Chapter 11

Photometric Detectors

Telescopes simply collect photons. To make useful observations requires some sort of gizmo to detect those photons and make a measurement or record of them.

11.1 Human Eye

The "natural" detecter of visible photons is the retina of the human eye. The human eye is an amazing thing. It does its intended job- provide us with panoramic views over a large solid angle with sub- second time resolution under an enormous range of lighting conditions, from full desert sun to the the darkness of the night sky in the same desert at midnight. However, the eye is not a very good photometric device, in the sense of being able to make quantitative measurements of flux. Also, the eye is not an integrating device, that is, the signal does not build up with time. The eye does not provide a permanent record of what it sees.

At present, the human eye is not used as the primary detecting device at professional research level telescopes telescopes. Amateur astronomers make some useful observations using their eye-balls. One field where the eye is useful is in seeing fine detail on planetary surfaces. The eye can can make use of the brief instants when the atmosphere steadies and the seeing, or image sharpness, becomes momentarily very good. Amateurs also use their eyes to make observations of the brightnesses of variable stars by comparing the brightness of the variable star to the brightnesses of comparison stars of known magnitude. The eye can make some useful, but limited, photometric measurements in this way.

The angular resolution of the human eye is limited by the small size of the lens. Under dark conditions, the lens is about 7 mm in diameter. The theoretical resolution is about 16 arcsec (about a quarter of an arcmin), but few people have eyesight that approaches this resolution.

The eye is sensitive to only a small slice of EMR wavelengths. The wavelength range of sensitivity is slightly dependent on whether the eye is observing in bright light conditions (photopic) or under very low light level conditions ("dark adapted" or scotopic). Figure 11.1 shows the relative response of the eye as a function of wavelength.

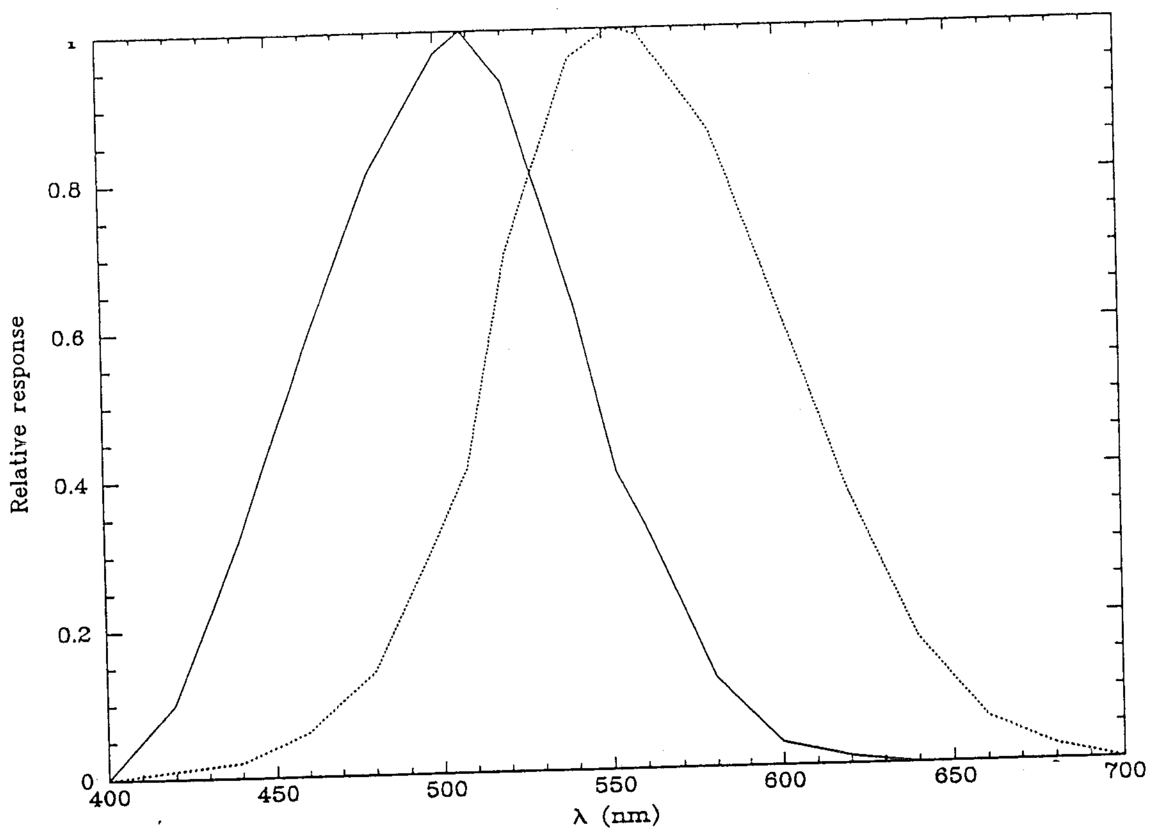

Figure 11.1: Wavelength sensitivity of the human eye. Solid line is dark-adapted (scotopic) response. Dotted line is for light- adapted (photopic) eye. The y axis is not calibrated in terms of any real units- it is simply relative to the peak of the response.

