

Signal to Noise: Understanding it, Measuring it, and Improving it

Part 2 - Understanding one pixel

Craig Stark

[All text and images Copyright Craig Stark, 2009. Material first appeared on Cloudy Nights (www.cloudynights.com) on May 26, 2009]

In [Part 1 of of this series](#), we went over the basic notions of what we mean by signal and what we mean by noise. Our images will always have noise and we're never going to get around that. The trick to making our pictures look pretty is to have enough signal to noise ratio (SNR) so that we can stretch well enough to resolve those faint bits without having the image look like a mess. SNR on the bright bits is going to be easy, so we'll spend most of our time talking about the low end of the intensity scale.

In this installment, we'll talk about the fundamental unit of our image – the pixel. While it may seem like a pretty low bar to set here, understanding the pixel is key to understanding our image and the quality of our data overall. Most of what comes down the pike hinges on understanding the pixel.

Before I get going, I do want to mention that it's not like I'm blazing a lot of new ground here. There have been a number of attempts to explain SNR in the pixel on the web. One well-known source is [Steve Cannistra's coverage](#) and another here on Cloudy Nights is [Charles Anstey's](#). I've also recently come across [Sam Fahmie's excellent article](#) on the topic. Each of these covers the math in great detail and these are great resources. What I hope to do here is to be faithful to the math (numbers don't lie – people lie with numbers!), but to try to cache this in a different format. I'll use a few equations, but nothing will get tougher than what the simple calculator program on your computer can handle. I'll toss numbers around (quite a bit actually), but the goal of this is to show you where they come from and how you can get to them yourself. To help out, I've also [provided a spreadsheet](#) that you can use in Excel, Open Office, Neo Office, or anything that can work with an Excel spreadsheet.

As we work through this and through the upcoming parts of this series, you may note a particular philosophy of mine. If you don't, I'll lay it out plainly here. It may seem odd coming from someone who makes a living as a scientist and researcher and someone who tries to do astrophotography, but the philosophy is to ***not stress and get hung up on being incredibly precise about this***. I've got a good buddy, we'll call him Michael (because that's his name), who was running some of the calculations to figure the optimal sub-exposure duration and was coming up with numbers like 3 minutes and 17 seconds. He had a nice library of 3 minute darks and was about to go and make a new one at 3 minutes and 17 seconds before I talked some sense into him. For starters, he'll never see the difference even if the extra 17 seconds puts him in a theoretically better place as the difference is trivial. Second, the calculations are all based on estimates of values and, if these estimates are a bit off, so are the calculations. So, don't sweat the details too much here. So, if we end up at a place in a future article saying that 1"/pixel is about as high a resolution as you should go and you're running 1.1" or 0.9", don't go out and swap cameras, scopes, reducers, etc. to get to exactly 1". Those differences are trivial and the suggestion is based on estimates of seeing, tracking,

etc. But, if you're at 0.1" per pixel... well, time to come up with a better solution.

Sources of Signal and Noise

Photons are hitting your CCD (or CMOS – it really doesn't matter here) and its job is to count those photons. You get an analog measure of the photon count out of each CCD well and this gets passed to an analog-digital-converter (ADC) whose job is to turn that analog measure into a number. That number is a scaled version of the number of photons that hit the CCD well, plus or minus something. That plus or minus is the noise or variability in our estimate of the true number of photons that hit the pixel.

As discussed last time, there are a few sources of signals and noise to contend with here. Let's briefly recap them in terms of what goes into that CCD well:

1. We have our photons from our faint fuzzy. For example, we may have 10 photons per second coming from the DSO. These are good. We like them. Let's call them the real signal.
2. We have photons coming from the sky but not from our DSO. These are skyglow-based photons. Perhaps we have 100 photons coming from skyglow into our image per second. So far, skyglow isn't really all that bad (but it will be in just a second). It's signal, but signal that we don't really care about.
3. We have photons coming from the dark. OK, they're not coming from the dark *per se*, but even when it's totally dark we have these getting picked up by our CCD well at a certain rate. The rate will vary a bit from pixel to pixel (and some pixels have a very high rate – hot pixels). But, if we consider just one pixel at a time, this rate is constant. Let's call it 0.03 per second for a cooled camera. Again, it's signal (it made the image brighter), but signal we'd just clip off in processing.

For what it's worth, these numbers aren't entirely made up. I just looked at a 1-minute image of M51 taken from my urban skies on a QSI 540 and a Borg 101 ED f/4. The background was ~4200 ADU, a spot of the arm was at ~4500 ADU, the average dark frame was ~215 ADU, and the average bias frame was ~214 ADU. The gain here is 0.8 e-/ADU and peak QE is 55% with 50% or so being a decent average value for the range recorded. So, if we remove the bias, that's ~4000 ADU of background, ~4300 of background plus target (~300 then of just the target when we remove the background as well), and ~1 ADU of dark current. Convert this into photons (multiply by 0.8 to get into electrons and then by 2 to get into photons that hit the detector, only half of which were recorded) we get 6400 photons per minute from skyglow, 480 per minute from my target, and 1.6 per minute from dark current (aka 107 per second, 8 per second, and 0.03 per second).

