Lab 3: CCD Lab

1 Introduction

1.1 Goals

In this lab you will investigate properties of the CCD, the most commonly used light detector at optical wavelengths (and beyond). You will use CCD images taken with the 1-meter *Nickel* telescope at Lick Observatory, and learn how to calibrate them. You will measure the *bias*, the *dark current*, the non-linearity, and the varying quantum efficiencies of each pixel with *flat-field* images.

1.2 Schedule

2 weeks - Report due Tuesday September 25

1.3 Recommended Reading

- Howell, *Handbook of CCD Astronomy,* Chapters 2, 3, & 4
- The lab webpage: `http://ugastro.berkeley.edu/optical/ccd/`

2 Background

Unlike the single element PMT used in the first lab, a charge coupled device or "CCD" is an array of millions of pixels each sensitive to photons. The CCD enables measurement of both the number of photons and the position where they hit. This spatial information allows CCDs to record an *image*, i.e. a picture.

Placing a CCD at the focus of a telescope or behind a camera lens allows it to detect images. CCD cameras can be found at every optical observatory and are ideal instruments for astronomical applications at optical wavelengths, from 300 - 1000 nm.

How do CCDs work? Photons enter the CCD and are absorbed by a silicon layer. This absorption excites an electron from the silicon's valence band to its conduction band in a process known as the photoelectric effect. These photoelectrons are then "captured" and stored by applying a positive voltage to the pixel to hold the electrons in a potential well. The varying number of electrons stored in each pixel produces different voltages across the pixel that is measured (by a fast voltmeter) and the voltage converted to a digital number (DN) that is presented as "counts" or ADUs (analog-to-digital units).

As you will be observing with the 1-meter Nickel telescope at Lick Observatory in a subsequent lab, this lab teaches you how professional astronomers analyze CCD data. You will be provided with a series of images that you will use to understand the properties and characteristics of the CCD camera.

3 Importing Archived FITS Files

A file format for astronomical images is called FITS (Flexible Image Transport System) with the extension, .fits for each file. FITS files contain text headers that can be viewed to quickly identify the contents of the file. The header generally contains information such as the date and time the image was taken, the exposure time (duration), the coordinates of RA and DEC of the object in the sky (if you don't know what RA and DEC is, please look it up with Google), the name of object observed, and many other pieces of information.

The first step of this lab involves locating the image files that you will use for this lab. The image files are in the directory,

```
/home/ay120/ccd_lab/
```

Copy these files to your own directory before working with the images.

```
$ mkdir ccd_lab
$ cd ccd_lab
$ cp /home/ay120/ccd_lab/*.fits .
```

Next, use IDL to open a FITS file. A convenient IDL function already has been written for this very purpose! The IDL Astronomy Library (`http://idlastro.gsfc.nasa.gov/homepage.html`) at NASA/Goddard Space Flight Center contains a number of helpful IDL utilities including programs to read and write FITS images. Fortunately, this library is already loaded when IDL begins. Use the program READFITS to read in a FITS file and explore the header. Plot the images with the DISPLAY procedure.

```
IDL> myimage = readfits('object0.fits', myheader)
IDL> display, myimage
```

The first line above opens the fits file *object0.fits* and stores the image in a variable called *myimage* and the FITS header in a variable called *myheader*. The second line plots the image on the screen. The readfits command also stores the FITS header. Take a look at the header to learn details about the image by printing it to the screen.

```
IDL> print, myheader
```

You may examine the image with a better gray scale (brightest and faintest pixel on the screen) by setting the min and max number of counts in the image. For example, to confine your gray scale to counts between 1000 and 1100:

```
IDL> display, myimage, min=1000, max=1100
```

You may also use the cursor to point to a given pixel and determine the column and row of that pixel:

```
IDL>cursor,x,y
Now click your mouse on a pixel
IDL>print,x,y
The value of x and y are the column and row of the cursor when clicked.
```

When you look at the FITS header, you will notice that the number of rows and columns is less than the 2048×2048 of the native CCD. All of the observatons in this lab were taken with 2×2 binning. That is, each "pixel" in the final FITS image is the average of four pixels with the added advantage of reducing the readout time as well as the noise in each pixel.

But why is the image somewhat bigger than 1024×1024 pixels? Each image from the Nickel telescope includes 32 "overscan" columns. These overscan pixels are not exposed to light and give an estimate of the bias voltage and dark current levels in the CCD. In future laboratory modules you may use the overscan region, but feel free to ignore these columns in this module.

During the course of this lab, it may be necessary to extract values from the FITS header. It is possible to read these values using the IDL command *sxpar* , another function of the pre-installed FITS library. As an example, the number of rows in the FITS image is stored as *nrows* with the following command.

```
IDL> nrows = sxpar(myheader, 'NAXIS2')
```

As with all measuring instruments, CCDs are not perfect, and they suffer from various effects and sources of error.

4 Bias and Dark Current

The bias and dark current are measures of the CCD output with no light input. Bias images are simply images taken by the CCD with no light input with a total integration time of 0.0 seconds. Dark current images also are made without any light input but have an exposure time greater than zero. These images allow measurement of the dark current of the CCD.

