4 ASTRONOMICAL DETECTORS

Human eyes respond only to the *rate* at which light from the source is reaching the retina. Once this rate falls to the threshold of sensitivity, the visibility of such a weak source cannot be improved by gazing at it for a long time. In fact, because of stress and tiredness, the very opposite tends to happen. Hence, although a telescope with a great light-gathering power undoubtedly helps to discover new weak sources, the sensitivity threshold of the eye is still a limiting factor.

4.1 Integrating detectors

If we replace the eye by an **integrating detector**, then by using long time exposures it is possible to detect sources that are several orders of magnitude weaker than can be detected with the eye. An integrating detector is so called because it can integrate (i.e. add up) the light it receives over a long period of time. The only limitation here is the background brightness of the sky itself. If the exposure is long, the detector will eventually record the intensity of the scattered light in the atmosphere, and the faintest astronomical sources will remain lost in this background. Ultimately, it is a combination of expert judgement, trial and error, and often a measure of luck, which leads to the most perfect images of the night sky.

In order to take images with long exposures, modern telescopes are equipped with sophisticated automatic guiding devices. As noted in Chapter 2, they make it possible to fix the field-of-view of the telescope on one particular object (or on a particular section of the sky) and to keep this field-of-view constant with such smoothness and precision that the images exhibit no loss of resolution, even though the Earth has been rotating around its axis and moving along its orbit during the exposure.

The importance of taking images through a telescope lies not only in the fact that images can record weaker sources than the eye can see; equally important is the fact that such images provide a permanent and accurate record of the observation.

An integrating detector is also capable of recording finer details in the structure of extended celestial objects, or of separating more closely spaced point-like objects, than can be immediately seen by the eye through the eyepiece. One reason for this is that the individual recording elements of a detector can be packed more closely than the receptors on the retina of the eye. Another reason is that the detector is often positioned in the focal plane of the primary mirror (or lens) and is therefore not affected by the residual aberrations of the eyepiece. The detector can therefore, in principle at least, make full use of the angular resolution of the telescope. Although, in practice, such limits of resolution are impossible to achieve when observing stars from the Earth's surface because of atmospheric turbulence, there is still a gap between the best angular resolution that can be achieved by the telescope and the acuity of the eye. An imaging detector can record details on scales of 1″ or better that can subsequently be enlarged to make them visible by eye.

Black-and-white photographic emulsion was the first, and still is the simplest, type of integrating detector used in connection with astronomical telescopes. However, for several reasons, photographic emulsion is *not* an ideal detector of the relative brightness of celestial bodies. Firstly, its response is non-linear. This means that the intensity recorded on the developed photograph is not directly proportional to the brightness of the light falling on it. (Although this is an inconvenience, the effect can often be allowed for and calibrated accordingly.) Secondly, and more importantly, photographic emulsions have a relatively low sensitivity to light when compared to electronic detectors. For more accurate photometric (i.e. brightness comparison) measurements of individual sources, it is preferable therefore to use some form of photoelectric detector.

4.2 CCDs

A photoelectric detector is essentially a device that responds to incoming photons of light by producing an electrical signal. This electrical signal is then detected, amplified and measured, and the resulting image is built up and processed using a computer. Several forms of photoelectric detector have been used over the years, such as *photomultipliers* and *photodiodes*, but nowadays, the commonest type of detector used in astronomy is known as a **CCD**, which stands for *charge-coupled device*.

A CCD is a two-dimensional, highly sensitive solid-state detector which can be used to generate, extremely rapidly, a pictorial representation of an area of the sky or a spectrum. Similar detectors are now routinely used in digital cameras. Figure 4.1 shows an example of an astronomical CCD. As you can see from the figure, physically CCDs are very small, typically only a couple of centimetres across. They are usually made from a silicon based semiconductor, arranged as a two-dimensional array of light-sensitive elements. The pictures generated from such detectors therefore consist of an array of picture elements, known as **pixels** for short, with one pixel in the image corresponding to each light-sensitive element in the CCD. Conventionally therefore, the light-sensitive elements of the CCD itself are also referred to as pixels.

