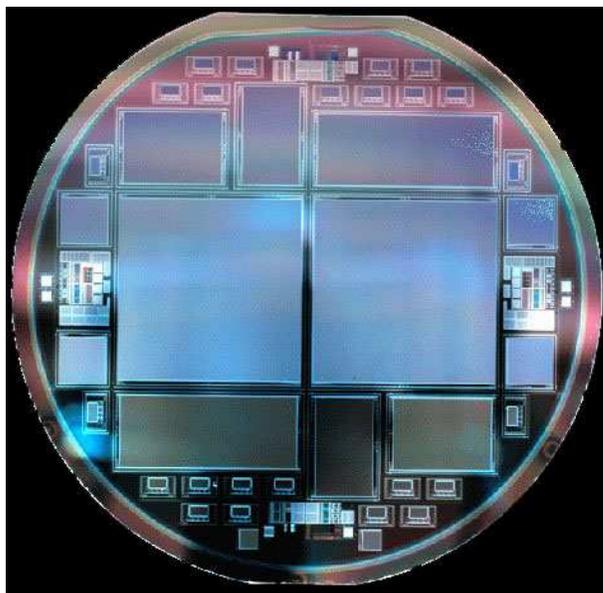


Signal-to-Noise in Optical Astronomy

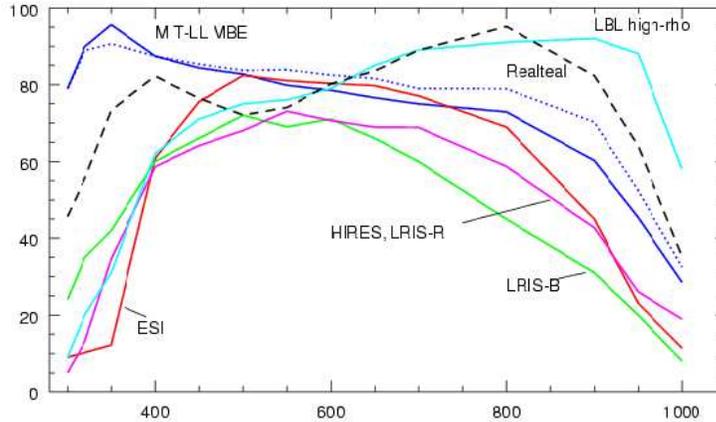
1 CCDs



1.1 CCD overview

Going to start with the easiest case: CCD photometry of point sources. First a bit about CCDs (charge-coupled devices). Ian Mclean’s book contains a fairly detailed account of the history off the development of CCDs. Electronic detectors came on the astronomical scene in the late 1970s. By the end of the 1980s, CCDs had won the battle of the various electronic devices and had replaced photographic plates for all applications except for very large field ($> 15'$) imaging. Early devices had 30μ pixels and were 400×400 to 800×800 pixels. In 2004, the first of the 4096×4096 devices are becoming available.

One of the principal reasons CCDs are so popular is they are very efficient with photons. The ratio of incoming photons to detected electrons is the quantum efficiency and for some wavelengths and some detectors it approaches 100%. For photographic plates the QE was generally less than 5%. Putting a CCD camera on your 1m telescope gave it the same capability of the 4m telescopes of the photographic era.



Real and imaginary QE curves for current and future CCDs in instruments at the Keck Observatory

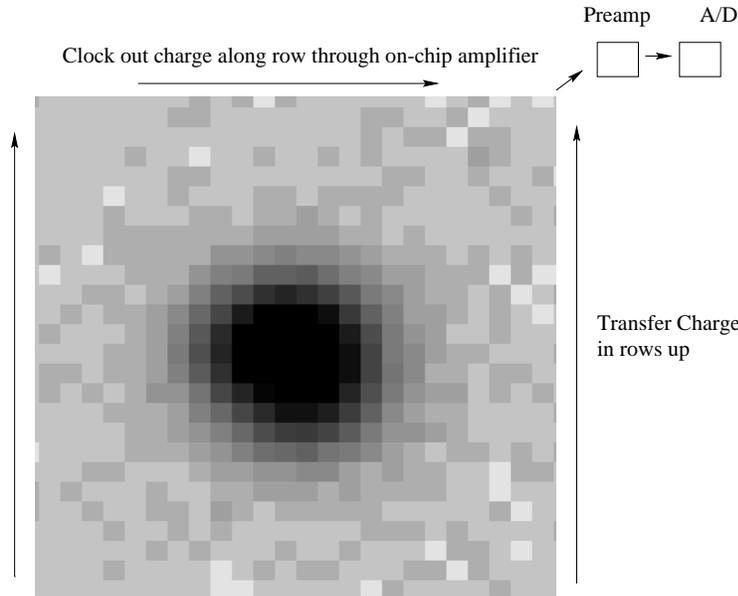
1.2 CCDs: Some Details

Photons \Rightarrow e^- : CCDs are silicon semiconductors. A structure is embedded in the device that creates localized electric corrals so that when a photon is absorbed in the silicon and pops an electron up into the conduction band, it is spatially localized in a “pixel”. When the device is to be read out, the top row is clocked out (serially) through a preamplifier, amplifier and analogue-to-digital converter, then all rows are shifted up and the new top row is clocked out serially. By keeping track of the timing the two-dimensional array can be reconstructed after the full array has been clocked out.

