Lab 2: Detecting Faint Signals

Andy Friedman [1,2]
September 25, 2001

ABSTRACT

In this lab experiment, we use our previous knowledge of photon statistics to analyze the detection of faint signals. We first do a set of experiments with the PMT (Photo-Multiplier Tube) where we produce both signal and noise by modulating the constant LED output with a flashing LED. We then do a set of experiments with a CCD (Charge Coupled Device) Camera, which in contrast with the single pixel (picture element) PMT, takes images on an array of 512×512 pixels and thus has more complex noise properties. We use the results of these experiments to illustrate some new statistical concepts, which include the signal, noise, signal variance, variance of the mean, and most importantly, the signal to noise ratio (SNR). For the CCD camera specifically, we measure the fundamental properties of dark current, bias, readnoise and gain. As it is relevant throughout this report, we employ our previous knowledge of Poisson Statistics, as well as the newly acquired techniques of error propagation, and least-squares fitting.

1. Introduction: How Low Can We Go?

Although the statistical analysis we do in this lab report is general for any type of signal in the presence of noise, here we deal specifically with detecting faint light signals, as such techniques are clearly quite relevant to astronomical imaging. In real life experiments that involve detecting light signals from astronomical objects, it is reasonable to ask quantitatively,"How faint a signal can one detect in the presence of noise?', or equivalently, "To what precision can one measure a given signal?" The specific details of the answer to that question will of course depend on the inherent properties of the equipment used. For example, the Hubble Space Telescope can detect fainter signals than the Keck Telescopes by virtue of its orbit in space. But the bottom line is that the techniques discussed in this lab report are completely general and can in theory answer the question for an arbitrary experimental apparatus.[3] Later, when we measure the infrared light from a star in the presence of a noisy infrared background, the techniques discussed here will be crucial.

[1]e-mail: friedman@ugastro.berkeley.edu

[2]Lab Group 3: Andy, John, Shane, Henry, (Unofficial Lab Group: Lee, Jim, Christina, and Lindsey)

[3]In this lab we assumed that the systematic errors for the PMT part of the lab were negligible, whereas we corrected many of the possible systematic errors for the CCD part.

2. PMT Grudge Match: The Flashing LED vs. the Constant LED

In this section, we will discuss four separate experiments performed with the PMT apparatus, where the PMT detects photons from the constant and flashing LED's. Via these experiments, we will discuss the concepts of signal, noise, signal variance, variance of the mean, and most importantly, the signal to noise ratio (SNR). For clarity, the experiments will be called, "On and Off", "Signal To Noise Ratio", "Smallest LED Pulse Width Prediction", and "Turning Up The Noise". The "Signal To Noise Ratio" subsection will be by far the most detailed, as it will include discussions of theoretical SNR plot, and how to interpret the measured SNR results in terms of error propagation and Poisson Statistics. Relevant Plots with detailed captions will be included throughout. IDL programming specifics will not be discussed, but the source code for all relevant programs and generated data files can be found in my home directory in the folder /lab2/snrdata.

2.1. On and Off

One way to increase your ability to detect a signal in the presence of a noisy background is to modulate the signal by turning it on and off. By doing this, we can isolate the background and subtract it to get the signal. For convenience, we define:

$$ON = SIGNAL + BACKGROUND$$
$$OFF = BACKGROUND$$
$$SIGNAL = ON - OFF$$

where ON = the mean of the "on" counts and OFF = the mean of the "off" counts.

In this experiment, we do this by modulating the constant LED (from Lab 1) with a flashing LED that emits a pulse of a specified width. We then take 1000 samples of data from the PMT apparatus at a rate of 100 Hz (settings we keep for the whole PMT experiment), and separate the data into "on" and "off" bins. The constant LED is always emitting photons and thus provides us with a constant source of background noise, whereas the flashing LED shoots out photons periodically and provides us with our signal. The flashing LED is driven by a square wave from the pulse generator which is triggered by a 100 Hz clock pulse from the PC. The flashing LED ("on") gets triggered either when the square wave is going from hi-lo or lo-hi depending on the pulse generator settings.

An ambiguity arises because we can never be sure what portion of the square wave is occurring when we will send the command to ask for data. At that time, it could just as easily be going either from hi-lo or lo-hi. Thus the first entry in our data file could be either "on" or "off" no matter how we define "on" and "off" in terms of hi-lo and lo-hi transitions in the square wave. To solve this problem, the data gathering program allows you to "mark" the data by adding 2^{14} to all the "on" data, (knowing, of course, whether "on" is triggered by hi-lo or lo-hi). At the bit level, this is equivalent to flipping the 14th bit in the number's binary representation from 0 to 1. (i.e 1×2^{14}) Incorporating this into our code, we can be sure which data entries are "on" and which are "off".

For a given constant background noise, we can take a data set and plot the signal, where signal is the mean "on" counts minus the mean "off" counts, or using our earlier notation, $SIGNAL = ON - OFF$. If we lower the LED pulse width, we get a smaller signal which makes it intrinsically more difficult to detect.[4] This is illustrated by the plot below, where we plot "on", "off", and "on - off" for data sets with LED pulse widths of 5ms and 1ms. In the second set, the flashing LED pulse is only 1ms long, which means the PMT counts photons over a time period 5 times shorter than for the 5ms data set. And since we know from lab 1 that the LED is emitting photons at a constant rate, a shorter LED pulse width means that the mean signal will decrease because the PMT measures photons over a shorter time period.

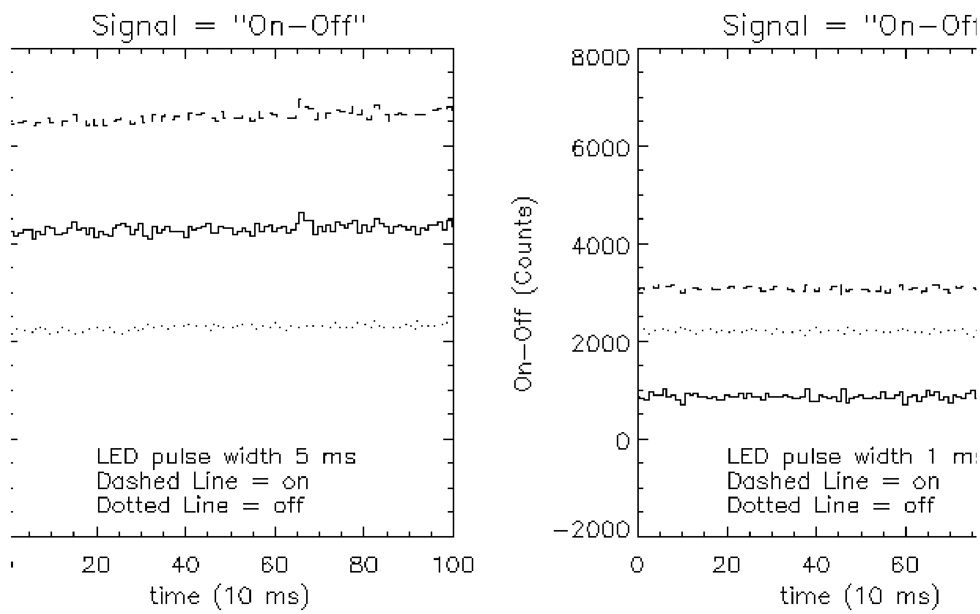

Fig. 1.— *Plots of "on" (SIGNAL + BACKGROUND), "off" (BACKGROUND), and "on-off" (SIGNAL) are shown for LED pulse widths of 5ms and 1ms. Notice that with a smaller LED pulse width, the mean signal counts drops (mean signal denoted by IDL pysm=10 symbols), since the PMT measures photons which the flashing LED emits over a shorter period of time.*