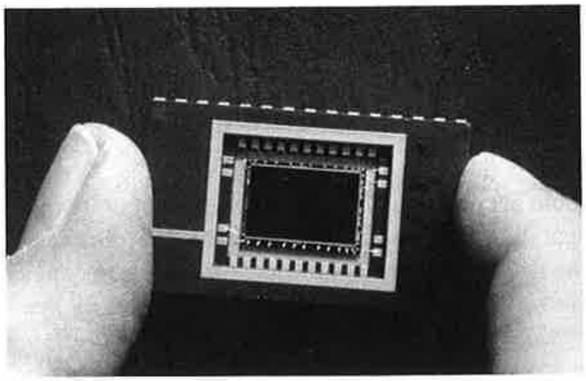

Figure 4.1 An example of a CCD used in astronomical imaging. (Courtesy of *John Walsh/ Science Photo Library.*)

Note that $4096 = 2^{12} = 4 \times 2^{10}$.

The individual pixels on the CCD can each be considered as tiny detectors in their own right. A modern CCD may contain up to 4096 × 4096 (referred to as 4k × 4k) pixels in an array, with each pixel typically of order 10 to 20 μm across. When light falls on a pixel, each photon generates one electron–hole pair in the semiconductor; the electron is called a **photoelectron**, since it is produced by a photon. Hence the number of pairs depends on the intensity of the radiation. Once an exposure is completed, the accumulated charges are transferred out of the array in a controlled manner, one row at a time. This is converted into a digital signal which can be displayed on a monitor screen and stored on computer for later processing and analysis.

CCDs have an extremely high efficiency at visible wavelengths, recording typically 70% of the photons that fall on them. This may be compared with the efficiency of photographic emulsion which is typically only a few percent.

A CCD consisting of 1024 × 1024 light-sensitive elements arranged in a square array is used to obtain an image of a star cluster under seeing conditions of 1″. In order to obtain an image that is 'well matched' to the seeing, it is reasonable to have the image of a point object stretching across about four pixels on the CCD. Hence, each pixel on the CCD must correspond to an angular size of 0.25″, to avoid an image falling on the dead area between pixels.

(a) If the physical size of the detector is 20 mm × 20 mm, what is the scale of the image formed on the CCD in arcseconds per mm?

(b) What is the field-of-view of the CCD in this case?

(c) What focal length telescope is required to match this performance?

(a) 1024 pixels occupy 20 mm on the CCD, hence each is only (20 mm/1024) = 19.5 μm across. Each pixel corresponds to 0.25 arcseconds on the sky image, hence the image scale is (0.25 arcseconds/19.5 μm) = 1.28×10^4 arcseconds m^{-1} = 12.8 arcseconds mm^{-1}.

(b) The whole CCD is 20 mm across, so the field-of-view is (12.8 arcseconds mm^{-1} × 20 mm) = 256 arcseconds along each side, or 4.3 arcminutes × 4.3 arcminutes.

(c) A field-of-view spanning 256 arcseconds corresponds to (256/3600) × (π/180) radians = 1.24×10^{-3} radians. This will extend over 20 mm when the focal length of the telescope is $f_o = 20\text{ mm}/1.24 \times 10^{-3} = 16\,100$ mm or about 16.1 m.

CCDs are now used for virtually all astronomical imaging and the images so obtained can then be used for astrometry (measuring the positions of objects), astronomical photometry (measuring the brightness of objects) and spectroscopy (measuring the spectra of objects).

4.3 Summary of Chapter 4 and Question

- *CCDs* have many advantages over the eye as a detector for use with an astronomical telescope. They are *integrating detectors* and so can detect *fainter* objects than the eye; they also enable a *permanent* record to be kept of the observation, and they allow *finer detail* to be investigated than possible with the eye alone.
- Unlike photographic images, those produced by CCDs have high efficiency, photometric linearity and the ability for images to be processed and analysed by computer techniques.

