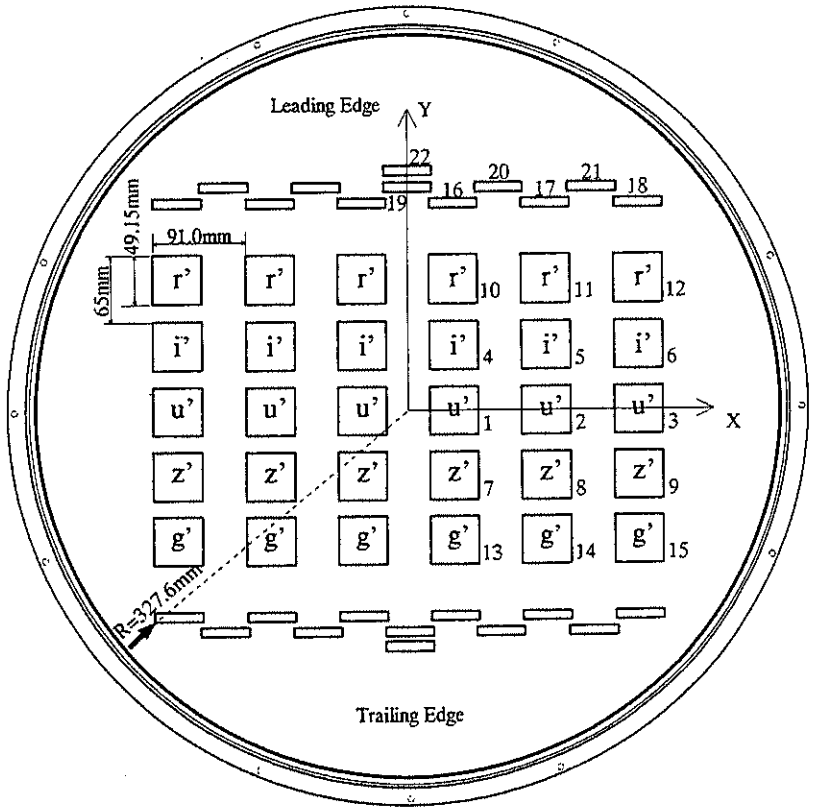

Fig. 4.8. The optical layout within the dewar of the Sloan digital sky survey CCD imager. The right side of the figure labels the CCDs as to their function; 1–15 are photometric CCDs, 16–21 are astrometric CCDs, and 22 (top and bottom) are focus CCDs. The left side gives the dimensions of the array. The labels r'–g' denote the five seperate intermediate band filters, each a single piece of glass covering all six horizontal CCDs. The scan direction is upward causing objects to traverse the array from top to bottom. From Gunn et al. (1998).

between colors caused by CCD crossing time and chip spacing. Complete details of the SDSS, too lengthly for presentation here, can be found in Gunn et al. (1998).

The SDSS has just begun operation and is probably the largest, most complex camera of its kind. Other projects of a similar nature are discussed in Boulade et al. (1998), Gunn et al. (1998), and Miyazaki et al. (1998).

5

Photometry and Astrometry

One of the basic astronomical pursuits throughout history has been to determine the amount and temporal nature of the flux emitted by an object as a function of wavelength. This process, termed photometry, forms one of the fundamental branches of astronomy. Photometry is important for all types of objects from planets to stars to galaxies, each with their own intricacies, procedures, and problems. At times, we may be interested in only a single measurement of the flux of some object, while at other times we could want to obtain temporal measurements on time scales from seconds or less to years or longer.

We start this chapter with a brief discussion of the basic methods of performing photometry when using digital data from 2-D arrays. It will be assumed here that the CCD images being operated on have already been reduced and calibrated as described in detail in the previous chapter. We will see that photometric measurements require that we accomplish only a few steps to provide output flux values. Additional steps are then required to produce light curves or absolute fluxes.

Some photometric output products, such as differential photometry, require fewer additional steps, whereas to obtain the absolute flux for an object, additional CCD frames of photometric standards are needed. These standard star frames are used to correct for the Earth's atmosphere, color terms, and other possible sources of extinction that may be peculiar to a given observing site or a certain time of year (Pecker, 1970).

As an introduction to the level of atmospheric extinction one might expect as a function of observational elevation and wavelength, Table 5.1 lists values of the extinction in magnitudes resulting from the Earth's atmosphere for an observing site at 2,200 m elevation. Note that for observations made at reasonable airmass and redward of 4,000 Å, the

Table 5.1. *Example atmospheric extinction values (magnitudes)*

Altitude	Airmass	3000 Å	3500 Å	4000 Å	4500 Å	5000 Å	5500 Å	6000 Å	6500 Å	7000 Å	8000 Å	9000 Å	10000 Å
90	1.00	1.2	0.65	0.4	0.3	0.2	0.2	0.2	0.1	0.1	0.1	0.1	0.1
75	1.04	1.2	0.65	0.4	0.3	0.2	0.2	0.2	0.1	0.1	0.1	0.1	0.1
60	1.15	1.3	0.75	0.5	0.3	0.3	0.2	0.2	0.2	0.1	0.1	0.1	0.1
45	1.41	1.6	0.9	0.6	0.4	0.3	0.3	0.2	0.2	0.2	0.1	0.1	0.1
30	1.99	2.3	1.3	0.8	0.6	0.4	0.4	0.3	0.3	0.2	0.2	0.1	0.1
20	2.90	3.3	1.55	1.2	0.8	0.6	0.5	0.5	0.4	0.3	0.2	0.2	0.2
15	3.82	4.4	2.5	1.6	1.1	0.8	0.7	0.6	0.5	0.4	0.3	0.3	0.2
10	5.60	6.4	3.65	2.3	1.6	1.2	1.1	1.0	0.7	0.6	0.4	0.4	0.3
5	10.21	11.8	6.7	4.2	2.9	2.2	1.9	1.7	1.4	1.1	0.8	0.7	0.6

effect of the Earth's atmosphere is, at worst, a few tenths of a magnitude. The details of photometric corrections for the Earth's atmosphere and extinction effects are not germane to the topic of this book and their discussion here would be beyond the allowed space limitations. The interested reader is referred to the excellant presentations in Dacosta (1992), Hendon & Kaitchuck (1982), and Young (1974). Further good discussions of photometric data handling are presented in Merline & Howell (1995), Stetson (1987), Howell (1992), Hiltner (1962), Walker (1990), Kreidl (1993), Howell & Jacoby (1986), and Howell, Everett, & Ousley (1999).

5.1 Stellar Photometry from Digital Images

Prior to the time when CCDs became generally available to the astronomical community, digital images of astronomical objects were being produced by detectors such as silicon intensified targets (SITs), video tube-type cameras, image tubes, and electronographic cameras. In addition, scanning of photographic plates with a microdensitometer resulted in large amounts of digital output. These mechanisms produced digital data in quantity and at rates far in excess of the ability of workers to individually examine each object of interest within each image. Today, the amount of CCD data greatly exceeds this limit. Thus, in the early 1980s, work began in earnest to develop methods by which photometry could be obtained from digital images in a robust, mostly automated manner.

One of the first such software packages to deal with digital images was written by Adams, Christian, Mould, Stryker, and Tody (Adams et al., 1980) in 1980. Numerous other papers and guides have been produced over the years containing methods, ideas, entire software packages that perform photometry, and specific detailed information for certain types of objects. I have tried to collect a fairly complete list of these in Appendix A. While details vary, the basic photometric toolbox must contain methods that perform at least the following primary tasks: (*i*) image centering, (*ii*) estimation of the background (sky) level, and (*iii*) calculation of the flux contained within the object of interest. We will assume below, for simplicity, that we are working with stellar images that to a good approximation are well represented by a point spread function of more-or-less Gaussian shape. Deviations from this idealistic assumption and nonpoint source photometry will be discussed as they arise.

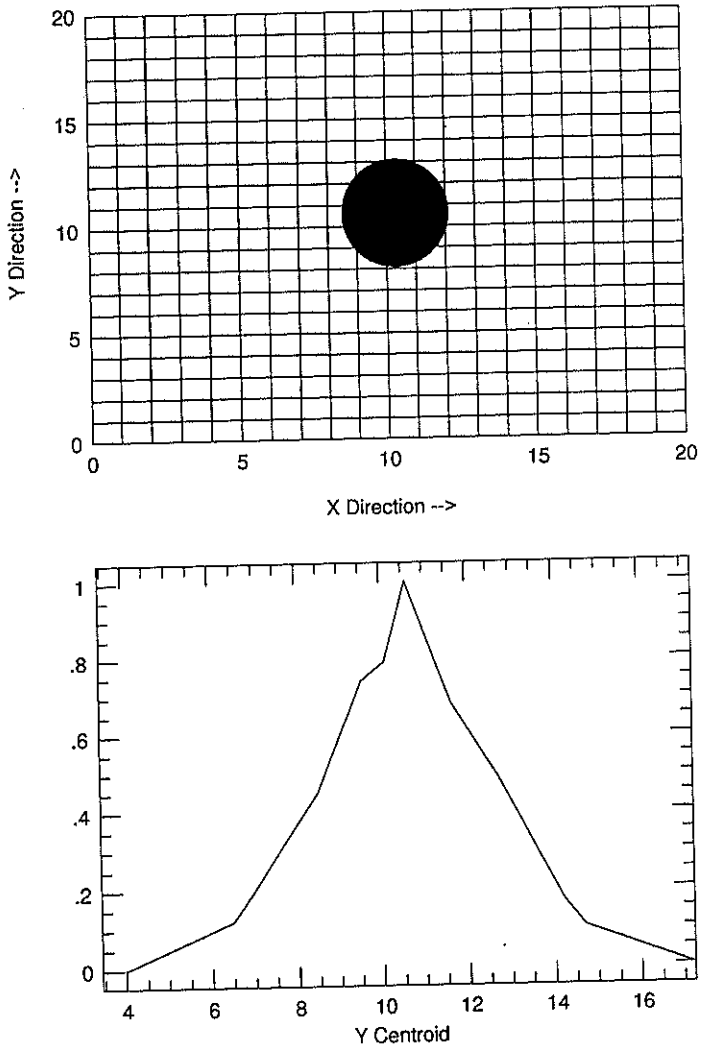

Fig. 5.1. An example of x, y centroiding. The idealized star image in the top box sits on a pixel grid with a center of $(x, y) = (10.3, 10.6)$. The two other plots represent the x, y centroids for the star image normalized to a maximum height of one. In this case, the star center is well approximated by the peaks in the x, y centroids.

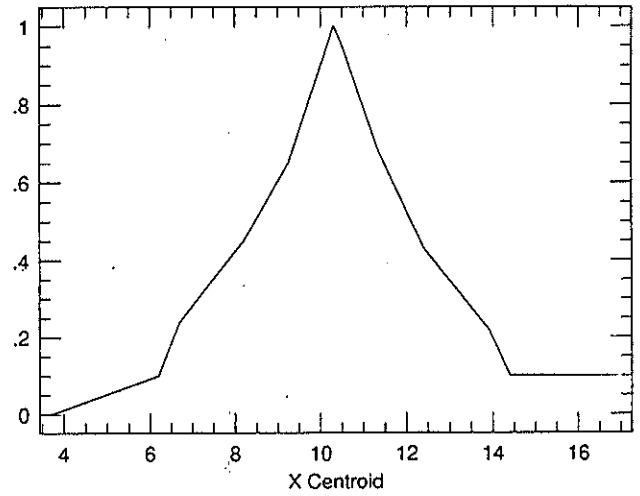

Fig. 5.1. (*Continued*)

5.1.1 Image Centering

Probably the simplest and most widely used centering approximation for a point spread function (PSF) is that of marginal sums or first moment distributions. Starting with a rough pointer to the position of the center of the star (e.g., the cursor position, reading off the x, y coordinates, or even a good guess), the intensity values of each pixel within a small box centered on the image and of size $L \times L$ (where L is comparable to the size of the PSF) are summed in both x and y directions (see Figure 5.1). The x, y center is computed as follows: The marginal distributions of the PSF are found from

$$I_i = \sum_{j=-L}^{j=L} I_{i,j}$$

and

$$J_j = \sum_{i=-L}^{i=L} I_{i,j},$$

where $I_{i,j}$ is the intensity (in ADU) at each x, y pixel; the mean intensities are determined from

$$\bar{I} = \frac{1}{2L+1} \sum_{i=-L}^{i=L} I_i$$

and

$$\bar{J} = \frac{1}{2L+1} \sum_{j=-L}^{j=L} J_j;$$

and finally the intensity weighted centroid is determined using

$$x_c = \frac{\Sigma_{i=-L}^{i=L}(I_i - \bar{I})x_i}{\Sigma_{i=-L}^{i=L}(I_i - \bar{I})}$$

for all $I_i - \bar{I} > 0$ and

$$y_c = \frac{\Sigma_{j=-L}^{j=L}(J_j - \bar{J})y_j}{\Sigma_{j=-L}^{j=L}(J_j - \bar{J})}$$

for all $J_j - \bar{J} > 0$.

For well-sampled (see Section 5.6), relatively good S/N (see Section 4.4) images, simple x, y centroiding provides a very good determination of the center position of a PSF, possibly as good as 1/5th of a pixel. More sophisticated schemes to provide better estimations of image centers or applications appropriate to various types of non-Gaussian PSFs are given in Davis (1994), Massey & Davis (1992), Lasker et al. (1990b), Stone (1989), Penny & Dickens (1986), Chiu (1977), and Howell et al. (1996).

5.1.2 Estimation of Background

The importance of properly estimating the background level on a CCD resides in the fact that the same pixels that collect photons of interest from an astronomical source also collect photons from the "sky" or background, which are of no interest. Remember that the background or sky value in a CCD image contains not only actual photons from the sky but also photons from unresolved astronomical objects, read noise, thermally generated dark current electrons, and other sources. All of these unwanted additional photons must be accounted for in some manner, estimated, and removed from the image before a final determination of the source flux is made. In order to determine this background level, a common technique is to place a software annulus around the source of interest and then use statistical analysis to estimate its mean level on a per pixel basis.

The background or sky annulus is usually defined by an inner and outer radius or by an inner radius and a width (see Figure 5.2). One simple, yet powerful, manner by which an estimation of the background

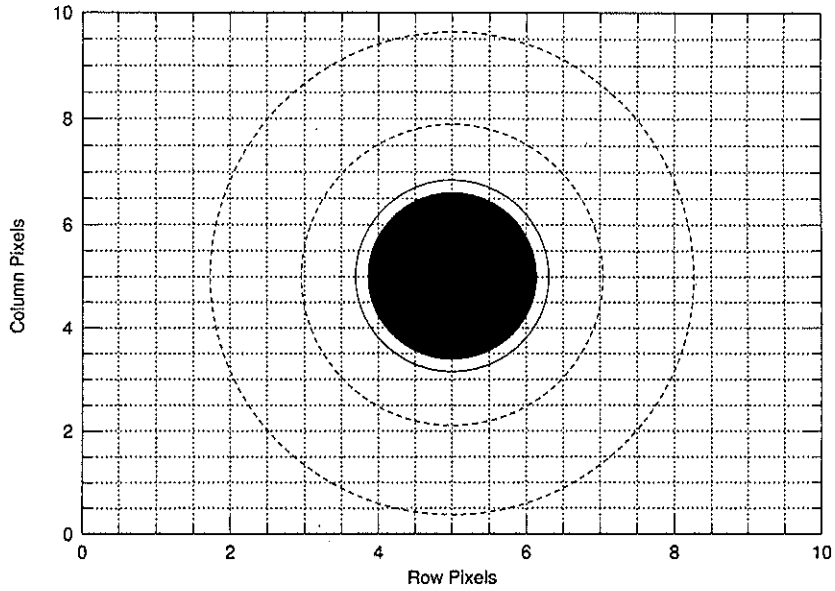

Fig. 5.2. Schematic drawing of a stellar image on a CCD pixel grid. The figure shows the location of the star, the "star" aperture (solid line), and the inner and outer "sky" annuli (dashed circles).

level can be made is to simply extract the values of all the pixels within the annulus, sum them, and divide by the total number of pixels within the annulus. This provides an average value per pixel for the background level of the CCD image. For a good statistical determination of the background level, the total number of pixels contained within this annulus should be relatively large, about three times the number contained within that of the source aperture (Merline & Howell, 1995).† A more robust estimator, requiring very little additional work, is to collect all the pixel values from inside the sky annulus, order them in increasing value, and find the median intensity, B_M.‡ A nice touch here is to reexamine

† Partial pixels arising from placing a circular annulus on a rectangular grid are usually not of concern here, as the number of annulus pixels is large. However, partial pixels cannot be so easily dismissed when we are determining the intensity within the much smaller source aperture.

