An Introduction to the CCD

Earth & Sky

6/28/06

1 Introduction

Astronomy has a considerable historical pedigree. Ancient Chinese astronomers kept detailed records of the manifestations of comets, and the ancient Greeks compiled the first known catalog of the stars. The advent of telescopic observations increased the astronomer's reach immensely, and scientists such as Galileo laid the foundations of our present understanding of the Universe.

In the mid-nineteenth century, astronomers began to take advantage of the new technology of chemical photography. Photography has two great advantages over naked eye observation. Firstly, it provides a lasting record of observations, allowing everyone access to the same data. Secondly, it makes possible extended exposures, allowing for the cumulative collection of light. Increased light-gathering power means that fainter objects can be observed. For example, spectra[1] of distant galaxies recorded on photographic plates led to such important discoveries as the expansion of the Universe.

Electronic instruments first saw use astronomical use in photometry, the accurate measurement of an object's luminous intensity. Photoelectric photometers use photocells[2] to produce sensitive measurements of the light falling on them. Photocells can be thought of as light-dependent resistors, meaning that the current flowing through them depends on the intensity of the light falling on them. This is a result of the photoelectric effect, first explained by Einstein. When light falls on matter, its energy can be transferred to electrons, exciting them to higher energy levels. In metals the excited electrons are often ejected entirely, while in semiconductors they often move into the material's *conduction band*. This causes an increase conductivity, making photocells possible.

---

[1] When light is passed through a prism or reflected off a diffraction grating, it spreads into a rainbow. This is called a spectrum, and scientists can learn much from the pattern of light and dark lines caused by the presence of various forms of matter.

[2] The same devices are used in simple motion sensors and automatic night lights.

2 Enter the CCD

The photoelectric effect is also employed in Charge-Coupled Devices (CCDs), which are the modern instrument of choice for photometry and all other forms of astronomical imaging in the near infrared, visual, and near ultraviolet wavelengths. A CCD chip is manufactured on a silicon wafer in much the same way as a a computer chip. Indeed, similar principles of semiconductor physics underlie the operation of both types of devices. A CCD has several advantages over photographic plates. One is ease and speed of use. CCDs do away with the preparation and developing of sensitive chemical emulsions; an image can be viewed within minutes of the completion of an exposure.

In addition, CCDs are much more sensitive than the eye or photographic film. Astronomers say that CCDs have a high *quantum efficiency* (QE), meaning that a large percentage of incoming photons are actually detected. While photographic plates might capture one photon out of every hundred, modern CCDs would capture eighty photons out of every hundred. This allows for a substantial decrease in exposure time.

CCDs are also *linear* in nature, meaning that the signal they produce is directly proportional to the amount of light collected. This makes it easier to calculate the number of photons that hit the detector in the time of an exposure. We can then calculate the apparent brightness of a star. The combination of high sensitivity and a linear response over varying light levels has made CCD detectors the premier astronomical instrument.

3 Dynamic Range

Since CCDs can detect light at both low and high levels, they are said to have a large *dynamic range*. In contrast, photographic film will under- or over-expose if light levels vary too much. The human eye also has a large dynamic range, but it is *logarithmic* rather than linear. This means that you can see over a broad range of light levels, but are less able to distinguish individual light levels. A linear CCD image thus contains more information than your eye can see at any one time. In addition, display devices such as computer monitors have a rather low dynamic range. As a consequence, screen display settings and image processing techniques are vital when extracting information from an image.

4 CCD Operation

A CCD chip is divided into pixels. Each pixel has a potential well that collects the electrons produced by the photoelectric effect. At the end of an exposure, each pixel has collected an amount of electrons (i.e., charge) proportional to the amount of light that fell onto it.

The CCD is then read out by cycling the voltages applied to the chip in a process called "clocking." Due to the structure of a CCD, clocking causes the charge in one pixel to be transferred to an adjacent pixel.

To understand how the whole chip can be read out in this way, consider the following analogy. 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[3]. This conveyor belt then pours these buckets one at a time into a graduated cylinder (the readout electronics). The volume of water from each bucket is measured and rounded to the nearest milliliter (corresponding to the digital output of a CCD, which reports the *counts*, or *analog-to-digital units* (ADUs), from each pixel). You can then reconstruct the distribution of rainfall on the array (an image). See Figure 1 for an illustration of this process.

Most CCDs have a setting called *binning* which causes them to read out in a slightly different manner. When binning is used, blocks of pixels are grouped together into "superpixels." Each superpixel acts as one large pixel, with several results. Read times are faster, since fewer actual measurements of charge are performed. (Water from a group of buckets is poured into the graduated cylinder, then the level is carefully read.) Each superpixel is also more sensitive, since it can collect more photons for a given exposure time, but this comes at the cost of resolution (see below).

