<u>10 bit, 12, bit, 14 bit, 16 etc?</u>

<u>Crisp 16 October 2017, rdcrisp@earthlink.net</u>

The dynamic range of a camera is defined as the full well capacity divided by the read noise. For electron based solid state imagers, both are usually expressed in units of electrons (versus ADU or DN, two names for the same concept). For example, if the well capacity is 50Ke- and the read noise is 10e- then the dynamic range is 50,000 / 10 = 5,000 discrete states that can be resolved by the camera in a single exposure.

If 12 bits are used, then 4096 states can be represented, which falls just short of 5000, so this will take 13 bits which allows up to 8192 states.

Mathematically, to calculate the required number of bits:

#bits = Ln(#states)/ln(2) [Ln is the natural logarithm transcendental function]

In our example that is Ln(5000)/Ln(2) = 12.28, again we use 13 bits: 12 is not quite enough.

Setting the camera gain is another design decision. A commonly used method for many cameras is to set the gain to equal the read noise. In our example, if we set the gain to equal 10 e-/ADU then with a 50000 e- signal we have 5,000 ADU and with 10 e- of signal we have 1 ADU.

If we have no added signal DC offset, then full well is reached with 5000ADU even though the 13 bits of the ADC can have full scale set to 8191 (0-8191, or 8192 states). You can apply an offset such that with zero electrons in the well, then the ADC output value is a non-zero number. For example, if full scale on the ADC is 8191 and full well takes 5000ADU to represent, then the difference (8191 – 5000) = 3191 can be used as the offset.

In that case full well is reached at ADC full scale (8191 ADU) and zero electrons are represented by a value of 3191 ADU.

For a standard DSLR one exposure = a photograph. They don't generally combine multiple images to make a single photo. So for that application there's little advantage having a camera gain that digitizes 10 e- of read noise with more than 1 bit.

<u>That may not be a good optimization for astro use, here's why:</u>

Faint signals like we constantly encounter may be less than the read noise of the camera for a given exposure time (what we call subexposures). You may only get 3 e- of signal but it is buried under 10 e- of read noise in our hypothetical example.

If your camera digitizes 10 e- as 1 ADU, then this 3 e- of signal is indistinguishable from 0, 1, 2, 3, 4 e-: all are digitized to the same value, which is 0.

Likewise, a signal of 5, 6, 7, 8, 9, 10, 11, 12, 13, 14 e- would be digitized as 1 (takes 10e- to record 1 ADU and we are assuming the ADC rounds to the nearest integer value).

This is a quantizing issue.

<u>For the astro use, it may be better to use a different camera gain and use more bits for digitizing "noise".</u>

One way to do this is to use a 16 bit ADC.

Since a 16 bit ADC can digitize 65536 states and we have 50,000 e- for full well we could use a gain of 50,000 / 65536 = 0.7629 e- per ADU. For several practical reasons let's not discuss right now, it may be a good idea to have an offset of 2000 ADU (meaning zero signal records a value of 2000 ADU) so we change the parameters a small amount:

50,000 / (65536 – 2000) = 0.7869 e-/ADU

In this case we are using a little more than 3 bits to digitize noise.

Let's see how that works in our favor.

Combining multiple exposures: random noise versus correlated signal

Theory teaches us that random noise is reduced by a factor of 2 for every quadrupling of signal when dealing with shot noise limited sampled data systems characterized by a Poisson distribution. That is the scenario we encounter when imaging astronomical targets.

The read noise is uncorrelated random noise. As a result, one can "beat down" read noise to arbitrarily low levels by simply combining images. Try it yourself: take 1024 bias frames and pick out a 100 x 100 box somewhere in a clean part of the frame and measure the Standard Deviation.