‡ Note that the statistical values determined from a CCD image for the median or the mode must pick their answer from the list of the actual CCD pixel ADU values, that is, values that are integers containing only digitized levels and thus digitization noise. The statistical mean, however, allows for noninteger answers. This seemingly subtle comment is of great importance when dealing with partial pixels, undersampled data, or high CCD gain values.

the list of annulus pixel values and toss out all those with values greater than $\pm 3\sigma$ from B_M. This last step will eliminate cosmic ray hits, bad pixels, and contamination from close by astronomical neighbors if they exist.

When applying a median filter and the 3σ cutoff technique to the list of background pixels, one can use the remaining annulus pixel values to construct a background histogram computed with a bin width resolution of say 0.5 ADU (Figure 5.3). The background histogram will be centered on the median background value with all contained pixel values within $\pm 3\sigma$. Since the detailed resolution of 0.5 ADU binning

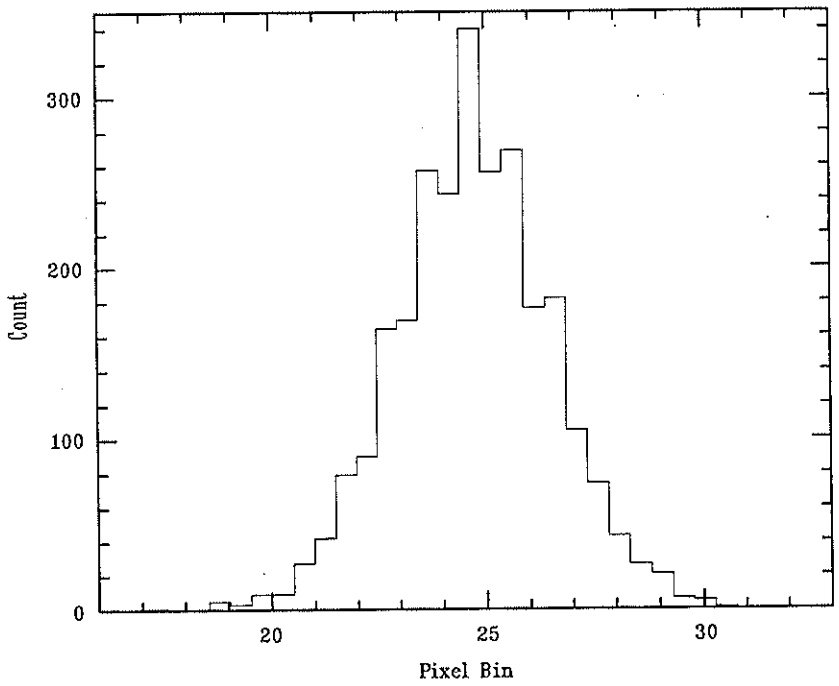

Fig. 5.3. Histogram of the "sky" annulus around a star in the CCD image shown in Figure 4.5. Notice the roughly Gaussian shape to the sky distribution but with an extended tail toward larger values. This tail is due to pixels that were not completely calibrated in the reduction process, pixels with possible contamination owing to dark current or cosmic rays, pixels with increased counts owing to unresolved PSF wings from nearby stars, and contamination of sky annulus pixels by faint unresolved background objects. The need for some histogram smoothing, such as that described in the text, is apparent, especially near the peak of the distribution.

will likely produce a ragged histogram (since only a finite number of background pixels are used), some smoothing of the histogram may be useful, such as that done by Lucy (1975). Lucy smoothing will broaden the histogram distribution slightly after one iteration but application of a second iteration will restore the correct shape and provide a smooth histogram from which to proceed. This final step produces a statistically valid and robust estimator from which we can now compute the mean value of the background, $\bar{B}$. A determination of the centroid of the final smoothed histogram, using equational forms for the centroid in one dimension such as those discussed above, can now be applied. Centroiding of this smoothed histogram is not influenced by asymmetries that may have been present in the wings of the initially unsmoothed values.

Correct estimation of the level of the CCD background on a per pixel basis is of increasing importance as the S/N of the data decreases and/or the CCD pixel sampling becomes poor. A background level estimation for each pixel that is off by as little as a few ADU can have large effects on the final result (Howell, 1989). Determination of the CCD background level has an associated error term $\propto (1 + n_{\rm pix}/n_B)^{-1/2}$, which should be included in the S/N calculation of the final result (see Section 4.4).

5.1.3 Estimation of Point Source Intensity

We now come to the pixel values of interest, that is, those which contain photons from the source itself. Using a software aperture of radius r centered on the x, y position of the centroid of the source PSF, we can extract the values from all pixels within the area A $(= \pi r^2)$ and sum them to form the quantity S, the total integrated photometric source signal. The sum S contains contributions from the source but also from the underlying background sources within A as well. To remove the estimated contribution to S from the background, we can make use of the value $\bar{B}$, discussed above. We can calculate an estimate of the collected source intensity, I, as $I = S - n_{\rm pix}\bar{B}$, where $n_{\rm pix}$ is the total number of pixels contained within the area A. There are some additional minor considerations concerning this procedure but these will not be discussed here (Merline & Howell, 1995).

A final step usually performed on the quantity I, which we will discuss further below, is to determine a source magnitude. The value of a magnitude is defined by the following standard equation:

$$\text{Magnitude} = -2.5 \log_{10}(I) + C,$$

where I is the source intensity per unit time, that is, the flux (universally given as per second), and C is an appropriate constant (usually ~23.5–26 for most earthly observing sites) and determined in such a manner so that the calculated source magnitude is placed on a standard magnitude scale such as that of the Johnson system or the Strömgren system.

As we mentioned above, when using circular apertures on rectangular pixel grids, partial pixels are inevitable. While we could toss them away for the large background area, we cannot follow a similar sloppy procedure for the smaller source aperture. Thus the question becomes, how do we handle partial pixels? This is not a simple question to answer and each photometric software package has its own methodology and approach. The three choices a software package can use are:

1. Do not use partial pixels at all. Any source intensity that falls into the source aperture but within a partially inscribed pixel is simply not used in the calculation of S.
2. Sum the values for every pixel within the source aperture regardless of how much or how little of the pixel area actually lies within A.
3. Make use of some computational weighting scheme that decides, in a predefined manner, how to deal with the counts contained within each partial pixel in the source aperture.

This last choice often uses the ratio of the pixel area inside the source aperture to that outside the aperture as a simple weighting factor. A computational scheme to handle partial pixels in a software package designed to perform digital photometry is the hardest of the above choices to implement, but it will provide the best overall final results. To know exactly how a certain software package handles partial pixels, the user is referred to the details presented within the documentation provided with the software. Many PC-type packages that perform photometry on CCD images do not detail their partial pixel and magnitude calculation methods and are therefore "black boxes" to be avoided.

There are two basic methods by which most observers estimate the total integrated signal within their source aperture: point spread function fitting and digital aperture photometry. The first method relies on fitting a 2-D function to the observed PSF and using the integrated value underneath this fitted function as an estimate of S. The second method, digital aperture photometry, attempts to place a software aperture about the source profile (as shown in Figure 5.2), centered in some manner (e.g., x, y centroids), and then simply sums the pixel values

within the source aperture to provide the estimation of S. We will discuss each of these methods in turn below and note here that it is unlikely that a single method of estimation will be the best to use in all possible situations. For example, for severely undersampled data the method of pixel mask fitting (Howell et al., 1996) provides the best solution.

5.2 Two-Dimensional Profile Fitting

The profiles of astronomical point sources that are imaged on two-dimensional arrays are commonly referred to as point spread functions or PSFs. In order to perform measurements on such images, one method of attack is profile fitting. PSFs can be modeled by a number of mathematical functions, the most common include Gaussian,

$$G(r) \propto e^{\left(\frac{r^2}{2a^2}\right)},$$

modified Lorentzian,

$$L(r) \propto \frac{1}{1+(r^2/a^2)^b},$$

and Moffat,

$$M(r) \propto \frac{1}{(1+r^2/a^2)^b},$$

representations, where r is the radius of the point source and a and b are fitting parameters (Stetson, Davis, & Crabtree, 1990). These types of functional forms can be used to define the PSF for each star within an image by the assumption that they provide a good representation of the data itself. For example, adjustment of the values of a and b within one of these functions may allow an imaged PSF to be matched well in radius and profile shape (height and width), allowing a simple integration to be performed to measure the underlying flux.

Generally, the above functions are only a fair match to actual PSFs and so a second method of profile fitting can be applied. This method consists of using an empirical PSF fit to the actual digital data itself, producing modified versions of the above functions. Depending on the application, PSFs may be evaluated at the center of a pixel or integrated over the area of each pixel. Even more general methods of allowing the data to produce completely analytic forms for the PSF functions have been attempted. The techniques and use of empirical PSFs could fill an entire chapter; we refer the reader to Stetson (1987), Diego (1985), and King (1971) for more details.

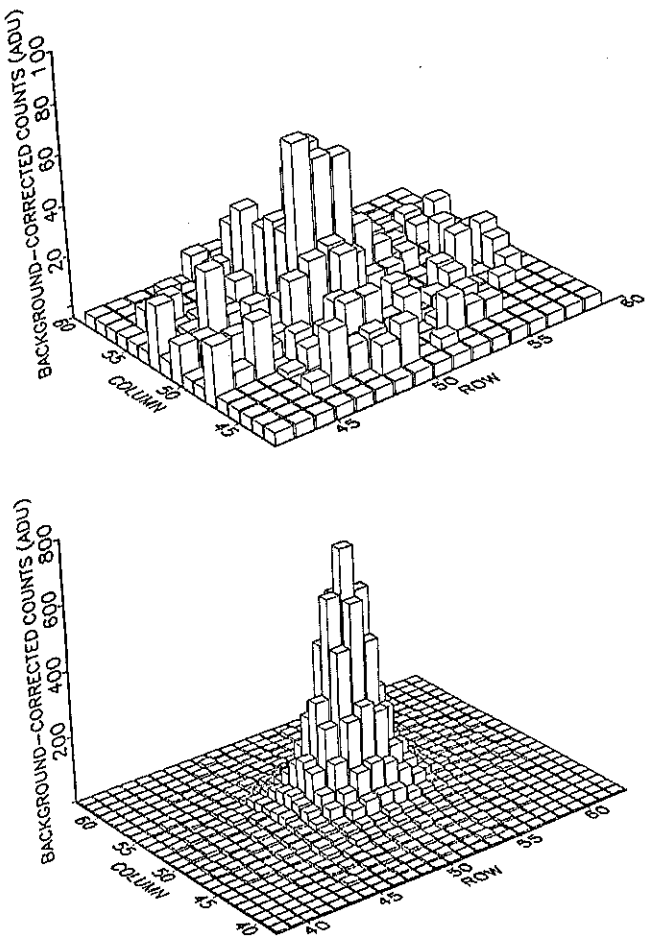

Fig. 5.4. Stellar PSFs are shown for various cases. The figure above shows two model PSFs, one for a bright star (S/N $\sim$ 125) and one for a faint star (S/N $\sim$ 20). The remaining two panels show similar brightness stars but are actual CCD data. Note that the models are shown as 3-D pixel histograms whereas the real data are represented as spline profile fits to the actual PSFs. The disadvantage of the latter type of plotting is that the true pixel nature of the image is lost.

Both techniques, the use of completely mathematical forms for a PSF approximation and the more empirical method, have their advantages and disadvantages. Model PSF fitting allows the necessary integrations and pixel interpolations to be carried out easily as the functions are well known, while the empirical method, which makes hardly any

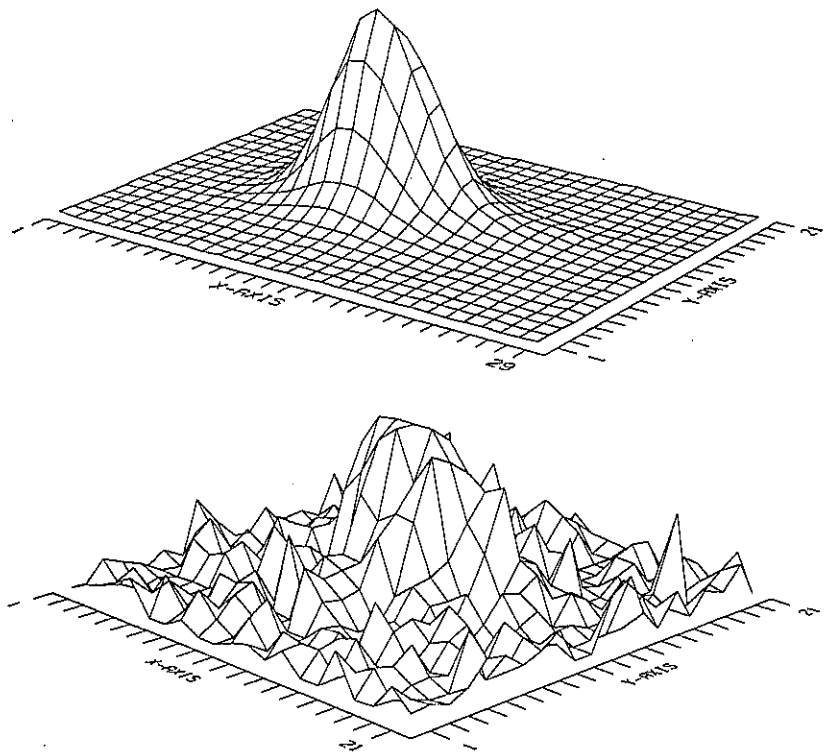

Fig. 5.4. (*Continued*)

assumptions about the actual shape of the PSF, is only defined on the CCD pixel grid and not in any general mathematical way. This latter complication can cause difficulties when trying to interpolate the empirical shape of one PSF somewhere on the CCD (say a bright reference star) to a PSF located somewhere else on the same image but that is likely to have a different pixel registration. For this reason, some implementations of empirical PSF fitting actually make use of the sum of an analytic function (such as one of those given as above) and a look-up table of residuals between the actual PSF and the fitting function. Figure 5.4 shows examples of some PSF models and some actual PSFs obtained with CCDs.

Procedurally, profile fitting techniques work by matching the implied PSF to the actual digital data in a 2-D fashion and within some radius, r, called the fitting radius. An attempt is then made to maximize some goodness-of-fit criteria between the assumed PSF and the observed one.

PSF fitting can be further optimized by fitting N point sources within the image simultaneously (usually one uses the brightest stars within the image) and using some combination of statistically weighted mean values for the final fitting parameters. PSF fitting can be very computationally demanding, much more so than the method of aperture photometry discussed below. However, for some types of imagery, for example crowded fields such as within star clusters for which some PSFs may overlap, PSF fitting may be the only method capable of producing scientifically valid results (Stetson, 1987).

5.3 Aperture Photometry

Aperture photometry is a technique that makes no assumption about the actual shape of the source PSF but simply collects and sums up the observed counts within a specified aperture centered on the source. The aperture used may be circular (usually the case for point sources), square, or any shape deemed useful. Aperture photometry is a simple technique, both computationally and conceptually, but this same simplicity may lead to errors if applied in an unsuitable manner or when profile fitting is more appropriate.

The basic application of aperture photometry starts with an estimate of the center of the PSF and then inscribes a circular software aperture of radius r about that center. The radius r may simply be taken as three times the full-width at half-maximum ($r = 3 \cdot \text{FWHM}$): the radius of a PSF that would contain 100% of the flux from the object (Figure 5.5) (Merline & Howell, 1995). Summing the counts collected by the CCD for all the pixels within the area $A = \pi r^2$, and removing the estimated background sky contribution within A, one finally arrives at an estimated value for I. We see again that partial pixels (a circular software aperture placed on a rectangular grid) are an issue, even for this simple technique. Using a square or rectangular aperture alleviates the need for involving partial pixels but may not provide the best estimate of the source flux. Noncircular apertures do not provide a good match to the 2-D areal footprint of a point source, thereby increasing the value of n_{pix} that must be used and decreasing the overall S/N of the measurement. Remember, however, that for bright sources, n_{pix} is essentially of no concern (see Section 4.4).

It has been shown (Howell, 1989; Howell, 1992) that there is a well-behaved relation between the radius of the aperture of extraction of a point source and the resultant S/N obtained for such a measurement.

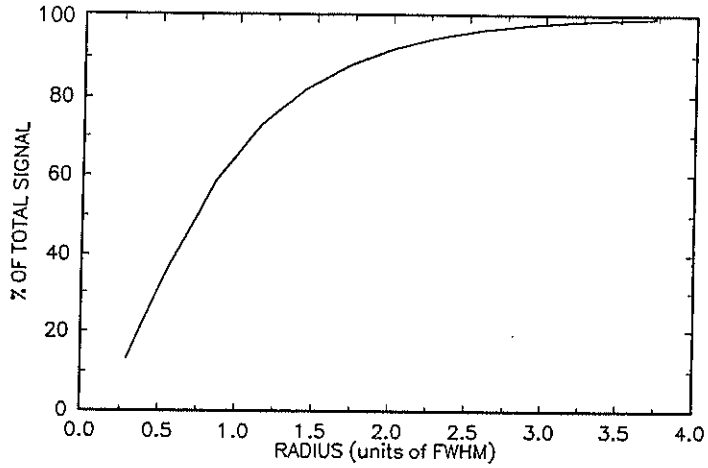

Fig. 5.5. For any reasonable PSF approximation, the figure above shows the run of the total encircled signal with radius of the PSF in FWHM units. Note that within a radius of $3 \cdot \text{FWHM}$ essentially 100% of the signal is included.