5 Nothing is perfect...

Like all instruments, a CCD has various limitations and sources of error. The *Charge Transfer Efficiency* (CTE) is the percentage of charge transferred from one pixel to the next. (In the bucket analogy, some water splashes out when one bucket pours into another.) Since some charges are transferred many times in good-sized CCD, the CTE must be quite high, usually in excess of 99.999%. Furthermore, the number and size of the pixels themselves limit the CCD's resolution (ability to detect detail) and field of view (image size). Smaller pixels mean higher resolution, and more pixels of the same size mean greater field of view.

Another limitation is the pixel *full well capacity*. This is the maximum number of electrons each pixel can hold, sometimes referred to as the *saturation count*. (In the bucket analogy, this is how much water each bucket can hold.) Saturated regions in an image usually appear to have a constant brightness. Saturation can also result in "blooming," where electrons overflow to neighboring pixels. This usually appears as vertical spikes emanating from bright stars. (In the bucket analogy, this occurs when a bucket overflows into neighboring buckets.)

---

[3] In an actual CCD, this is another row of pixels that is never exposed to light

Problems can also arise due to nonlinear effects. Some CCDs depart from a linear response well before saturation occurs. This effect is not immediately obvious, so astronomers must know the linear range of their CCDs. For accurate measurements (such as photometry), this knowledge is absolutely vital.

CCDs have several identifiable sources of error. *Read noise* is produced in the camera electronics and during the readout process, and is independent of exposure time. In addition, thermal energy in the CCD causes electrons to be excited over time. The electrons due to this process are referred to as *dark current*, since they accumulate even when the shutter is closed. Dark current can be reduced by using shorter exposures, or by cooling the CCD with thermoelectric coolers or liquid nitrogen (see Figure 2). Professional instruments cooled in the latter manner can have negligible dark current. Finally, individual pixels have slightly different responses to light of the same intensity, introducing another source of error.

6 Basic Image Processing

There are a couple of simple techniques used to compensate for the effects outlined above. A *bias frame* is an image exposed for zero seconds (or as short as the camera allows) with a closed shutter. It contains information about pixel-to-pixel variations in the read noise, often visible as a gradient in the bias frame, as well as any defects in the chip. Subtracting a bias frame from a *light frame* (i.e., a normal image) removes these effects.

A *dark frame* is an image exposed for a finite amount of time with a closed shutter. It contains the level of the dark current and any pixel-to-pixel variations in it. Subtracting a dark frame from a light frame corrects for the dark current. Note that the bias information is also in the dark frame, so bias-subtraction is subsumed by dark subtraction.

A *flat frame* is an image exposed to a uniform light source[4]. It contains information about pixel-to-pixel variations in quantum efficiency (i.e., response to light) and light variations due to the instrument configuration. Dividing a dark-subtracted light frame by a dark-subtracted flat frame corrects for this effect.

All of the above processes increase the overall noise level of the image. Astronomers often compensate for this by using libraries of bias, dark, and flat frames. Averaging many frames together then reduces their noise levels. As a result, the amount of noise added by image calibration[5] is minimized. Astronomers quantify the presence of noise with the signal-to-noise ratio (SNR), where the signal is from the object of interest (e.g., a star), and all possible noise sources are considered. One way to achieve a high SNR is to use longer exposures, since the signal from the object often (but not always!) accumulates faster than the noise.

---

[4] Astronomers use illuminated dome screens or the twilight sky.

[5] Making the corrections described in his section is called calibrating the image. More processing is often done later on, depending on the scientific application.

7 Questions

1. What makes a CCD a better astronomical tool than photographic plates?

2. The Hirsch Observatory's camera has a CCD chip that is 1530×1020 pixels. What is the maximum number of times a charge would have to be transferred? After this many transfers, what percentage of the original charge makes it out, assuming a CTE of 99.9%? How about for a CTE of 99.999%?

3. A dark frame must be taken at the same temperature and for the same exposure time as the light frame it is meant to correct. Explain why this is so.

4. As noted, a sufficiently cooled camera has virtually no dark current. Which calibration steps must you still perform, and which can you skip?

5. While watching a video taken with a digital camcorder, you will sometimes see vertical blue streaks extending from bright objects. Using what you now know about CCDs, can you explain this phenomenon?

6. The high sensitivity of CCDs allows for much shorter exposure times, but they still must be carefully planned. Given what you now know about CCDs, what factors must be taken into account when planning exposure times?