Now this isn't it. If it were, we'd be dancing. But we also have:

4. Shot noise from the target. Photons don't arrive at a nice, constant rate. So, there is an uncertainty here of \sqrt{N} in our signal (where N is the photon count). In our example, in one second this is about 3.2 photons of variability (noise).
5. Shot noise from the skyglow.
6. Shot noise from the dark current.
7. Read noise

The Pixel SNR Equation

Since SNR is a ratio (signal to noise ratio), we've got two parts to consider. First, we've got the "signal". That's got the first three parts laid out above. The real signal we want from our DSO, the skyglow, and the dark current.

$$\text{Total_Signal} = \text{Duration} * (\text{Target} + \text{Skyglow} + \text{Dark})$$

We can think of this in terms of photons or electrons (recall only about half my photons were captured) and here, we can even think of it in ADU. All this says is we've got three sources of "flux" (streams of photons coming into our detector at certain rates). Add them up you've got a total rate per second. Multiply that by your exposure duration and you've got the total signal. No rocket science here.

$$\text{Shot noise} = \text{sqrt}(\text{Total_Signal})$$

Recall from above and from [Part 1](#) that we've got this uncertainty in the actual photon count that goes along with the the square root of the expected number. So, we've got that as part of our noise.

$$\text{Read noise} = \dots \text{well} \dots \text{read noise}$$

Every time you read off an image, your camera will inject some number of electrons of read noise.

Now, we need to combine these two sources of noise. We do that by taking the square root of each of them squared. So, we have $\text{sqrt}(\text{sqrt}(\text{Total_Signal})^2 + \text{Read_Noise}^2)$ which simplifies to:

$$\text{Noise} = \text{sqrt}(\text{Total_Signal} + \text{Read_Noise}^2)$$

For now, this is all the math we need. Let's stick some numbers in here and let's work in electrons instead of photons. Based on what I recorded from that image of M51 we have:

| Signal source | e-/sec |
|---------------|--------|
| Target | 5 |
| Skyglow | 50 |
| Dark current | 0.02 |

This makes our equation for the signal be:

$$\text{Total_Signal} = \text{Duration} * 55.02$$

If we plug in 60 seconds here, we get a total signal of ~3300 electrons. (Recall my camera has a gain of 0.8 e-/ADU so if we divide this by 0.8 we get a bit north of 4000 ADU, nicely back where we started give or take the rounding I did for nice, easy numbers).

Now, a typical value for a good camera is about 8 e- of read noise. If we figure our noise in this 60 second exposure we get:

$$\text{Noise} = \text{sqrt}(3300 + 8^2) = \text{sqrt}(3364) = 58$$

At this point, you might be tempted to just say the SNR is $3300 / 58$ or about 57 and be dancing as that's a very good SNR for a single frame. But, remember that a good bunch of this "signal" is signal we don't care about. We're going to take our Levels tool or our B slider

and “reset the black point” to remove all that skyglow (and dark current). We're throwing that away as we don't care to know just how nicely bright our skies are. Our real signal is the signal from the DSO. That is 60 seconds * 5 e-/second or 300 electrons. Our SNR here is now 5.2. This is better than an SNR of 1 (the proverbial “you've got squat” line), but it's not 57.

Here's the shot (raw, no pre-processing, just simple linear stretching) so you can see what this level of SNR looks like. Something is there. You can see M51's arms riding above the noise. But this is no prize winner as it stands. Our total SNR will have to be a good bit better than this in the end.

Another case


The data from that M51 shot could leave you thinking that dark current and read noise don't mean diddly (technical term). The read noise was 8 e- and the shot noise from the dark current was nil. The shot noise from skyglow was almost 55 e-. Since the total noise was 58 e- (remember, the way to add noise isn't just to add them but to square each, sum, then take the square root), that 55 e- of shot noise from the sky is pretty daunting. Let's consider another case though and look at some data shot through an H-a filter using the same camera and scope. Here, the background sky near the Cone Nebula reads about 440 ADU in a 10 minute exposure and the nebula itself is up around 505 ADU. Dark frames are at about 216 ADU still. If we pull the 214 ADU bias signal out we're left with 2 ADU coming from the dark current (1.6 e-), 224 ADU coming from skyglow (179 e-) and 65 ADU (52 e-) coming from the target (or at least one spot on the target).