An optimum radius aperture, that is, one that provides the optimum or best S/N for the measurement, can be determined for any given PSF and generally has a radius of near $1 \cdot \text{FWHM}$. This optimum radius is a weak function of the source brightness, becoming smaller in size for fainter sources. Figure 5.6 illustrates this idea for three point sources of different brightness.

To understand the idea of an optimum radius and why such a radius should exist, one simply has to examine the S/N equation given in Section 4.4 in some detail. To obtain a higher S/N for a given measurement, more signal needs to be collected. To collect more signal, one can use a larger aperture radius, up to the maximum of $3 \cdot \text{FWHM}$. However, the larger r is, the more pixels that get included within the source aperture and the larger the value of n_{pix}. As n_{pix} increases, so does the contribution to the error term from noise sources other than the source itself. Thus, a balance between inclusion of more signal (larger r) and minimizing n_{pix} in the source aperture (smaller r) leads to an optimum extraction radius for a given source.

We saw in Figure 5.5 that if extracted at or very near an aperture radius of $3 \cdot \text{FWHM}$, 100% of the light from a point source would be collected. However, to obtain the maximum S/N from your measurement, extraction at a smaller radius is warranted. If one extracts the source

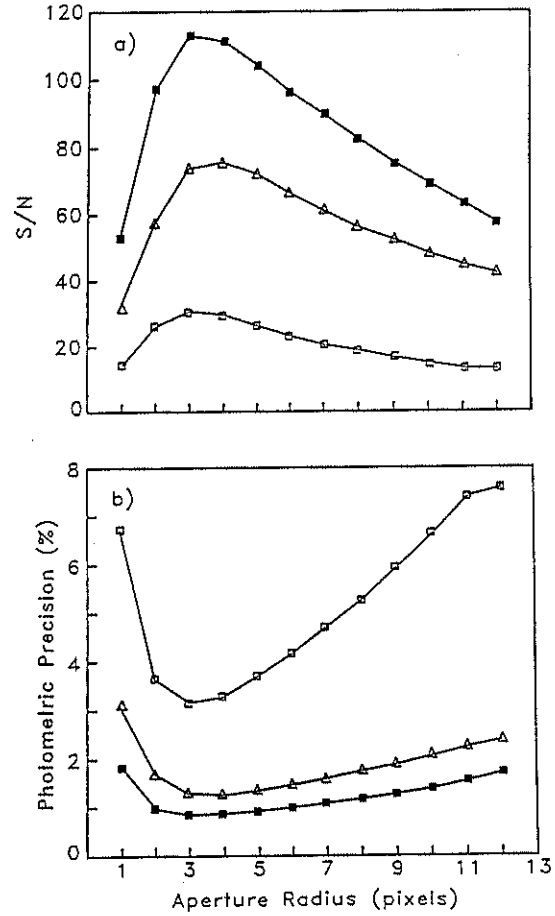

Fig. 5.6. The S/N obtained for the measurement of a point source is not constant as a function of radius. There is an optimum radius at which the S/N will be a maximum. The top panel shows this effect for three point sources that differ in brightness by 0.3 (middle curve) and 2.0 (bottom curve) magnitudes compared with the top curve (filled squares). The bottom panel presents the same three stars as a function of their photometric precision. The image scale is 0.4 arcsec/pixel and the seeing (FWHM) was near 1.2 arcsec. From Howell (1989).

signal using an aperture that is smaller then the actual PSF radius itself, some of the source light that was collected by the CCD is not included in the summing process and is thus lost. This sounds like an incorrect methodology to use, but remember that inclusion of many pixels lying far from the source center also means inclusion of additional noise

contributions to the aperture sum. Therefore, while one may wish to obtain the maximum S/N possible through the use of a smaller aperture (i.e., summation of less than the total collected source counts), for the final result it is often necessary to correct the answer obtained for this shortcoming.

In order to recover the "missing light," one can make use of the process of aperture corrections or growth curves as detailed by Howell (1989) and Stetson (1992). Growth curves do not make any demands on the underlying PSF except through the assumption that any bright stars used to define the aperture corrections are exact PSF replicas of any other (fainter) stars that are to be corrected. As we can see in Figure 5.7, growth curves for the brightest stars follow the same general shape, leading to minor or no necessary aperture corrections at a radius of $3 \cdot \mathrm{FHWM}$. The fainter stars, however, begin to deviate from the canonical growth curve at small radii, resulting in 0.5 up to 1.5 magnitudes of needed aperture correction. In general, as a point source becomes fainter, the wings of the source will contain pixels that have an increasingly larger error contribution from the background, leading to greater deviations from a master growth curve at large r, and thus a larger aperture correction will be needed.

As we will see below, if differential photometric results are desired, the aperture corrections used to bring the extracted source signal back to 100% are not necessary. This is only strictly true if all point sources of interest (those to be used in the differential measures) are extracted with the same (optimum) aperture† and have identical PSFs.

5.4 Absolute versus Differential Photometry

Whether an observer needs to obtain absolute or differential photometric measurements depends on the objectives of the scientific and observational program. Absolute photometry results in a measurement of a given astronomical source leading to the true level of the flux received, say in ergs sec^{-1}, or the total luminosity of a source in ergs, each within the specific bandpass used (e.g., a Johnson R filter). Differential photometry

† Note, it is likely that on a given CCD image all stars of interest will not be of exactly the same brightness and will therefore not all have exactly the same optimum aperture radius (see Figure 5.6). Thus, a compromise is usually needed in which the extraction radius used for all sources of interest is set to that of the optimum size for the faintest stars. This procedure allows the faintest sources to produce their best possible S/N result while decreasing the S/N for bright stars only slightly.

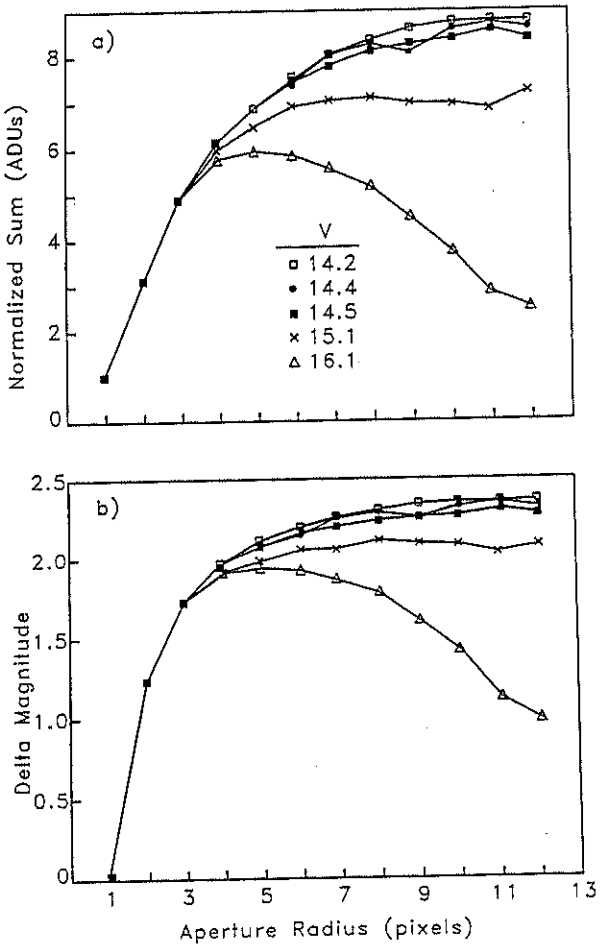

Fig. 5.7. Growth curves for five stars on a single CCD frame. The three brightest stars follow the same curve, which is very similar to the theoretical expectation as shown in Figure 5.5. The two faint stars start out in a similar manner, but eventually the background level is sufficient to overtake their PSF in the wings and they deviate strongly from the other three. Corrections, based on the bright stars, can be applied to these curves to obtain good estimates of their true brightnesses. The top panel presents growth curves as a function of normalized aperture sums while the bottom panel shows the curves as a function of magnitude differences within each successive aperture. The relative magnitudes of the point sources are given in the top panel and the image scale is the same as in Figure 5.6. From Howell (1989).

is just that: The final result is only known with respect to one or more other sources. This is a relative measurement, which, if the other source(s) are known in an absolute manner, can be placed on an absolute scale. All scientific measurements are really differential ones. The difference between absolute and differential measurements is simply that differencing from a known (in an absolute sense) source allows for an absolute result, whereas differencing from an unknown (in an absolute sense) source can only produce a relative result.

In photometric work one can view this in the following way. If the star Vega has its absolute flux known as a function of color (Tüg, White, & Lockwood, 1977), then a comparison of an unknown star with Vega would allow one to calibrate the unknown in an absolute sense. This differential measurement would produce an absolute result. However, if two unknown sources are compared over time, one may be seen to vary periodically allowing a scientific result to be produced, but its absolute flux is unknown. This is a differential measurement that produces only a relative result. If you are concerned about the true brightness of an object, say within a standard photometric system, or its color (e.g., B) in order to place it on an H–R diagram or to obtain its age or metallicity, then absolute photometry is needed. However, if you are after relative variability, say for periodicity analysis, or color indices (e.g., $B-V$) of a large group of stars such as a cluster in which you are looking for outliers, then differential photometry is likely to suffice.

A discussion of the observational and reduction techniques involved in absolute photometry is beyond the scope of this section. The reader is referred to any of the numerous books and articles that discuss the conversion of digital photometric data to (absolute) magnitudes and standard photometric systems (e.g., Dacosta (1992), Hendon & Kaitchuck (1982), Hiltner (1962), and Young (1974)). The observational methodologies and equational transformations used to convert from CCD instrumental magnitudes to standard photometric system magnitudes are essentially identical to those used for similar observational programs with photomultiplier tubes (PMT). Absolute photometry from CCD data has been shown to be just as reliable and every bit as good as PMT measurements (Walker, 1990; Kreidl, 1993; Young, 1974). At the present, typical very good values for errors in the source magnitude for a absolute photometric result from CCD observations are $\pm 1\%$ or less.

Differential photometry concerns itself with the measurement of the difference in brightness of one or more astronomical sources when compared with one or more reference sources. With a two-dimensional CCD

array, unlike the situation when using a PMT, it is often the case that multiple sources are imaged simultaneously and can therefore be used as references by which to compare the source(s) of interest. The assumption that is made in this simple approach is that the reference source(s) is not variable (at least over the time period of the observation) and the object(s) of interest can be compared with the reference source(s) on a frame by frame basis. The last step is important as it cancels out any seeing or atmospheric effects that may change with time. Proper statistical comparison of the object of interest and the reference(s) must take into account all sources of error as well as photon statistics in order to use differential photometry correctly (Howell & Jacoby, 1986; Howell, 1992; Howell, Mitchell, & Warnock, 1988; Honeycutt, 1992). Differential techniques allow one to obtain incredible precisions from CCD data. For example, photometric errors of near ±0.001 magnitude are easily obtainable with differential techniques and even higher precisions, approaching or reaching the theoretical limit, have been realized (Howell, 1993; Gilliland et al., 1993; Howell et al., 1999; Howell, Everett, & Ousley, 1999).

5.5 Astrometry

The science of the exact location and possible movement of astronomical sources has also been greatly advanced through the use of CCDs. Differential astrometric measurements of 1 milliarcsecond (mas) or better are achievable today, yielding new information on positions, proper motions, and parallaxes of astronomical objects. In differential astrometry, one measures the position of an object with respect to other objects within the same CCD field of view. This type of measurement is applicable to point source PSFs (e.g., stars) as well as moving objects (e.g., comets) and fuzzy objects (e.g., galaxies); however, the two latter types of object have additional complexities not discussed here (Monet et al., 1991; Lasker et al., 1990a; Monet, 1992).

Astrometric information is one method by which distance determination may be made for an object of interest. Of course the desire for a distance does not always translate into the ability to obtain such information. In general, in one night of obtaining a few good CCD observations, the determined position for a "well-behaved" star can be known to an astrometric accuracy of about 3 mas. A few tens of CCD frames taken over the course of a year (at a minimum of three different epochs) will allow astrometric accuracies near ±1 mas to be reached, while the best astrometric positions to date have errors of ±0.5 mas or less.

Astrometry performed with front-side illuminated CCDs will suffer from the effects of the overlying gate structures in ways such as we have discussed previously; that is, locating the exact center of the PSF is not without bias. The front-side gates must be traversed by incoming photons leading to a decreased (or no) blue response, intrapixel variations in photometric (QE) response (Jorden, Deltorn, & Oates, 1994), and the need to use color terms to adjust the measured PSF centroid position to that of the actual source position. We will discuss an interesting CCD effect caused by intrapixel QE variations in the next chapter.

Back-side illuminated CCDs present a different situation for astrometric work entirely. As mentioned earlier, their response to blue light is much improved and no major source of intrapixel deviation exists. The physical flatness of a back-side illuminated CCD over its surface can be of concern and can introduce an additional term to be corrected for when measuring source positions and relative source offsets. Additionally, thinned CCDs may allow very long wavelength light to pass entirely through the device, be reflected back into the CCD itself, and collected at a slightly different location (e.g., a neighboring pixel) than its original incoming path location. This long wavelength reflection effect can cause a slight blurring of the measured light from the source, a particularly worrisome issue for astrometry of very late type stars. The best astrometric measurements to date, that is, the ones with the smallest rms error and greatest repeatability, have all been made with thinned, back-side illuminated CCDs.

As with any CCD imaging or photometric observation, the use of a particular type of CCD and filter combination will produce a different effective wavelength of the imaged scene and thus a change in the relative positions of different sources as a function of the zenith distance of the observation and the color (spectral type) of each individual source. There is no simple or even complete solution to this issue of differential color refraction but there are some items to note.† The use of narrow band filters in front of your CCD is of help in this situation because such filters greatly restrict the bandpass of the observation, thereby reducing the effects of differential refraction and color terms. However, the use of narrow band filters is usually impractical owing to the large loss of incoming signal from the astronomical sources of interest. Astrometric observations made in long wavelength (red) bandpasses have merit as they eliminate much of the refractive effects to start with. Finally, CCD observations obtained near source meridian passage are also a plus given

† While reading the remainder of this paragraph, the reader may wish to skip ahead and glance at Table 6.2.

that for an airmass of 1.0, refractive changes are essentially zero within the entire CCD field of view. Details of astrometric observations and a discussion of such effects is presented in Monet et al. (1991) and Monet & Dahn (1983).

Issues of related concern for precision astrometry involve image focus, seeing, use of different filters, PSF stability, telescope collimation, and many others (Monet, 1992). Detailed information concerning one's CCD is vital when attempting precision astrometry. For example, the standard pixel size value available in most of the literature for a 15-micron pixel TI CCD states that the pixel size is 15 microns across. However, the true pixel size is 15.24 microns, a difference of one quarter of a pixel, a value that can cause significant errors in precise astrometric measurements. Finally, even more subtle effects such as nonuniformly spaced pixels and improperly produced pixels that differ slightly (~1%) in size must be considered. The bottom line is, as in all cases where highly precise results are desired, one must know thy CCD in great detail.

Data reduction methods of CCD images from which astrometric measures are to be obtained are similar to those discussed above for general CCD imagery. The differences in the process occur when the final reduced frames are to be used to produce output astrometric information. The interested reader is referred to the discussions and results given in Monet et al. (1991), Lasker et al. (1990a), Monet (1992), and Monet & Dahn (1983).

5.6 Pixel Sampling

An important consideration in photometric and astrometric measurements made with a CCD is how well the PSF is sampled on the two-dimensional array. PSFs that are well sampled by a CCD observation will lead directly to the result that the center and shape of the PSF will be known to higher precision, and thus one will obtain a final answer that will be of higher accuracy. We can define a sampling parameter, r, as follows (Howell et al., 1996; Buonanno & Iannicola, 1989):

$$r = \frac{\text{FWHM}}{p},$$

where FWHM is the full-width half-maximum value of the source PSF and p is the pixel size, both values given in the same units. For r less than about 1.5, digital data are considered undersampled. As can be seen from the above expression, r will be small for the case of a CCD

with large pixel sizes compared with the total areal coverage of the imaged PSF. The other possible case of small r values is if the CCD image contains very tight PSFs such as those that might be obtained at observing sites with very good seeing, if using adaptive optics systems, or for CCD images obtained outside the Earth's atmosphere (i.e., space-based telescopes). Real life examples of cases that will produce undersampled images (i.e., small r values) are a typical wide-field telescope outfitted with a large-format CCD, such as a Schmidt telescope or a camera lens, or a space-based telescope such as the *Hubble Space Telescope* wide-field planetary camera (WFPC) (Howell et al. (1996), Holtzman (1990), and Section 7.1).