$e^- \Rightarrow$ **DN:** The accumulated e^- in each pixel gets represented by a number in “units” of a DN (Data Number) or ADU (Analogue-to-Digital Unit). Until recently, the typical A/D converter used 15-bits to represent the data numbers and 1-bit for the sign. The dynamic range was therefore $\pm(2^{15} - 1) = \pm 32,767$. Since there are no sources of anti-photons, it makes more sense to use all 16 bits for positive data numbers and many systems now have a dynamic range for data numbers of: $2^{16} - 1 = 65,535$.

As a general rule, the number counts in any one pixel is NOT the number of photo-electrons detected but related to e^- through the e^- /DN conversion:

Inverse-Gain: (almost everyone will refer to this as Gain and use something like “G” to represent it) is this conversion factor in strange but dimensionless units of e^- /DN. Will use the symbol \mathcal{G} for inverse gain in units of e^- /DN



CCD image of star

This is often something that you have the option to change. How do you pick an inverse-gain? Depends on the nature of the data you intend to acquire. The tradeoff is usually between dynamic range and effective read noise although there are with some chips other subtle issues. For most broadband photometry, the read noise makes only a tiny contribution to the total noise. In this case it makes the most sense to set the gain to:

$$\mathcal{G} = (\text{full well capacity of pixel}) / (\text{ADC saturation})$$

For LRIS, the CCD pixels “fill up” at around $300,000e^-$, the highest data number is 65,535 so a sensible inverse-gain is $4.58e^-/\text{DN}$. Using the inverse-gain of around 2 that is available, the ADC saturates long before the pixels well have filled, giving up a factor of 2.3 in dynamic range per exposure.

2 Noise

2.1 Random (Statistical) Noise

For the discussion that follows be thinking about an aperture with radius r centered on a point source on a CCD.

Here is the key to the symbols I’ll be using:

R_*	count rate from star	e^-/second
R_{sky}	count rate from background	$e^-/\text{second/pixel}$
t	exposure time	seconds
r	radius of aperture	pixels
n_{pix}	number of pixels in aperture	$\pi \times r^2$
\mathcal{G}	inverse-gain	e^-/DN
D	dark current	$e^-/\text{pixel/sec}$

Noise in photon counting: “Poisson” or “shot” or “statistical” noise associated with counting statistics (appropriate to counting the number of photons that arrive each time interval). This comes from a binomial distribution with $N \rightarrow \infty$ and “p” $\rightarrow 0$. The events (detections of photons) in one interval are independent of the events in the previous or subsequent time intervals. For the Poisson distribution, the standard deviation is just $\sigma = \sqrt{N}$ where N is the number of detected events.

Don’t forget that this applies to the detected e^- – not counts (unless $\mathcal{G} = 1$).
 In e^- , $\sigma = \sqrt{N \times \mathcal{G}}$. In DN, $\sigma = \sqrt{N \times \mathcal{G}}/\mathcal{G} = \sqrt{N/\mathcal{G}}$.

Don’t forget that if you have averaged frames together, the inverse-gain changes. If n frames are averaged together, the effective inverse-gain of the averaged frame is $n \times \mathcal{G}$.

2.2 Statistical noise sources in CCD images

2.2.1 Read Noise (RN):

For many CCDs/dewar electronics systems the read noise can be as low as a few e^- . There is a tradeoff between readnoise and “dwell time” such that you can decrease the time to clock a chip out at the expense of higher read noise. For broadband imaging – go for the faster readout. For the case of n frames averaged together, RN will decrease by a factor \sqrt{n} .

The way RN is defined, it implies that $(RN)^2$ readout e^- were measured and within the aperture of radius r the noise will:

$$\sigma_{RN} = \sqrt{\pi \times r^2 \times RN^2} = \sqrt{n_{pix} \times RN^2}$$

In some cases, there can be a component of “digitization noise” (undersampled readout noise). If the RN is $5e^-$ and the inverse-gain is $\mathcal{G} = 10e^-/\text{DN}$ then the read noise is undersampled or poorly resolved. The effective read noise will be $[RN^2 + (0.5 \times \mathcal{G})^2]^{1/2}$. For broadband imaging this is rarely important because usually observations are very quickly sky-noise or source-noise dominated.