One might also ask what the uncertainty in the $SIGNAL = ON - OFF$, σ_{signal} is. Since the measurement of $SIGNAL$ depends on the measurement of both ON and OFF, we must propagate the errors in those two measurements to calculate σ_{signal}. Since we know that the uncertainty in the measurements of ON and OFF are independent (they come from different samples), the theory of error propagation tells us that we can add their variances in quadrature to get the variance in the signal, $\sigma^2_{ON-OFF} = \sigma^2_{ON} + \sigma^2_{OFF}$. From this we can define the variance of the mean or

[4] If we increased the background noise, the signal would also become harder to detect, but this will be discussed in detail later in the "Turning Up The Noise" section of the lab.

$VOM = \frac{\sigma^2_{ON}+\sigma^2_{OFF}}{N}$, where N is the number of samples. $N = 500$ in our case.[5]

2.2. Signal To Noise Ratio

Extending the concept of the VOM, in the next section, we define possibly the most important concept in this lab, the signal to noise ratio, or SNR. As expected, the SNR quantitatively tells you how strong your signal is in comparison to the noise. We define it as:

$$SNR = \frac{SIGNAL}{NOISE} = \frac{SIGNAL}{\sqrt{VOM}} \tag{1}$$

where we can clearly see that $NOISE = \sqrt{VOM}$.

Because this experiment is a counting experiment and thus obeys Poisson Statistics, we can use the theoretical prediction that $mean = variance$ or $\mu = \sigma^2$ (as defined in lab1) to describe even more precisely what we mean by ON and OFF. To be clear, we write:

$$\begin{gathered}\text{mean "on" counts} = ON = \sigma^2_{ON} \\ \text{mean "off" counts} = OFF = \sigma^2_{OFF} \\ SIGNAL = \sigma^2_{ON} - \sigma^2_{OFF} = ON - OFF\end{gathered}$$

Rewriting the SNR expression in terms of the above definitions for $SIGNAL, ON, OFF$, and our earlier definition of VOM, we see that

$$SNR = \frac{SIGNAL}{\sqrt{VOM}} = \frac{\sigma^2_{ON} - \sigma^2_{OFF}}{\sqrt{(\sigma^2_{ON} + \sigma^2_{OFF})/N}} = \sqrt{N}\frac{ON - OFF}{\sqrt{ON + OFF}} \tag{2}$$

We will use this equation to define the SNR for future parts of the lab.

In addition, we have an expression for an estimate of the SNR, which we can denote here by SNR_2. In IDL code, it looks like

$$SNR_2 = \frac{total(on - off)}{\sqrt{total(on + off)}} = \frac{N mean(on - off)}{\sqrt{N mean(on + off)}} = \sqrt{N}\frac{mean(on - off)}{\sqrt{mean(on + off)}} \tag{3}$$

where on average, $SNR_2 \approx SNR$ because in Poisson Statistics, $mean = variance$ and the variances for measurements with independent errors always add. In other words, $\sigma^2_{ON-OFF} = \sigma^2_{ON+OFF} = \sigma^2_{ON} + \sigma^2_{OFF}$. In the experiment to follow, we will utilize definitions for both the measured SNR and the estimated SNR_2.

[5]Confusion arose over whether N should be 1000 or 500. It is true that we take 1000 total ON and OFF bins (500 bins each), but when we take the VOM, we need to consider the number off samples of $ON - OFF$, which in this case happens to be $500 = 1000/2$

2.2.1. *Experiment: Varying the LED Pulse Width*

After defining the SNR, we then attempted to measure it for varying LED pulse widths, while keeping the background a constant ≈ 2300 counts/sample for our data. We took data for LED pulse widths of 0.1, 0.3, 0.5, 1.0, 3.0, and 5.0 ms, calculated the SNR for each data set, and endeavored to explain the results quantitatively based on propagation of errors and Poisson Statistics. Here are plots of our signal for each LED pulse width. *We display the data in a table on the next page.*