Anytime the value of r approaches the limiting case of undersampling, standard software methods and techniques of astrometric and photometric data analysis will begin to produce increasingly larger errors and poorer fits as r decreases further (see Section 7.1). Photometric and astrometric errors obtained from CCD observations are related in that the analysis techniques for each type of measurement are very similar. Both photometry and astrometry require intimate knowledge of the centroid position of the source PSF. However, it has been shown (Stone, 1989; Howell & Merline, 1991; King, 1983) that for undersampled data the photometric error is least for source PSFs that are centered on the edge of a pixel or exactly at its center, whereas for astrometric data, the resulting error is least when the source PSF is centered midway between a pixel's edge and its center.

The rule of thumb for pixel sampling on a CCD follows directly from the statistical result of Nyquist sampling. That is, sampling of the PSF of an astronomical source will be optimal in terms of S/N, error rejection, data analysis, and so on for a source PSF that has its FWHM value sampled over about two pixels (i.e., FWHM $\sim 2 \cdot$ pixel size). For example, if the average seeing at an observing site produces source PSFs with FWHM values of 2 arcsec, then an ideal (optimal) CCD pixel size to use would be one for which each pixel within the array has a projected image size of 1 arcsec across.† A rigorous mathematical definition of undersampling, based on the Nyquist theorem, identifies critical sampling as the sampling interval that is equal to the width (i.e., standard deviation) of the PSF. For a Gaussian PSF this corresponds to a FHWM equal to 2.355 pixels. Of course an ideal image scale is hard to meet in reality as seeing and telescope focus change with time, source PSFs generally

† The determination and measurement of CCD pixel size or plate scale was discussed in Section 4.1.

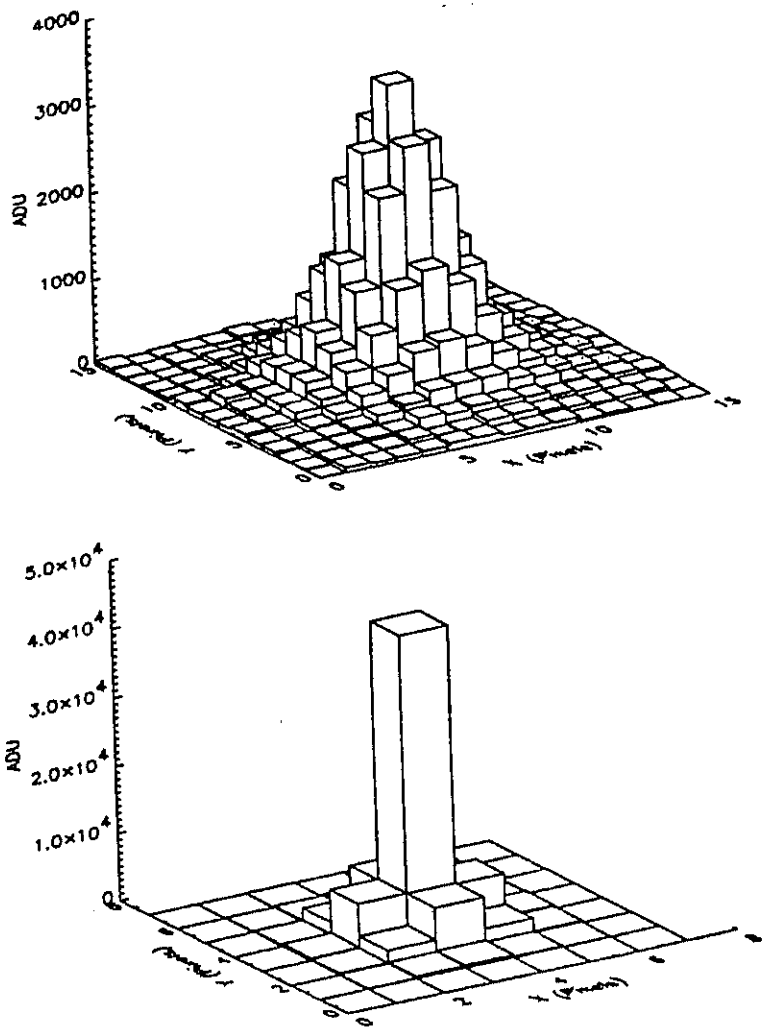

Fig. 5.8. The effects of pixel sampling are shown. The top PSF is a well-sampled star image with a S/N of near 230. The bottom panel shows the same PSF but now severely undersampled and centered at the middle of a pixel and (next page) at the corner of four pixels respectively. Note that the undersampled profiles are not well represented by a Gaussian function. From Howell et al. (1996).

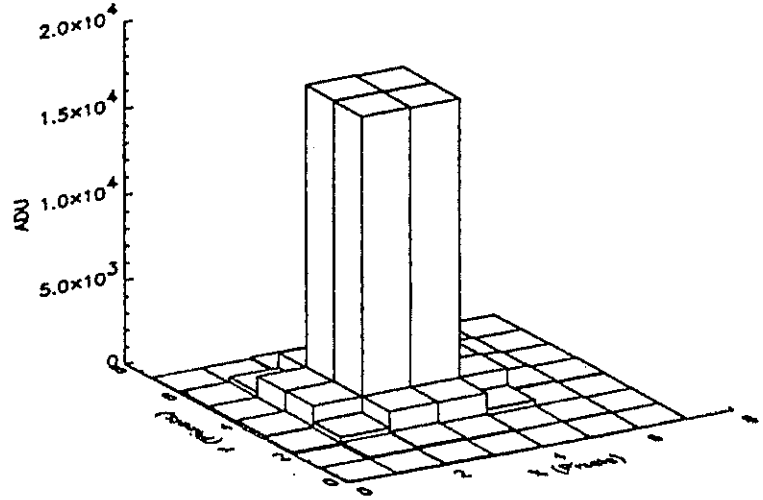

Fig. 5.8. (*Continued*)

do not fall onto the CCD pixel grid exactly on a pixel boundary, and one generally has only a limited number of available CCD cameras with fixed pixel sizes.

Since CCD detectors do indeed sample astronomical sources in a quantized manner, pixel sampling within the array will cause even identical PSFs to change their detailed appearance slightly, even within the same observation. The effects of such sampling differences become worse as the sampling parameter (r) becomes smaller (Merline & Howell, 1995; Howell & Merline, 1991). Figure 5.8 illustrates some examples of various values of r caused by CCD pixel size. The top panel shows a well-sampled source PSF that appears to be more or less Gaussian in shape. The remaining two panels in Figure 5.8 show the same model PSF but now imaged as a poorly sampled ($r = 1$) source. The undersampled cases are for a source PSF with a pixel-centered centroid and a corner-centered centroid respectively. Notice in Figure 5.8 that the undersampled PSFs are not well represented by a Gaussian function (Howell et al., 1996; Buonanno & Iannicola, 1989; Howell & Merline, 1991; Holtzman, 1990).

6

Spectroscopy with CCDs

Although imaging and photometry have been and continue to be mainstays of astronomical observations, spectroscopy is indeed the premier method by which we can learn the physics that occurs within and near the object under study. Photographic plates obtained the first astronomical spectra of bright stars in the late nineteenth century, while the early twentieth century saw the hand-in-hand development of astronomical spectroscopy and atomic physics. Astronomical spectroscopy with photographic plates, or with some method of image enhancement placed in front of a photographic plate, has led to numerous discoveries and formed the basis for modern astrophysics. Astronomical spectra have also had a profound influence on the development of the fields of quantum mechanics and the physics of extreme environments.† The low quantum efficiency and nonlinear response of photographic plates placed the ultimate limiting factors on their use.

During the 1970s and early 1980s, astronomy saw the introduction of numerous electronic imaging devices, most of which were applied as detectors for spectroscopic observations. Television-type devices, diode arrays, and various silicon arrays such as Reticons were called into use. They were a step up from plates in a number of respects, one of which was their ability to image not only a spectrum of an object of interest, but simultaneously, the nearby sky background spectrum as well – a feat not always possible with photographic plates. Additionally, and just as important, were the advantages of higher obtainable signal-to-noise ratios, higher quantum efficiency, and very good linearity over a large dynamic range. These advances permitted spectral observations of much fainter

† Extreme physical environments, such as very high temperatures, magnetic fields, and gravitational fields, are not often available within earthly laboratories.

sources than were previously available. Two-dimensional spectroscopy allowed the large error contribution from the often unknown sky background to be removed during data reduction procedures. However, the above electronic devices still had their problems. Use of high voltage, which caused distortions in the image, and the often low dynamic range both limited the ability to numerically resolve strong spectral lines from a weak continuum or to resolve weak lines in general.

The introduction of CCD detectors for use in astronomical spectroscopy was quick to follow their introduction into the field of astronomy itself. One of the first devices put into general use for astronomical spectroscopy was a Fairchild CCD placed into service at the Coude' feed telescope on Kitt Peak circa 1982. This author remembers that particular device and the remarkable advance it made to astronomy at the time. Observing without photographic plates was amazing.† The introduction of CCDs allowed test observations to be made, no chemical development was needed, you could view your data almost immediately after taking it, and mistakes caused you very little in lost time. In addition, fainter sources and better spectral resolution were easily obtained and caused a renaissance in astronomical spectroscopy. The fact that the Fairchild CCD had a read noise of 350 electrons seemed unimportant at the time.

We will begin our discussion of astronomical spectroscopy with point source observations. The term "point source" in generally taken to mean a star, but under various conditions, many objects can be observed as a point source. For example, an active galaxy does indeed often show extended structure in terms of spiral arms, but short exposures or observations intended to study only the nuclear regions are essentially point source measurements. A more formal definition might be that point sources are objects for which their angular size is determined by the seeing disk or instrumental resolution. We will follow point source observations by introducing extended object spectroscopy. The major difference in these two types of spectroscopy is the type of output data product you end up with and the science obtained from the collected data. Our discussion here will concentrate more on the CCD aspects of astronomical spectroscopy and not as much on the actual observational techniques

† For those of you interested in a bit of nostalgia, remember how one needed to cut the photographic plates completely in the dark, attempt to fit them into the plate holder licking one side along the way to find the emulsion, and then suffering the agony of defeat when you discovered that your 1 hour integration was made with the dark slide closed or the plate had been placed in the holder backwards!

and data reduction procedures. Various types of spectrographs and other related topics are discussed in detail in the excellent reviews given by Walker (1987), Wagner (1992), and Pogge (1992).

6.1 Review of Spectrographs

Spectroscopic observations can be thought of as a method by which one samples the emitted energy distribution from an astronomical source in wavelength bins of size $\Delta\lambda$. Broad-band filter photometry, for example, is a form of spectroscopy; it is merely one with extremely poor spectral resolution. To use spectral information to learn detailed physics for an astronomical object, one must be able to differentiate specific spectral features (lines) from the continuum within the observed spectrum and be able to make quantitative measurements of such features. Generally, this type of analysis requires a spectral resolution of at least 20–40 Å or better. Keep in mind, however, that various scientific objectives can be accomplished with varying amounts of spectral resolution. Schmidt telescope observations using an objective prism and imaging each spectrum onto a CCD have fairly low spectral resolution, but the imaged spectra are indeed useful if the purpose is to identify objects that have blue color excesses (see Section 6.7).

Figure 6.1 illustrates a typical astronomical spectrograph with the common components identified. An entrance slit, onto which the telescope focuses the incoming light from the source of interest, is used to both set the spectral resolution and to eliminate unnecessary background light. An internal light source for the production of a flat field (called a projector flat in spectroscopy) and various wavelength calibration emission line sources are also included. These lamps usually consist of a quartz projector lamp for the flat fielding and a hollow cathode or arc lamp for the calibration sources. Both types of calibration lamp are included in the spectrograph in such a way as to attempt to make their light path through the slit and onto the CCD detector match as closely as possible that of the incoming telescope beam from an astronomical object. Some type of grating (commonly a concave reflection grating) is needed as the dispersive element, although a prism can also be used. Various camera optics, to re-image the slit onto the CCD detector and provide chromatic and field flatness corrections, finish the suite of standard components. Numerous variations on this standard theme have been and will continue to be used as cost, complexity, and purpose of the instrument are always issues.

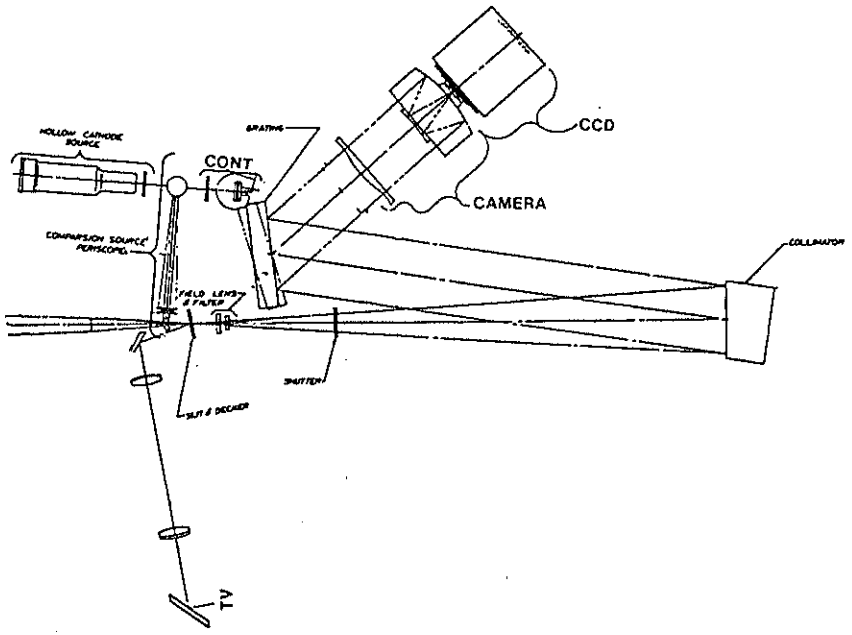

Fig. 6.1. Schematic diagram of a typical astronomical spectrograph. The major components are the CCD detector, the continuum and comparison line calibration sources, the TV slit viewer, and the grating. From Wagner (1992).

Let us define a few useful quantities in CCD spectroscopy. Table 6.1 lists the various definitions needed for use when discussing the optical properties of the telescope, collimator mirror (or lens), and the camera itself. Using the definitions in Table 6.1, we can define the magnification of the spectrograph as $M = F_{\rm cam}/F_{\rm col}$, the projected width of the slit

Table 6.1. *Spectrograph definitions*

D	Diameter of Telescope
F	Focal Length of Telescope
f	Focal Ratio of Telescope
w	Width of Entrance Slit
$d_{\rm col}$	Diameter of Collimating Mirror
$F_{\rm col}$	Focal Length of Collimator
$f_{\rm col}$	Focal Ratio of Collimator
$d_{\rm cam}$	Diameter of Camera Mirror
$F_{\rm cam}$	Focal Length of Camera
$f_{\rm cam}$	Effective Focal Ratio of Camera

at the CCD will be wM, and the slit width will subtend an angle on the sky of $\Theta = w/F$. The projected width of the slit at the CCD detector is then given by

$$r = wM = \Theta FM = \Theta f D \frac{F_{\rm cam}}{F_{\rm col}} = \frac{\Theta f D F_{\rm cam}}{f_{\rm col} d_{\rm col}}.$$

To avoid loss of efficiency within the spectrograph, the collimator focal ratio should match that of the telescope. To make this clear, we can write the projected slit width at the CCD detector as $\Re = \Theta D f_{\rm cam}$. We assume here that the dispersing element does not change the collimated beam size.† If the slit is opened wide enough to allow all the light from a point source to pass through, the projected image size of the point source at the CCD detector is simply $\Re$.

The ability to separate closely spaced spectral features is determined by the resolution of the spectrograph. Spectral resolution is defined as $R = \lambda/\Delta\lambda$, in which $\Delta\lambda$ is the difference in wavelength between two closely spaced spectral features, say two spectral lines of equal intensity, each with approximate wavelength λ. Optical light spectral resolutions of a few hundred thousand to one million have been obtained for the Sun, whereas for typical astronomical spectra, R is much less, being near a few thousand down to a few hundred for very faint sources. For comparison, current R values in the infrared are typically less than about 1,000.

For further details of the actual components of astronomical spectrographs, various types of spectroscopy for different applications, and the way in which these components are used to produce spectra at various wavelengths and resolutions, see Wagner (1992), Pogge (1992), DeVeny (1990), Stover et al. (1995), Queloz (1995), Corbally (1995), Cochran (1995), and Robinson (1988b).