11.2 Photographic Emulsions

The first technological advance in astronomical detectors was the photographic emulsion. Photographic emulsions can make a permanent record of astronomical objects imaged by telescopes. However, photographic emulsions are not a very good photometric devices for astronomy because of several drawbacks. Briefly, these are: photographic emulsions only record a small fraction (around 1%) of the photons that hit the emulsion. Because of the analog (rather than digital) nature of the image record on an emulsion, it is difficult to make quantitative measurements of star brightnesses. Photographic emulsions are also nonlinear- twice the input light does not produce twice the output on the film. This feature is called **reciprocity failure**. One advantage of the photographic plate is that it can be made larger than the largest CCDs, at least in the year 2000 (CCDs are getting bigger as time goes on). However, this may be a moot point. The demand for astronomical plates has pretty much gone to zero, and there are rumors that Kodak (which never really made any money from the small specialized astronomical market) has or will soon quit making astronomical emulsions.

11.3 Modern Detectors - PMT and CCD

Modern detectors are inherently **digital**. They detect individual photons and output a number which is directly and linearly related to the number of photons that were incident on the detector.

The first modern detecter in this sense is the photomultiplier tube (or PMT). A PMT consists of an evacuated glass tube, on one end of which is deposited a film of a material (such as indium antimonide) called a **photocathode**. This material has the property that when it is struck by a photon, an electron is often liberated from the material. Each electron liberated from the cathode is directed away from the cathode by an electric field, and is amplified into a **pulse** of electrons by a series of metal plates (called dynodes) and an accelerating electric field in the tube. Electronics coupled to the PMT counts these pulses. Thus, a **single** photon hitting the cathode results in an easily counted pulse of many electrons. (See Figure 11.2 for a diagram of a PMT and Figure 11.3 for a figure of a photometer, the gizmo that holds the PMT and associated equipment onto the telescope.)

The drawback of a PMT is that it is essentially a single channel device, meaning there is **no** positional information in the signal. The output signal does not depend on where on the cathode the photon hit, so we get only a measure of all the light that fell on the photocathode. However, the PMT, unlike a photographic plate, has a digital output, meaning we can easily make quantitative measurements. The fraction of the photons which hit the cathode that are actually detected by the PMT is set by the fraction of photons that hit the cathode that liberate an electron. This fraction is typically 20% or so, so the efficiency of observing is much higher for a PMT than for the best photographic emulsion.

The detector of choice for optical astronomy is now the CCD (charge coupled device). The CCD has many advantages- it is a linear, photon counting device which records a large fraction of the photons that fall on it. It is far better than a PMT because it can record a two dimensional image- i. e. there is positional information. CCDs have truly revolutionized astronomy in the last

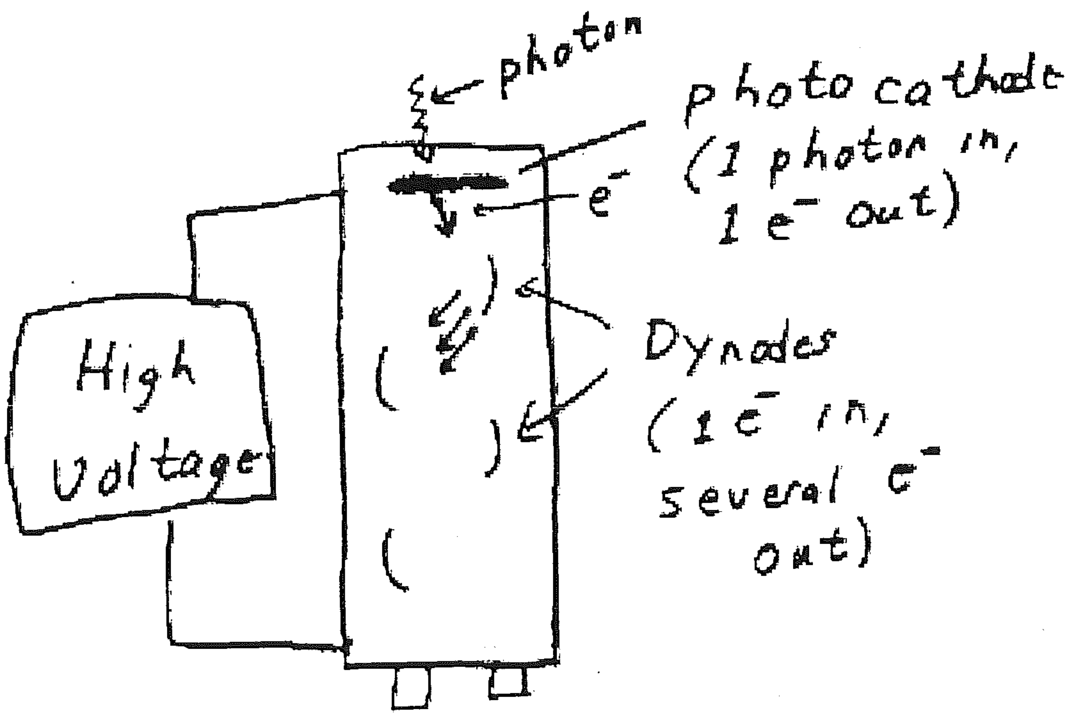

Figure 11.2: Schematic diagram of a photomultiplier tube (PMT). The high voltage supply creates an electric field that accelerates electrons along the tube. At each dynode, an impinging electron releases several electrons, which are then accelerated towards the next dynode, where each of them knock loose several more electrons. Through this cascade, a single photon hitting the photocathode releases an easily counted **pulse** of many electrons.

