CCD Theory

The CCD is a semi-conductor device, very similar to a chip on a digital camera. The manufacturing process for both is essentially the same — complex circuits etched into a silicon wafer — but the eventual use is somewhat different. The CCDs used in astronomy are sophisticated, low-noise variants of the imaging devices used in digital cameras. The primary difference is that astronomical CCDs are designed to detect extremely faint objects.

A CCD is essentially an array of light sensitive elements. These elements are arranged in a grid, and each is referred to as a pixel. A typical modern CCD may have 2048×2048 pixels — just over 4 million in total. The pixels are typically 15 to 25 microns in size. A photon hitting a pixel knocks loose an electron, and hence deposits a charge on the pixel. The charge on each pixel is thus a measure of the number of photons which struck it. A typical modern CCD is 90% efficient at this capture — that is, on average, roughly 90 out of every 100 photons which strike the CCD produce a measurable signal. This is called the *quantum efficiency* (QE) of the device. The extremely high QE of CCDs is why they are so much better than photographic film — the QE of the best film is only about 1%!

CCDs have another great advantage over photographic film: linear response. At a fixed photon energy, one photon equals one electron. Period. Photographic film isn't nearly as easy to interpret — the strength of an image is not linearly related to its brightness. Also, the images produced by CCDs are digital, which allows one to analyze CCD data more directly on a computer. Photographic film must be laboriously scanned and digitized before such analysis can be done[1].

1 Preprocessing CCD Data

CCDs do, however, have some complications. A raw image from a CCD cannot be analyzed directly; several steps must be taken to pre-process the data prior to analysis. The basic pre-processing steps are: overscan correction and trimming, bias removal, dark current removal, and flat-fielding. Each is described in some detail below. [2]

1.1 Overscan Correction and Trimming

The value of the pixels on a CCD image of zero duration are considered to consist of three components[3]. The first is a bulk offset, simply the mean level of all pixels in the image. The bulk component is known as the 'overscan level', and varies from image to image unpredictably with the CCD temperature and time. It is purely an offset, and does not have an associated Poisson noise. The second component of a zero duration image is called the bias. It represents the individual pixel-to-pixel variations of the offset level. The bias often has some sort of pattern across the CCD, and does not appear random. The bias is generally fixed, despite variations in the mean offset. That is, apart from noise effects, the difference between the value of a given pixel and the overscan level of the entire image is fixed, regardless of the value of the overscan level. The third component of a zero duration exposure is read noise, which occurs because the amplifier electronics attached to the CCD are not noiseless; the simple act of reading the pixels generates a small amount of noise. Removal of the overscan level and bias are usually treated as separate pre-processing steps. The read noise, of course, cannot be removed from a single image, though other noise sources, such as the Poisson noise associated with photons, often dominate.

---

[1] There is another advantage to CCDs often forgotten by modern astronomers. Using a CCD is just plain easier than using photographic film — no fumbling around in the dark loading the cameras, no long hours in the dark-room, inhaling noxious fumes and squinting at barely visible equipment.

[2] N.B.: the following applies mainly to photometric data. For spectra, calibrations are somewhat different. For example, arc lamp exposures must be used to determine the dispersion (i.e., the conversion between pixels and wavelength along a spectrum).

[3] Indeed, all raw CCD images, whether zero duration or otherwise, contain these components, though only zero duration exposures have no additional components such as signal and noise due to photons.

To remove the overscan level we must first know what it is. This is done by analyzing the signal recorded in a set of pixels known as the overscan region. These pixels (which usually appear in a strip along one or more sides of a CCD image) are extra pixels generated by the CCD electronics when the CCD is read out. They are not in any way connected to real physical pixels on the CCD — they correspond to locations on the pixel grid which don't really exist. The mean level of the pixels in the overscan region thus gives a measure of the average signal introduced by reading the CCD, i.e. the overscan level. Once this level has been computed, it should be removed from all pixels in the image; an image with this level removed is considered to be 'overscan corrected'. We can then trim off the overscan region (keeping only those pixels which correspond to real light-gathering pixels on the CCD). This is called 'trimming' the image. This process of overscan correction and trimming should be done to ALL images prior to any other pre-processing.

And, now, after such heavy emphasis on always doing this, I should tell you that neither the FLI IMG1024-S CCD (Scarborough) nor the SBIG ST7 CCD (DA) which you'll be using for your observing produce an overscan region, and so you won't need to overscan correct any images. However, the above is certainly true for all CCDs used at all the large observatories, and serves to illustrate the process in general.

1.2 Bias Removal

After the overscan has been removed from an image and the image has been trimmed, the next step is to remove the bias. As described above, the bias is the pixel-to-pixel structure in the the read noise on an image. Since the bias varies across the CCD, we must use a bias image (a.k.a. bias frame) to remove the bias structure from *other* images. A bias frame is acquired by simply taking a zero second exposure. In this case the CCD is not exposed to any light, so the measured signal is merely the bias (+ overscan, which should be removed prior to using the bias frame to de-bias other images). In practice, one usually acquires multiple bias frames (typically 6-10) and then uses an average of these frames (a so-called 'master bias frame') to de-bias all the other images. This averaging process ensures that the signal to noise ratio (S/N) of the master bias frame is good enough that the process of de-biasing does not introduce significant extra noise into the corrected images.