6.2 CCD Spectrographs

Current-day spectrographs almost exclusively use CCDs as their detector. The major reasons for this choice are the large free spectral range of modern CCDs (covering roughly 3,000 to 10,000 Å), the linearity of the devices (better than 1% from zero counts to saturation) over a large dynamic range (allowing, for example, detection of absorption or emission lines as well as the continuum), and the large areal format of

† While not always true for diffraction gratings, this condition is realized for prisms.

modern CCDs (2,048 up to 4,096 pixels or more in extent). This latter property is especially important for applications such as wide-field objective prism work with Schmidt Telescopes, multiobject fiber spectroscopy (Robinson, 1988c) in which the fiber fed spectra are placed in row order on the CCD, and for Echelle spectroscopy for which many spectral orders are two-dimensionally imaged at once (Vogt & Penrod, 1988). The free spectral range obtained could, in principle, be as large as the detector's quantum response, but in practice limits in optical and grating design and CCD size restrict a single spectrograph coverage to somewhere near 4,000 Å or less in the optical bandpass. Some observatories have solved this limited spectral coverage issue by designing and using double spectrographs, with each instrument having two separate spectrograph arms (one red and one blue) with a CCD for each (Oke, 1988). A double spectrograph almost always uses a different type of CCD detector in each arm; each CCD is customized for the best possible detection properties for their particular wavelength coverage. Both high resolution ($R \sim$ 30,000–80,000 or more) and low resolution ($R \sim$ a few thousand or less) spectroscopic applications are well suited to using CCDs as the detector; one simply has to match the CCD pixel size to the particular spectrograph and resolution being used.

Optimal sampling of a spectral line that is just unresolved occurs when the FWHM of the line is twice the physical pixel size (Nyquist sampling criteria). It can be assumed in this discussion that a spectral line has an approximately Gaussian shape for which a formal FWHM value can be determined. Note, however, that real spectral lines are not always this well behaved. In addition to matching the spectral line width to the pixel size, CCDs used for astronomical spectroscopy must also have very good charge transfer efficiency (CTE) in order to reduce smearing of spectral lines during readout, which would lead to a loss in spectral resolution. Also, small pixel sizes such as 15 or 9 microns are often desired to meet the Nyquist criteria discussed above.

Let us look at an example for a spectrograph that uses a CCD with 9-micron pixels as the detector. With this setup, the projected slit width size, $\Re$, must be near 18 microns to achieve optimal sampling. For an observing site with typical seeing of 1.5 arcsec, and using a 2- to 5-m telescope, we find (using the formulations given in Section 6.1) that the spectrograph camera must have a focal ratio near unity. This is a very fast focal ratio and requires excellent optical design and near perfect optical surfaces. For the CCD itself, this requirement means that its

physical surface must be extremely flat throughout the entire extent of the chip (less than 0.5% rms for accurate spectrophotometry), in order to allow all parts of the spectrum to be in focus simultaneously. As we have seen, this level of flatness can be a difficult requirement to fulfill for certain types of CCD (e.g., thinned devices).

The above example for a CCD spectrograph informs us that, for large-aperture telescopes (say 8–12 m), optimum spectral sampling can only occur if some combination of the following conditions are met. As the telescope diameter increases, the camera focal length must decrease, the seeing disk must decrease, and the detector resolution element (2 CCD pixels) must increase in size. Currently, the hardest requirement to meet in this list is the design and construction of very fast focal ratio cameras. Increasing the CCD pixel size while retaining the large range of total wavelength coverage is the major driving force behind producing even larger format CCDs.

Exceptional seeing, less than 1 arcsec for example, would seem to be the dream of any spectroscopist. However, let us look at an example when very good seeing can cause unexpected results in CCD spectroscopy. The problem is as follows: CCDs are mounted in dewars and attached to the end of a spectrograph in some manner. The dewars are then aligned in an attempt to have the observed spectrum fall onto the detector along either its rows or columns. Perfect alignment of the CCD across the entire spectrum is rarely achieved and thus the imaged spectrum centroid must cross pixel boundarys in the dispersion direction (i.e., as a function of wavelength).

If the object seeing disk becomes less than the projected pixel size, the position of the spectral centroid falls within the pixel itself, alternately occurring at the center of some pixel and then at the pixel boundaries themselves. Wavelength-dependent QE effects within the pixels, due to their gate structures and intrapixel "dead" spots, will cause apparent flux variations that can be as large as ±10% in amplitude. Additional complexities, such as which type of CCD is used, telescope focus, and tracking changes, are harder to quantify and correct for but can have similar effects. The problem described here, that of undersampling, can also occur in CCD imaging applications as well, when the majority of the source PSFs fall within a single pixel. Optimum sampling in CCD imaging also occurs at the Nyquist sampling limit, that is, a point source FWHM should be imaged across about two CCD pixels (see Section 5.6). A several percent error in CCD photometry can occur for images that are undersampled (Howell et al., 1996; Holtzman, 1990).

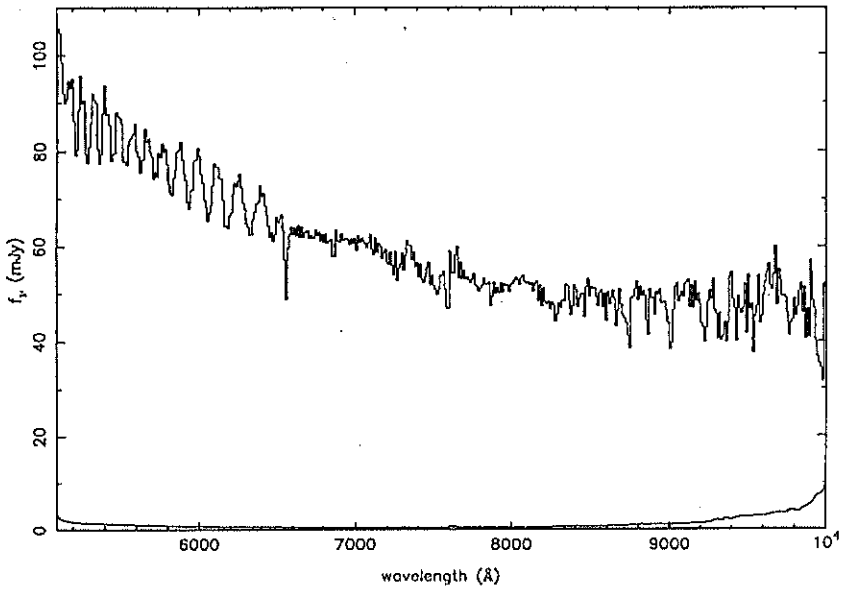

Fig. 6.2. Spectrum of the standard star Feige 92 taken during a time of excellent seeing. This 30-second exposure, obtained with the Faint Object Spectrograph on the Issac Newton telescope, has a pixel scale of 1.2 arcseconds/pixel and this spectrum was obtained when the seeing was sub-arcsecond. The ripple seen in the continuum near 6,000 Å and blueward is due to intrapixel QE differences, which are especially prominent for short wavelength photons in the EEV CCD used. From Dhillon, Rutten, & Jorden (1993).

Figure 6.2 shows an example of the ripple that can occur in a spectrum obtained under conditions of excellent seeing and for which intrapixel QE effects are present (Rutten, Dhillon, & Horne, 1992). Possible methods of correction for spectral ripple include de-focus of the telescope or slightly trailing the spectrum up and down the slit during each integration,† neither of which are desirable. Corrections to the spectrum can also be applied after the fact using some empirical model during the reduction phase (Dhillon, Rutten, & Jorden, 1993; Jorden, Deltorn, & Oates, 1993). Spectral ripple, as well as our discussion of pixel sampling in the chapter on photometry, indicates that while very poor sampling is not ideal (as the collected image is spread out over many pixels, each of which adds unwanted noise) undersampling is not necessarily better. Optimum sampling is the best but is not always possible with a given instrument and CCD combination. In addition, the conditions of

† This particular solution may remind us long-time observers of a Wobble Plate.

optimum sampling can change with time owing to effects such as seeing, telescope focus, the wavelength of light imaged, and other more subtle effects.

The two-dimensional nature of a CCD allows one to place the spectrum anywhere on the array, targeting good regions of the CCD that avoid manufacturing flaws such as areas of poor sensitivity or dead pixels. Their 2-D design also provides the ability to simultaneously image the (nearby) sky spectrum with the object spectrum. Accurate sky subtraction from a source spectrum increases the S/N of the final result and allows much fainter sources to be observed. It also increases the reliability of the flux measurements and is probably the greatest factor making CCDs superior to other detectors as spectroscopic imagers.

Although large-format CCDs for spectroscopy (2,048 × 2,048, 1,024 × 3,072, etc.) are useful in many ways, they have their drawbacks as well. Large CCD formats take longer to read out, make critical areal flat fielding (important for accurate fluxes and line profiles) difficult, and provide more stringent restrictions on mounting and dewar design and overall cost. Physically nonflat CCDs under vacuum (such as thinned devices), shifting of the CCD within the dewar, or movement of the larger, heavier dewar itself can all cause slight movement of spectral features, leading to errors in the analysis. Loss of LN2 throughout the night has been identified as one cause of spectral movement (Smith, 1990a).

Long readout times for large-format CCDs (a few minutes) do not seem like such a big deal compared with typical integrations of 30–60 minutes or more. However, the longer the readout time, the lower the overall observational duty cycle one experiences, and, after obtaining all the necessary calibration spectra, flat fields, comparison arcs, and standard stars, this extra time can become costly. In addition, to produce the final spectrum one often wishes to co-add shorter object exposures to avoid numerous cosmic ray events that can hinder accurate flux and line profile measurements. Obtaining fast temporal coverage with large-format CCDs, such as required for phase-resolved spectroscopy, also gets thwarted, again by these long readout times. In a likewise manner, CCDs that have a limited dynamic range (those with small pixel sixes) force the user to make multiple (shorter) exposures for each type of needed image, especially calibration data for which the high S/N desired can easily saturate the detector.

The finite readout time of large-format CCDs has led to numerous attempts to obtain high-speed spectroscopy via some technique that takes

advantage of the two-dimensional nature of the detector. For example, one could step the detector, or the spectrum, along in a direction perpendicular to the dispersion at some predetermined rate, reading out the CCD only after a entire frame of individual spectra is collected. A better idea might be to use the fact that the collected charge can be moved back-and-forth electronically on the CCD from row to row, without actually reading it out. This movement is quite fast, say 200 rows in a few milliseconds, and suffers no readout noise penalty but only a slightly increased dark current. Spectra placed on the top 50 rows of a device will allow faster readout times or one could use the bottom 50 rows with electronic charge movement used to shift the spectrum upward between each integration time. Both of these processes have been tried with success (Robinson, 1988b).

One could even slosh the charge back-and-forth periodically in order to build up the signal in two spectra for say a given phase of a binary star. The spectrum not being exposed gets hidden under a mask and the final image produces two or more "phased" spectra. An idea such as this was developed for spectropolarimetry (Miller, Robinson, & Goodrich, 1988) with some success and led to the discovery of some previously unknown, yet interesting CCD effects (Janesick et al., 1987a; Blouke et al., 1988).

6.3 CCD Spectroscopy

Two common measurements from an astronomical spectrum are those of the flux level as a function of wavelength and the shape and strength of spectral lines. Currently, the best determined absolute fluxes are near 1%, usually for spectral data covering a small wavelength range and for bright (high S/N) sources (e.g., Tüg, White, & Lockwood (1977), Furenlid & Meylon (1990), Hayes & Latham (1975), Oke & Gunn (1983), and Schectman & Hiltner (1976)). Absolute flux observations must make use of a wide slit or hole (10 arcsec or larger in size is typical) as the entrance aperture to the spectrograph. Use of such a wide slit assures that 100% of the source light is collected and avoids problems related to telescope tracking or guiding errors and differential refraction (see below).

In principle, relative flux measurements should be determinable to even better levels than those quoted above and the shape of a spectral line should be almost completely set by the instrumental profile of the spectrograph. Spectroscopy desiring relative fluxes generally makes use of a narrow slit (matched to the seeing) to preserve the best possible spectral resolution. For example, for an observation requiring

accurate radial velocity measurements a narrow slit would be preferred to achieve the highest possible spectral resolution. After all, the imaged spectral features (lines) are merely images of the slit itself focused onto the CCD.

Calibration of observed object spectra is normally performed by obtaining spectra of flux standard stars with the same instrumental setup including the slit width. During the reduction process, these standard stars are used to "flux" the object data, that is, assign relative fluxes to your object spectrum counts as a function of wavelength. It is assumed that during each integration, the object of interest and each standard star sent the same relative color and percentage of their PSF light per time interval through the slit and onto the CCD. Seeing changes, nonphotometric conditions, guiding errors, and differential refraction can all negate this assumption.

Three factors unrelated to the CCD detector itself – tracking and guiding errors, seeing changes, and spectrograph slit angle – are important to your final spectral result in the following ways. Observations of an astronomical source made through a spectrograph slit are often obtained such that the slit size matches the object seeing disk. That is, the slit width is set to allow most of the PSF (for a point source) to pass through but kept small enough to eliminate as much of the sky background as possible from entering the spectrograph. Therefore, changes in telescope guiding or image seeing will cause slightly more or less source light to enter through the slit. These effects can cause large noticeable effects as well as more subtle possible unknown effects to be part of your final image. These issues, as well as items such as the use of a red sensitive slit viewing/guiding camera when observing very blue sources, will not be discussed further here (Wagner, 1992).

The angle on the sky of the spectrograph slit is very relevant to the use of CCDs as detectors. Observations of an astronomical source at some angle away from the zenith (an airmass of 1.0) can cause differential refraction to become an issue. Differential refraction is the variation of the angle of refraction of a light ray with zenith distance. All objects in the sky become slightly prismatic owing to differential refraction by the Earth's atmosphere. If the spectrograph slit is not placed parallel to the direction of atmospheric dispersion, nonchromatically uniform light loss may occur at the slit since the image will have an extended, color-dependent PSF. Atmospheric dispersion is caused by the variation of the angle of refraction of a light ray as a function of its wavelength, and

the direction parallel to this dispersion is called the parallactic angle. Using spherical trigonometry, the parallactic angle can be determined from the following:

$$\cos(\text{object declination}) \times \sin(\text{parallactic angle}) = \text{sign}(\text{hour angle}) \times \cos(\text{observers latitude}) \times \sin(\text{object azimuth}),$$

where sign $= +1$, if the hour angle is positive and -1 if it is negative.

Table 6.2 presents an example of the amount of differential refraction that a point source will experience, as a function of airmass, for an observatory at an elevation of 2,200 m. Note that all values in the table are relative to 5,500 Å and are of significant magnitude at essentially all wavelengths and airmasses. CCD observations are particularly sensitive to an effect such as differential refraction as their very good red or blue QE could produce spectral data with highly inaccurate flux values. Even more troublesome is the case of using, for example, a blue flux standard star when observing red sources. Differential refraction can cause the light entering the slit, and therefore imaged by the CCD itself, to record incorrect intensities as a function of wavelength. Aligning the slit at the parallactic angle solves this problem.

A fun example to consider concerning differential refraction is the possible use of a direct CCD imager at very high airmass. The imager will act like a spectrograph and record spectra of sources using the Earth's atmosphere as the dispersive element. Some simple calculations will reveal that while possibly a good idea in principle, at the very high airmasses needed, the dispersion is extremely nonlinear and changes very rapidly with time. Extremely short exposures would be required, thereby allowing data to be gathered for only the brightest of stars. In addition, most astronomical telescopes can not point to the needed positions to obtain these sorts of very high airmass observations.

Some additional considerations related to the use of CCDs in astronomical spectrographs are those of order sorting filters, stray light, fringing, and economics. The first three of these cause undesired light to be imaged onto the CCD, some of which falls directly on top of the object spectrum of interest. This extra light increases the apparent background, decreases the resulting S/N, and is often difficult or impossible to remove or even measure. As an example, using a diffraction grating to observe blue light in second order necessitates removal of the first-order red spectral light from the incident beam before detection by the CCD. Since CCDs are generally very red sensitive, this is a critical step to perform.

Table 6.2. *Example Differential Refraction Values (in arcsec)*

Altitude	Airmass	3000 Å	3500 Å	4000 Å	4500 Å	5000 Å	5500 Å	6000 Å	6500 Å	7000 Å	8000 Å	9000 Å	10000 Å
90	1.00	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
75	1.04	0.6	0.35	0.2	0.1	0.1	0.0	0.0	−0.1	−0.1	−0.1	−0.1	−0.2
60	1.15	1.3	0.8	0.5	0.2	0.1	0.0	−0.1	−0.1	−0.2	−0.3	−0.3	−0.4
45	1.41	2.3	1.45	0.8	0.4	0.2	0.0	−0.1	−0.2	−0.3	−0.5	−0.6	−0.6
30	1.99	4.0	2.45	1.4	0.7	0.3	0.0	−0.2	−0.4	−0.6	−0.8	−1.0	−1.1
20	2.90	6.3	3.9	2.3	1.2	0.5	0.0	−0.4	−0.6	−1.0	−1.3	−1.5	−1.7
15	3.82	8.5	5.3	3.1	1.6	0.7	0.0	−0.5	−0.9	−1.3	−1.7	−2.1	−2.3
10	5.60	12.9	8.0	4.7	2.4	1.1	0.0	−0.8	−1.3	−2.0	−2.6	−3.1	−3.5
5	10.21	26.1	16.2	9.5	4.9	2.2	0.0	−1.6	−2.7	−4.0	−5.2	−6.3	−7.0

A $CuSO_4$ filter is probably the best order sorting filter to use in this circumstance but is far from ideal. $CuSO_4$ filters have poor UV and blue transmission, are inconvenient to use (they are either a crystal or a liquid filter), and have a long wavelength red leak. Further discussion of such intricacies can be found in books concerning spectrographic observations of astronomical sources; also see Wagner (1992), Pogge (1992), and Robinson (1988b).