So, the signal we care about is 52 e- and the Total_Signal is 233 e- (target + skyglow + dark). Our noise is then:

$$\text{sqrt}(233 + 8^2) = 17.2$$

The SNR then in this image is 3.0 (a bit worse than before, but much of this depends on where you look in the image). The real point here, though is to consider where the noise is coming from. Overall, there is **a lot less noise here than in the earlier case**. In fact you can see it in the image. Take a look at the background of the Cone here and of M51. Notice all the extra noise in the M51 background? That noise is from the skyglow's shot noise.

In the Cone shot here, we have 17.2 e- of noise and before we had 58 e- of noise. Last time, if we ignored read noise, we had 57 e- of total noise. Removing the shot noise from the target, we still had 55 e- of noise just from the skyglow. This, out of a total of 58 e- of



noise. In this case, who cares about the read noise? Our SNR considering the read noise was 5.2. Pull it out and the SNR is 5.3.

Here, though, we have a different story. Here, our total noise term is 17.2. If we just look at the shot noise overall it's 15.2 e-. Shot noise from the sky is 13.4 e- and the read noise is 8 e-. These are a lot closer. SNR with the read noise in there is 3 and pulling it out, it's 3.4. Clearly, it's more of a factor now. If we could eliminate it, our SNR would go up 13% whereas previously it would go up under 2%.

This is why you may have heard people say that, oddly enough, a pristine camera, low read noise, low dark current, etc. are more important at dark sky sites. This is also why the various calculators for optimal exposure durations say you should use longer exposures at dark sky sites than under urban skies. You might think that the urban skies present more of a challenge (they do) and therefore need a cleaner camera. The truth is that the urban skies themselves are giving you so much noise that you'll be hard-pressed to take advantage of the lower noise offered by better cameras. (Rather large caveat – all this assumes that, when comparing cameras, both have “well-behaved” noise. Many don't and this will be covered in a later article).

Playing with the numbers

At this point, I hope you're sitting there scratching your head and saying “Hmmm...”. Perhaps there are questions like: What kind of SNR do I have in that galaxy arm? What kinds of SNRs give me shots I like? What would happen if I moved your scope to a darker site? Is that filter really helping me? Does my buddy who claims his light pollution is worse than mine have a leg to stand on or is his SNR no worse than mine? What would happen if I got a camera with lower dark current?

Playing with the numbers lets you get at these issues. I call it playing because it should be fun. Again, don't stress over the details. If you sample a spot in the image and see it's at 1017 ADU in one place but 1024 right next door, don't sweat it. 1020 is a fine number. The goal is to get a handle for how some of these things behave and to see where the big wins and losses are and what may not matter. I've put together a spreadsheet in Excel format to let you start to get a handle on this. You don't actually need Excel and, in fact, I didn't put it together in Excel. OpenOffice is a perfectly fine stand-in (and is free) and Google Docs can do this as well (although the graphs need some love if you use Google Docs). Anything that can work with a .xls file should do.

Spreadsheet basics

To use the spreadsheet, start by finding a light, a matched dark, and a bias. Open up the light and you'll be looking for two values. One is just the intensity of the skyglow background and the other is the intensity of the target (pick some area that's not the galaxy's core, but rather is that arm or bit of nebulosity you want to pull out). These should both be from a raw light frame (no pre-processing). Enter in the intensity of the dark frame and of a bias frame and enter in the exposure duration for your light and your dark (make sure your dark frame is at least as bright as your bias frame). Finally, enter two parameters about your camera that you should be able to get from the camera maker (or any of my reviews): the gain and the read noise.

The bottom half of the first sheet shows the calculations I've performed for the samples here.

By default, it loads up with sample data from my urban back yard. You'll see that the SNR of the part of the arm in M51 I sampled is about 4 (the numbers used in the example above were rounded a bit, but the spreadsheet has the actual values from my yard). Supplied in there are also values from the L channel of a [shot of the Veil nebula I took](#). This was out in Julian at my buddy Chuck Kimball's ["Big Cat Cabin" B&B](#). He's got much darker skies as you'll see. It'd take 20 minutes of exposure to reach the same level of background glow out there that it took me to hit in only 1 minute!

Fun with graphs

The fun doesn't end with the simple calculations. Included in the spreadsheet are several other sheets with graphs to show various "what if" scenarios. The first extra sheet (*Skyglow effect*) shows what would happen to your single-frame SNR if you changed to a brighter or darker sky site. As you can see, skyglow has a pretty bad effect on SNR.

The next (*Duration effect*) shows the effect of exposure duration on SNR in a single frame. Certainly, the more data you get the better as the SNR goes up with exposure duration. Recall that read noise is always there and the skyglow's shot noise goes up with the square root of the duration. The target's signal goes up with duration so that's why the single-frame SNR continues to rise. Of course, you could have taken four shots at 5 minutes and stacked them rather than one at 20 minutes. That effect is plotted in another sheet (*Sub-exposure length*).

