$

Summary

Proper use of the Photon Transfer Curve removes the guess-work from camera selection, saving both time and money for the imaging system integrator. For this reason, the Photon Transfer Curve is used by Jet Propulsion Labs, Air Force Research Laboratories and leading high-performance camera manufacturers.

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.