Finally, we conclude this section with our wish list for the ideal CCD to use within an astronomical spectrograph. The CCD should have low read and dark noise, high QE over a large wavelength range, very good CTE, small pixel size, and large dynamic range, and it should be "tuned" to the application desired. Tuning a CCD simply means that one should use the device that is best suited for the job at hand. Properties such as high red sensitivity, very deep pixel wells, small pixel size, back-side illumination and coating, etc. are items worthy of consideration.

6.4 Signal-to-Noise Calculations for Spectroscopy

Calculation of the signal-to-noise ratio for spectroscopic observations is performed in a manner similar to that described earlier in this book (Chapter 4). As in photometric measures, we find that for bright sources $S/N \propto \sqrt{N}$ while for faint sources we must use the entire S/N expression (see Section 4.4). For spectroscopic observations obtained with a good CCD system, the largest noise contributors that will degrade the resulting S/N are the background sky contamination and how well the data can be flat fielded. The value of the S/N of a spectroscopic observation can have a few different meanings.[†] For CCD spectroscopy, one can calculate the S/N for the continuum or the S/N for a given spectral line.

In the continuum case, the number of pixels, n_{pix}, used in the S/N calculation will be determined by the continuum bandpass range over which the S/N is desired times the finite width of the spectrum on the CCD. For example, a typical CCD spectrograph might have an image scale of 0.85 Å/pixel and the imaged spectrum may have a width of 3 pixels on the array in the direction perpendicular to the dispersion. To calculate the S/N for the spectral continuum over a 100 Å bandpass

[†] Note here that this is also true for CCD imaging or photometric applications as well. Depending on the user's choice of parameters such as n_{pix}, the final S/N value will change. Likewise, reported S/N values without comment on the exact choice of specific parameters are, at times, difficult to interpret.

in this example, one would use the value of 353 for $n_{\rm pix}$. In contrast, a narrow emission line with a full width at zero intensity (FWZI) of 40 Å would use an $n_{\rm pix}$ of 141. The S/N of the emission line will therefore be higher in value owing to the smaller overall error contribution (approximately 3 to 1) not to mention the higher flux values per pixel within the line itself. As an exercise, the reader might consider how one would calculate the S/N in the case of an absorption line.

Signal-to-noise calculations are also useful in predicting observational values such as the integration time needed to obtain a desired scientific result from your spectroscopic observations. This type of calculation can be performed using the formulae presented in Section 4.4. Spectroscopic S/N values in the continuum of near 10 are often sufficient for gross spectral classification, whereas values in excess of 500 are needed for detailed abundance analysis or weak absorption line measurements. When making predictions of the S/N to expect in a spectroscopic observation, keep in mind that spectrographs are much less efficient overall than direct imaging cameras (2–4% vs. 30–40%) and that seeing effects and slit losses can be considerable in terms of the actual flux collected by the spectrograph.

6.5 Data Reduction for CCD Spectroscopy

This section discusses the basics of astronomical CCD spectroscopic data reduction. Figure 6.3 presents a raw, unprocessed CCD spectroscopic image of a point source. The initial reduction steps for CCD spectroscopy are exactly the same as previously discussed for imaging applications. Bias (or dark) frames and flat field calibration images are needed and used in the same way. After performing these basic reduction procedures for the CCD images, there are additional steps one must take that are specifically related to spectroscopy. These extra steps involve the use of spectra of spectrophotometric flux standards and wavelength calibration (arc) lamps. Figure 6.4 illustrates the five types of CCD images needed for complete reduction and calibration of spectroscopic observations. We will not discuss some minor, yet important processing steps such as cosmic ray removal, bad pixel replacement, and night sky emission line complexities. Detailed instructions for CCD spectroscopic data reduction can be found in Pogge (1992) and Massey, Valdes, & Barnes, (1992). The article by Wagner (1992) is particularly useful, being the best review of the subject to date.

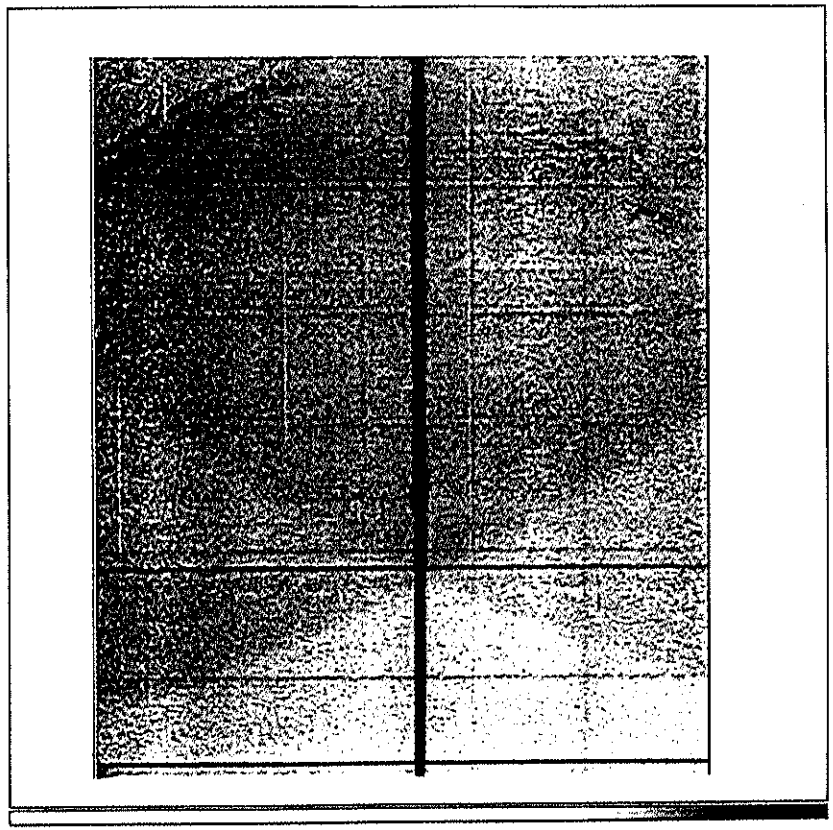

Fig. 6.3. Raw CCD image of a point source. The spectrum runs vertically up and down near the middle of the CCD with wavelength increasing toward the top. Night sky emission lines are clearly seen running horizontally across the entire frame. The sky lines due to OI at 5,577, 6,300, and 6,363 Å are visible at the bottom, and numerous OH ariglow lines are seen near the top of the image. The large-scale gradients are due to nonuniform CCD illumination within the spectrograph. From Wagner (1992).

For a point source, the final product desired is a 1-D spectrum consisting of wavelength versus relative or absolute flux. It is assumed that you have in hand CCD spectra of your object(s) of interest, at least one observation of a flux standard, and wavelength calibration spectra of arc lamps (see Figure 6.4). The first step with the reduced two-dimensional CCD image, after the standard processing with bias and flats is performed, is to extract the spectrum itself and collapse it into

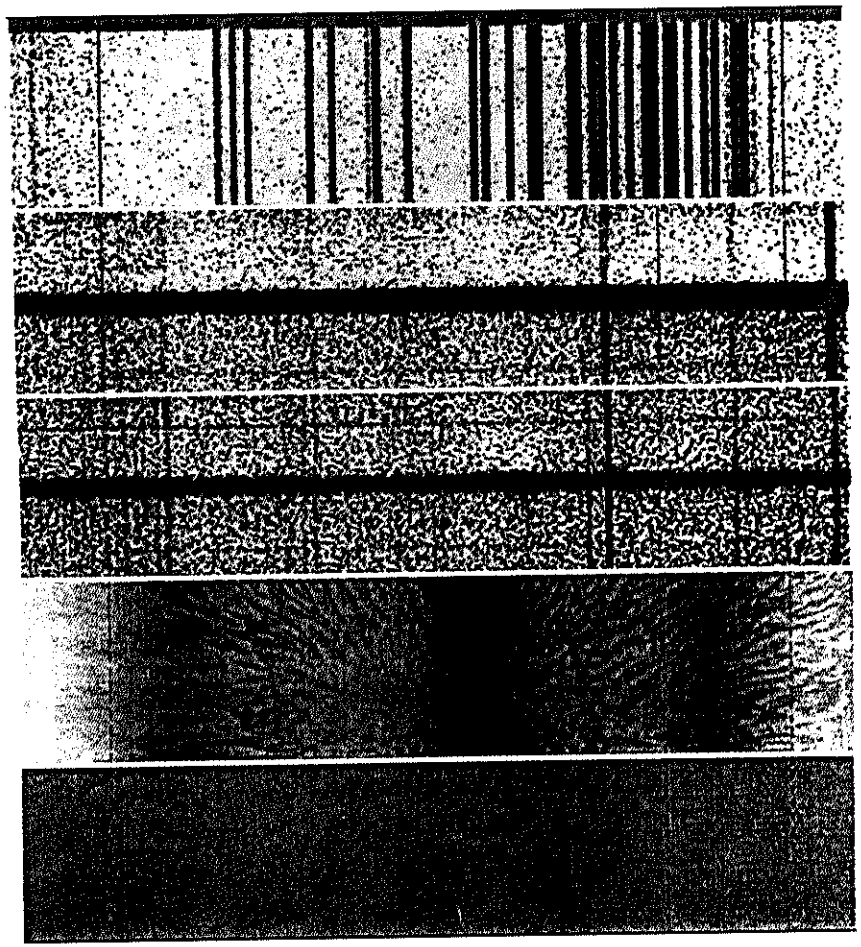

Fig. 6.4. The five necessary types of CCD image needed for spectral reduction. From bottom to top we see a bias frame, a flat field frame, the object spectrum of interest, the spectrum of a flux standard, and finally a FeNe arc spectrum. The flux standard will be used to turn observed counts in the object into relative flux, and the arc spectrum will be used to set the wavelength scale. From Wagner (1992).

a 1-D image of pixel number versus ADU/pixel. In the simplest (and unrealistic) case, the imaged spectrum lies exactly along one row (or column) of the CCD and one can extract it simply by extracting the single row from the final image. Generally, the imaged spectrum covers two or three or more rows (or columns) and the extraction process

involves some manner of summing a few adjacent rows perpendicular to the dispersion (an "after-the-fact" pixel binning) at each point along the dispersion. Furthermore, in practice it is often found that CCD spectra are not precisely aligned with the CCD pixels and are curved on the detector as the result of the camera optics, instrumental distortions, or CCD flatness issues. Fainter sources present additional complexities, as centroiding the spectrum (for, by example, using cuts across the dispersion direction) in order to extract it is often difficult or impossible. A typical example might be a spectral observation of a faint continuum source that contains bright emission lines. Details of the spectral extraction process sky subtraction and optimal extraction techniques are discussed in Wagner (1992), Pogge (1992), Robinson (1988b), Robinson (1988c), Schectman & Hiltner (1976), and Horne (1986).

At this point, your extracted spectrum will have an x axis that is in pixels and we wish to convert this to wavelength. The procedure to perform such a task involves observations of calibration arc spectra obtained often during your observing run. The idea is to match the x-axis pixel scale of the calibration arc lamps with their known wavelengths and then apply this scaling procedure to your object data. Calibration arc spectra of sources such as hollow cathode Fe lamps or He-Ne-Ar lamps contain numerous narrow emission lines of known wavelength. Use of these emission lines during the reduction procedures allows a conversion from a pixel scale to a (possibly rebinned linear) wavelength scale. Correction for atmospheric extinction (Hendon & Kaitchuck, 1982), similar in manner to photometric corrections already discussed but generally a bit more complex owing to the larger wavelength coverage, can now be applied to all obtained spectra. Since you are relying in this step on the hope that the arc emission lines fall onto specific CCD pixels and define a wavelength scale that is identical to that of your object, you want to obtain arc data in as similar a manner to that used to gather your object spectra as possible. Instrument flexure caused by telescope motion throughout the night or movement of the CCD within the dewar are two of many possible effects that will invalidate the wavelength to pixel scaling procedure.

We now have to deal with our collected spectra of flux standards. These stars are observed solely for the purpose of converting the collected pixel count values into absolute or relative flux values. Application of all of the above steps to your spectra of spectrophotometric flux standards will produce 1-D data with a wavelength scale (x axis) versus

counts on the y axis. We now wish to have the y axis of ADUs or counts converted into flux units such as ergs s^{-1} cm^{-2} $Å^{-1}$. Most observatories and data reduction packages (such as IRAF and MIDAS) contain lists of spectrophotometric flux standards appropriate to observe and use for fluxing of your spectroscopic data. Within the reduction software, tables of wavelength versus flux are kept for a large number of spectrophotometric flux standards. To understand and appreciate the details involved in setting up even a single spectrophotometric standard star, see Tüg, White, & Lockwood (1977), Jacoby, Hunter, & Christian (1984), and Massey et al. (1988). In a similar manner to the method by which we took the known arc wavelengths and converted their pixel scale into a wavelength scale, we can now take the known fluxes of the standard stars and convert pixel counts into relative or absolute fluxes. The difference between relative and absolute is essentially the difference between a narrow or large slit width as mentioned above. The conversion of counts to flux is performed under the assumption that slit losses, color terms, transparency, and seeing were similar between the standard star observations and the object(s) of interest.

One can never have too many calibration data and must always trade off time spent collecting CCD frames of standard stars and arcs with collection of data for the objects of interest. Instrument flexure, non-photometric conditions, color terms, and accurate wavelength calibration are crucial to the production of accurate final spectroscopic results such as that shown in Figure 6.5.

6.6 Extended Object Spectroscopy

The definition of extended object spectroscopy follows from the fact that we wish to obtain spectra, not from an unresolved point source, but from a region of the sky for which we desire simultaneous wavelength and spatial information. Examples might include galaxies, nebulae, and planets within our solar system. While there is no fundamental difference between this type of spectroscopy and point source observations such as those described above, their are differences in the instruments used and the reduction techniques involved. We will present here a very basic introduction to the subject and refer the reader to the more detailed review given by Pogge (1992).

One method of obtaining spectra of an extended object is by using long-slit spectroscopy. While sounding like an entirely new method of observing, long-slit spectroscopy is very basic. When we discussed

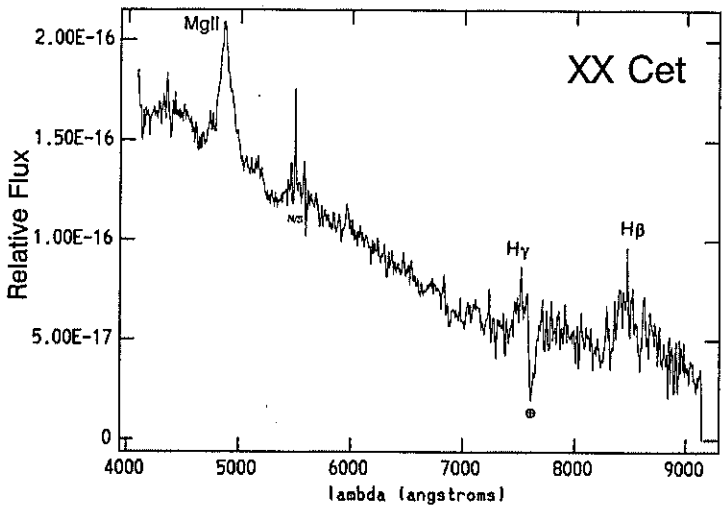

Fig. 6.5. Final reduced CCD spectrum of the quasar XX Cet. This object, long suspected to be a cataclysmic variable, shows redshifted spectral lines (redshift = 0.736), the telluric A band due to the Earth's atmosphere, and false very narrow spectral lines (N/S) due to incomplete night sky emission line subtraction during the reduction process. The flux is in units of ergs s^{-1} cm^{-2} $Å^{-1}$. From Howell & Usher (1993).

point source observations, we concerned ourselves with the details of the spectrograph and final spectrum as they related to the incident light from a point source focused on the spectrograph slit and imaged on the CCD array. We mentioned in that discussion that we desired to use the slit to keep out as much background light as possible. However, all objects that lie along the slit, say in the x direction as presented in Figure 6.6, will produce a spectrum and be imaged on the 2-D CCD array. This last statement makes some assumptions about the ability of the spectrograph optics and optical path to support such an endeavor.