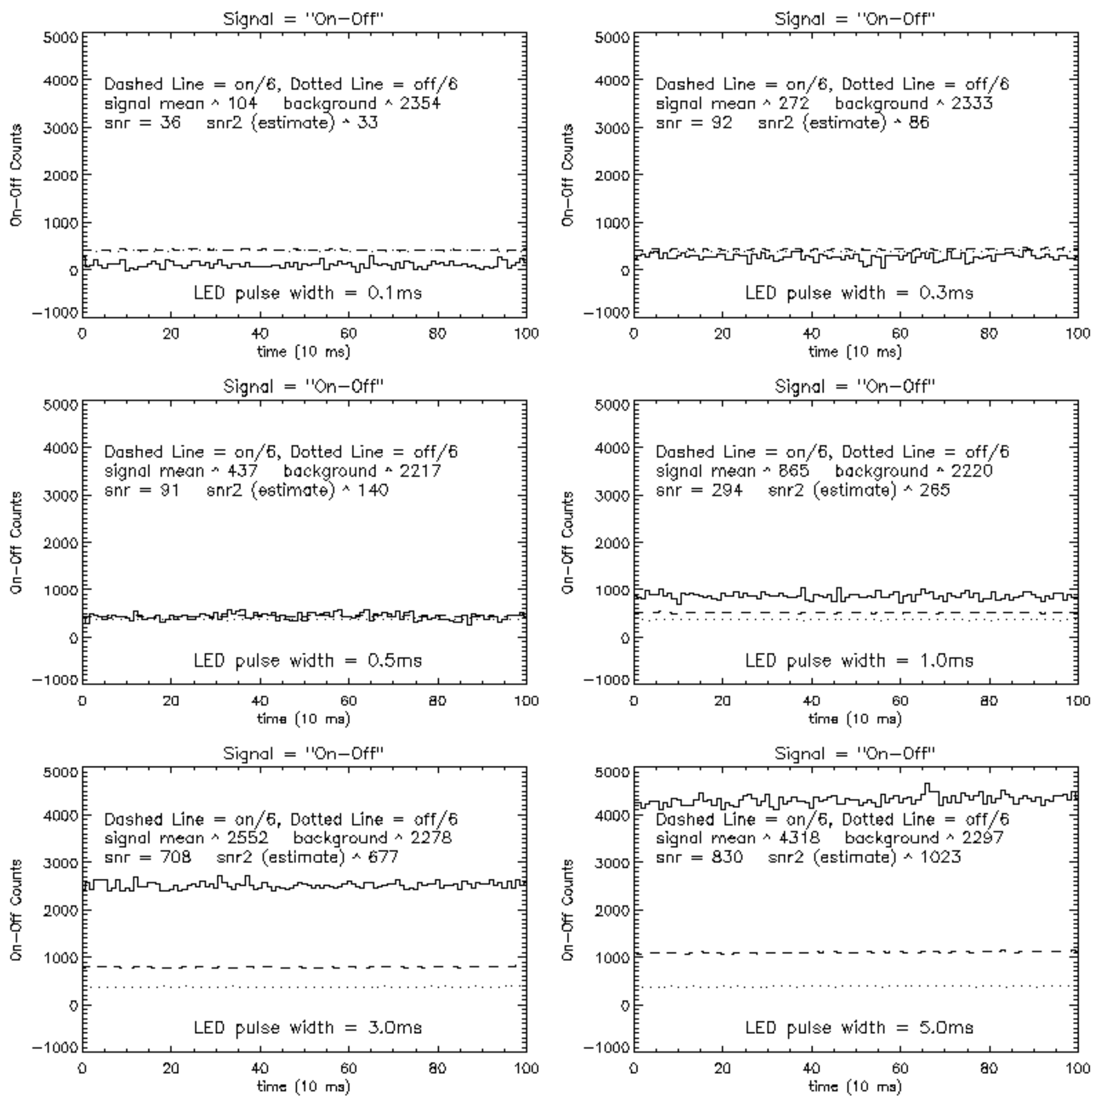

Fig. 2.— *6 runs of 1000 samples taken at 100 Hz. We bin the data into on off bins and calculate the SNR for each LED pulse width using Eq. (1). On the graph we display our best values for the mean signal, background counts, and values for both SNR and SNR_2 for each LED pulse width.*

LED pulse width t (ms)	mean signal $ON \pm \sqrt{ON}$	$SNR + \delta_{SNR}$	$SNR_2 + \delta_{SNR_2}$
0.100	100 ± 10	40 ± 20	30 ± 20
0.300	270 ± 20	90 ± 20	90 ± 20
0.500	440 ± 20	90 ± 20	140 ± 20
1.00	870 ± 30	290 ± 20	270 ± 20
3.00	2550 ± 50	710 ± 20	680 ± 20
5.00	4320 ± 70	830 ± 20	1020 ± 20

Table 1: *All data sets have mean background counts ≈ 2300. We ignore the error in the LED pulse width, trusting that the pulse generator can specify it to 3 significant figures. The error in the mean signal* $ON = \sqrt{ON}$ *from Poisson Statistics. Note that the* SNR *increases as t increases, and that the data for the estimate* SNR_2 *is comparable to the data for the* SNR.

2.2.2. *Deriving the Theoretical SNR curve from Poisson Statistics*

Before we display the plot of our data as a graph of SNR vs. LED pulse width, we must discuss how to derive the Poisson Statistics theoretical formula for the SNR, which we overplot and compare with our data. This is a rather complicated process, but we feel it is necessary to explain what we did to arrive at our plot before presenting it.

Starting with equation (2), , We use the previous definitions of $SIGNAL = ON - OFF$ and $BACKGROUND = OFF$ to rewrite equation (2) in the form:

$$SNR = \sqrt{N}\frac{ON-OFF}{\sqrt{ON+OFF}} = SNR = \sqrt{N}\frac{SIGNAL}{\sqrt{SIGNAL+2\times BACKGROUND}}.^{6}$$

Noting that the $SIGNAL$ should be proportional to the LED pulse width t, we can write $SIGNAL = \alpha t$, where α is measured in counts/ms. We get our best estimate of α by taking $\alpha = SIGNAL/t$ for the largest LED pulse width, which for our data was $t = 5ms$.[7] Similarly, we can define our mean background counts as β, which finally allows us to write a Poisson Statistics theoretical equation for the SNR.[8]

$$SNR = \sqrt{N}\frac{\alpha t}{\sqrt{\alpha t + 2\beta}} \tag{4}$$

Once we measure the parameters α and β, we can overplot this curve on top of our actual data.

[6] i.e for the argument of the square root in the denominator, $ON + OFF = ON - OFF + OFF + OFF = SIGNAL + OFF + OFF = SIGNAL + 2 \times OFF = SIGNAL + 2 \times BACKGROUND$

[7] For a given background, α should be a constant, but we will get our best value of α when t is largest, since we will get a larger signal, which means the percent error in the signal will be a smaller than for other values of t.

[8] it was ultimately derived based on the Poisson Statistics prediction $mean = variance$.

2.2.3. *Error Propagation*

In addition, we must use the techniques of error propagation to explain how we determined the error bars on our data, without which, the comparison to the theoretical overplot would be much less meaningful. So we now must first calculate the error in each of our measured values of SNR and SNR_2 via the methods of error propagation, and place error bars on the plot. As it turns out, the calculation of the errors δ_{SNR} and δ_{SNR_2} are identical, so we need only present one of them here. Starting with Eq. (2), $SNR = \sqrt{N}\frac{ON-OFF}{\sqrt{ON+OFF}}$, we can rewrite the equation by substituting $q = SNR$, $\sqrt{N} = \kappa$, $x = ON$, and $y = OFF$. And since we know from Poisson Statistics that $x = ON = \sigma_{ON}^2$ and $y = OFF = \sigma_{OFF}^2$, we recognize that the errors in x and y are given simply by $\delta x = \sqrt{x}$ and $\delta y = \sqrt{y}$, which we know to be independent. Making the substitutions, we get:

$$q = \kappa\frac{x-y}{\sqrt{x+y}} = \kappa z \tag{5}$$

Now in general for a function of 2 variables $q(x,y) = \kappa z(x,y)$, when the errors in x and y, δx and δy are independent and random, we can use the general formula for error propagation to calculate $\delta q = |\kappa|\delta z$. [9]

$$\delta q = |\kappa|\sqrt{\left(\frac{\partial z}{\partial x}\delta x\right)^2 + \left(\frac{\partial z}{\partial y}\delta y\right)^2} \tag{6}$$

Sparing the reader the differentiation and algebra, but noting again that $\delta x = \sqrt{x}$ and $\delta y = \sqrt{y}$, we arrive at:

$$\delta q = |\kappa|\sqrt{1 - \frac{3}{4}\left(\frac{x-y}{x+y}\right)^2} = \sqrt{N}\sqrt{1 - \frac{3}{4}\left(\frac{\sigma_{ON}^2 - \sigma_{OFF}^2}{\sigma_{ON}^2 + \sigma_{OFF}^2}\right)^2} = \delta_{SNR} = \delta_{SNR_2} \tag{7}$$

As expected, this quantity is unit-less since it must have the same units as the SNR, which itself is unit less. Note that since the subtracted variances $\sigma_{ON}^2 - \sigma_{OFF}^2$ are smaller than the added variances $\sigma_{ON}^2 + \sigma_{OFF}^2$, when we take their ratio and square it we will get a very small number. Thus inside the square root we have 1 minus a small number which tells us that $\delta_{SNR} = \delta_{SNR_2} \approx \sqrt{N}$. Now we're done with the error propagation for the SNR and we're ready to plot our data with error bars and compare it to the theoretical curve!

[9] Taylor Ch 3.