QUESTION 4.1

Figure 4.2 shows a (schematic) CCD spectral image, obtained using a single-slit grating spectrograph, containing the spectra of two stars. The CCD comprises 800 × 800 pixels arranged in a square array and has a linear size of 10.0 mm along each side. The telescope used to obtain the image has an effective focal length of 4.00 m.

(a) What is the image scale in the plane of the telescope, in arcseconds per mm?

(b) What is the field-of-view at the detector in arcminutes?

(c) What is the angular scale of the image in arcsec per pixel?

(d) What is the angular separation of the two stars?

(e) If the spectral scale of the image is 0.4 nm per pixel, what is the difference in wavelength of the two emission lines (labelled H_α and H_β) in each spectrum?

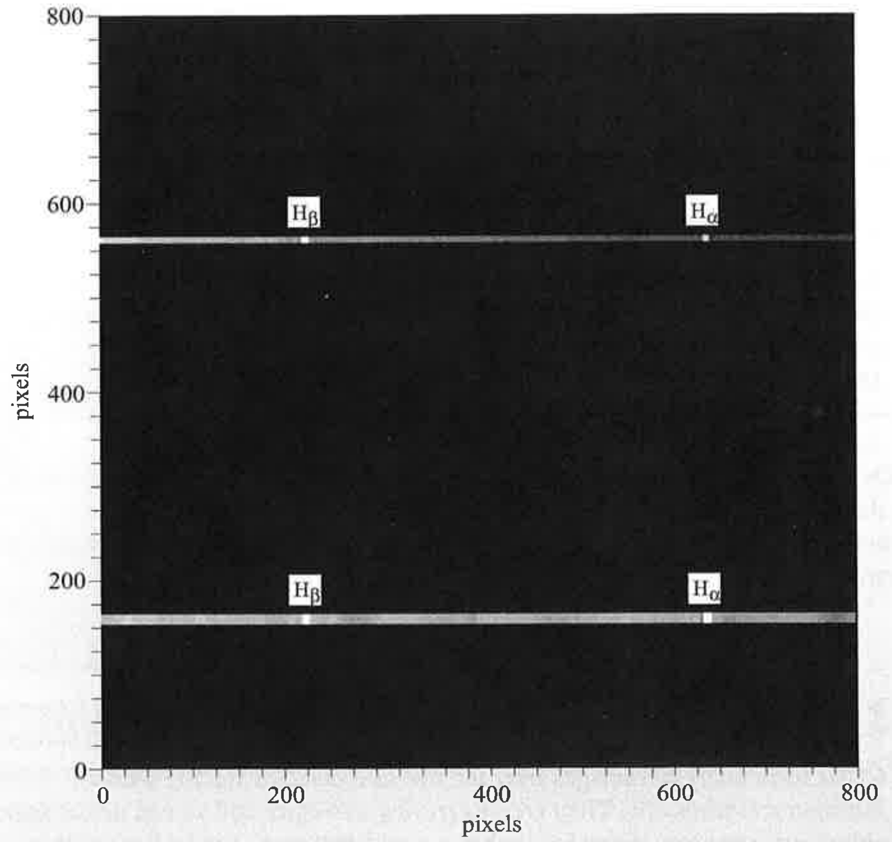

Figure 4.2 The CCD spectral image referred to in Question 4.1.

5 REDUCING CCD DATA

A CCD is just a semiconductor chip. To make it into a useful astronomical device, it must be connected up to electronics that power it, control it, and allow its data to be read out. The charge that has accumulated in each detector pixel is initially read out as a tiny electric current. This current is amplified and converted into a number expressed in so-called analogue data units (ADU). The ADU value is therefore a measure of the charge that was read out from the detector pixel in question, but is on an arbitrary scale. In order to quantify the number of photoelectrons that the pixel held, and therefore the number of photons that were incident on the pixel during the exposure, an analogue-to-digital conversion (ADC) factor is applied to the number in ADUs. The ADC factor is essentially the number of photoelectrons per ADU. The result of this process is a digital image where the value in each image pixel is the number of photons incident on the detector pixel during the exposure (subject to some calibration which we discuss below). Bear in mind however, that the CCD itself does *not* count photons – a distinction which is crucial in assessing uncertainties.