Next up is one called *Stacking*. It shows what will happen to your SNR as you take multiple exposures (at your current settings) and stack them (again, assuming well-behaved noise on your camera's part). Here, you'll see just why we stack! You can also, though, do things like look at the SNR of a single 10 minute exposure on the *Duration effect* sheet and compare it to 10 frames at one minute here (or whatever works out given the exposure time from your data – the default data are from 1-minute shots).

Finally, the last one, called *Sub-exposure length* lets you enter a fixed total imaging time (default is 180 minutes) and examine the SNR in a stack as a function of the duration of individual sub-exposures.

Note, in each of these, there are a number of simplifications. Most notable is the fact that we're assuming either no dark subtraction or infinite SNR on your darks. Likewise, your flats are assumed to have no noise. Each of these steps introduces noise into your image in the process of correcting artifacts. Despite these limitations, they can give you a good feel for how things affect your SNR.

Sample "What if?" Experiment: Effect of Skyglow

Above, I gave the example of the Veil shot at Chuck's site in Julian and my yard here. We can use this sheet to see just what would happen if I tried this same M51 shot that I did from my yard out in Julian. This serves as a sample of the kind of thing you can do when you start playing with the numbers here and it lets you see just what you've traded off.

If we use my original parameters, we see that my skyglow flux is 53 e-/second and the target flux there on M51's arm was 4 e-/second. Let's transport my rig to Chuck's site where the skyglow flux is only 2.36 e-/sec. We can't just paste the parameters from the Veil in there and be done as that's the Veil and we're shooting M51. So, we need to twiddle a bit. If we first adjust the skyglow background down to 390, we get close to Chuck's skyglow flux. Of course,

the target + skyglow won't be up at 4500 where we started. We want 300 ADU of target signal (where we were before) which should lead to the target flux of 4 e-/sec (it's not like M51 is any brighter out there than here – it's just that his skies are darker than mine). So, if we set the target signal in the top there to 690 (390 + 300) we get back to the right flux from M51.

Take a look at the SNR how here. You'll see we have an SNR of 11.38 in our section of M51's arm there from the dark site. Recall that this same spot had an SNR of 4.06 from my urban yard. We're talking almost 3x the SNR here and recall that you need 4x the data to get 2x the SNR ($1/\sqrt{N}$). Paste back in the original values and then head on over to the *Stacking* sheet and you'll see it takes about 8 stacked frames to get to this SNR (look at the graph or the table on the right). So, this says we need 8 times as much data from the urban site to equal the dark sky site.

Paste back in the dark-sky values and head over to the *Sub-exposure length* sheet and you'll see now that we're over 160 on our SNR with most sub-exposure lengths. That's what 3 hours at Chuck's site will do. Now, what kind of total exposure time do I need from home to hit this? Paste back in my yard's parameters. In 3 hours, we're below 55 on the SNR. Now, up the *Total duration* field here. Start upping my total exposure time and look at the effect. You'll need to get to about 1550 before my SNR matches Chuck's cabins. That's about 26 hours. Ouch! When doing LRGB work, I really should get out there to his B&B (or to my club's dark sky site). (Note, how 26 hours is about 8x the 3 hours and recall the 8x factor we came up with just above?) Armed with this we can now start to ask other questions. Hmm... what about when I shoot Ha? Get the idea?

Conclusions

Coming down the pike are going to be articles on things like sampling and how you can view this as a trade-off between SNR and resolution. Along for the ride in that one will be a discussion of f-ratio and just what it does and doesn't mean to us. Color vs. mono cameras will be covered as well. We'll go a bit deeper into things like how pre-processing can add noise into your images as well. Each of these requires a handle on SNR in one pixel. Along the way, we'll also cover how to measure things like read noise or system gain in your camera without any fancy equipment. Hopefully, the first two articles have filled in any gaps in your knowledge here and you're well on the way.

A few things to note before closing. First, is that one of the great things about [Cloudy Nights](#) is that we've got such great discussion boards. There are threads there for each of the articles and I'm reading them. **If something doesn't make sense, speak up!** I'm a professor by day and try to convince my students that if they've done their job by coming to class prepared and paying attention and still don't get something, it's my fault not theirs. The same kind of logic applies here even though the roles aren't exactly the same. If you're reading along and trying to get it but things still don't compute, odds are I skipped a step and/or wasn't as clear as I could have been. So, speak up!

Second, Mike and Mike here have given the A-OK for me to post PDFs of these up on [my personal website](#). So, if you're looking to print out hard copies for some offline reading (I won't ask where), that's a good place to start.

Craig