

CCD Characteristics Lab

Observational Astronomy

6/6/07

1 Introduction

In this laboratory exercise, you will be using the Hirsch Observatory's CCD camera, a Santa Barbara Instruments Group (SBIG) ST-8E. The operation of CCDs, or *Charge Coupled Devices*, can be understood on the basis of the photoelectric effect and semiconductor physics¹. A CCD chip is based on the Metal² Oxide Semiconductor (MOS) capacitor, which is also a basic component of the MOSFET (MOS Field Effect Transistor). Indeed, CCD and computer chips are manufactured on a silicon wafer in much the same way.

The MOS capacitors in a CCD create an electric potential maximum in the n-type region of a p-n junction. This is a potential well for electrons (why?). When light falls on the CCD, the incoming photons excite electrons into silicon's conduction band in a classic example of the photoelectric effect, thereby creating electron-hole pairs. The photoelectrons collect in the potential well region, which prevents them from recombining with the holes. A structure of channel stops and biased gate electrodes divides the CCD chip into pixels; the ST-8E has 9×9 micron pixels in a 1530×1020 array.

In a basic *three-phase* device, each pixel has three gates; during an exposure the center gate is biased at a higher potential. Photoelectrons then collect in each pixel's potential well, as shown in Figure 1. When an exposure is completed and the camera shutter has closed, the gate voltages in the device are cycled in a process called "clocking the gates." A side gate is biased to a higher potential, allowing the collected charge to spread over. The center gate is then biased to a lower potential, leaving the charge centered on the side gate (see Figure 2). Repeating this process eventually results in the transfer of collected charge from one pixel to the next. The *Charge Transfer Efficiency* (CTE) is the percentage of charge actually transferred from one pixel to the next. This value is quite high in modern CCDs ($> 99.999\%$), which is necessary considering how many times some of the charges are transferred (think about it!).

The CCD reads out by shifting all the pixel columns down one row, with the row at the edge being shifted into the *serial register*, a row of pixels that is never exposed to light. The gates in the serial register are then clocked to transfer these pixels one at a time into

¹You remember Physics II, right?

²In modern devices, polycrystalline silicon is used instead of metal.

the output amplifier and other readout electronics, where the voltage associated with the collected charge is amplified, measured, and converted into a digital number. The number associated with each pixel is referred to as its *counts* or *ADUs* (analog-to-digital units).

(If the explanation above leaves you confused, there exists a helpful analogy for the operation of a CCD. The incoming photons are rain, and the CCD chip is a 2D array of buckets. Each bucket is a pixel, and the water it collects is the charge accumulation due to photoelectrons. Once the rain has stopped (the shutter is closed), conveyor belts move the columns of buckets down one row (the gates are clocked). The water in the buckets at the edge of the array pours into more buckets on a horizontal conveyor belt (serial register). This conveyor belt then pours these buckets one at a time into a graduated cylinder (output amplifier). The volume of water from each bucket is rounded to the nearest milliliter (counts), and you can then reconstruct the distribution of rainfall on the array (an image). See Figure 3 for an illustration.)

However, you should keep in mind that these “counts” are not direct counts of incoming photons. The *quantum efficiency* (QE) is the fraction of photons that excite photoelectrons and produce a count; it is a function of wavelength for each type of CCD³. Furthermore, the gain is the number of photoelectrons the analog-digital (A/D) converter needs to produce one ADU (more on this below). Other factors contribute various types of noise to the signal. In this lab, you will explore the following CCD characteristics⁴.

read noise A CCD camera has several inherent sources of noise collectively referred to as read noise. The output amplifier and other camera electronics generate both thermal and flicker noise. Thermal (white) noise results from the thermal motion of atoms and electrons and has a \sqrt{T} temperature dependence. Flicker (pink, or $1/f$) noise results from trapped states in the semiconductor. There is also noise associated with the process of clocking the CCD, which typically has a \sqrt{f} dependence on clocking frequency. Note that the levels of all these noise sources are independent of the signal level.

dark current Over time, thermal excitation of electrons in the silicon builds up a signal referred to as dark current. This is dealt with by taking dark frames with the same exposure times as the light frames, but with the shutter closed. The dark frame is then subtracted from the light frame, eliminating pixel-by-pixel variations due to dark current. Dark current is also a source of Poisson noise. This noise cannot be subtracted; dark subtraction in fact increases the overall noise level (why?). This effect can be minimized by averaging several dark frames together. Still, the optimal solution is to make the dark current itself (and therefore its Poisson noise) as small as possible. Since dark current is extremely temperature dependent, CCDs are either

³CCDs can have QEs of 40%–90% at visual and near-IR wavelengths, compared with around 1% for the human eye and 5% for photographic film.