2.2.4. SNR vs. LED Pulse Width

So what should we expect for a plot of SNR vs. LED Pulse Width? Well, clearly, for a constant background, as we increase the LED pulse width t, we increase the signal, and thus increase the SNR. This can easily be seen from equation (4) $SNR = \sqrt{N}\frac{\alpha t}{\sqrt{\alpha t + 2\beta}}$ where as t increases, eventually $\alpha t >> 2\beta$ and the entire equation goes like $SNR \approx \sqrt{N\alpha t}$, or equivalently $SNR \propto t^{1/2}$. So we expect the SNR to scale as a power law, which happens go as the square root of t in the limit of large t. This tells us that it will be convenient to plot the data on a log-log scale, since if the curve behaves like a power law in certain limits, it will appear linear on a log-log scale, allowing for easier extrapolation to smaller SNR's and LED pulse widths.

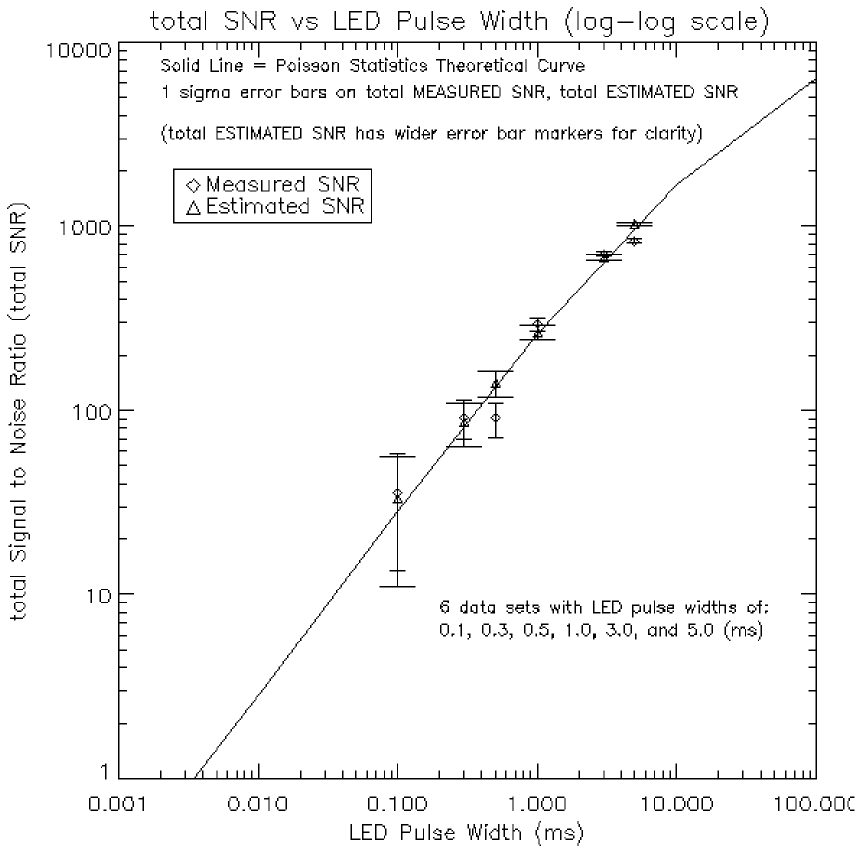

Fig. 3.— *Plot of our SNR (MEASURED SNR) and SNR_2 (ESTIMATED SNR) vs. LED pulse width plotted on a log-log scale. We overplot the theoretical curve derived from Poisson Statistics, and include error bars on both SNR and SNR_2. Note that although the logscale distorts the error bars along the y-axis, in reality they are all comparable in size since $\delta_{SNR} = \delta_{SNR_2} \approx \sqrt{N}$. The error bars for SNR and SNR_2 clearly overlap to 1-σ, (4 out of 6 points overlapping gives 4/6 $\approx$ 67% - almost exactly the 68% for 1-σ) showing that the estimate is a very good one. All told, the data for both SNR and SNR_2 clearly fit the theoretical curve to within the 1-σ error bars.*

2.3. Smallest LED Pulse Width Prediction

Noting that the theoretical curve fits the SNR data to 1-σ, we can use it to extrapolate the smallest LED pulse width in which we can detect a signal, and from that, the smallest signal we can detect given a background level comparable to that of our data ($\approx$ 2300 counts). Quantitatively we define the limit of signal detection as occurring when the $SNR = 1$ (i.e when the signal and the noise are comparable). We can then extrapolate that the smallest t occurs where the theoretical curve intersects the line $SNR = 1$. For our data, we calculated the best value for our intercept as $t_{best} = 3.51\mu s$. So what's the best thing to do next? Well, how about do the experiment? We tuned the background on the PMT apparatus to approximately our old level of $\approx$ 2300 counts, we set the pulse generator to give us a pulse of 3.51 μs and we took some data. We then measured the SNR for that data set and added it to our original curve, where we now predict that the SNR for that data will be 1. So what did we get?

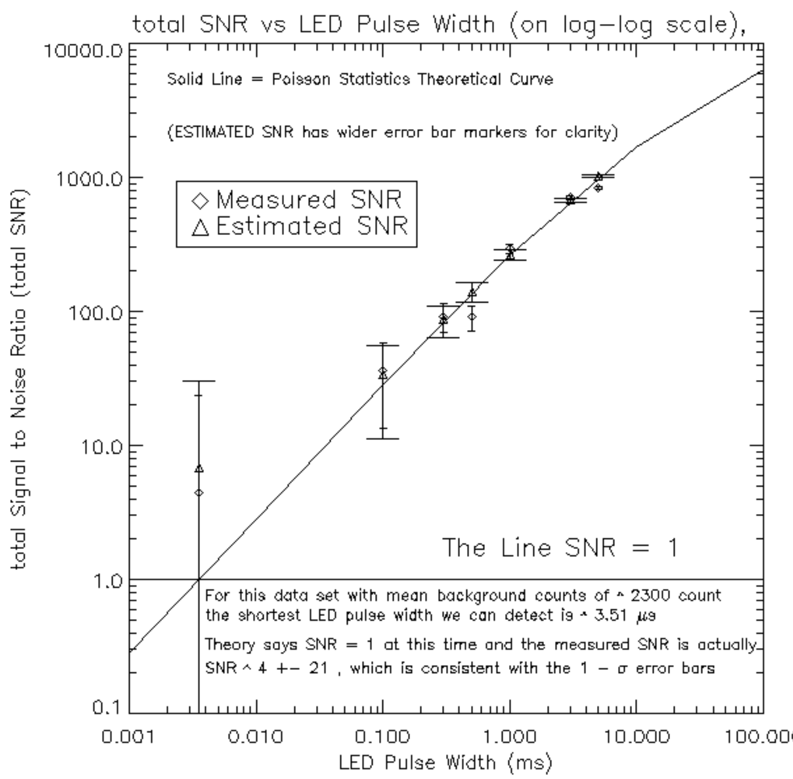

Fig. 4.— *Extrapolating our theoretical curve, we predicted we would measure $SNR = 1$ at an LED pulse width of $t = 3.51\mu s$. What we measured was $SNR = 4 \pm 20$, which is clearly consistent with our prediction to within 1-σ. This shows us the power of the theoretical fit. Note also that the error in SNR, δ_{SNR} is basically independent of the SNR, and is roughly $\sqrt{N} \approx 20$ for all values. This tells us that the errors in measuring small SNR's will always be large compared to the SNR, as is clear for the predicted data point. Note again that the error bars appear skewed due to the log scale.*

As far as detecting the smallest signal, remember that $SIGNAL = \alpha t$ and that we get our best measurement of α from taking $SIGNAL/t$ for the data point with the largest LED pulse width t. For our data, we found that the best value was $\alpha_{best} = 900$ counts/ms. This value tells us that the smallest signal we can detect at this background level of ≈ 2300 counts is ≈ 3 counts. The mean signal per sample we measured at this SNR turns out to be ≈ 20 counts per sample. This differs from our prediction of ≈ 3 counts, but the error bars on it are likely to be so large that the numbers may even be consistent to 1-σ. I will not go through the calculation of the error here since it is not the major point of this particular experiment. The bottom line is that we made a prediction based on the theoretical curve of the smallest LED pulse width in which we could detect a signal, we performed the experiment, and the measurement was consistent with the theory to within the error bars.