2.2.2 Dark Current noise:

For CCDs, dark noise is usually not an issue. For near-IR detectors it can still be significant.

For a dark current of D in e^- /second/pixel, the noise in the aperture is simply:

$$\sigma_{dark} = \sqrt{D \times n_{pix} \times t}$$

2.2.3 Shot noise in source:

This is the Poisson noise based on the number of detected photons from the source. If source provides R_* photo-electrons per second, then the total number collected during the integration is simply $R_* \times t$ where t is the integration time and the standard deviation is:

$$\sigma_* = \sqrt{R_* \times t}$$

Where do you find values for R_* ? Hopefully at the Web site describing the properties of your instrument. You could calculate this value given a true flux in energy units of a star of given brightness, the size of your telescope mirror and the total effect throughput of your mirrors, filter and CCD. But, it is simpler and more accurate to measure it with the system you will use. The usual way it is presented is normalized to a U=B=V=R=I=20 mag star in counts per second at an airmass of 1 in the available filters. It is then simple to figure out the expected count rate for an object of magnitude m :

$$R_*(m) = R_*(20) \times 10^{0.4(20-m)}$$

As an example, here are the approximate numbers for LRIS on Keck:

Count Rates for B=V=R=I=20	
B —	1470 e^- /second
V —	1521 e^- /second
R —	1890 e^- /second
I —	1367 e^- /second

2.2.4 Shot noise in the sky:

The uncertainty in the sky brightness per pixel is

$$\sigma_{sky} = \sqrt{R_{sky} \times t}$$

and within the aperture simply

$$\sigma_{\text{sky}} = \sqrt{\pi \times r^2 \times R_{\text{sky}} \times t} = \sqrt{n_{\text{pix}} \times R_{\text{sky}} \times t}$$

What is the sky brightness? It is different at different sites, in different bandpasses and of course varies with the phase of the moon and distance from the moon. From NOAO Newsletter #10, Alistair Walker measured for the skies above Cerro Tololo:

Sky Brightness (mag/arcsec ²)					
lunar age (days)	U	B	V	R	I
0	22.0	22.7	21.8	20.9	19.9
3	21.5	22.4	21.7	20.8	19.9
7	19.9	21.6	21.4	20.6	19.7
10	18.5	20.7	20.7	20.3	19.5
14	17.0	19.5	20.0	19.9	19.2

This is an interesting table. First, the color of moonlight is clearly pretty blue. There is a tremendous dependence on moon phase in the U-band (5 magnitudes = factor of 100), and a little less than a factor of 2 between new and full moon in the I-band. Also, the sky at new moon is relatively faint in the blue filters, and quite bright in R and I.

How does this translate into R_{sky} (photo-electron rate per pixel)? You need to know the plate scale of your detector and the effective QE of the system – most straightforwardly in units of e-/magnitude/filter as discussed above.

Let’s call the numbers from the table up there SB_{sky} (surface brightness of the sky) and the linear pixel scale of the detector SCALE ("/pixel) which for LRIS is 0".218/pixel.

So, the area of one pixel is $SCALE^2 = 0.0475(\text{arcsec})^2$ for LRIS. This is also the ratio of the flux in a pixel to the flux falling on a square arcsecond. In magnitudes this will be:

$$m_{\text{pixel}} = m_{\text{arcsec}^2} - 2.5\log(SCALE^2)$$

For LRIS, add 3.308 to each of the entries in the sky brightness table and you get the sky brightness in magnitudes for LRIS pixels. To convert to counts per pixel from the sky, use this value in

$$R(m) = R(20) \times 10^{0.4(20-m)}$$

For LRIS, the counts per second in the sky are pretty high. For example, in the R-band at new moon:

$$R_{\text{sky}} = 1890 \times 10^{0.4(20-24.21)} = 39.1e^- / \text{sec/pix}$$

So, the sky noise, $\sqrt{39.1} = 6.25e^-$, is equal to the readout noise in just 1 second of integration time. In a 10 minute-long exposure you have 1/5 of the dynamic range of the system filled with sky photons.

2.2.5 An interesting side-trip – sky brightness:

For ground-based sites much of the sky brightness shortward of 7200Å is due to emission lines from city lights. Redward of there it is mostly OH emission in the upper atmosphere.