1.3 Dark Current Removal

The signal recorded at each pixel on a CCD may, in some cases, have an additional component which has nothing to do with the number of photons which struck it. The signal is essentially thermal noise, or 'dark current': the motion of atoms (due to heat) in the material of the CCD itself causes some charge deposition in the pixels. To mitigate this effect, all CCDs used in astronomy are cooled to very low temperatures. The best CCDs are cooled with liquid nitrogen to about -110°C, and have negligible dark current (this is true, for example, for the CCD in use at the DDO). The less expensive CCDs (such as the ones you'll be using in this course) are cooled thermo-electrically (to typically -30° to -50°C), and have a significant dark current which must be removed.

To remove the dark current image, one must, as in bias removal, acquire a separate frame which has only the dark current signal. This is done by taking an exposure of the same duration as the images to be processed, but keeping the shutter to the CCD closed. This ensures that no light reaches the CCD, so the only recorded signal is the dark current. *It is extremely important that the dark current frames be of exactly the same duration as images to which they will be applied.* Dark current is a time dependent phenomenon; a long duration exposure will have more dark current than a short one. The dark frames must also be taken at the same temperature as the images to which they will be applied.

Like bias frames, the typical procedure for using dark frames is to acquire several (6-10 is again sufficient) and average them to construct a master dark frame. Unlike bias frames, one may end up with more than one master dark frame, with each corresponding to a different exposure time. This situation can easily arise in multi-colour photometry work; the required exposure time in different filters can vary quite dramatically for the same target. Note that individual dark frames should be overscan corrected, trimmed, and de-biased *prior* to construction of a master dark frame.

1.4 Flat Fielding

The final correction applied to a CCD image is flat-fielding. This correction accounts for the fact that the sensitivity of individual pixels on the CCD is not constant. This variation has two major sources. First, the optics of the telescope may not transmit light uniformly across the entire field of view. This may be due to vignetting (a reduction in throughput near the edges of the field) or the presence of dust on the CCD itself or the glass CCD cover (this typically appears as broad dark rings on the image). Second, even in the presence of uniform illumination, the QE of individual pixels varies across the CCD — causing a variation in sensitivity.

To correct for this effect, one must acquire flat-field images. This requires a source of even illumination — usually either a diffusing screen and a light within the telescope dome itself, or the twilight sky. Flat fields taken with a dome diffusing screen are known as 'dome flats' and those taken using the twilight sky as the illumination source are known as 'sky flats'. The exposure time of flat field images should be set so that a significant signal level is recorded (say 10,000 counts per pixel or so). Several flats should be taken to allow for the construction of a master flat, *and a separate master flat must be constructed for each filter used.* The flat field correction is made with this master flat field, by dividing the image to be corrected by the flat field. The individual flat field images should be overscan corrected, trimmed, de-biased and corrected for dark current effects prior to combining. Always be sure to get dark current frames of appropriate duration for the flat field images.

1.5 Required Calibration Frames for the DA 14”/16” + S-BIG CCDs

The above discussion on CCD pre-processing is relevant for 99% of the CCDs used in astronomy. However, the SBIG and FLI CCDs you’ll be using are somewhat simplistic devices, and the amount of calibration that needs to be done is thus somewhat less. As mentioned above, the SBIG and FLI do not produce an overscan region, so there is no need (and no way) to overscan correct and trim the images. Because of this, de-biasing and dark current effects can be corrected in one step by simply removing a dark frame which has *not* been bias corrected (we can only get away with this because there is no overscan removal step). So, there is no need to acquire separate bias frames. The required images (in addition to ’science’ frames of the target objects) are then simply:

1) Dark current frames with durations appropriate to all other images, and taken at the same temperature as all other images. In other words, if you have a set of science images which are 300 seconds each, another set of science frames which are all 100 seconds each, you will require two sets (of 4 each) of dark current frames (100sec and 300sec). Though 10 darks would be preferable, 4 is often a reasonable number given time constraints. (I’d rather have you taking science frames than spending most of your time taking darks!)

2) A set (3) of flat field frames, properly exposed, for each filter used in the science observations. By properly exposed, I mean that the signal level on each pixel should be no less than 10,000 counts, and shouldn’t be more than 30,000 counts. You’ll typically be taking dome flats, since properly exposed sky flats are somewhat difficult to take, due to the rapidly varying sky brightness at sunrise/set.

Acquiring these data can take 0.5-1.0 hours each night of observing, so be sure to give yourself enough time to get them. I recommend showing up around sunset and getting the calibration frames prior to starting the science observations. This makes packing up and going home at the end of observing a lot quicker, and doesn’t waste dark sky time, as the sky will take on order of 0.5-1.0 hour to get fully dark after the sun has set. Also, in general you must get dark and flat field images each night of observing, as neither of these are guaranteed to be constant from night to night, particularly if the CCD is taken on and off the telescope, or the temperature changes significantly. As the course progresses we will assess the stability of the whole system; if it is sufficiently stable, we may dispense with taking darks and flats every night. Until further notice, however, you must do so to make the rest of your data useful.