The numbers in a digital image are stored to a certain numerical precision. For instance, each pixel value may be stored as a 16-bit binary number. In this case, the maximum value that may be stored in any pixel of the image is $2^{16} - 1 = 65\,535$, i.e. there are 65 536 unique values possible, ranging from 0 to 65 535. If the light intensity falling on any single detector pixel is such that it would generate a value larger than this limit, the analogue-to-digital conversion will **saturate**. In such a case, the value stored in an image pixel would no longer be representative of the number of photons falling on the detector pixel at that location. Exposure times need to be carefully selected so that the objects of interest do not saturate the CCD. If necessary, a number of short exposures may be combined to overcome this.

■ What is the maximum value that can be stored in an image pixel produced by a CCD whose analogue-to-digital converter operates with 15-bit precision?

☐ The maximum value is $2^{15} - 1 = 32\,767$.

A raw CCD image is shown in Figure 5.1. Reducing, or processing, CCD data consists of taking the array of values stored on the pixels of an image, and manipulating it appropriately to produce an image in which the numerical value in each pixel is directly proportional to the number of photons falling on the detector at that location. In order to do this, a series of corrections need to be applied to each raw image, and we consider these below.

5.1 Bad pixels and cosmic rays

In any CCD, some of the detector pixels will be faulty and will return values that are misrepresentative of the light falling on them. Such pixels are referred to as *hot* or *cold* or *bad*. Sometimes an entire column or row of the CCD may contain **bad pixels**. Software to process CCD images will generally have the facility to either ignore bad pixels, or to replace them with an interpolated value based on the values in adjacent non-bad pixels.

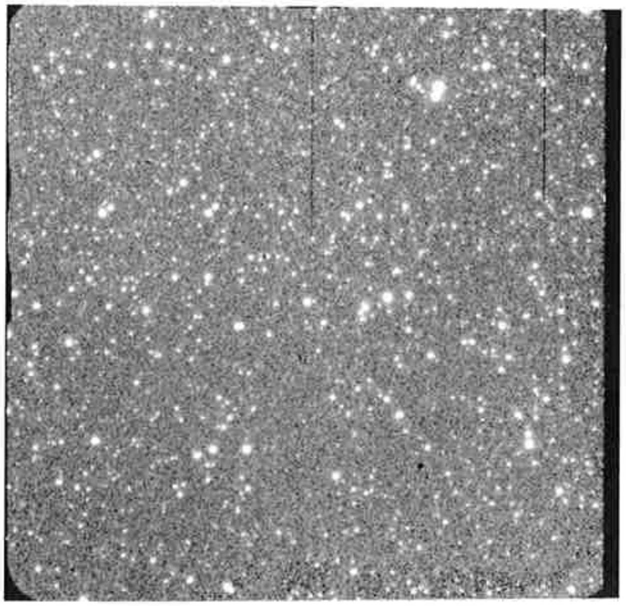

Figure 5.1 A raw CCD image showing bad pixels, bad columns and varying intensity from one region to another.

When ionizing radiation, either from local, naturally occurring sources of radioactivity or from cosmic rays, hits the CCD it releases charge in a pixel that is similar to that caused by light falling on the chip, though often many times greater. These spurious signals are usually confined to a single detector pixel or a few adjacent pixels, and any individual image may have several dozen or several hundred **cosmic ray events**, with the number increasing with exposure time. However, as cosmic ray events have abnormally high values in single pixels they are usually easy to distinguish from genuine 'point sources' such as stars whose light will be spread over a few pixels with a characteristic distribution, namely the point spread function of the telescope as discussed in Chapter 2 Section 4. Automated software to remove cosmic ray events from individual images operates in the same way as for bad pixels by interpolating using adjacent pixel values. Note, however, that this is only a partial correction, since the interpolated value is only an estimate of the real one.