⁴For more information, refer to <http://www.ccd.com/ccdu.html>;
<http://micro.magnet.fsu.edu/primer/digitalimaging/concepts/concepts.html>;
<http://spiff.rit.edu/classes/phys559/lectures/ccd1/ccd1.html>.

equipped with thermoelectric cooling (as the ST-8E is) or cooled with liquid nitrogen (as are most scientific instruments).

gain The number of collected electrons (really, the analog voltage their charge produces) needed to read out one ADU is referred to as the gain of the system. A low gain means less digitization noise, but also a lower effective well depth. This is due to the limitations of the readout electronics; the ST-8E has a 16-bit A/D converter and can therefore represent ADU counts from 0 to $2^{16} - 1 = 65,535$. A very low gain may result in A/D saturation before the actual full well capacity of the CCD (in electrons) is reached, though full well saturation almost always happens first. Telltale signs of saturation are “flat” image regions (where pixels have constant value) and streaks from “blooming” stars (where charge has overflowed into neighboring pixels).

linearity A chief attraction of CCDs is their linear response to varying light levels, ideally from the limits of detection to full well saturation. CCDs have a much wider range of sensitivity (larger *dynamic range*) than photographic film, typically by one or two orders of magnitude. In addition, their wide range of linearity means that the brightnesses observed correspond to actual brightnesses. In a CCD, full well or A/D saturation produces obvious effects, but nonlinearity does not. It is therefore very important to note the range of your CCD’s linear behavior, or your scientific results may turn out to be useless! It is interesting to note that the human eye is not linear but logarithmic. This allows for a large dynamic range at the cost of the ability to distinguish varying light levels. A linear CCD image contains much more information than can be displayed on a computer screen or seen by the eye at one time; this is why screen display settings (“screen stretch”) and image processing are so important in bringing out various details.

2 Setup

After you have successfully connected to the camera, set the cooling to a reasonable value (try -5°C) and wait for the temperature to stabilize. Meanwhile, you can do some setup that will become useful later in the lab. Place a piece of plain white paper over the large aperture attachment. It is best to place the paper on a desk, put the attachment down on this, use black electrical tape to seal the paper to the attachment, and then trim off any excess paper. Insert the covered end into the light box; you may need to press firmly to do this. Now attach the camera, using the wooden blocks to support it. Use electrical tape on the inside of the box to make a light-tight seal where the camera enters. Put the cover on the box.

Throughout this lab, you want to avoid light leaks in the camera, which are possible even with a closed shutter. Keep the room as dark as possible and point light sources, such as computer screens, away from the camera. Also, be sure to check for temperature stability throughout the lab. When analyzing your data, it is advisable to use Excel or MATLAB to save time and effort.

3 Read Noise

Record the CCD temperature. Take two *bias frames*, or images of zero seconds exposure⁵ with the shutter closed. After saving these images, subtract the second image from the first. This process eliminates any overall gradients in the bias frame, leaving just the structure of the noise. Since negative pixel values are set to zero, you must add a constant (try 1000) before (CCDOps) or during (MaxIm DL) the subtraction. If the differenced image has pixel values of zero, try again with a larger constant; otherwise, save the image.

Examine the differenced image to identify any hot pixels or other unusual features; adjusting the zoom and display settings can help bring out various details here. Then, measure and record the standard deviation of the counts in a typical region (record the size of your region as well). The central region from (510, 340) to (1020, 680) usually works well. How does the added constant affect this measurement? The read noise is then found by dividing the standard deviation by $\sqrt{2}$ (why?). Keep in mind that this is the read noise for an individual pixel.

For future measurements, you may want to set the camera to read out only part of the full frame; CCDOps has the ability to only read out the central half or quarter of the chip. This will shorten image download times dramatically but still provide a relatively large region for measurement. You should continue to check for hot pixels, unusual regions, etc. Measure in a smaller subframe to avoid these, if necessary.