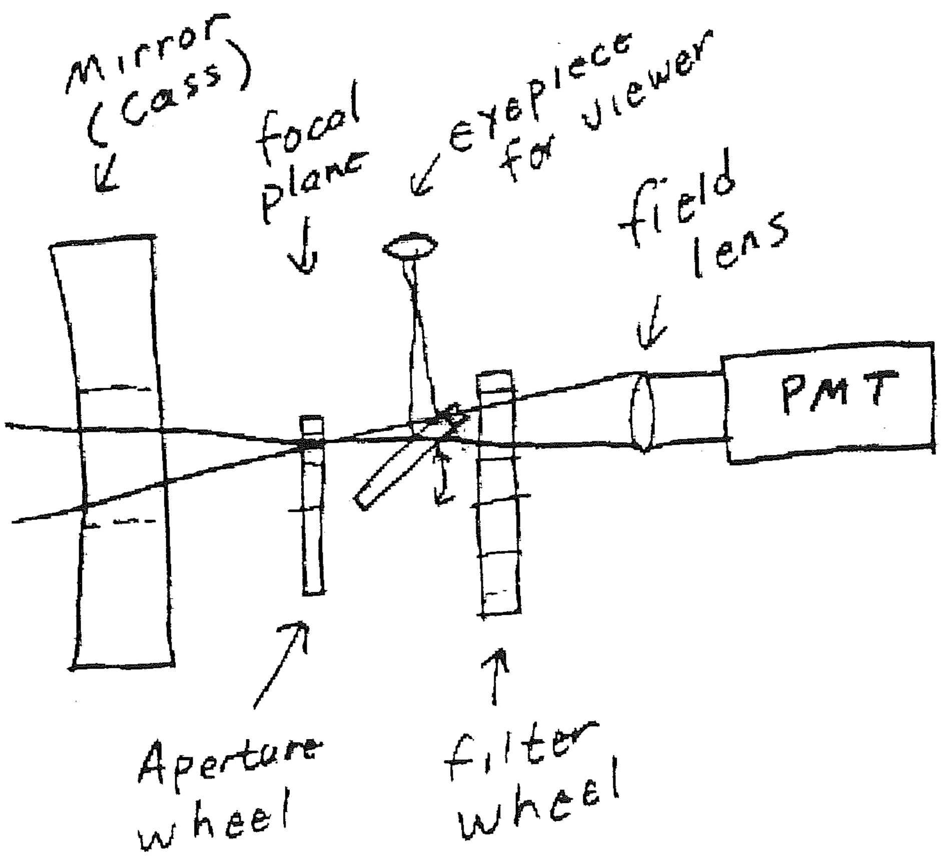

Figure 11.3: Schematic diagram of a simple photometer. The aperture wheel contains holes of various sizes that define the angular size of the spot measured by the photometer. Behind the aperture wheel is a mirror that can be flipped into the light path to direct light to an eyepiece. This "behind the aperture viwer" allows the astronomer to visually center the star in the aperture. The mirror is flipped down to allow light to pass to the detector. The filter wheel holds various filters that define the passband measured by the PMT. The field lens directs the light to the photocathode of the PMT. The PMT is usually cooled by dry ice.

two decades. The next chapter is devoted to further discussion of these amazing devices.

Further Reading

Sky on a Chip: The Fabulous CCD Sky and Telescope Sept. 1987. Dated in spots (the first sentence is no longer true!), but still an excellent introduction to CCDs.

Chapter 12

CCDs (Charge Coupled Devices)

12.1 Basic Concepts

A CCD is a light sensitive silicon "chip" which is electrically divided into a large number of independent pieces called **pixels** (for "picture elements). Present day CCDs have 512 x 512 (262144) to at least 4096 x 4096 (16,777,216) individual pixels, and are from about 0.5 cm to 10 cm in linear size (typical sizes of each pixel are 10 to 30 micron square). For astronomical use, we use the CCD as a device to measure how much light falls on each pixel. The output is a **digital image**, consisting of a matrix of numbers, one per pixel, each number being related to the amount of light that falls on that pixel. Of course, one of the beauties of the CCD is that the image, coming out in a digital form, is readily manipulated, measured, and analyzed by computer. Research astronomers spend **FAR** more time sitting in front of computers than anywhere near telescopes!

Several concepts are basic to CCD use as a low light level detector in astronomy. The following should give you enough information that you can understand the reasons for the various steps in the computer data reduction, which we will do later in the course:

Quantum efficiency (QE)- A CCD detects individual photons, but even the best CCD does not detect every single photon that falls on it. The fraction of photons falling on a CCD that are actually detected by the CCD is called the quantum efficiency (QE), usually expressed as a percentage. (Note: There is another way of defining QE that involves the signal to noise ratio of the input and detected signal. For CCDs in which photon noise is the dominant noise source, the fraction of photons detected and the signal to noise definitions of QE are equivalent.)

QE is a function of wavelength. For optical detection, there are two basic styles of chip: **thick chips**, or "frontside illuminated chips", in which the light passes through some of the electronic layers of the CCD before hitting the silicon detecting level, and **thin chips**, ("backside illuminated") in which the silicon layer is mechanically or chemically thinned and the light enters the silicon directly (see Figure 12.1). Thick chips have low QE in the blue, because the electronic layers absorb much of the blue light. Thin chips have better blue QE. Thin chips and thick chips have more similar red QE, but the thin chips usually have higher QE at all wavelengths than the thick chip. Figure 12.2 shows the QE vs. wavelength for several CCDs. The Steward/Loral chip and the

NURO chips are thinned devices, and have blue response. The original Kodak chip is a thick device that has little blue sensitivity. Kodak has recently released an enhanced "E" chip that has greatly improved blue QE. The "E" chip is still a thick chip, but uses more transparent gate structures and materials than the original chip. (Unfortunately, when we bought our ST8 CCD, there were no "E" chips.)