Alternatively, since the locations of cosmic ray events are entirely random, they may be removed by combining together several individual images of the same field, obtained with the same exposure time. By taking the **median** value in each pixel from such a stack of images, the anomalously high values in pixels affected by cosmic rays may be rejected.

A stack of nine CCD images contains the following values in a particular single image pixel: 4356, 4421, 4324, 4309, 4401, 4967, 4397, 4391, 4364. Which image contains the cosmic ray? What is the median value in the pixel?

The sixth image (with a pixel value of 4967) contains the cosmic ray. The median value is the fifth value when the pixel values are arranged in order. The pixel values in order are 4309, 4324, 4356, 4364, 4391, 4397, 4401, 4421, 4967. So the median value is 4391.

Notice that the median value is *not* equal to the *mean* value: the mean value is much higher due to the influence of the anomalous value in the image affected by the cosmic ray.

5.2 Bias and dark-current subtraction

When the signals from a CCD are digitized in the analogue-to-digital conversion process, an offset or **bias** signal is intentionally introduced into the digital value to prevent the signal from going negative at any point (which it may otherwise do, due to random fluctuations or noise). The value of the bias may vary depending on the position on the CCD and can vary with time over the seconds or minutes it takes to read out the array. There are two techniques for correcting for this bias signal.

Overscan strips are narrow regions of the CCD image, usually running down either side of the image, a few tens of pixels wide. They are virtual pixels created by continuing to read out the CCD even after the last real pixel has been read out. That is, they do not correspond to a piece of the detector, and hence contain no photoelectrons of their own. They therefore indicate how the CCD electronics, and the analogue-to-digital converter in particular, responds to a genuine zero signal, and how that response varies with time. For each row in the CCD, the values of the signal in the overscan strip corresponding to that row may be averaged and subtracted from the value in each other pixel in that row. After this stage of processing, the overscan strips may be discarded, thus reducing the sizes of images.

(Sometimes overscan strips are called bias strips, but this can lead to confusion with the bias frames described below.)

Bias frames are entire images created by reading out the CCD following a zero second exposure, which means there are no photoelectrons stored in its pixels. (A CCD 'image' is often called a **frame** and the two terms should be seen as interchangeable.) The bias frame enables the average noise across the chip to be measured and accounted for. Bias frames are usually obtained by taking zero-length exposures with the shutter closed at the beginning and end of each night's observing. A master bias, see Figure 5.2, is generated by creating the median of a stack of many such frames, and this can then be subtracted from every other image obtained during the night.

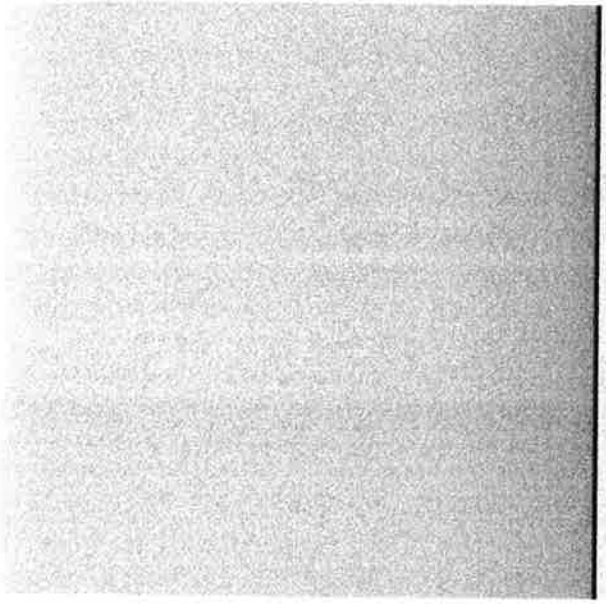

Figure 5.2 A master bias frame showing small-scale structure in the noise across the CCD. There is an overscan strip down the extreme right-hand side.