4 Dark Current

Record the CCD temperature. Take a bias frame and save it. Then take and save a series of *dark frames* (i.e., with the shutter closed), starting with an exposure time of 30 seconds and proceeding to 300 seconds in 30 second increments. Subtract the bias frame from each of the dark frames; you may need to add a constant to avoid negative pixel values. Once this is done, measure and record the mean of the counts for each exposure time. Use the same measurement region in all images, and remember to subtract the constant you added from the means.

Make a plot of mean counts versus exposure time and fit a linear function. Sketch your plot and record your fit parameters. The fit slope is the dark current in ADU/sec. You will calculate the dark current in electrons/sec once you know the gain.

5 Linearity & Gain

For this part of the lab, you need to achieve a uniform illumination of the CCD. This is why you put the paper over the aperture attachment earlier. Put a frosted bulb in the ceramic stand and place it at the far end of the light box. Plug the light bulb stand into the VARIAC

⁵The ST-8E has a minimum exposure time of 0.12 seconds, which it automatically uses if you tell it to take any shorter exposures.

and set it to a low voltage (15 volts seems to work well; record what you actually use). Put the cover back on the box; the cord goes up, over, and down at a corner of the box.

Record the CCD temperature. Now take and save a series of frames starting close to the minimum exposure time and increasing in even increments until the camera saturates. (At a VARIAC setting of 15 volts, 0.2 seconds to about 6.4 seconds in increments of 0.2 seconds seems to work well.) Saturation should be apparent from both the image (flat regions) and its statistics (mean pixel value will stop increasing). Since the camera is cooled, you can ignore the effects of dark current on these fairly short exposures. For each image, measure and record the mean and standard deviation of the counts. Then calculate and record the count variance.⁶

You now have the data needed to calculate the CCD's gain. Because the gain relates the number of collected electrons to ADUs, you know that

$$electrons = gain \cdot \mu_{ADU} \tag{1}$$

and that

$$\sqrt{electrons} = gain \cdot \sigma_{ADU} \tag{2}$$

The second equation follows from the fact that electron counts (*not* ADU counts) follow Poisson statistics. Squaring this equation, you find that

$$electrons = gain^2 \cdot \sigma_{ADU}^2. \tag{3}$$

Setting Eqns. 1 and 3 equal and canceling a factor of gain, you conclude that

$$\mu_{ADU} = gain \cdot \sigma_{ADU}^2. \tag{4}$$

Make a plot of mean counts versus count variance and sketch it on the worksheet. Note where the linear regime ends; it may be earlier than you expect. Truncate your data to remove the nonlinear regime and make (and sketch) another plot. Now fit a linear function to this data; record your fit parameters. Following the derivation above, the slope gives the gain, which you should record. The nonzero intercept is due to the read noise.

Speaking of read noise, you can now convert your measurement of it from ADUs to electrons; record this value. You should also convert your measurement of the dark current to electrons/sec and record that as well. You now know the rate of thermal electron excitation in each pixel on the CCD chip (for a particular temperature).

Next, make and sketch a plot of mean counts versus exposure time. Note where the curve flattens out; this is the CCD saturation count. Record your estimate for both this value and the CCD's full well capacity. You previously identified the linear regime in the process of finding the gain; notice that it is not easy to identify the linear regime directly from this plot. Estimate and record the size of the linear regime as a fraction of full well capacity. Make and sketch another plot of mean counts versus exposure time, omitting the nonlinear regime. You should fit a linear function to this data and record your fit parameters.

When you are all done working, be sure to properly warm up and shut down the camera.

⁶Remember, variance is just the square of standard deviation.

6 Application: The Signal-to-Noise Ratio

Some of the quantities measured in this lab are useful when calculating the signal-to-noise ratio (SNR) for astronomical observations. For a signal with no background or instrument noise, the collected photoelectrons follow Poisson statistics. This results in a SNR of $N_*/\sqrt{N_*} = \sqrt{N_*}$, where N_* is the number of electrons from the the signal source (i.e., the number of incoming photons times the quantum efficiency). Taking into account the other main sources of noise, we can construct the so-called “CCD Equation,”

$$SNR = \frac{N_*}{\sqrt{N_* + n_{pix} \cdot [N_{sky} + N_{dark} + N_{read}^2]}}, \quad (5)$$