I am not sure how seriously to take the numbers in CCD QE curves for various CCDs. You all remember YMMV - "your mileage may vary" (but somehow always seemed less than the number advertised). Just remember YQEMV. (We at OU will be getting a new CCD soon, an Apogee Ap7. This will allow us to directly compare the QE of that CCD with the Kodak chip we now have. It will be interesting to compare the results with the QE curves.)

To make a thin chip, you start with a standard thick chip, then thin the silicon layer. (The silicon layer must be very thin in a back-side illuminated chip, or else electrons will never make it to the counting structures.) "Thinning" the silicon layer to the required thickness in a uniform manner without destroying the device is a real Black Art!

The OU ST-8 CCD camera uses a thick chip made by Kodak. The camera is made by SBIG (Santa Barbara Instrument Group: www.sbig.com). The chip has 1530 x 1020 pixels, each 9 microns square. (As we will see, we usually bin the pixels togther in groups of 2×2 or 3×3 when reading the device out to make the chip appear to have 765×510 or 510×340 pixels). The price of such a chip depends on the quality- our chip is not the best grade (defined by the number of defects in the chip) and the price of the chip itself was about \$2500. ("Perfect" chips of this type are about \$10,000). Large, thin chips are very expensive- many hundreds of thousands of dollars. In the last few years, several astronomical camera manufacturers (notably Apogee Instruments: www.apogee-ccd.com) have offered cameras with thinned CCDs at prices low enough for small observatories and dedicated amateurs. Major observatories have their own labs where they take commercial (thick) chips and thin them to get good blue QE (using ancient alchemistic incantations and virgin sacrifice!- Just kidding about the virgin sacrifice, at least in Arizona where its against the law.) These labs also "package" the devices. This doesn't mean sticking them in cardboard boxes, but rather the difficult and delicate job of mounting the thinned CCDs so that these very fragile devices are optically flat (they tend to curl like potato chips when thinned) and securely anchored to their base. (Most professional CCDs are mounted in vacuum chambers, so that frost doesn't form on them when they are cooled. Improperly packaged CCDs have been sucked into the vacuum pump when it was first turned on!)

Counts - So, are those numbers that we read out of the CCD the actual number of photons that fell on each pixel? Well, no. First off, part of the number is an electrical offset called the **bias** (see below) and part may be due to **dark current** (see below). After we subtract these components, the signal is related to the number of electrons liberated by photons in each pixel. Only a fraction QE of photons generate electrons, so that the number of electrons is: (number of photons) × QE. For several technical reasons, the numbers the CCD give out are related to the number of electrons by a divisive number called the **gain** (the gain is usually a small number greater than 1). So, basically, the number of photons that fell on a pixel is related to the output number (sometimes called "DN" for data number) as follows:

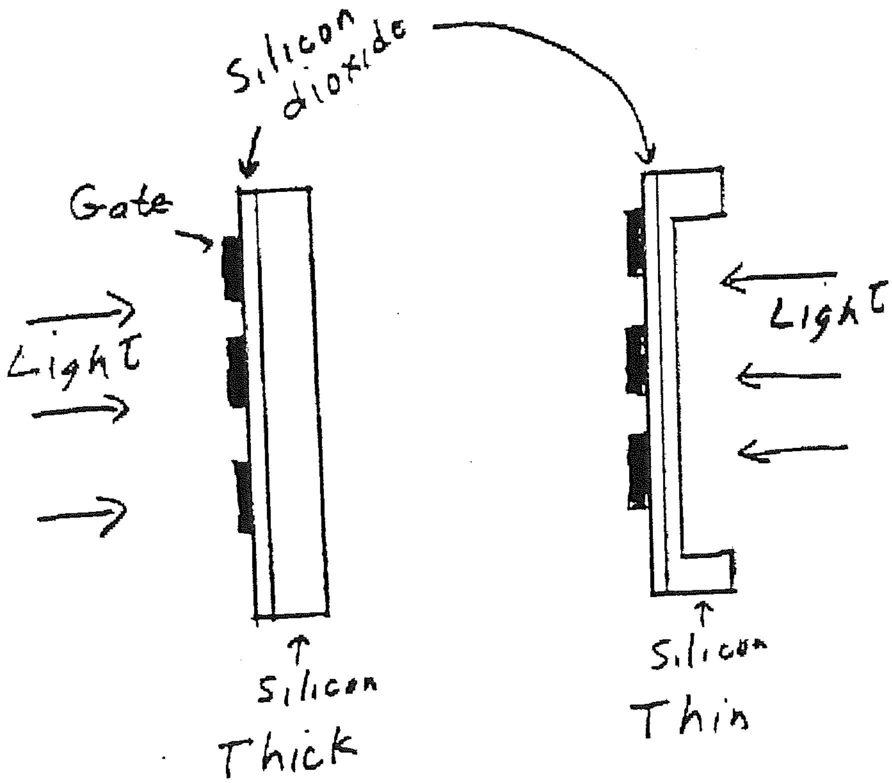

Figure 12.1: Thick or thin at Pizza Inn. The thick chip (left) has a silicon layer about 500 microns thick. Light enters the silicon from the left, after passing through the gate structures. Thin chip (right) has silicon layer only 10 or so microns thick. Light enters from the right, so does not have to pass through the gate structures.