2.4. Turning The Noise Up

As mentioned earlier, turning the background up will clearly decrease the signal to noise ratio.[10] If we again revisit the theoretical equation for SNR Eq. (4) $SNR = \sqrt{N}\frac{\alpha t}{\sqrt{\alpha t + 2\beta}}$, we can easily see that increasing $BACKGROUND = \beta$ will decrease the SNR. in the table below, we display the data for the high background set and compare it with the same low background data displayed previously in Table 1. When we actually perform the experiment with increased background, the fact that the SNR decreases is easily confirmed, as seen below and in the plot on the next page.

BACKGROUND	High	High	Low	Low
t t (ms)	$SNR + \delta_{SNR}$	$SNR_2 + \delta_{SNR_2}$	$SNR + \delta_{SNR}$	$SNR_2 + \delta_{SNR_2}$
0.100	20 ± 20	20 ± 20	40 ± 20	30 ± 20
0.300	30 ± 20	30 ± 20	90 ± 20	90 ± 20
0.500	30 ± 20	50 ± 20	90 ± 20	140 ± 20
1.00	90 ± 20	90 ± 20	290 ± 20	270 ± 20
3.00	300 ± 20	250 ± 20	710 ± 20	680 ± 20
5.00	400 ± 20	390 ± 20	830 ± 20	1020 ± 20

Table 2: *The high background data has mean background counts $\approx$ 6400 whereas the low background data has mean background counts $\approx$ 2300. Again, we ignore the error in the LED pulse width, trusting that the pulse generator can specify it to 3 significant figures. Note that for the same LED pulse widths, the High Background data clearly has lower SNR and SNR_2 values than the low background data. In short, upping the background decreases the SNR.*

[10] To dispel some confusion the quantities $BACKGROUND$ and Noise are not the same, even though we casually speak of "Background Noise". The $BACKGROUND$ is simply the mean of the "off" counts, whereas the noise is defined as $\sqrt{VOM}$. Although noise and $BACKGROUND$ are related in the sense that raising either of them decreases the SNR, the two quantities are distinct.

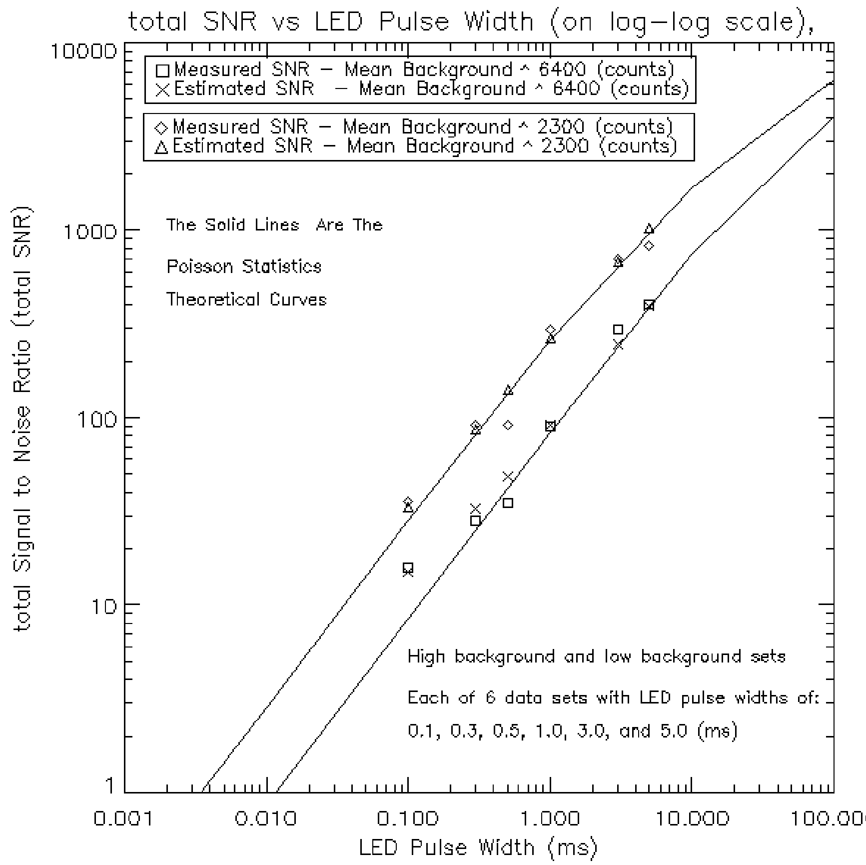

Fig. 5.— *Here we have high background and low background data sets, plotting SNR vs LED pulse width t over the same pulse widths as in previous plots (also listed in the plot). Clearly the theoretical curve for the high background data (mean background $\approx$ 6400 counts) has shifted to the right of the plot, increasing the size of the smallest t for which we could still detect a signal. Error bars are neglected as they are not relevant to this plot. It is sufficient to see that the data follows its associated theoretical plot, and that the SNR does indeed decrease for the higher background data.*

A relevant question for the high background data set is this. "Can we detect a flashing signal that is 1% of the constant LED?" Quantitatively, this question just asks, for the smallest LED pulse width in which we can detect a signal t_{best}, is the quantity $p < 0.01$, where p is defined as:

$$p = \frac{SIGNAL}{BACKGROUND} = \frac{\alpha t_{best}}{\beta} < 0.01 \ ? \tag{8}$$

If $p < 0.01$, then the answer is yes, we can detect the signal. For the high noise data set, we measured $\alpha = 430$ counts/ms, $\beta = 6400$ counts, and $t_{best} = 0.0117$ ms. Plugging in, we get a best value for p of 0.00078, which is clearly less than 0.01. So yes, we can detect a signal that is only 1% of the constant LED! (assuming the error in p is less than an order of magnitude)