where n_{pix} is the number of pixels under consideration, N_{sky} is the number of electrons from the sky, N_{dark} is the number of electrons from the dark current, and N_{read} is the number of electrons due to the read noise (these three N values are per pixel). Note that N_* , N_{sky} , and N_{dark} will depend on exposure time, while N_{read} will not. For high signal levels, the signal’s Poisson noise dominates over the other noise sources, and the SNR is again approximated by $\sqrt{N_*}$. For low signal and/or high noise levels, a more complete version of this equation can be found in §4.4 of Steve B. Howell’s *Handbook of CCD Astronomy*. Be sure to answer the worksheet’s questions about the CCD Equation.

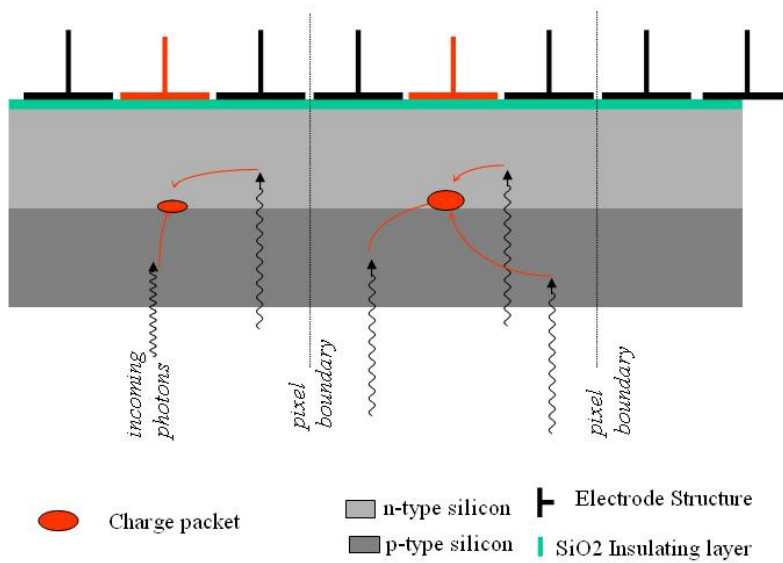


Figure 1: When a CCD is exposed to light, incoming photons excite electrons, which are then trapped in the potential wells of individual pixels. Image from http://www.ing.iac.es/~smt/CCD_Primer/CCD_Primer.htm.

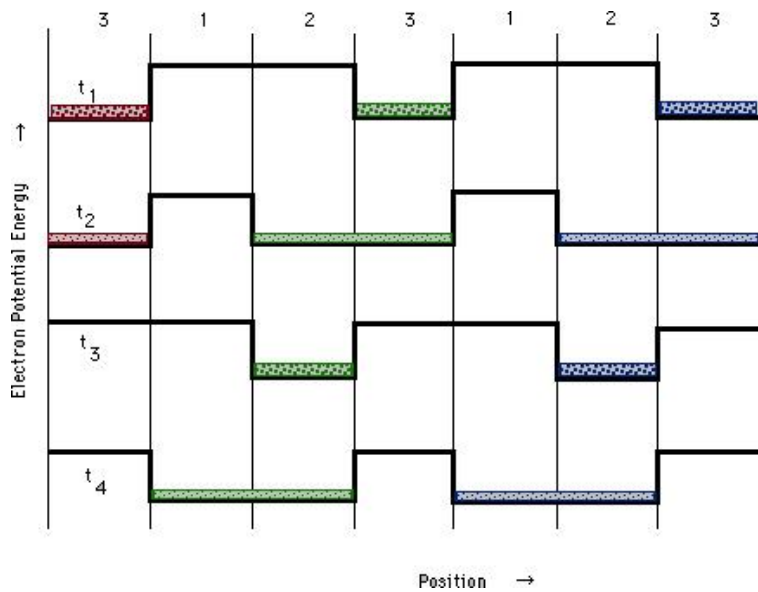


Figure 2: By changing the gate bias potentials, charge is transferred laterally in the CCD. Image from <http://oak.cats.ohiou.edu/~piccard/radnotes/detectors.html>.




Figure 3: The "bucket brigade" analogy for CCD operation. Image from http://www.ing.iac.es/~smt/CCD_Primer/CCD_Primer.htm.