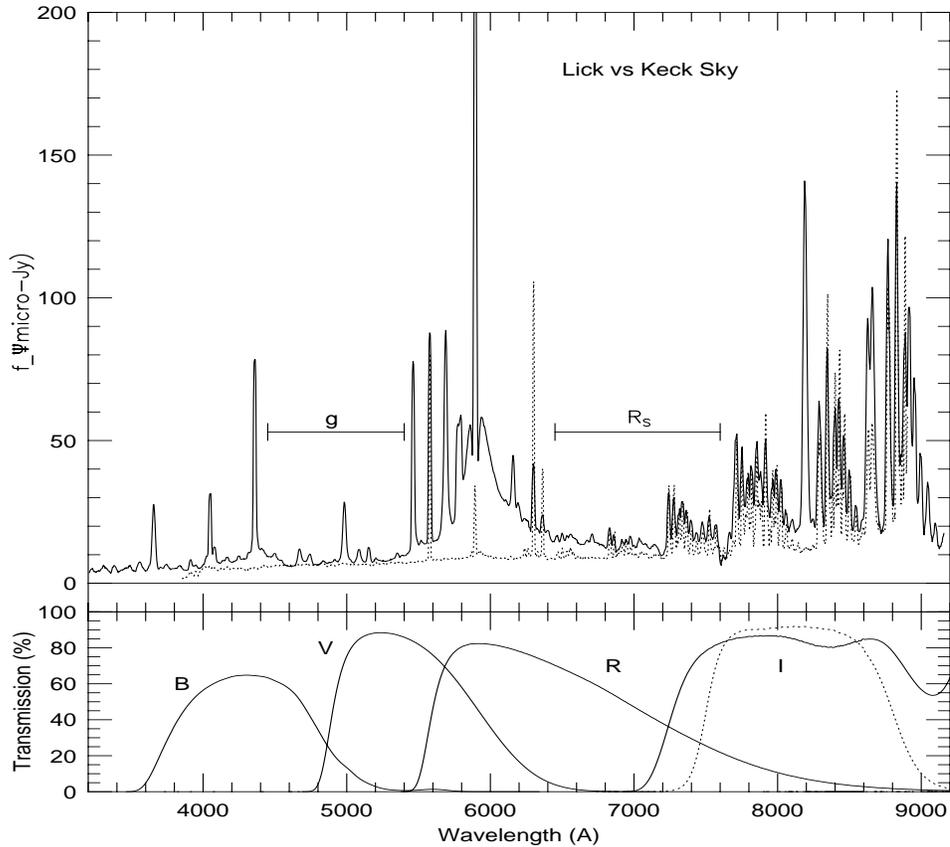


Figure 1: Spectra of the night sky from Mt. Hamilton (dotted line) and from Mauna Kea (solid line). The transmission curves in the lower panel are measured for the Harris B,V,R,I filter set from KPNO. The dashed line is the Thuan-Gunn I filter.

The standard *B, V, R* filters don't make a lot of sense for a site like Mt. Hamilton or any ground-based site in the future. The Thuan-Gunn set was designed to miss the worst of the usual emission lines from city lights, although when they were designed, the high-pressure sodium lights were not available and the wings of the Na "D" lines were not so extended. The T-G "r" filter works less well than the "Spinrad" R.

3 The Standard S/N Equation

All these sources of noise are independent so they get added in quadrature. The general S/N equation is:

$$S/N = \frac{R_* \times t}{[(R_* \times t) + (R_{sky} \times t \times n_{pix}) + (RN^2 + (\frac{G}{2})^2 \times n_{pix}) + (D \times n_{pix} \times t)]^{1/2}}$$

3.1 Conversion from S/N to uncertainties on the magnitude scale

Usually errors in astronomy are expressed on the logarithmic magnitude scale. To go from S/N to magnitudes errors consider:

$$\begin{aligned} m \pm \sigma(m) &= C_0 - 2.5\log(S \pm N) \\ &= C_0 - 2.5\log[S(1 \pm \frac{N}{S})] \\ &= C_0 - 2.5\log(S) - 2.5\log(1 + \frac{N}{S}) \\ \sigma(m) &= \pm 2.5\log(1 + \frac{1}{S/N}) \end{aligned}$$

Note, that magnitude errors are often taken to be fractional errors and measurements like $V = 31.9 \pm 0.1$ are claimed to be 10% photometric accuracy. This is close to but not quite correct:

$$\begin{aligned} \sigma(m) &= \pm 2.5\log(1 + \frac{N}{S}) \\ \sigma(m) &= \pm \frac{2.5}{2.3}[\frac{N}{S} - \frac{1}{2}(\frac{N}{S})^2 + \frac{1}{3}(\frac{N}{S})^3 - \dots] \\ \sigma(m) &\approx \pm 1.0875(\frac{N}{S}) \end{aligned}$$

3.2 Systematic Errors: What has been left out of this equation?

The S/N equation above is really the **best** you can ever hope for. There are many other opportunities to degrade the S/N.

There are also many “analysis errors” that contribute to uncertainty in the final value derived for a measured quantity. In many cases these lead to systematic errors rather than random ones. Some of these could be the source of an entire lecture on their own.