Another effect is that some signal may be generated in CCD pixels even when no light is present. This is referred to as **dark current** and is due to the motion of electrons that arises from the thermal energy of the CCD and defects. Like the bias signal, it varies from pixel to pixel and also changes with time. Dark current can be minimized by cooling the CCD to liquid-nitrogen temperatures. However, if this is not possible, or not sufficiently effective, then dark current may be accounted for by taking long exposures with the shutter closed, removing the bias, and then dividing by the exposure time to obtain the dark current per second in each pixel. This may then be scaled by the exposure time of every other image, and subtracted off. Dark current is often insignificant for many visible-light CCDs, but is more important when working in the infrared.

5.3 Flat-fielding

The sensitivity to light of the many pixels in any CCD will vary slightly with position, by a few percent. This is due to irregularities introduced by the manufacturing process. In order to calibrate for this relative variation in pixel-to-pixel sensitivity, a CCD is exposed to a uniformly illuminated light source, such as the twilight sky, or the inside of an illuminated observatory dome. The images obtained by such a process are known as **flat fields**. Target images may then be corrected, using flat fields, to the values they would have had if all the detector pixels had the *same* sensitivity to light. This process is known as flat-fielding. It is important to note that the pixel-to-pixel variation will also be a function of wavelength, so when observing through filters, flat fields must be obtained through the same filters as the target observations.

Flat-fielding also corrects for effects such as dust particles on the CCD itself; dust on the filters which cast ring-shaped (out of focus) shadows; and the dimming of objects observed towards the edge of the telescope field-of-view (this is known as *vignetting*) caused by obstructions in the light path or just the change of angle. Two types of flat fields commonly used are as follows.

Dome flats are images of the inside of the observatory dome, illuminated by a continuum spectral source such as a tungsten-filament light bulb. The dome will necessarily be out of focus, and the images will be featureless. Dome flats may be taken during the day, before or after an observing session. They have two disadvantages though. First, light reflected from the dome enters the telescope at a different angle from that at which light from the sky enters. This can affect the response to vignetting and dust on the filters or CCD. Second, the spectrum of a tungsten-filament light bulb is not the same as that of the night sky and can make it more difficult (or even impossible) to correct for an effect known as **fringing**. Fringing is caused by interference between rays from multiple reflections, within the CCD or filters, of light at a single wavelength. It can give rise to wave-like patterns of intensity variation across the CCD, known as **fringes**. In order to correct for fringes, we must use a flat-field source whose spectrum closely matches that of the image in question.

Fringing can usually be accounted for by using **sky flats**. These are images of the sky taken in twilight, either before or after an observing session. The sky needs to be brighter than any stars that are in the field-of-view, but not bright enough to saturate the CCD. Disadvantages of sky flats are that it can be difficult to judge the appropriate exposure times and that sunlight reflected from the inside of the observatory dome can also reach the CCD, affecting the vignetting response as with dome flats.

The procedure to 'flat-field-correct' the images is that several flat fields are de-biased and dark current subtracted, and then combined (median stacked) to produce a single master flat (in each filter); see Figure 5.3. The value in each pixel of the master flat is then divided by the mean value of all the pixels. This has the effect of normalizing the mean value of the master flat to unity (i.e. a value of one). De-biased, dark-subtracted target images in each filter are then *divided by* the normalized master flat image in the appropriate filter to remove the pixel-to-pixel sensitivity variation, dust images, and vignetting effects (Figure 5.4).