Then combine 4, 16, 64, 256, 1024 frames and make the same measurement. If you plot the Std Deviation versus # frames combined using logarithmic axes for the chart, you will see a straight line with a slope of -1/2 (Figure 1)

Noise Reduction by Stacking

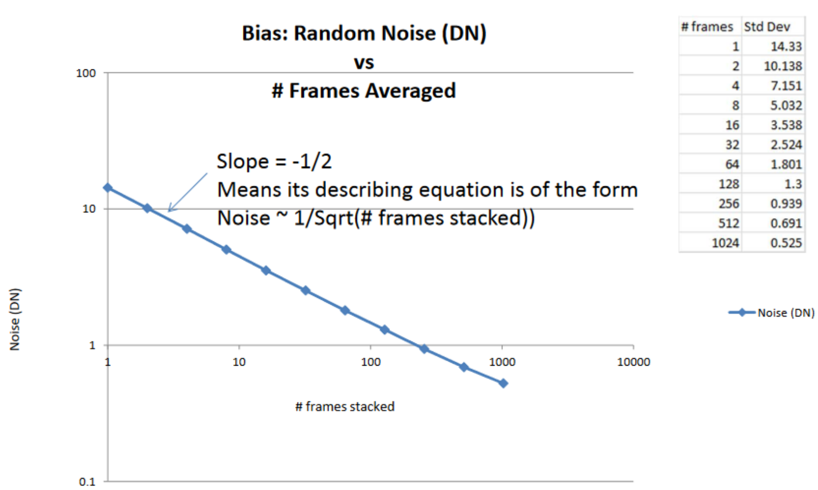

# frames	Std Dev
1	14.33
2	10.138
4	7.151
8	5.032
16	3.538
32	2.524
64	1.801
128	1.3
256	0.939
512	0.691
1024	0.525

Figure 1

What the graph is showing is that you can combine an arbitrarily large number of exposures that are filled with read noise and drive down the read noise to any value you so desire.

If you have signal buried under that read noise, that signal will improve as the noise is decreased: the random noise is halved for every quadrupling of signal and for the correlated signal, it's random noise (shot noise) is decreased by a factor of two for the same quadrupling of signal collected. You see the combined effects below in Figure 2.

ABG Sensor (KAI29050) with 20,500 e- saturation signal on HDR object (M42)

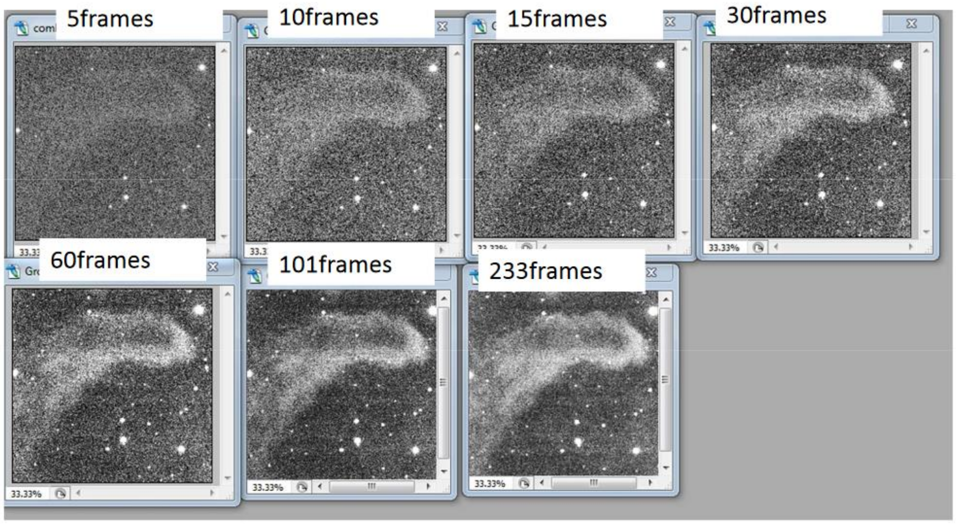

Very faint parts of image are buried in read noise: use large numbers of frames to improve SNR (ie 90 second exposures in Orion Nebula)

Figure 2

<u>Conclusion:</u>

It is wasteful of storage and transmission bandwidth to encode noise using more than one bit for applications that involve single images to make a digital photograph. Optimum encoding suggests setting camera gain to numerically match the read noise in order to use the least possible bits when using linear encoding without compression.

For cameras used to image very faint objects that have RMS magnitude less than the read noise, a different optimization is beneficial to better resolve the faint images: many subexposures are combined to make a final product.

Because random uncorrelated noise is reduced by the square root of the number of frames combined it is beneficial to use several bits of the ADC to digitize signals that are smaller in magnitude in a single exposure than is the read noise for cameras used in such applications. This is another way that cameras optimized for astronomy are different than ones used for terrestrial photography.