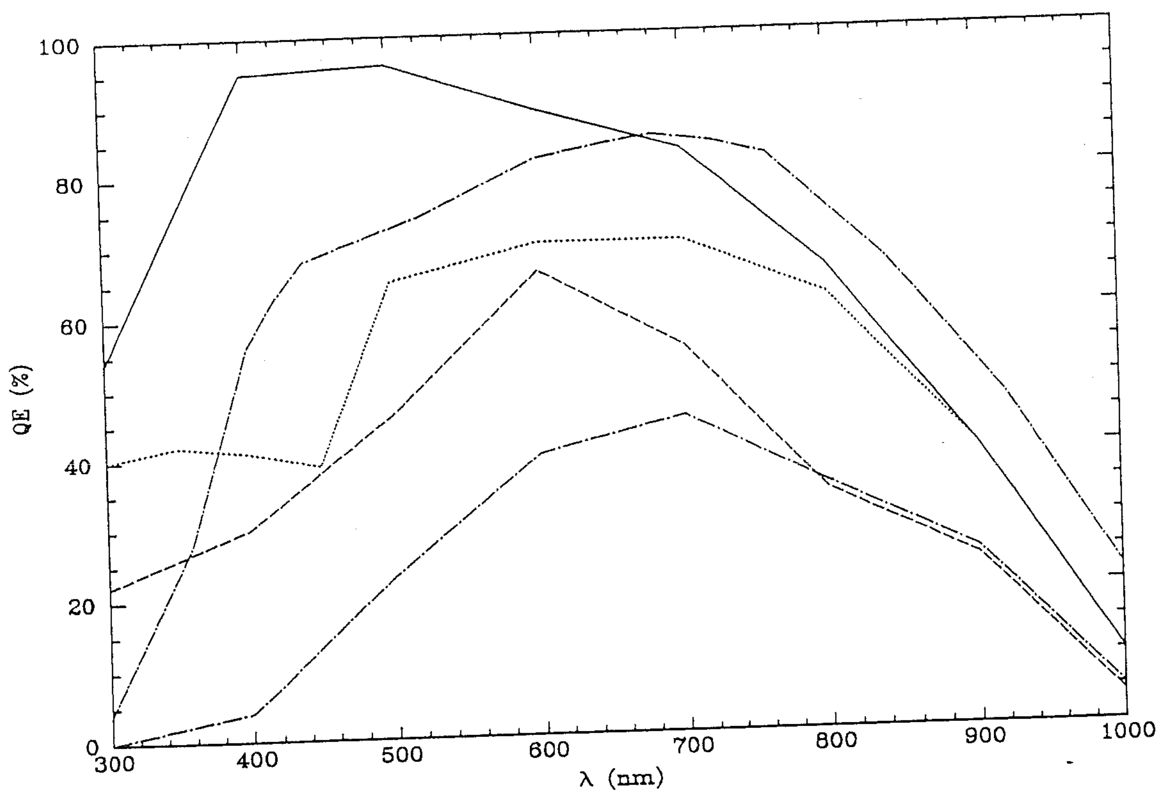

Figure 12.2: QE curves for 5 CCDs. From top to bottom at 500 nm: Solid curve: Steward Observatory thinned CCD ; dot-dash curve: Site back-illuminated (in AP7 camera) ; dotted curve: NURO TEK thinned CCD ; dashed curve: Kodak "E" (for Enhanced) chip. This is a thick chip, but one with more transparent gate material than usual ; dot-dash curve: old (non-"E") Kodak thick chip

$$\text{Photons} = \frac{\text{number of electrons}}{\text{QE}} = \frac{\text{gain} \times \text{DN}}{\text{QE}} \tag{12.1}$$

here the number of electrons is only those that came from photons- i.e. the bias and dark contribution have been subtracted off.

When doing astronomical photometry, we don't usually calculate the actual number of photon per pixel, as we make our measurements by ratioing the DN for our objects to the DN for stars (called **standard stars** whose flux we know. (More about this later, of course)

Integration time- The CCD (unlike the human eye but like a piece of film) is an **integrating** device. The signal (electrons knocked loose from the silicon by impinging photons in each pixel) builds up with time. The integration time (or exposure time) is controlled by a mechanical shutter (like in a camera) or electrically (changing voltages in CCD).

Read noise- After an integration (exposure), the CCD must be "read out" to find the signal value at each pixel - because the signal may be as low as a few electrons per pixel, this step involves some very sophisticated amplifiers that are part of the CCD itself ("on chip" amps). Unfortunately, but inevitably, the read out process itself generates some electronic noise. The average noise per pixel is called the **read noise**. Modern CCDs typically have a read noise of 5 to 20 electrons per pixel per read out (read noise is the same whether exposure is 0.1 sec or 3 hours).

Bias frame- If we simply read out the CCD, without making an integration, (or think of a zero second integration), there will be a signal called the bias signal. (You might think that the bias would be identically zero, but it isn't. Think of it as an electrical offset or background.) This bias signal must be measured (it changes somewhat with things like CCD temperature) and subtracted from the images we take. Since there is read noise associated with ANY readout of the CCD, even bias frames have read noise associated with them. To minimize noise introduced when we subtract the bias, we take many bias frames and them combine them to "beat down the noise".