One can imagine a case in which a number of point sources are lined up in an exact east–west manner and a spectroscopic observation is performed of these stars using an east–west slit alignment. The output image will contain a number of parallel spectra, one for each point source, lined up on the CCD. Obvious extensions of this simple example to real cases include placing the long slit of a CCD spectrograph along the axis of a spiral galaxy or alternately across and alongside of bright knots of ejecta within a recently erupted classical novae. The uses of such

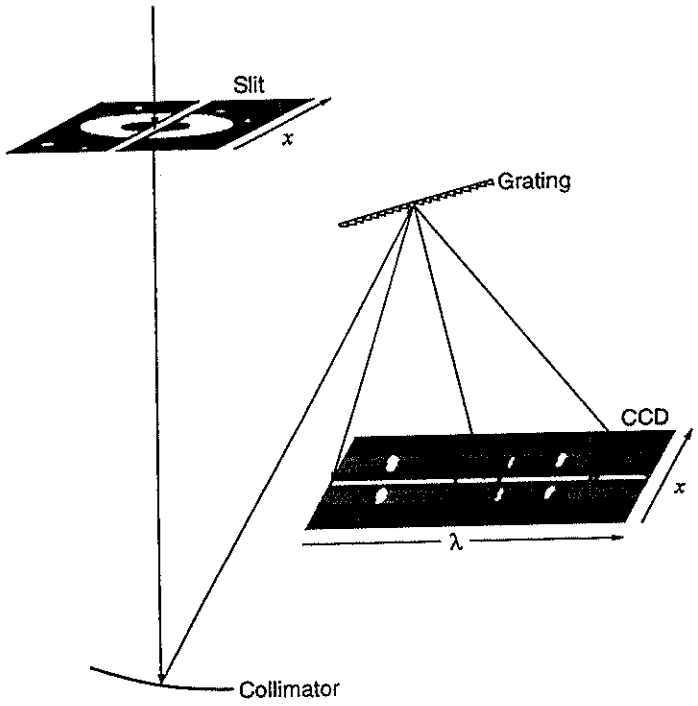

Fig. 6.6. Schematic view of a typical long-slit CCD spectrograph. Positions along the slit are mapped in a one-to-one manner onto the CCD detector. A number of optical elements in the camera, used to re-image and focus the spectrum, have been omitted from this drawing. From Pogge (1992).

spectrographic techniques are many and yield information on both the spectral properties of an object (radial velocity, line widths, etc.) and the spatial properties (spiral arms versus the bulge, etc.). This type of observational program is indeed another that benefits from large-format CCDs.

A more complex and upcoming type of two-dimensional CCD spectrograph is the imaging Fabry–Perot interferometer. Within this type of a spectrograph, 2-D spectral imaging over a narrow bandpass can be collected on the CCD, and with a change in the Fabry–Perot etalon spacing, a new narrow bandpass can be isolated. These types of instruments are often quite complex to master and setup, produce complex data cubes of resultant imagery, and tend to be less efficient than conventional spectrographs owing to their many optical surfaces. However,

to its benefit, Fabry–Perot interferometry has a unique capability to provide tremendous amounts of scientific data per observational data set and to obtain fantastic results for the objects observed. The details of this type of spectrograph and its use can be found in Pogge (1992), Bland & Tully (1989), and Roesler et al. (1982).

A final method of obtaining spatial as well as spectral information is that of narrow-band imaging. The spatial aspect comes about through the use of a CCD as a wide-field imager and the spectroscopic information is provided at selected wavelengths by the use of narrow-band filters. This is no different than observing a given area of the sky in B, V, and R filters, except that here one usually is imaging one or more extended objects and selects carefully two or more restrictive filters in order to obtain specific results. Simultaneous imaging in all colors is lost as one has only limited spectral coverage with this technique, but one obtains full spatial sampling over the entire the field of view and the instrumental setup and data reduction are quite simple. In addition, this type of 2-D spectrophotometry can be performed with any CCD imager with the addition of the appropriate filters.

There are numerous other complexities associated with extended object spectroscopy, some in the observational setup procedures needed and some in the data handling and reduction procedures. However, the wealth of information available from such observations far outweigh the difficulties of these more complex types of CCD spectroscopy.

6.7 Slitless Spectroscopy

The above discussion of astronomical spectroscopy was based on the use of a spectrograph that allowed an image of a slit to be dispersed and then re-imaged onto the CCD detector. Another method, slitless spectroscopy, has also been a workhorse for astrophysical science (Bowen, 1960a) and continues to be useful today (MacConnell, 1995). The idea of slitless spectroscopy is to take advantage of a simple direct imaging system to provide a mechanism for producing and recording astronomical spectra. Some form of dispersive element is placed either before the telescope entrance aperture (e.g., an objective prism) or just before the detector (e.g., a grism), allowing a spectrum to be formed and imaged on the detector.

Work with objective prisms placed in the optical path of a Schmidt telescope and the resulting slitless spectra being imaged on large-format

photographic plates has been used in astronomy for over seventy years. This type of an arrangement allows one to obtain spectra of all imaged objects in the field of view with the integration time setting the limit for the faintest usable spectra (Walker, 1987). The resulting spectra are generally of low dispersion and used for gross classification purposes (MacConnell, 1995), but even radial velocities can be measured (Fehrenbach, 1967). The dispersions available with an objective prism arrangement are near 200–300 Å/mm and the spectra obtained generally cover a total bandpass of 400–1000 Å. The limiting magnitude for a usable spectrum with an objective prism setup is approximately given by

$$m_{\rm lim} = 18.5 + 5\log_{10}T,$$

where T is the focal length of the telescope in meters. Generally, this limit is 8–10 magnitudes brighter than that available at the same integration time with the same optical/detector system when used as a direct imager.

Objective prisms have the disadvantage of being massive and producing nonlinear dispersions. Transmission gratings have replaced objective prisms in some applications although they can introduce coma into the final image. To counter the effect of coma, grating-prism combinations or grisms were developed (Bowen, 1960b). A variant of this idea is the grens or a grating-lens combination. Gratings can be blazed to produce high throughputs in a single order, but the zeroth-order image is always present and leads to field crowding, background light issues, and possible confusion as an emission line in another spectrum (Schmidt, Schneider, & Gunn, 1986).

Current work with slitless spectroscopy of interest here are those applications that use a CCD as a detector. Such examples are the PFUEI (Schmidt, Schneider, & Gunn, 1986; Gunn & Westphal, 1981) and COSMIC (Kells et al., 1998) instruments used at the 200″ Hale telescope and the QUEST instrument (Sabby, Coppi, & Oemler, 1998) used on a 1-m Schmidt telescope. These instruments use a transmission grating, a grism (preceded by a multislit aperture mask), and an objective prism respectively to form their spectra, which are then imaged onto a CCD. The above formula for the limiting magnitude of a useful spectrum should be revised downward by 1–2 magnitudes for a high throughput, high QE, low noise CCD slitless spectroscopy system.

All types of slitless spectroscopy must cope with three similar issues in terms of their image acquisition and data reduction. First, it is often useful to obtain a direct image of the field without the dispersing element in place. This image allows accurate source location independent of intervening spectral overlap or zero-order confusion (in the case of using a grating). Slitless spectra can cover many hundreds of pixels on the CCD with the zeroth-order image being separated by 100–200 pixels. Second, data extraction from the final image can often be complex as field crowding or seeing changes during an exposure will cause broadened, ill-defined spectra. Procedures here generally use a rectangular extraction box containing the spectrum, yet being small enough to avoid large sky contributions. Finally, there is the issue of calibration of the spectrum in flux and wavelength. Wavelength calibration is usually accomplished by one of two methods: Observation of an astronomical source with well-defined, known spectral (usually emission) lines or the use of a calibration lamp shone through pinholes placed at the focal plane. Other methods include use of known spectral lines in one or more of the imaged spectra or the centroid of the zeroth-order image[†] and a knowledge of the plate scale to calculate the dispersion per CCD pixel (Schmidt, Schneider, & Gunn, 1986).

Of the above three concerns, the calibration of the obtained spectra is the most important and the hardest to perform. It is often the case that slitless spectra are presented as unfluxed, having a spectral shape dominated by the instrument and detector response, and often with a wavelength scale that is only approximate. The latter is usually sufficient, as dispersions of 10–50 Å/pixel or more provide wavelength resolutions of only 10 to 200 Å or so. Thus, precise wavelengths are often unimportant. Providing fluxes for the spectra is a difficult and often ignored issue (see below), although not always (Schmidt, Schneider, & Gunn, 1986; Schmidt & Gunn, 1986). Spectra classification, identification of (redshifted) emission lines, or separation into blue and red objects are typical goals, none of which require more than relative flux values.

Two considerations in the determination of the flux from an object imaged via slitless spectroscopy are background subtraction and image flat fielding. The determination of the sky background in a CCD image

[†] The zeroth-order image, while possibly a nuisance, can be used to obtain photometric information such as an instrumental magnitude or, for time-series slitless spectroscopy, a light curve.

of slitless spectra can be confusing. While all the light from the astronomical objects is passed through some sort of dispersing element, so is the light from the background sky. Thus, at any point on the CCD detector, the background is a combination of spectrally dispersed but unresolved sky light. Thus, a "sky" section of the CCD lying adjacent to an object spectrum is not a true representation of the (dispersed) background level imaged with the spectrum. Sky subtraction is not a simple procedure and is often not performed.

Flat fielding a slitless image also presents challenges. We discussed above how the color terms in a flat field image can have large effects on the outcome of the calibration procedure. You can imagine that color terms are even more important in this circumstance, as the flat field light is imaged and dispersed across the CCD in a hopefully similar manner to that of your data. As expected, sky flats provide a far better flat field than dome flats (Schmidt, Schneider, & Gunn, 1986).

Let us examine a few slitless spectroscopy systems in detail. The prime focus universal extragalactic instrument (PFUEI) (Schmidt, Schneider, & Gunn, 1986; Gunn & Westphal, 1981; Schmidt & Gunn, 1986) uses a TI CCD with 15-micron pixels, an image or plate scale of 0.4 arcsec/pixel, and a field of view of 30 arcmin. A transmission grating with 75 lines per mm, blazed at 6,500 Å, is used to provide a spectral dispersion of 35 Å/pixel. The transmission grating can be easily removed from the optical path allowing a direct CCD image to be obtained. When placed in the beam, spectra of each object within the field of view are recorded on the CCD. Figure 6.7 shows a portion of a typical slitless spectroscopic CCD image obtained with the PFUEI.

An innovative setup can combine CCD drift scanning techniques with slitless spectroscopy. The QUEST instrument [Sabby, Coppi, & Oemler (1998) and Section 4.6] is such a camera operating on a 1-m Schmidt telescope and using a CCD array of 16 2,048 × 2,048 Loral CCDs. A 3.2 degree objective prism is employed to provide 401 Å/mm dispersion, yielding a spectral resolution of ~10 Å. A cutoff filter is used to stop light longward of 7,000 Å. This cutoff filter reduces the increased red-ward sky background as well as shortening the length of the imaged spectra, thus decreasing the chance of spectral overlap. Since a prism is used as the dispersing element, no zero-order image concerns exist; however, the spectra do have a nonlinear dispersion.

Figure 6.8 shows a portion of a QUEST drift scan with the prism in place. The spectra are about 400 pixels long and have the dispersion direction aligned with the scan direction to avoid spectral clipping at

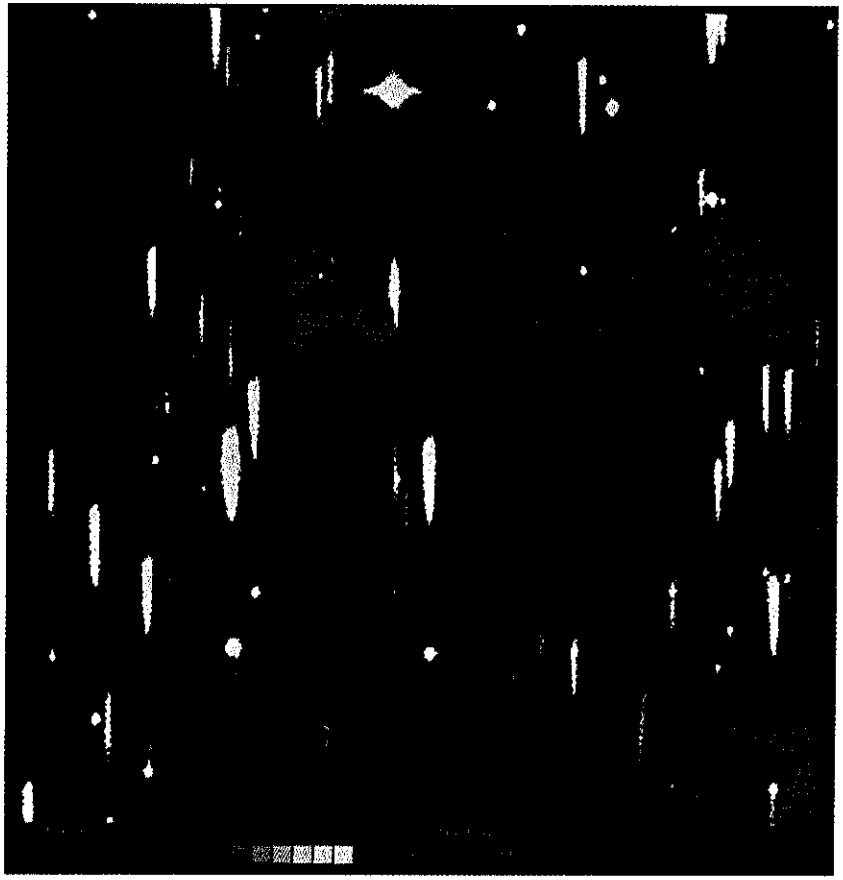

Fig. 6.7. A portion of a PFUEI spectroscopic CCD image. The figure is 5.5 arcmin on a side with east at the top and north to the left. Note the zeroth-order (pointlike) images for each spectrum that appear about one spectral length to their west. The large starlike object near the top middle of the frame is a saturated zeroth-order image from an object whose spectrum is off the top of the image. The zeroth-order images are at the positions of the sources as seen in a direct image and are used in the PFUEI data reduction to set the wavelength scale. The arrow points to the C IV 1550 Å line in a redshifted 19.5 magnitude quasar. From Schmidt & Gunn (1986).

the CCD edges. No attempt is made to flux calibrate the data but initial wavelength calibration is accomplished by identification of known spectral lines detectable in stars.

Slitless spectroscopy is a wonderful tool for obtaining spectra of many objects simultaneously and to very faint limits. The simplistic nature

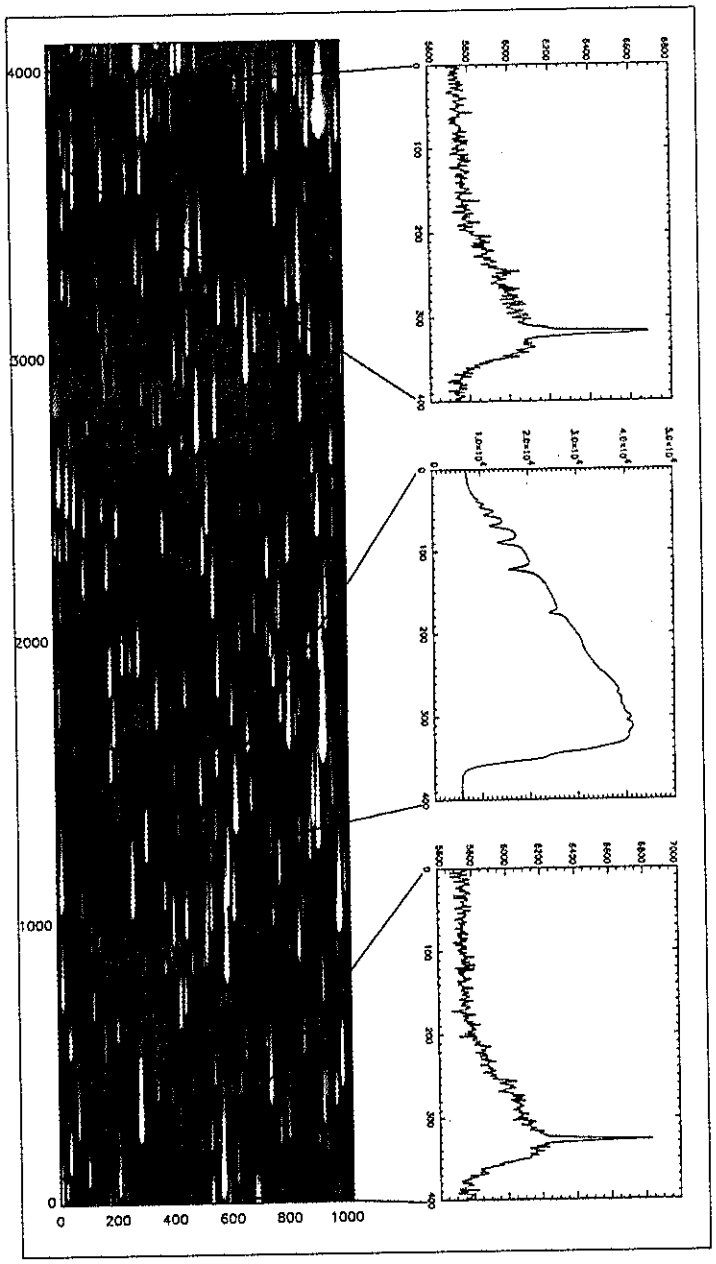

Fig. 6.8. A 1,000 × 4,000 pixel section of a QUEST CCD drift scan obtained with the use of an objective prism. The dispersion is parallel to the scan direction and each spectrum is 200–400 pixels in extent. Examples of some extracted spectra are shown on the right. From Sabby, Coppi, & Oemler (1998).

of the technique makes for an ease of use given that a direct imaging CCD camera is available. The addition of slitless spectroscopy with large-format CCDs and CCD drift scanning go even further, allowing spectroscopic identification of many thousands of objects with little additional effort beyond the direct imaging already in place. Slitless spectroscopy promises to be an evolving CCD application, yet to have reached its peak potential.