1. **Bias level correction:** There is a positive constant added to all pixels that is usually determined from the mean of an overscan region. You want to determine this value and subtract it from your data, but the consequences of getting it wrong by a little bit are small.
2. **Bias structure correction:** Once upon a time this was an issue with some chips showing a distinct, stable structure that needed to be subtracted out or, for some chips that were clocked in sync with 60 hz there was a stable pickup noise pattern that needed to be subtracted off. As this is an additive correction, you can get it pretty wrong and it won't make much difference
3. **Flat-fielding:** This is a tricky business and can never be done perfectly. For many widely used systems, it is hard to flat-field to better than 1 part in 100 on large spatial scales. So, if a given star is measured on different parts of a chip, your answer will change by this amount.

This can often be a significant, unquantified source of error.

4. **Interpolation errors:** When calculating the number of detected e^- within a circular aperture, you have to do something about the fact that you have square pixels...
5. **CTE problems:** Charge transfer efficiency is measured in “nines” – “five nines” means 0.99999 efficiency per pixel shift. Sounds pretty good, but think about how many shifts are required to get the last pixel clocked out. If a fraction f of the charge is lost per shift, then after clocking out the fraction of the initial charge than remains is (for the last pixel in a 2048-row CCD):

$$(1 - f)^{2048}$$

For “five nines” – for the worst case pixel loss is 2%.

6. **Non-linearity when approaching full well values:**

A good thing to check with your first images at the telescope is the magnitude difference for stars of different brightness measured on frames with different exposure times. Take a long and short exposure of a field with at least several stars and difference the magnitudes through a fixed aperture. The deviations from a straight line in the plot of magnitudes derived from the short vs long exposures will tell you where the response becomes non-linear.

7. **Charge skimming:**
8. **Shutter Vignetting:**
9. **an unending list of possible complications...**

3.3 The Limiting cases

3.3.1 Bright sources

for bright sources:

$$S/N \simeq \sqrt{R_* \times t} \propto t^{\frac{1}{2}}$$

3.3.2 Background Limited

The usual case for broadband imaging is to be “sky limited”. This means that $\sqrt{R_{\text{sky}} \times t} > 3 \times RN$ (note these are both per pixel).

For faint sources in this case,

$$S/N \simeq \frac{R_* \times t}{\sqrt{n_{\text{pix}} \times R_{\text{sky}} \times t}} \propto t^{\frac{1}{2}}$$

3.4 Example: LRIS V-band imaging:

For LRIS, the readnoise is well-sampled and there is essentially no dark current. The basic S/N equation is:

$$S/N = \frac{R_* \times t}{[(R_* \times t) + (R_{\text{sky}} \times t \times n_{\text{pix}}) + (RN^2 \times n_{\text{pix}})]^{1/2}}$$

First, what is a sensible aperture to use? For making estimates for observing proposals and first-order planning of observing runs you can assume a Gaussian profile for point sources so

$$I(r) = I_0 \times e^{\frac{-r^2}{2\sigma^2}}$$

and the count rate within an aperture R (sorry for using R’s for radius and also for count Rates) can be found from:

$$R_*(r) = \int_0^R I(r) 2\pi r \, dr$$

Good exercise for the reader: For the case of a Gaussian profile, find the radius R that maximize S/N for a faint source in the sky-limited case. The answer is $R_{\text{optimal}} = 1.7\sigma$. Will use this. For a Gaussian, the Full-Width-at-Half-Maximum (FWHM) = 2.35σ so the best radius is at around $\frac{2}{3}$ FWHM.

If points sources really could be described as gaussians, you would be using only $\sim 90\%$ of the light, but to get the other 10% you pay too much in added sky noise.

It's easy to plug in and derive S/N as a function of integration time for an $V = 26$ magnitude point source at new moon:

$$R_*(V = 26) = R_*(V = 20) \times 10^{0.4 \times (20 - 26)} = 6.1e^-/\text{sec}$$

The counts in the sky are:

$$R_{\text{sky}} = R(V = 20) \times 10^{0.4 \times (20 - 25.1)} = 13.8e^-/\text{sec/pixel}$$

LRIS readnoise is $RN \sim 8e^-$.

Figure 2 shows the S/N vs exposure time for different FWHM values for the images. This assumes that you stick with the “optimum” aperture. Note the importance of image quality by comparing the exposure time at fixed S/N.