Dark frame- If we allow the CCD to integrate for some amount of time, WITHOUT any light falling on it, there will be a signal (and more importantly **noise** associated with that signal) caused by thermal excitation of electrons in the CCD. This is called the **dark signal**. The dark signal is very sensitive to temperature (lower temperature = lower dark signal), and that is why CCDs used in astronomy are cooled (often to liquid nitrogen temperature). Even with cooling, some CCDs have a non negligible dark current. This must be measured and subtracted from the image. As for the bias, we want to take many dark frames and combine them to beat down the noise. (The dark frame and bias frames are *NOT* the same thing!)

Flat frame- All CCDs have non-uniformities. That is, uniformly illuminating the CCD will NOT generate an equal signal in each pixel (even ignoring noise for the moment). Small scale (pixel to pixel) non-uniformities (typically a few percent from one pixel to next) are caused by slight differences in pixel sizes. Larger scale (over large fraction of chip) nonuniformities are caused by small variations in the silicon thickness across the chip, non-uniform illumination caused by telescope optics (vignetting) etc etc. These can be up to maybe 10% variations over the chip. To correct for these, we want to shine a uniform light on the entire CCD and see what the signal (image) looks like. This frame (called a flat) can then be used to correct for the non-uniformities (we divide our images by the flat).

Data (object) frame- To take an image of an astronomical object, we point the telescope at the right place in the sky, and open a shutter to allow light to fall on the CCD. We allow the signal to build up (integrate) on the CCD for some length of time (anywhere from 1 second to 1 hour) and then read it out. The exposure time used depends on many things. The basic goal is to get an image of the source with the best signal to noise ratio (S/N) possible in a given amount of available telescope time. Now, since the signal is composed of photons, there is an unavoidable noise associated with photon counting statistics ("root N" noise, where N is the number of photons collected, to be discussed later in detail). There is NO WAY to get rid of this noise. HOWEVER, by collecting more photons, we can improve the S/N (the signal goes linearly with time, while the noise goes as the square root of time). During the integration, the dark signal is also building up. We also have to worry about other sources of noise- readout, dark, and also the effects of cosmic ray particles, which give a spurious signal. The first thing to insure is that these other sources of noise are much less than the photon noise, so that we are not limiting ourselves unnecessarily. This argues for a long exposure. However, cosmic rays argues for several shorter exposures which can then be combined (as you might imagine, the CCDs aboard the HST have real problems with cosmic rays!). Getting the optimum exposure time is very complicated!

** NOTE: the following applies to "professional" CCD systems (like at NURO)- see following for some changes for "amateur" systems**

SO the basic steps in observing and reducing a CCD image taken with a telescope are as follows:

Collect a number of bias frames - median combine them to a single low noise bias frame

Collect a number of dark frames (no light, finite integration time equal to the data frame integration). If dark current is non- negligible, combine dark frames into a single low-noise frame (after subtracting bias frame, of course)

Collect a flat frame in each filter- flat frames can be made by pointing the telescope at the twilight sky , or by pointing at the inside of the dome. Bias frames (and dark frames, if the dark current is non- negligible over the time interval covered by the flat exposure) must be subtracted from the flat frame. The signal level in the flat is arbitrary- it is related to how bright the twilight sky was etc- all we need is the information on the differences of the signal across the chip. Thus, we normalize the flat so that the average signal in each pixel is 1.00 (we do this simply by dividing by the average signal).

Subtract low - noise bias frame and low noise dark frame from object frame. Then divide this by the normalized flat frame.

Symbolically:

$$Reducedframe = \frac{(Rawobjectframe) - (lownoisebiasframe) - (lownoisedarkframe)}{(normalizedflatframe)} \tag{12.2}$$

(For most modern professional grade CCDs the dark current is so low that we can ignore it- We always take dark frames, as one test just to make sure the system is operating properly.)

12.2 Amateur vs. Professional CCDs

CCD systems come in a range of sophistication and price levels. “Professional” systems, employed by major observatories, are usually cooled by liquid nitrogen to an operating temperature of −100 C or colder. "Amateur" or "advanced amateur" systems usually use thermoelectric cooling systems. These thermoelectric systems typically cool the CCD to 20 to 40 C below the ambient temperature. These systems thus operate at temperatures of 0 to −40 C, depending on the temperature in the dome. At these temperatures, dark current can be a problem, and is a major problem with some chips. Also, the temperature stability of the CCD is typically much less prescise in these systems, as opposed to the more carefully engineered professional systems, so that changes in the dark current are also a source of problems.

Becuase of these considerations, amateur systems are sometimes operated in a slightly different mode than the professional systems. Often, dark frames (of length equal to the object frames) are taken before and after each object frame. These are then avearged and subtracted from the data frames. Here the "dark" frames are really "dark+bias" so that separate bias frames are not taken. The problem with this type of procedure is that it is inefficient at using telescope time- much time is devoted to dark frames during the night.