Read the bias image (bias.fits) and display the image to the screen. What do you see? Just by looking at this image, can you identify imperfections in the CCD? Is so, what are they? How will these imperfections affect your results? One way to deal with these bad pixels is to ignore them in your anaylsis. With the bias image as a guide, identify the pixels to be excluded from your future analysis. What was your criterion?

Now, plot a histogram of the pixels of the bias image. Is the histogram nearly Gaussian in shape? What is the mean count level (the bias level) in the CCD? What is the standard deviation of the counts in the bias image? Why do observers need to determine this mean bias level?

The dark current in the CCD is a measure of the thermal agitation of the electronis in the CCD. That is, occasionally electrons will be expelled from their valence electrons due to the thermal energy of the CCD. Once free, these electrons are captured by the CCD and are indistinguishable from those electrons freed by astronomical photons. Given the histogram results from the zero-second bias image, estimate the lower limit (in counts/pixel/hr) of the dark current that could be measured with a dark image having an exposure time of 64 seconds. Now, open the FITS file *dark.ccd* and plot a histogram of the pixels. How different is this histogram from the bias histogram? What does this tell you about the dark current of the Nickel CCD. Search the Lick Observatory website to record the measured dark current of the CCD (in counts/pixel/hr) at the Nickel Telescope. Does your measured value agree with that on the Lick Obs. website?

5 CCD Saturation

The number of counts recorded in a pixel is proportional to the number of photoelectrons. The ratio of photoelectrons (i.e. detected photons) to counts is called the *Gain* of the CCD. The Nickel CCD has a gain 1.7 e^-/count. Every count represents 1.7 photons detected.

The number of counts is linearly proportional to the number of photons that hit the pixel. But a CCD's response to light is linear *only* up to a specific light intensity. When a large amount of light is shined on the CCD, the number of counts is slightly less than it should have been, i.e. *non-linear*. When even more light is shined, the CCD will *saturate*: photons hitting the CCD no longer produce any counts. Measuring how the CCD becomes non-linear and eventually saturates involves shining uniform light on the CCD repeatedly with increasing exposure times, 1 sec, 2 sec, 4 sec, etc.

Use the files named *flatfield0-11.fits* to calculate the mean number of counts as a function of exposure time. Plot the results and determine how linear the CCD is. At what count level does the linearity breakdown? Show in your plot the maximum counts that can be recorded by a 16-bit Analog-to-Digital (A/D) converter. Remember that 16 bits represents 16 values of either 0 or 1 using base-2 to represent any number. What is the highest achievable number in a 16-bit A/D converter?

6 Measuring Quantum Efficiency: Flat Field Exposures

Each pixel in a CCD has a different sensitivity to photons. The fraction of photons hitting a pixel that result in a photoelectron is called the *Quantum Efficiency* of the pixel. The typical QE for CCD pixels is 0.6 - 0.8, i.e. 60% to 80% of the light is detected. But you must determine the different QE for each pixel, to correct the CCD image for these pixel-to-pixel differences.

A *flat-field* exposure is made by shining a uniform smear of light on the CCD. Every pixel receives the same amount of light. The resulting *flat-field* image has different numbers of counts in each pixel because of their differing QE. The ratio of counts in neighboring pixels is equal to the ratio of their QE.

For example, if one pixel has 10,000 counts and the neighboring pixel has 11,000 counts, the second pixel has a QE that is 10% higher.

All images of the night sky must be corrected for the differing QE of the pixels by dividing the image by a flat-field image, thereby "dividing out" the QE variations. Ideally a master flat-field image is the sum of many flat-field images, so that the Poisson variations (square-root of the number of photons) is minimized.

Determine the typical fractional variations in the QE of the pixels in your CCD. To do this, consider one flat-field image, such as the 4-second or 8-second exposure used in the determination of the non-linearity of the CCD. Create a sub-image consisting of 100 x 100 pixels centered on pixel, (500,500), i.e., near the center.

```
IDL> subim = im(450:550,450:550)
```

- Compute the mean count level of that subimage.
- Compute the standard deviation of that subimage
- Make a histogram of that subimage

From those three efforts, compute the variation in the QE, as a fraction of the typical QE. (You can't figure out the mean QE itself.) That is, does the QE vary by 0.1%, 1%, 10% or by how much from pixel to pixel?

7 A CCD Reduction Package

From your understanding of the CCD behavior and using the bias, dark current and flat field images, construct corrected *reduced* image. The object fields can be found in the *object0-9.fits* files. Describe how you created the corrected image step by step, using the bias, dark, and flat-field images.

8 Supporting Material

- UC Lick Observatory homepage - `http://mthamilton.ucolick.org/`
- Overview of the 40-inch Nickel telescope - `http://mthamilton.ucolick.org/techdocs/telescopes/Nickel/nickel.html`
- Overview of the Nickel Direct Imaging Camera - `http://mthamilton.ucolick.org/techdocs/instruments/nickel\_direct/intro/`
- IDL Astronomy Library - `http://idlastro.gsfc.nasa.gov/homepage.html`
- FITS I/O with the IDL Astronomy Library - `http://idlastro.gsfc.nasa.gov/fitsio.html`