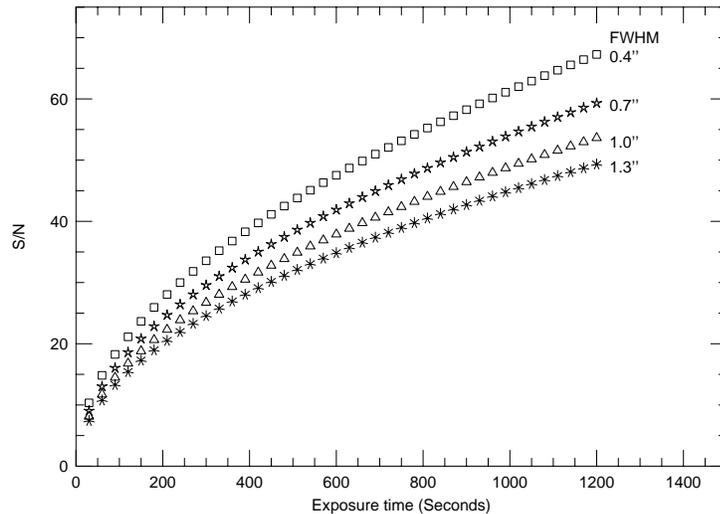


Figure 2: S/N vs exposure time for $V=26$ mag star using LRIS on KeckII plotted for different values of the FWHM of point sources.

For planning observations, the more typical case is to have a target S/N and solve for the integration time. Good exercise #2 is to fill in the easy steps to derive:

$$t = \frac{[SN^2(R_* + n_{\text{pix}} \times R_{\text{sky}})] + [(SN^4(R_* + n_{\text{pix}} \times R_{\text{sky}})^2 + 4 \times R_*^2 \times SN^2 \times n_{\text{pix}} \times RN^2)]^{\frac{1}{2}}}{2 \times R_*^2}$$

In Fig 3 it you can see that for bright stars (with LRIS this means $V < 21.5$) readnoise contributes to the noise budget. At around $V = 24$ the shot noise in the star and sky

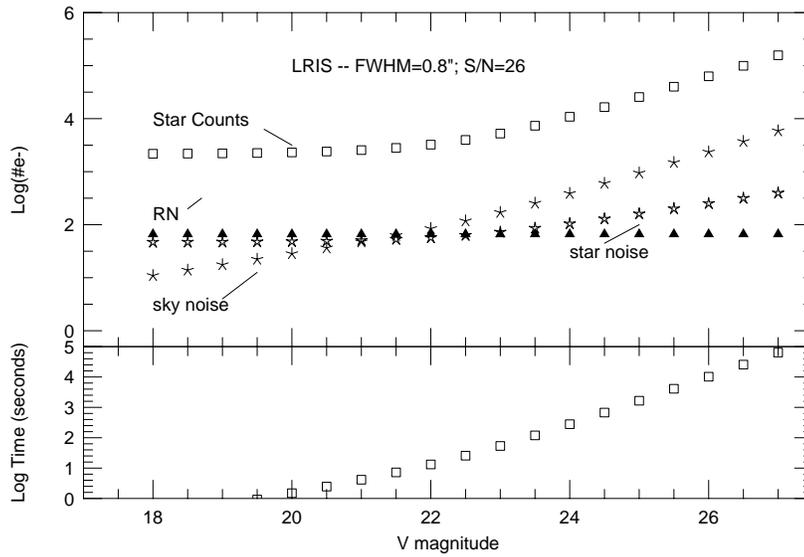


Figure 3: **Upper Panel:** The contribution of each noise source in $\log(e^-)$ at fixed S/N over a range of magnitudes. **Lower panel** Exposure time vs mag.

are about equal, fainter than that the sky dominates. This crossover depends on the filter (sky brightness) and on FWHM.

3.5 S/N in spectra

For spectra, the basic rules are the same but there are differences in the effective aperture. For each pixel in a spectrum, there are counts due to the object, the sky and counts due to the readout noise. At the wavelengths where there are sky emission lines, the fractional noise goes up significantly. In the final extracted spectrum (much more about this later), you will have combined several rows containing object counts and often the S/N is calculated after averaging in the spectral direction to be not S/N per row-summed pixel, but rather S/N per resolution element. The goal for the wavelengths that include sky emission lines is to determine the value of the sky line at the position of the object spectrum accurately enough so that after subtraction, only the shot noise from the line intensity remains behind.