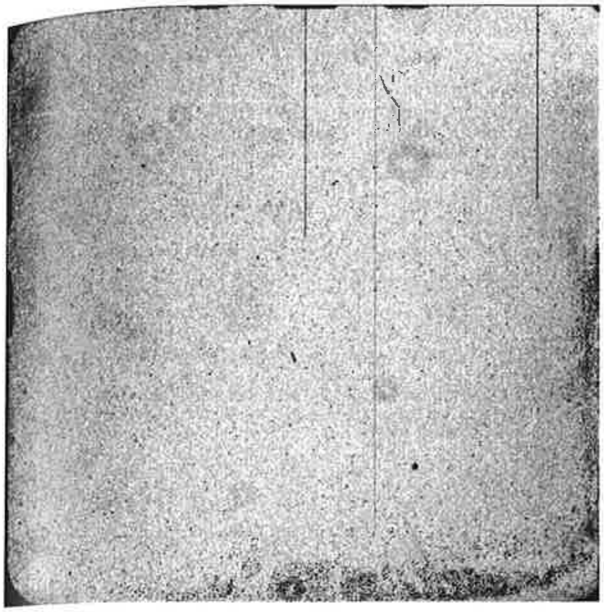

Figure 5.3 A master flat field showing spurious images due to dust particles, vignetting effects and pixel-to-pixel sensitivity variations.

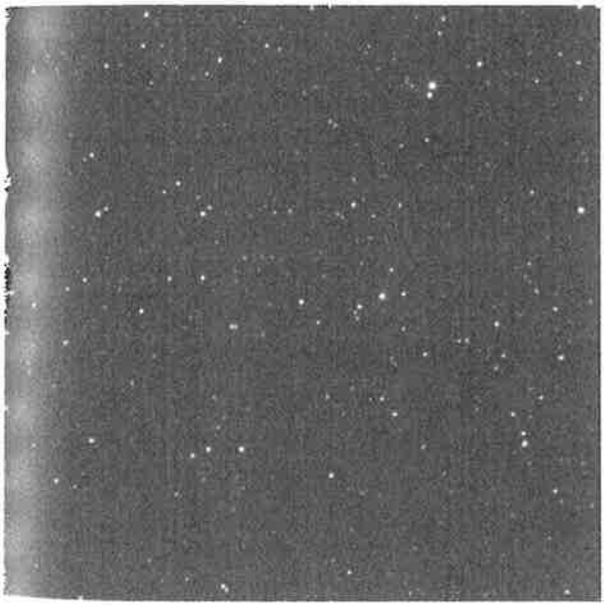

Figure 5.4 A de-biased, flat-fielded image of Figure 5.1. Comparison of Figures 5.1 and 5.4 as 'before' and 'after' illustrates why we say the data have been reduced.

5.4 Summary of Chapter 5 and Questions

The stages in reducing CCD data are as follows:

- Bad pixels and cosmic rays are removed by interpolating across the affected pixels or by taking the median of a stack of images.
- A master bias frame is created and subtracted from each image, or overscan strips are used to correct for the bias level in each image.
- The dark current is removed if necessary, scaled to the exposure time of each image.
- Dome flats or sky flats in each filter are combined and normalized to a mean of unity, then divided into each target image to correct for pixel-to-pixel sensitivity, dust particle shadows and vignetting.

QUESTION 5.1

Figure 5.5 shows cross-sections through the images of two stars on a CCD frame. One of the two stars is saturated. Which one is it and why? How could you avoid this problem?

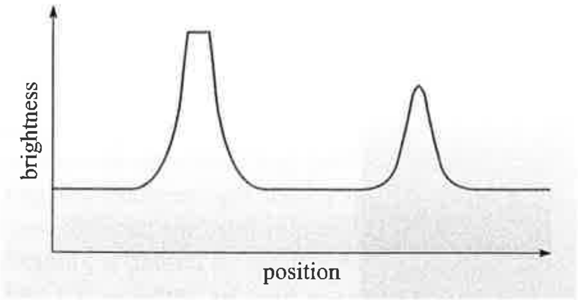

Figure 5.5 Cross-sections through the images of two stars on a CCD frame. One of the stars is saturated, the other is not. See Question 5.1.

QUESTION 5.2

If a series of exposures are made of targets through different filters, why must flat fields also be obtained through *each* filter too? Conversely why is a single set of bias frames and dark frames sufficient?