3. Getting To Know Your CCD

The second half of the lab consisted of acquiring several CCD images, creating a web page to display them on, then manipulating and analyzing them using IDL code and statistical techniques in order to determine some of the CCD's data storage methods and fundamental noise properties. In general, the noise of the CCD is much smaller than the noise in the PMT because the CCD camera is cooled to a temperature where the pixels have a low probability of ejecting stray electrons due to shot noise, which could be incorrectly interpreted as photon counts. Despite the comparatively low noise, we attempted to detect and correct for the major noise sources which would produce systematic errors if left alone. We first attempted to measure the CCD's dark current in order to subtract it to arrive at the actual photon count rate. We also measured the bias of the CCD using dark current frames of different exposure times. We then fully illuminated the CCD with light from the ceiling in order to measure the gain and readout noise, each of which was extracted from a least squares fit to our data using error propagation techniques. Our final results have the $BIAS$ = 78.6 $\pm$ 0.2 ADU (Analog-Digital Units), the gain as $G = 0.0487 \pm 0.0003$ counts/electron, and readout noise ${\sigma_{RON}}^2$ = 9.2 $\pm$ 0.2 electrons. For clarity, we will discuss each of these measurements in sections named "Web Page", "Dark Current subtraction", "Bias Measurements", and "Measuring the Gain and Readnoise". As relevant we will include the equations for least squares fitting and error propagation.[11]

3.1. Web Page

The first thing we did was to acquire CCD images of ourselves or other lab members to display on our Astro 122 webpage. We saved these first as .fits files then used xv to convert them to JPEG's for easy display on the web. We experimented with the CCD camera's focus, how far to stand from the camera, but left the exposure time as the default of 0.1s. This step is there to get us familiar with how to use the camera and get data from it to later be read into IDL using the readfits command. Once in IDL, we then plotted color coded surface plots and histograms of the data. Bright points of light in the CCD image appear as tall spikes in the surface plot whose z-axes correspond to photon counts. The relevant CCD images, surface plots, histograms and a few possibly witty comments can be found on my Astro 122 Home page at http://www.UGAstro.Berkeley.EDU/ friedman/ For anyone interested, the HTML code can be viewed by clicking on the VIEW PAGE SOURCE option in Netscape.[12]

---

[11] The data used here was shared by myself, James Brennan, Lee Huss, Christina Lee, and Lindsey Pollack. We worked together on the analysis as well. Thus our measured results are the same.

[12] As an aside, I would like to note the similarities between LaTeXand HTML, which both are essentially formatting languages that work from the top down and display corresponding images in real time either in xdvi or on the web, respectively. I've know HTML for a while but so I was excited to learn that aside from esoteric syntax, the structure of the two programming languages are quite similar.

3.2. Dark Current

Dark Current is the phenomenon where the CCD will measure "fake" counts even when the shutter is closed and no photons from your light source are actually hitting the pixels. This comes from the inherent nature of the pixels themselves and can be reduced by cooling the CCD, but not removed altogether. Thus one must take the data and subtract the dark current after the fact to arrive at the true signal. As one might expect, the longer the dark exposures are, the more dark current you can measure. This results in a fundamental limit to exposure times for particular signals because no matter how bright a source is, at some point, the dark current noise will drown out the signal. Luckily, for the relatively short exposure times used here, this is not an issue, but it would be for someone wanting to take a long snapshot exposure of say, a cluster of distant galaxies.

In this lab in particular, we took 2 frames each for dark exposure times of 0.01, 0.1, 1, 10, 20, 30, and 50 seconds. The second exposure will be used later in measuring the variance of the CCD pixel ensemble, using techniques from lab 1. For now, we plot a histogram of the counts as shown below.

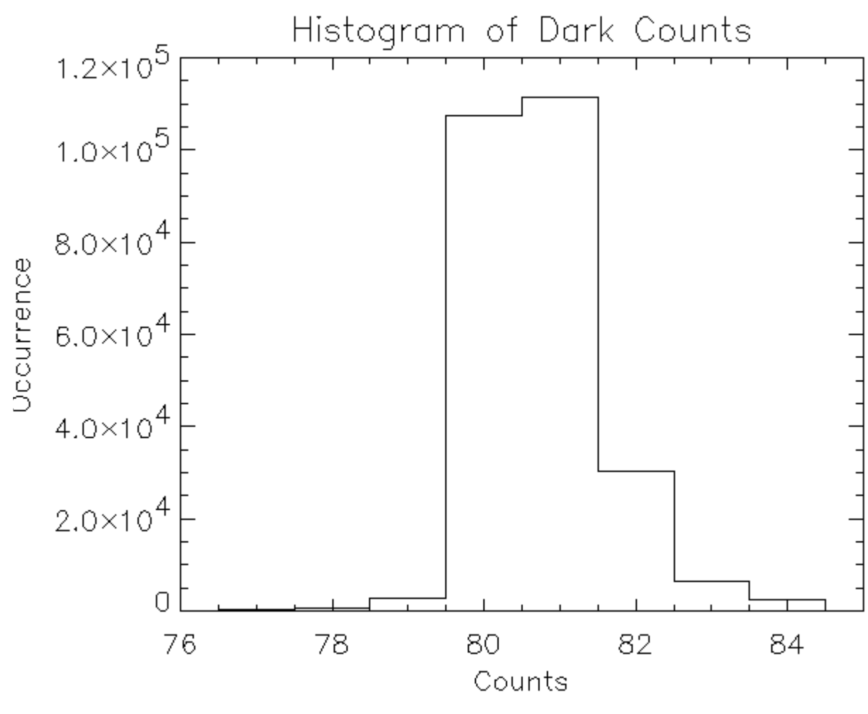

Fig. 6.— *The mean dark counts in this frame are ≈ 80 counts, and the variance ≈ 1. They appear strongly clustered in almost a narrow spike around the mean value, possibly indicating that we can not get a good fit to the data by overplottoing the Poisson Distribution on this histogram. At first this seems to indicate that the dark count electrons are not emitted at a definite average rate, possibly invalidating a description of them using Poisson Statistics. On the other hand, the counts could be subject to a systematic error, a bias possibly....hmmmm*

3.3. Bias Measurements

As it is, the CCD array does have a bias where an arbitrary offset is added to the counts in order to keep their numerical range within the bounds that can be accommodated by the 16 bit unsigned positive integer data type with range 0 to 65,535. The idea here is to avoid negative numbers. The CCD measures the counts as positive integers and if for some reason, the light source being imaged emits photons which hit individual pixels more times than the range of the integer data type, the bias is there to keep the recorded number within the range and thus prevent a program crash. Therefore, we must subtract off the bias after acquiring data to arrive closer to the true signal. That means we first have to measure the $BIAS$.

Using a single image for each of our dark exposure times, we plot the CCD output signal (in ADU) as a function of the dark exposure time t_{exp}. We then perform a least squared fit on the data to extract a value for the $BIAS$. The details are discussed in the figure caption of the plot below. The theory and relevant equations of least squares fitting will follow.