8 Investigating Linearity & Saturation

Now, it is time to test the linear response and full well capacity of the Hirsch Observatory's SBIG[6] ST-8E CCD camera. After you have successfully connected a computer to the camera, set the cooling to a reasonable value (try −5°C) and wait for the temperature to stabilize. Meanwhile, there is some setup that needs to be done (or just checked over, since it's probably already set up).

You need to achieve a uniform illumination of the CCD, so 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. You can use electrical tape on the inside of the box to make a light-tight seal where the camera enters. Then 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 and set it to a low voltage. 15 volts seems to work well; record what value you actually use.

volts

Put the cover back on the box; the cord goes up, over, and down at a corner of the box. The box should block out all external sources of light. Since it is possible to have light leaks in the camera, you should keep the room as dark as possible and direct light sources, such as computer screens, away from the camera. Also, be sure to check that the camera's temperature is stable as you work. Record the temperature at which you are working.

°C

Take one bias frame, an image of zero[7] seconds exposure with the shutter closed. Adjust the screen display setting to bring out the detail; describe what you see in the space below this paragraph. One thing you should notice are "hot" pixels and other defects in the CCD. Pick a region that seems to avoid most of the problem areas; you will make your subsequent measurements in this area. The central quarter of the image usually works well. In fact, you can set the camera to only read out this portion of the frame, saving you from long image download times.

---

[6] Santa Barbara Instruments Group

[7] The ST-8E actually has a minimum exposure time of 0.12 seconds, which it will automatically use if you give it any shorter time.

Now you will take and save a series of flat frames, gradually increasing the exposure time until the CCD saturates. You should start close to the minimum exposure time and increase it in even increments until the CCD noticeably 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 quite 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 pixel value (use Table 1). This is in counts, but you can (and should) convert this to electrons by multiplying by the camera's *gain*. For the ST-8E, the gain is approximately 2.3 electrons/count. This means that the readout electronics report one count for every 2.3 electrons worth of charge. (Using Microsoft Excel may be useful in expediting your calculations here.)

Now make a plot of mean pixel value in electrons versus exposure time and sketch it below. (Again, you should use Excel or another program to do this.)

Mean Pixel Value vs. Exposure Time

Note where the graph levels off; record this below as the full well capacity.

electrons

You should also notice that the theoretically linear part of the graph is "bent." The ST-8E apparently deviates from true linearity before its full well capacity is reached. Estimate the size of the true linear region as a percentage of full well capacity.

%

9 Taking Pictures!

Now you will get to take and calibrate an actual picture. Remove the camera from the light box setup (be sure to turn off the VARIAC). You can form images without an attached telescope or any other lenses by making a pinhole camera. Fit black electrical tape over the small aperture attachment and use a paper clip or pin to make a clean hole in its center (this has probably been done for you ahead of time). Attach this to the camera, and you are ready to go. Record the temperature at which you are working.

°C

Make sure the camera is set to take full frames at low resolution (3 × 3 binning), then try taking images of yourself or the room. With normal room lighting, exposure times of up to several seconds (try 5 seconds) should produce good pictures. Avoid aiming at windows and lights, as they will cause part of the image to saturate. When you have saved an image you like, take and save several dark frames corresponding to the exposure time you used. You should also take and save several flat frames. This can be done by holding a piece of white paper over the pinhole and exposing the camera to a uniform light source, such as a white ceiling or distant light. Use the same exposure time as your image and dark frames.

Examine the dark frames and describe their appearance below. Average the dark frames together, save the result, and remark on the appearance of the combined image. This averaging procedure reduces the noise present in the dark frames, so less noise will be added to the final processed image.

Next, examine the flat frames and describe their appearance below. Average the flat frames together, save the result, and remark on the appearance of the combined image. Then subtract the averaged dark frame from the averaged flat frame, save the result, and comment on its appearance.

Now you can calibrate your image. Subtract the averaged dark frame from your image and save the result. Then divide by the dark-subtracted averaged flat frame and save the final result. In the space below, explain how the image's appearance changes after each calibration step. Is there significant improvement over the raw image frame?

Table 1: Record your linearity data here.