There are many possible variations in these procedures. We will have to determine the optimum procedure for our particular CCD. The Kodak chip we are using has a reputation for very low and stable dark current (there is a tradeoff for this- the QE is not as high as some other "amateur" chips). We should be able to get away with making a "library" of dark frames at different temperatures, instead of taking frequent dark frames during a night.

so, for a system where before and after dark+bias frames are needed:

$$\mathrm{Reducedframe} = \frac{(\mathrm{Rawobjectframe}) - [\frac{((\mathrm{dark+biasbefore})+(\mathrm{dark+biasafter}))}{2}]}{(\mathrm{normalizedflatframe})} \tag{12.3}$$

where the dark+bias frames are the same exposure time as the object frame

12.3 Flat Field Frames

There are a number of ways of getting flat field frames. I have found the twilight flat method to work well, at least for small chips on large telescopes, where the field of view is small. After the Sun sets, we take images of the twilight sky, which should be a reasonably uniform light source. But getting good twilight flats can be hard, particularly if a number of filters are involved. If you are waiting to observe, twilight sems to last too long- if you are trying to get good twilight flats, the sky seems to get dark too quickly, particularly is you have a number of different filters to get flats for! One good practice is to make a number of flat field exposures in each filter, moving the telescope between exposures. Then if stars appear in the frames, you can get rid of them by appropriately scaling and combining the images. To combine the images and remove the stars, we would use a median combine algorithm (discussed later).

As chips are getting larger, and field of views are also getting larger, we have to worry about the fact that the twilight sky is not uniform in brightness- for example it is obviously brighter in the west, near the setting Sun than overhead. However, by picking the right position in the sky, we can minimize the gradient in the twilight sky brightness (see article listed in references to this chapter).

Another way to get flat fields is to use a screen in the dome, and illuminate the screen with artificial lights. With effort, this also works well, although there are a number of potential concerns. One is that the electric lights used to illuminate the spot are quite a bit redder than the sky. This concern can be partially addressed by using very hot lamps, or using filters that decrease the red light from the lamps. For exacting work, particularly in the blue, you must also worry about the reflectivity of the screen. A screen may look fine and uniform with your naked eye, but be quite non- uniform at wavelengths less than the eye can see. Special paint concoctions with high UV/blue reflectivity have been formulated.

For any flat field system, one often overlooked concern is **scattered light** in the optical system. If one is trying to do photometry of stars , then one should only use light that has been imaged by the complete optical system. Scattered light is light that reaches the CCD without going through the complete optical system. For instance light may come through the front of the telescope and bounce off of the inner surface of the baffle tube found in most cassegrain systems an directly on to the CCD. If you use a flat field image that includes scattered light, then it will **NOT** properly correct the images of stars, where all the light in the image of a star passes through the optical system.

One way to test for scattered light problems in flat fields is to image a bright star near zenith on a clear nite at a number of positions across the CCD. For each image, look at the brightness of the star on a field that has been bias and dark subtracted, but not flat fielded. Then do the flat fielding and again look at the constancy of the star brightness across the chip.

For "pretty pictures", most any flat field will do. For photometry, one must be vary careful to eliminate all forms of scattered light. Don't just assume that because your images "look nice" that they are photometrically correct.

Further Reading

Special Considerations for Flat Fielding by F. R. Chromey. CCD Astronomy Fall 1996.

$\frac{0.000014}{15} = 0.000000 9333$

Chapter 13

Computer Image Processing

Computer image processing is a general name for using a computer to manipulate images or make measurements of something on the image. Commercial image processing programs (e.g. Photoshop) are usually concerned solely with the visual appearance of the image (e.g. editing out that jerk Larry from the family photos after what he did to your sister.) Scientific image processing programs are usually more concerned with making some sort of **quantitative** measurement of something on the image. There are a number of image processing programs available that are tailored to astronomical images. These range from small homebrewed systems for low end computers to staggeringly complex systems for the fastest machines available. We will use a program called IRAF that will discussed in the next chapter.

No matter which program we use, there are some basic ideas and concepts of image processing as applied to astronomical CCD images. Here I give a very brief introduction to some common image processing opertaions. You will get to use IRAF to do these operations in a series of computer projects.

13.1 Image Format

The word pixel can refer to an area of silicon on a CCD or to one tiny piece of the picture. A CCD is a collection of pixels arranged in rows and columns, and a picture the CCD produces is an array of pixels arranged in rows and columns. Each pixel in the image is represented by a number that is related to the amount of light that fell on each pixel on the CCD. Pixels on the CCD have a finite size. The area of sky imaged on each pixel also has a finite angular size, set by the (linear) size of the pixel on the CCD and the plate scale of the telescope. A typical CCD has square pixels with sides of length 24 microns (s = 2.4E−5 m). On a 4 meter, f/2.7 telescope, with a focal length of f = 10.8 meters, each pixel subtends an angle of $\Theta = s/f$ or 2.22E−6 radians, or about 0.46 arcsec.

A CCD image can be thought of as a 2 dimensional array of numbers. We specify a single pixel with an x and y value from the origin. Instead of x and y, we often refer to **rows** and **columns**. The origin is (usually) at the lower left corner of the image. Rows run parallel to the floor, and columns perpendicular to the rows. In x,y notation, a row is all the pixels with the same y value, and a column all the pixels with the same x value.

The brightness of each pixel is stored as a number. The size of the computer file representing the image depends on the type of number used to store the image. A raw image, the image as read out of the CCD, is often stored in an integer number format. Such images are often stored as 2 byte (or 16 bit) integers. Thus, the value of each pixel can only take on one of $2^{16} = 65536$ different values. (This "discreteness" is not a problem at this stage, because the output of the CCDs is in this integer form.)

However, when we go to do various processing steps to the images, such as flat fielding, we find that representing the value of pixels as integers is a bad idea. This is because integer arithmetic **truncates** numbers - divide 27 by 7 and the (integer) answer is 3, not 3.857.., or even 4, which is the rounded answer. So, one of the first steps in reducing the images is to convert the integer numbers into real numbers. Real numbers are usually stored in 4 bytes on most computers.