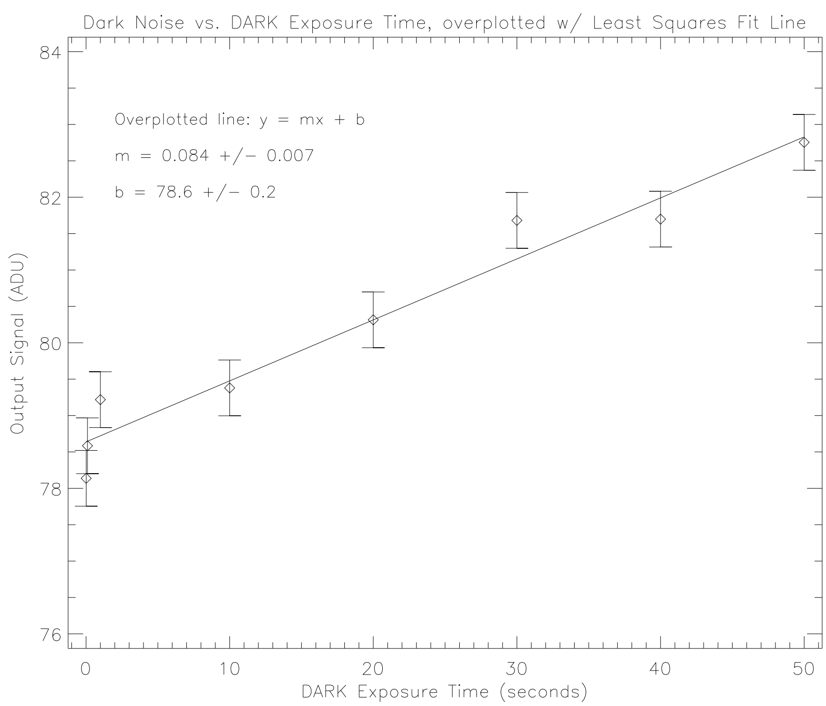

Fig. 7.— *Here we plot the Dark counts vs.* t_{exp}, *along with the least squares best fit straight line to the data. Notice that the y-intercept of the line corresponds to* $t_{exp} = 0$, *which tells us that not only does the CCD measure counts when the shutter is closed, but also when its not even taking an image! Thus the intercept corresponds to our* $BIAS$, *which we measure here as* $BIAS = 78.6 \pm 0.2$ *ADU. As discussed earlier, we see that the Dark counts rise linearly with time. From this we can extrapolate to the longest exposure time one can use before the dark noise overwhelms the signal.*

The best fit line and error bars in the previous plot were determined here by using the statisitcal techniques of least squares fitting and error propagation. Rather than stating the measured results without comment, it will be useful to discuss the relevant equations in some detail, following the derivations in Taylor.[13]. The general formula for straight line is of the form $y = mx + b$, and thus your best fit straight line is determined entirely by the parameters $m = slope$ and $b = y_{intercept}$. Given a data set of N corresponding points $\{x_i, y_i\}$ where $i = 1, 2, ...N$, we calculate our best values for m and b with the following formulae:

$$m = \frac{N \sum xy \sum y - \sum x \sum y}{\Delta} \tag{9}$$

$$b = \frac{\sum x^2 \sum y - \sum x \sum xy}{\Delta} \tag{10}$$

where

$$\Delta = N \sum x^2 - \left(\sum x\right)^2 \tag{11}$$

and in each case $\sum$ implies a sum over i from $i = 1$ to N

By assuming that the error in the time axis is negligible, we can safely apply the method of Least Squares Fitting. To obtain the error in our y-axis (output signal) we used the equation:

$$\sigma_y = \sqrt{\frac{1}{N-2} \sum_{i=1}^{N} (y_i - b - mx_i)^2} \tag{12}$$

In addition, the values of m and b are not as useful without corresponding errors σ_m and σ_b, so we can calculated them using the two formulae:[14]

$$\sigma_b = \sigma_y \sqrt{\frac{\sum x^2}{\Delta}} \tag{13}$$

$$\sigma_m = \sigma_y \sqrt{\frac{N}{\Delta}} \tag{14}$$

Where Δ is still given by Eq. (11). For our data, we after performing the least squares calculations, our measured values end up as $b \pm \sigma_b = BIAS = 78.6 \pm\ 0.2$ in ADU and $m \pm \sigma_m = 0.084 \pm\ 0.007$ counts/second.

[13] Taylor, pg. 182-89

[14] The $\frac{1}{N-2}$ in the denominator of Eq. (12) represents the fact that we already used 2 degrees of freedom in calculating the values of b and m.

3.4. Measuring The Gain and Readnoise

When we aquire normal image data from the CCD and hope to interpret it in terms of the positive integer number of photons that actually hit each pixel, we will have to account for the systematic errors that could be caused by the dark current and the bias. But in addition, there are conversion factors that converts photo-electrons (or net pixel charge) to counts (the gain) and there is a phenomenon called readnoise that adds another source of error having to do with simply reading out the count from each pixel. It is therefore the goal of this section to measure the gain g and the readnoise σ^2_{RON} from our data.

This time, we got the data by openeing the shutter and taking two uniformly illuminated images (white light reflecting off the ceiling), for exposure times of 0.001, 0.01, 0.1, 0.25, 0.3, and 0.35 s. We also took two dark frames for the same exposure times. We then calculated the variance per pixel by subtracting the two illuminated images with corresponding exposure times and calculating the variance of the subtracted image. This removed the dark current and bias, but not the variance in the dark current since, for measurements with inependent errors, variances add. We then took the variance, got the SDOM by squaring it and dividing by the number of samples, which is 2 in this case. The signal, on the other hand, is given by the average of two same-exposure illuminated frames minus the average of two same-exposure dark frames.

So how do we actually measure g and σ^2_{RON} from our data? Again, we use least squares to fit it to a straight line and calculate our errors. We can summarize what happens in the process of converting photo-electrons to counts and derive our theoretical best fit line via the following equations:

$$COUNT = BIAS + (I_\gamma)(t_{exp})\left(\frac{g}{c}\right) + (I_{dark})(t_{exp})\left(\frac{g}{c}\right) \tag{15}$$

where I_γ = intensity (counts/s falling on pixel), I_{dark} = (dark counts/s falling on pixel), t_{exp} = the exposure time, c = the capacitance per pixel, and g is the gain. Subtracting off the Dark Current I_{dark} and the $BIAS$ our signal is just,

$$SIGNAL = (I_\gamma)(t_{exp})(\frac{g}{c}) \tag{16}$$

and the variance V is:

$$V = (I_\gamma)(t_{exp})\left(\frac{g}{c}\right)\left(\frac{g}{c}\right) + \left(\frac{g}{c}\right)^2 \sigma^2_{RON} \tag{17}$$

The extra factors of (g/c) are there to put V in units of $counts^2$. Notice that we can write the above equation in the form $y = mx + b$.

$$V = \left(\frac{g}{c}\right) X + \left(\frac{g}{c}\right)^2 \sigma^2_{RON} \tag{18}$$

where $X = SIGNAL = (I_\gamma)(t_{exp})(\frac{g}{c}), m = \left(\frac{g}{c}\right)$, and $b = \left(\frac{g}{c}\right)^2 \sigma^2_{RON}$

We can then find our best values for m and b using the previously discussed Least Squares Fit line equations. In the end, we find that $g = 0.0487 \pm 0.0003$ *counts per electron*, $m = 0.2 \pm 0.5$ *counts*^{2}, and $\sigma^2_{RON} = 9.2 \pm 0.3$ *electrons*. Using the gain value and the intercept value, the error in the readout noise was calculated using the error propagation formula:

$$\frac{\delta_{RON}}{\sigma^2_{RON}} = \sqrt{\left(\frac{\delta_{\frac{g}{c}}}{\frac{g}{c}}\right)^2 + \left(\frac{\delta_b}{b}\right)^2} \tag{19}$$