Exposure Time (seconds)	Mean Pixel Value (counts)	Mean Pixel Value (electrons)

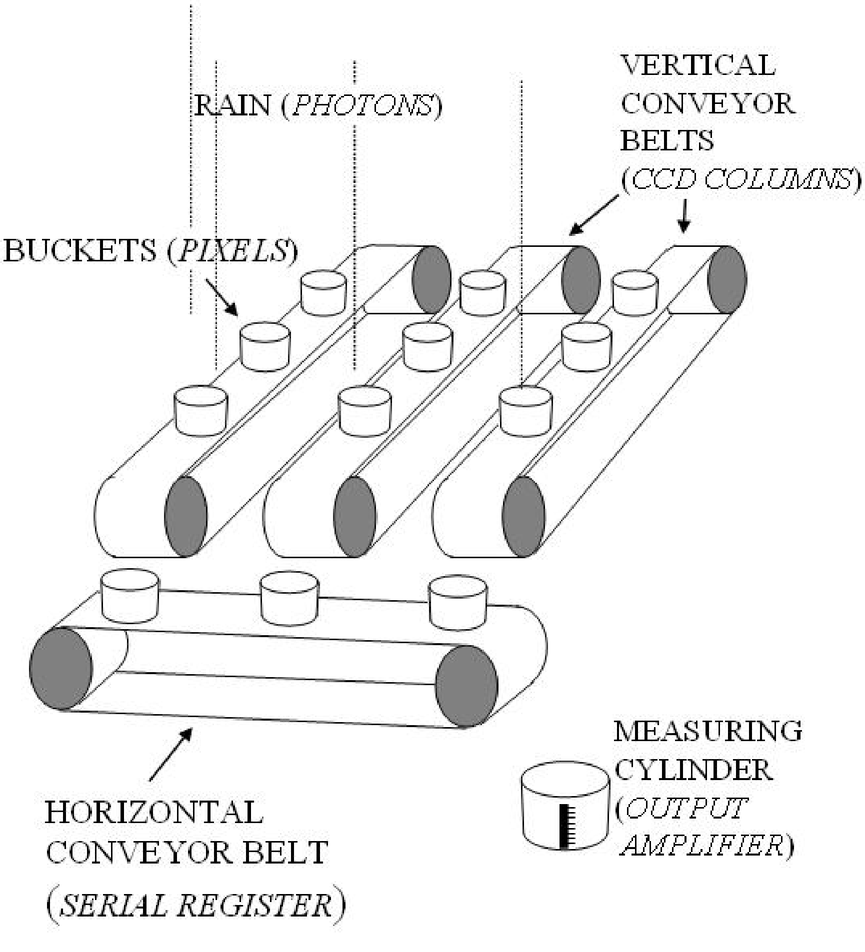

Figure 1: The "bucket brigade" analogy for CCD operation. Image from http://www.ing.iac.es/˜smt/CCD_Primer/CCD_Primer.htm.

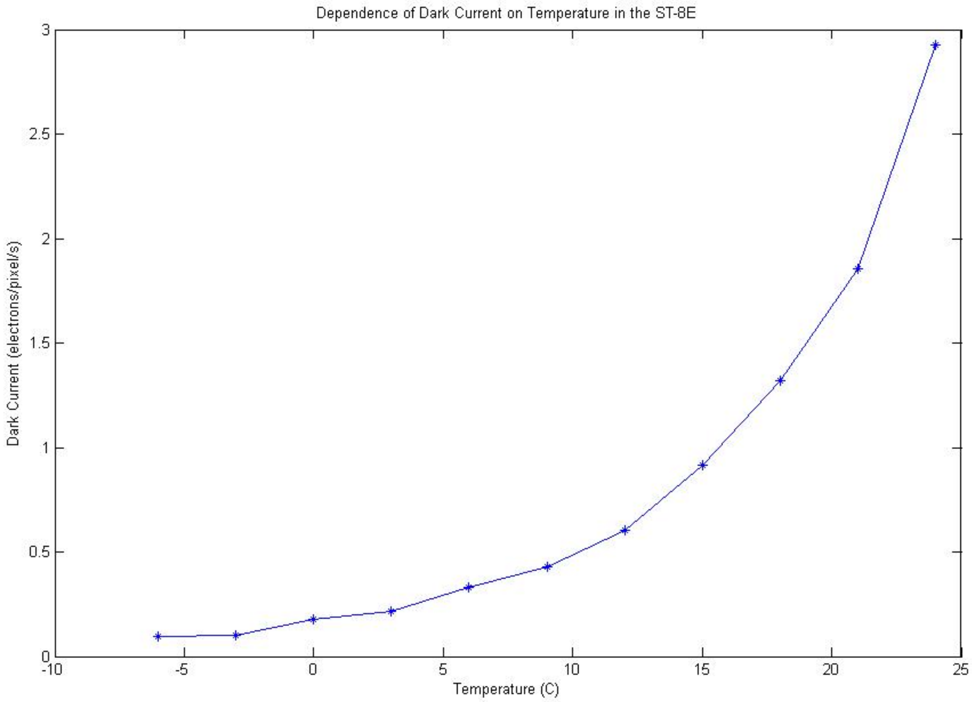

Figure 2: The variation of dark current with temperature appears to follow an exponential pattern. The actual relationship is a bit more complicated.