Thus, an integer image from a 1024 $\times$ 1024 CCD will require 1024 $\times$ 1024 $\times$ 2 = 2,097,152 bytes (or 2MB) of storage. The same size image in real format will require twice as much storage (4 MB). A night of observing can produce hundreds of images- you can see why astronomers are always looking for bigger and faster hard disks and tape backup units!

Computer files of images usually contain some header information (with info like when and where the picures was taken). but the space taken up by the header information is usually trivial compared to the space taken up by the pixel data.

13.2 Image Format - FITS

There are a number of different image processing programs used by astronomers, and astronomers use many types of computers that store files in different ways. To allow images to be easily transfered between computers, astronomers have developed something called the flexible image transport system (**FITS**). FITS is an image **interchange** format. Each image processing program has a task to read FITS images, converting from FITS to the internal format required by the program and computer, and each program has a task to write images into FITS files. So if you want to send an image to someone, you don't need to know what kind of computer or program she is using- simply write a FITS file, send it, then the person at the other end will read the FITS file, converting it to the internal format required by her program and computer.

13.3 Basic Image Arithmetic and Combining

We can combine two input images into one output image using basic arithmetic operations. The usual case is that the input images and the output image have the same dimensions in pixels. To add (subtract, divide, multiply) two images, we simply add (subtract, divide, multiply) the values at each pixel. We can also do these operations with a constant replacing one image- e.g. we can create a new image by dividing each pixel of an image by a number.

Another important task is to **average** two or more images into one image. If we make an unweighted average of n images, the output image at each pixel is simply the sum of the values at that pixel divided by n. Sometimes we wish to weight the images differently. Say we have 2 images of an object, but one has higher noise than the other. An unweighted average would give equal weight to the two images. The lower S/N image would still add some useful information, but obviously not as much as the higher S/N image. To make an optimal combination, we would want to weight the higher S/N image more than the lower S/N. IRAF allows many several different weighting schemes- in the above example, we might want o weight the images by the inverse of the noise in each image, so that the higher S/N image was given more weight than the lower.

Another useful method of combining images is the **median combine**. To find the median of a set up numbers, we order them and pick the middle value. To median combine 3 images, look at the pixel values of the 3 images at each pixel, find the middle value, and put that value in the output image. The median combine is particularly useful to get rid of noise spikes such as cosmic rays.

13.4 Smoothing Images

We sometimes want to smooth an image. This helps to reduce noise in the image, but at the expense of reducing spatial resolution. A **boxcar smooth** replaces the value at each pixel with the average of a rectangular region around the pixel. The number of pixels in the smoothed image equals the number in the original image. A somewhat related operation in the **block average.** Here a new image is created by averaging pixels which are in a box in the original image. This does change the size (in pixels) of the image. Say we start with an image which has 1024 × 1024 pixels, and block average the image using a 4 × 4 pixel box. The resulting image will have 256 × 256 pixels.

Another way to smooth images is with a **median smooth**. Here, each pixel is replaced with the median of the pixel values in a rectangular region of the original image. The median smooth is different from the median combine- the median smooth works on a single image, the combine on 3 or more images. Median operations are useful when we want to get rid of unwanted signals- say cosmic rays, or stars in a twilight flat field exposure. You have to always be **extremely** careful that you don't get rid of or change of the signal you want to measure!

13.5 Image Flipping and Transposing

Image flipping is just changing the order of the pixels. E. g. to flip an image left to right, we would just replace the first pixel in each row by the last, the 2nd by the next to last etc.

Transposing an image is just like transposing a matrix. We interchange the rows and columns.

13.6 Image Shifting and Rotating

Image flipping or transposing result in the same image, just viewed from a different coordinate system. There are also operations which result in a different image by shifting or rotation. For any of these operations, we always have to worry about edge effects. We do not know what is beyond the edge of an image.

We often have to shift images. One common example is when we want to combine several images of the same field which are slightly offset, due to drift in the telescope tracking. If the images are not undersampled, we can make shifts of a fraction of a pixel. There are several mathematical techniques to do this. One that works well is called **bicubic splines**, which you can roughly think of as a "numerical French curve". Unless an image is grossly oversampled (very small angular size pixels) a linear or bilinear shifting algorithm does not work well, as such an algorithm cannot get local maxima (e. g. tops of stars) right.

We can rotate images, by arbitrary angles around arbitrary centers. We can use one of a number of image interpolation schemes, such as bicubic splines.

13.7 Image Subsections

Often we want to deal with only a portion of an image. IRAF has a very nice way to deal with this. Say we want an image consisting of the first 700 rows and columns of an image originally 800 x 800 pixels. We could make a new image by

```
imcopy big[1:700,1:700] small
```

The stuff in the square brackets is called the image section.

13.8 Mosaicing

Another way of combining images is to make several small images into one larger image covering a larger piece of the sky. If the image were of precisely adjacent parts of the sky, with no overlap and no "gaps" on the sky, you can just think of putting them together like pieces in a jigsaw puzzle. In reality, getting images aligned like this is hard, so usually when one wants to image a piece of sky bigger than the telescope field of view, one makes a number of images with considerable overlap. A mosiacing program can help align the images (by finding stars in common between images in the overlap region) and do the necessary image trimming and combining.