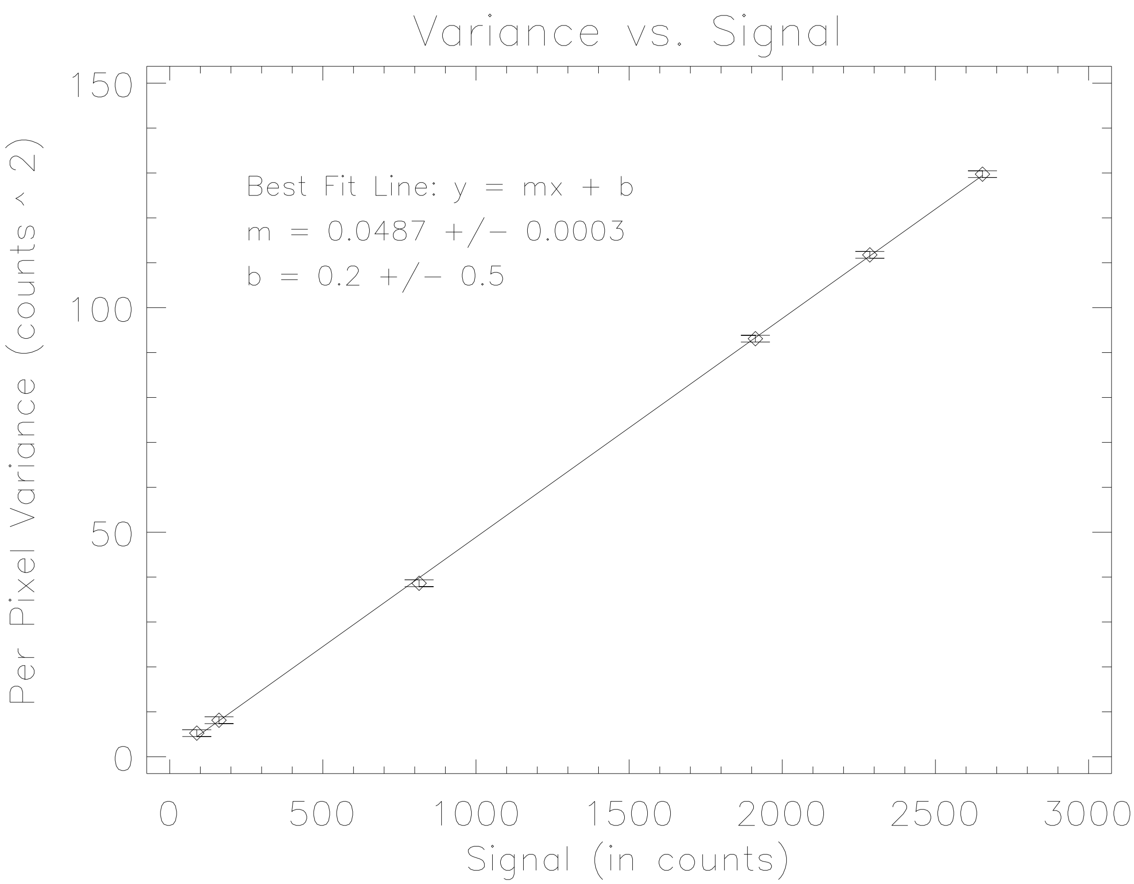

Fig. 8.— *Plot of Variance (Per Pixel) vs. Signal. The error in each pixel variance is $\approx \pm$ 1 counts^2. We assume the error in the x-axis is negligible. We overplot a Least Squares Fit line to our data an it is extremely consistent with the 1-σ error bars. The slope of the line represents our gain factor.g, which we extract the readnoise σ^2_{RON} from the $y_{intercept}$.*

4. Conclusion

This lab was undoubtedly a step up from the first lab, as I sit here having stayed up for about 33 hours straight trying to finish everything up, and not because I was behind, mind you. The sheer level of difficulty and the number of details hidden within are belied by the brief text in the lab handout. For example, take my favorite sentence: "Explain quantitatively your result in terms of error propagation and Poisson Statistics". Who would have thought that 12 words could be such an epic saga. And by the way, this is not complaining as much as concluding, in a way that seems quite reasonable to me in this state of mind. But I can say that I'm getting a lot smarter at using IDL code, UNIX and emacs are becoming second nature, and I've become a bit of a high stress perfectionist with LaTeX. That aside, I've found that the best way to conclude is by summarizing the most important concepts I now leave with, in bullet form.

1. For measurements with independent errors, the variances add. (i.e covariance = 0)
2. The general formula for error propagation (Taylor) works when the covariance = 0.
3. One can detecting faint signals or strong signals to high precision by modulating the signal with the noise.
4. The Poisson Statistics prediction of $mean = variance$ applies to all counting experiments where the signal is random but possesses a definite average rate.
5. The limit of signal detection occurs when the signal to noise ratio = 1.
6. Increasing the background noise decreases the signal to noise ratio.
7. Increasing the sample time increases the signal to noise ratio.
8. HTML and LaTeX are structurally similar programming languages.
9. Dark current and bias must be subtracted from CCD images to avoid systematic errors.
10. Least squares fitting is generally applicable to data sets where the error is all in one axis.
11. CCD gain and readnoise must be taken into account to convert some pixel charge to real photon counts in each pixel.
12. CCD cameras have much lower noise than Photo-Multiplier Tubes, mainly due to their being cooled.

All in all, this lab extended our statistical knowledge from the previous lab, explored further facets of the PMT, and introduced us briefly to the CCD camera, which has quite a number of advantages to the PMT regarding signal detection, and indeed has become the default standard for astronomical imaging. Right now, I am now anticipating how difficult things will get when we stray away from controlled experiments where we generate our own noise in a black box or take serendipitous, uniformly illuminated pictures. I have a feeling that several balls of fire in the distant sky will have their way with us.

5. Acknowledgments/References

ACKNOWLEDGMENTS

I would like to thank in particular James Graham for showing me how to derive and overplot the Poisson Statistic theoretical curve for the total SNR, and for helping me find the cool LEGEND program online. I would especially like to give credit the fellow Astro 122 juggernaughts Lee, Jim, Christina, Lindsey, John, Shane, Lauren, and Kirsten. We ended up doing much more of a practical division of labor for this lab, as it involved quite a step up from the first lab in regard to application of the tools we've learned. Specifically, Lee and I focused on the initial part of the lab dealing with the PMT, the signal to noise ratio, overplotting the theoretical curve, and grinding through the error propagation, whereas, Jim, Lindsey, and Christina focused more on the second half of the lab with the CCD, dealing with the dark current, gain measurement, and performing the relevant least squares fits. In many cases, we shared data and traded our insights into the sections we individually focused on, becoming mini experts and teaching one another what we had learned. But overall, all participants made an effort to gain a thorough understanding of all the important facets of the lab, even though much of the data gathering, analysis, and especially the idl coding was a shared effort. Such teamwork will undoubtedly be crucial for future labs as they grow exponentially in complexity.

REFERENCES

1. Taylor, John R., *An Introduction to Error Analysis.* University Science Books, Sausalito, CA 1997, 2nd ed.

2. Graham, James R., *ERRORS & STATISTICS: Ay 122 Lab Notes*, 2001