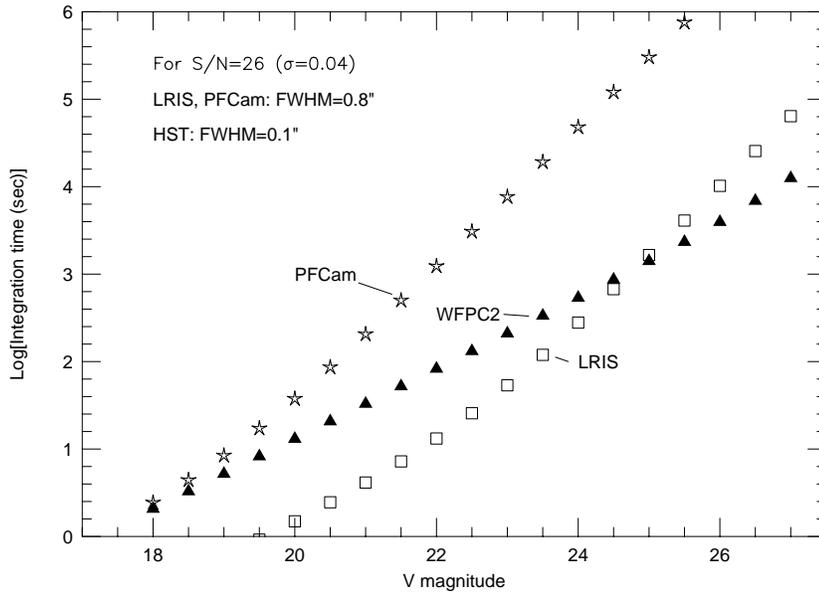


Figure 4: Comparison of the relative exposure times for the prime-focus camera at the Lick 3m, LRIS on Keck and HST/WFPC2. Don't take the HST numbers too seriously.

3.6 Sampling and Binning

You will occasionally run into claims that for $\text{FWHM} < 2$ pixels, images are undersampled because the Nyquist theorem. The Nyquist criterion doesn't really have much to do with CCD data as it is based on the idea of sampling a function in discrete intervals with delta functions (unlike CCDs which integrate over contiguous intervals). Nyquist was concerned with aliasing problems.

But, the $\text{FWHM} = 2$ pixels is a pretty good rule of thumb anyway. For smaller images (wrt pixels), it becomes noisy to do point-spread-function (PSF) fitting photometry, interpolation errors grow and you become more sensitive to intra-pixel QE variations.

It also doesn't make sense to significantly *oversample* your PSF mostly because of the added RN. There are two options for binning CCD data. The smart way is to recognize at the telescope that you have a poor match of image quality to pixel size and use "on-chip" binning. If you bin 2×2 *before* readout then you double (or is it halve?) your scale but the RN of the large pixel is the same as the individual smaller pixels. You can also bin the data *after* the fact in the analysis procedure, but this is not as smart as in this case you are only averaging over RN from the four pixels and the effective read noise is $RN_{\text{binned}} = \frac{4 \times RN}{\sqrt{4}} = 2 \times RN$.

With spectra, it is much more common to set work with the CCD binned in at least one direction.

If this CCD was binned in *columns*, then spectral resolution would be lost. On the other hand, binning by *rows* can make sense from the point of view of reducing the readout noise (assuming nothing will be made of the spatial resolution). The same rule of thumb works here as for point sources. Making a plot of intensity vs row number will show a not-too-different-from Gaussian distribution of counts along a column and a good goal is to try and make that distribution have a FWHM of between 2 and 3 pixels.

Some references:

1. Much of this was cribbed from Bill Harris's notes for a McMaster University Graduate Student Workshop/
2. Harris, W. E., 1990, PASP, 102, 949.
3. Howell, S. E., 1989, PASP, 101, 616.
4. King, I. R., 1983, PASP, 95, 163.
5. ASP Conf. Series Vol 8: "CCDs in Astronomy".
6. ASP Conf. Series Vol 23: "Astronomical CCD Observing and Reduction Techniques".
7. <http://www.noao.edu/kpno/docs.html> Check the Direct Imaging Manual.

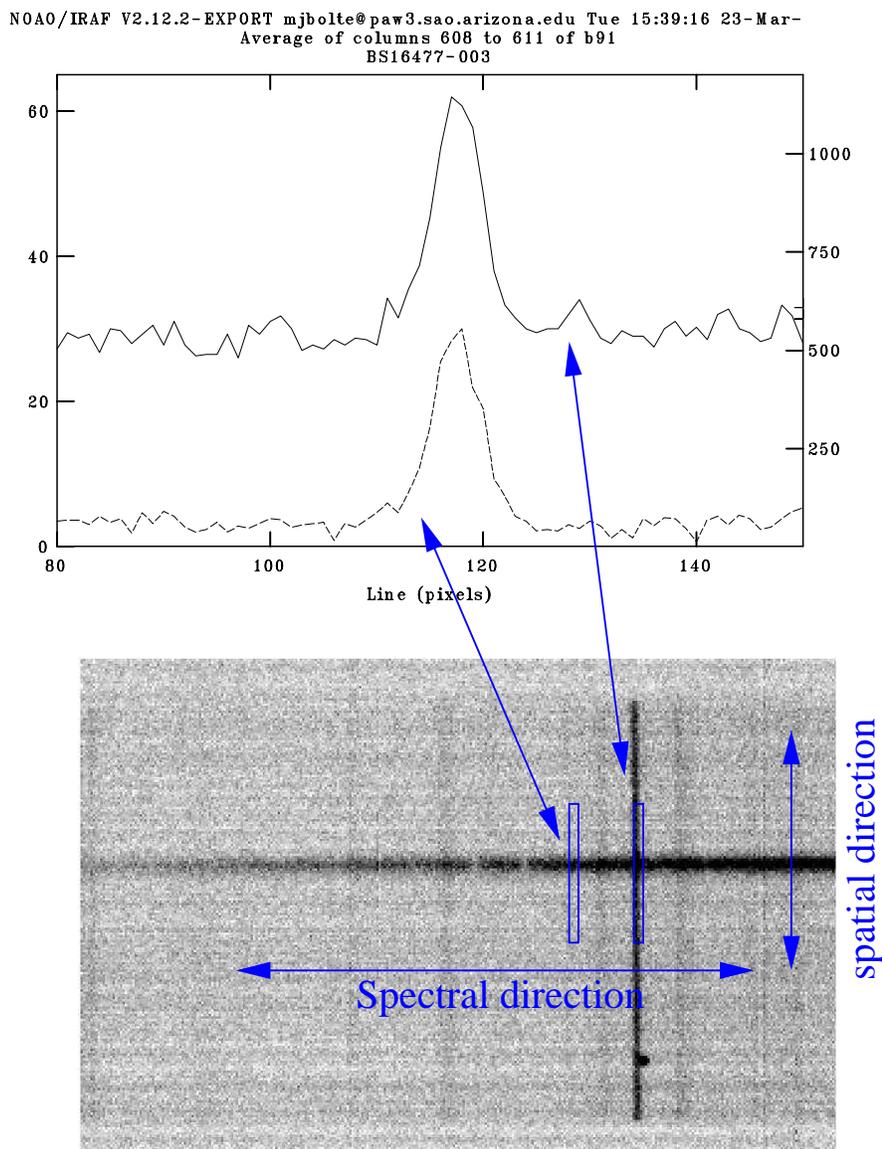


Figure 5: Two cuts across a spectrum. Under a sky emission line and not.

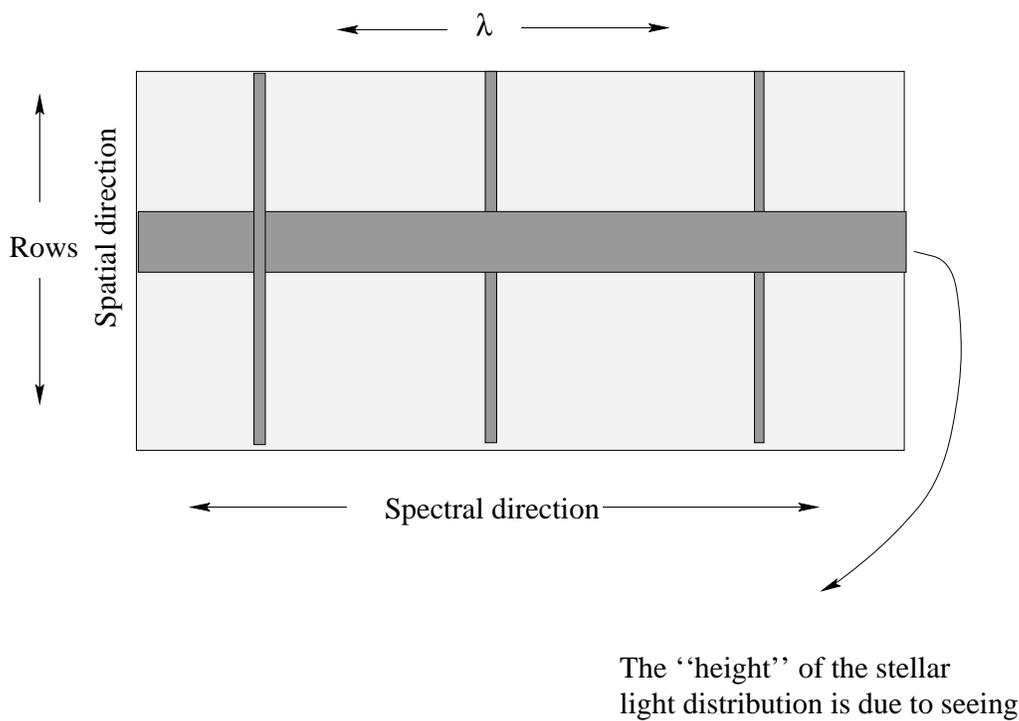


Figure 6: An example of a 2-D spectrum on a CCD. The dark rows down the center are supposed to be the object spectrum, the vertical lines are night-sky emission lines.