7

CCDs Used in Space and at Short Wavelengths

The current high level of understanding of CCDs in terms of their manufacture, inherent characteristics, instrumental capabilities, and data analysis techniques make these devices desirable for use in spacecraft and satellite observatories and at wavelengths other than the optical. Silicon provides at least some response to photons over the large wavelength range from about 1 to 10,000 Å. Figure 7.1 shows this response by presenting the absorption depth of silicon over an expanded wavelength range. Unless aided in some manner, the intrinsic properties of silicon over the UV and EUV spectral range (1,000–3,000 Å) are such that the QE of the device at these wavelengths is typically only a few percent or less. This low QE value is due to the fact that for these very short wavelengths, the absorption depth of silicon is near 30–50 Å, far less than the wavelength of the incident light itself. Thus, the majority of the light (~70%) is reflected with the remaining percentage passing directly through the CCD unhindered.

Observations at wavelengths shorter than about 3,000 Å involve additional complexities not encountered with ground-based optical observations. Access to these short wavelengths can only be obtained via space-based telescopes or high altitude rocket and balloon flights. The latter are of short duration from only a few hours up to possibly hundreds of days and use newly developing high-altitude ultra-long duration balloon flight technologies. Space-based observations in the high energy regime from UV to shorter wavelengths usually require detectors to be "solar blind." The term solar blind means that the detector must be completely insensitive to visible light photons. This is generally accomplished by using a nonoptically active type of detector or through the use of various types of filters. The majority of astronomical objects emit 10^4–10^6 visible light photons for every UV or shorter wavelength photon; thus even

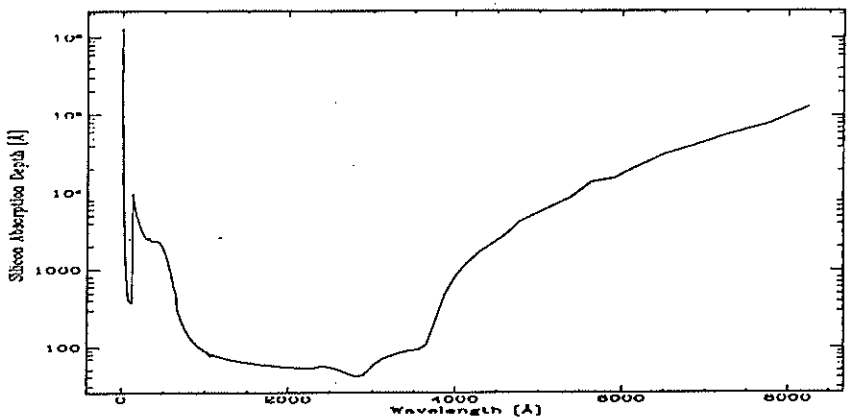

Fig. 7.1. Silicon absorption depth (in Å) from 1.2 to 8,500 Å. The vertical axis is a log scale with major tick marks starting at 100 and ending at 10^6. From Bonanna (1995).

a visible light blocking filter with a 1% UV transmission is not nearly sufficient to remove optical contamination. In addition, most common filters used to block visible light also absorb some of the incident higher energy radiation as well. Use of such absorbing filters cause even a high QE CCD at UV wavelengths (say 20%) to be reduced to a low effective QE near 2%.

Long-term exposure to high vacuum can cause contamination of the dewar in which the CCD resides. This contamination can be through outgassing of various materials such as vacuum grease or AR coatings, normally not a problem in well-produced ground-based systems. Exposure to high energy radiation can cause changes in the QE of the CCD or cause permanent damage to the pixels and electronic structures within the array. Studies of the effects of high energy radiation and space environment observations on CCDs are ongoing at a number of laboratories, the most famous probably being the Jet Propulsion Laboratory (Janesick & Elliott, 1992; Janesick, Hynecek, & Blouke, 1981; Janesick, Elliott, & Pool, 1988; Holtzman, 1990).

Before we discuss the details of observations at wavelengths shorter than the optical, we need to make a brief detour to look into some special issues related to space-based observations with CCDs. The more notable of these issues are the calibration of the CCD throughout the instrument or mission lifetime, the fact that the point spread function is much smaller than generally obtained with ground-based data, and continual degradation of the CCD with time as the result of radiation damage.

7.1 CCDs in Space

Space-based CCDs have a number of special problems associated with them that are often not considered for ground-based systems. Once launched, human intervention is unlikely and the CCD and instrument package can never be retrieved for fault correction or calibration purposes. Even simple procedures, such as bias calibration, take on new meaning as CCD evolution or changes in the gain or other CCD electronics mean new calibration images are needed. Damage to the array (see Section 7.2), or the possibility that the primary circuits fail and the backup electronics or even a different clocking scheme must be used, means that new calibration images must be produced. Also, each observer does not have the ability to obtain all the needed calibration data and the project must provide the finest and most up-to-date calibration images for each CCD, instrument and mode of operation. All issues have to be thought out completely prior to launch or dealt with through analysis of downloaded data during the mission.

One such example of a significant change in CCD operation is provided by the *Hubble* WFPC2 instrument (Holtzman et al., 1995b). After operating in space for only about three months, it was noticed that the CCDs developed an odd sort of CTE effect. The effect caused stars to appear fainter if imaged in higher numbered rows. The apparent cause was the development of a large number of traps within the CCDs not seen during preflight tests. Photometric gradients of 10–15% were present along CCD columns and, even worse, the effect was highly dependent on the brightness of the imaged star, being only about 5% for bright stars.

Using ground-based laboratory tests with similar electronics and CCDs, it was determined that changing the operating temperature from −76°C to −88°C would cause a sharp decrease in the CTE effect. Such a change caused the CTE variations to almost disappear, leaving only a 3–4% gradient. A further temperature decrease would probably improve the situation but in-flight hardware did not allow the CCDs to be operated at colder levels. Thus, considerable effort has been put into the development of a semi-empirical software model that can be applied to data obtained with the WFPC2 in order to correct for the remaining effect (Holtzman et al., 1995a; Whitmore & Heyer, 1998).

As a consequence of the CCD operating temperature being lowered, a decreased dark current was observed. However, on-orbit hot pixel development was greater than expected with many of these hot pixels "fixing"

themselves after dewar warming (Section 7.2). Calibration dark frames are therefore required often to monitor the dark current and to provide the best dark frames to use given any set of observational circumstances. Hot pixels are especially important to understand in space-based CCD imagery as the very small PSF of imaged scenes and the appearance of numerous cosmic rays with a plethora of shapes, including single pixel events, must be distinguished from the collected flux of interest.

We alluded above to the importance of cosmic ray identification in order to avoid misinterpretation of imaged scenes. For a sample of 2,000 second dark images taken with the WFPC2 it was found that 5–10% of the cosmic ray events were single pixel events of 5 sigma or greater above the bias level. Fully one half or more of these events showed consistent pixel positions from frame to frame and thus could not be identified with true cosmic rays or local radioactivity from the dewar and surroundings. Typical signal levels for true single pixel cosmic ray events were near 200 electrons while multiple events peaked near 700 electrons (Holtzman et al., 1995b). Multiple pixel cosmic ray hits (averaging 7 affected pixels per event) are much more common than single pixel events, and a rate of almost 2 events per CCD per second was seen.

CCD dewars, once sealed, evacuated, and chilled, are often seen to produce contaminants owing to outgassing of grease or other coatings used in their construction. When at operating temperatures of −80°C or so, the dewar window is a good site for condensation of such contaminants. These small particles of material are very good absorbers of light, particularly UV radiation, because of their characteristic sizes.

One simple calibration test that allows monitoring of this effect is to obtain fairly regular observations of a bright UV star. If the dewar window does indeed get fogged with material, careful measurements of the UV throughput of the observed flux will show a slow degradation. Even in the best space-based instruments, small amounts of material out gas, and after several weeks UV performance can be noticeably lower. One solution that seems to work, at least for the *Hubble Space Telescope* WFPC2 CCDs, is to warm the dewar to near 20°C for about 6 hours approximately every month.

Flat fields, as we have discussed before, are very important to have in one's calibration toolkit. Once in orbit, either as a satellite or a spacecraft destined for another world, the CCDs aboard generally have little ability to obtain flat field calibration images. High S/N flats made prior to launch in the laboratory are often the best available. These usually

provide overall correction to 5% or a bit better, but small effects, such as illumination or instrument changes, limit the accuracy of the correction. Sometimes, the space-based CCD has significant changes, and large corrections are needed or new flats have to be generated in some manner.

The original WFPC camera aboard *Hubble* could obtain on-orbit flats through observation of the bright earth (Holtzman, 1990; Faber & Westphal, 1991). These were not elegant flats, having streaks and nonuniformities, but were all that was available. WFPC2 used Loral CCDs, which have an increased stability over the original TI CCDs, allowing preflight laboratory flats to work very well, even after the reduction in operating temperature. Numerous other small effects, such as color dependence, radiation damage, hot pixels, CCD illumination, and optical distortions, seen in the on-orbit WFPC2 flats are discussed in detail in Holtzman et al. (1995b). The effects of flat fielding, CTE, and the other issues discussed above on the photometric performance of the *Hubble* WFPC2 are described in Holtzman et al. (1995a), Whitmore & Heyer (1998), and Faber & Westphal (1991).

The *Galileo* spacecraft, currently providing impressive imagery of the planet Jupiter and its satellites, was one of the first public CCD cameras to be launched into space. Its CCD camera is described in detail in Belton et al. (1992) and can be used as an example of the details of space-based observations, their calibrations, properties, and difficulties. CCD and instrument stability and processes for their calibration after launch are major effects to consider as well as proper treatment of the photometric calibration images in lieu of the much reduced PSF.

The solid-state imager (SSI) aboard *Galileo* consisted of a single 800 × 800 TI CCD with a read noise of 40 electrons, gains of 38 to 380 electrons per DN, and a pixel size of 15 microns yielding 2.1 arcsec per pixel. The SSI, like the WFPC2, developed a CTE problem after about 8–12 months in space. Detailed study of SSI images taken during periods of cruise science (Howell & Merline, 1991) revealed that the CTE problem resulted in a readout tail containing 400 electrons, independent of the brightness of an imaged star or its location within the CCD. The cause was attributed to a trap, not in the active CCD array, but in the output register. Radiation damage (see next section) was the most likely cause. Owing to the constant number of trapped electrons, photometric correction was possible to a high degree of accuracy.

Point sources imaged in space are free from the blurring effects of the Earth's atmosphere and have a very small PSF compared with those commonly obtained with ground-based telescopes. A theoretical

diffraction-limited image formed through a circular open aperture will have a FWHM (of the Airy disk) in radians of

$$FWHM = \frac{1.03\lambda}{D},$$

where λ is the wavelength of observation and D is the diameter of the aperture (Born & Wolf, 1959). Note that if we were to use the radius of the first Airy disk dark ring as our definition of image size, we would have the traditional formula

$$r = \frac{1.22\lambda}{D}.$$

Figure 7.2 shows theoretical Airy disk PSFs expected to be imaged by the SSI at three representative wavelengths and five different possible slight defocus values.

The FWHM of the SSI images (being obtained without any atmospheric or other seeing effects) were predicted to be about 0.55 arcsec at 4,000 Å and 1.2 arcsec at 9,000 Å. These PSF sizes correspond to 0.25 and 0.6 pixels respectively, making the SSI images severely undersampled ($r \sim 0.2$). This level of undersampling makes it impossible to directly determine the true FWHM or profile shape of a PSF. Using multiple images with slight offsets, images containing multiple stars with different pixel grid placements, and model CCD images, one can re-construct the true PSF imaged by an undersampled space-based CCD camera. In the SSI case, the PSF was found to be slightly larger than predicted and attributed to a slight camera focus problem.

As we have seen, undersampled images will lead to astrometric and photometric error, as the lack of a well-sampled PSF makes it hard to determine the true image center or the actual flux contained in the image. For the SSI, astrometric error amounted to about 0.8 arcsec even for bright stars, or about half a pixel. Observations of bright guide stars are a common occurrence for spacecraft and are used for navigation and course correction. Large astrometric uncertainties are hazardous and can lead to spacecraft orbital trajectories with inaccurate pointings, having the potential of producing spacecraft course corrections that could cause it to miss a target or, even worse, come too close. In the *Galileo* case, it was determined that a large number of guide star images were needed and careful analysis of these could be used to determine the path and navigation of the spacecraft within acceptable limits.

Photometrically, the nature of the undersampling manifests itself in two ways. First is the way in which one extracts the data and how a flux

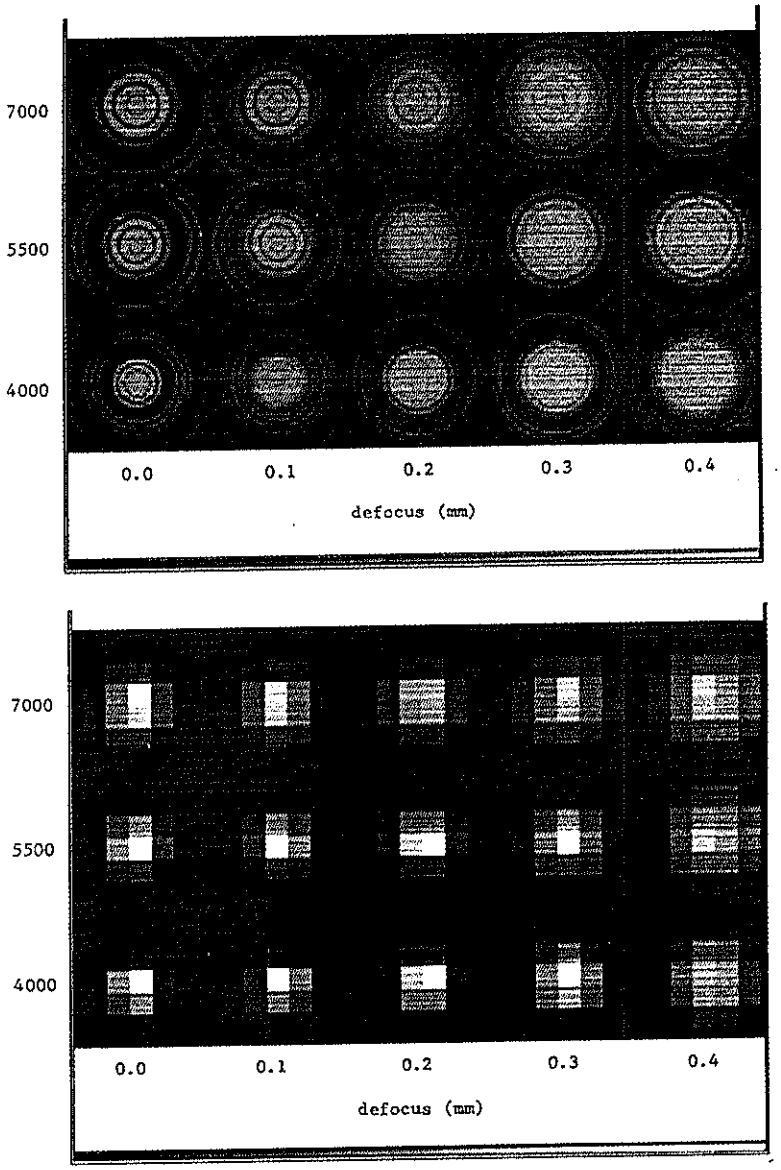

Fig. 7.2. Modeled Airy disk patterns imaged by the *Galileo* SSI. The top panel shows the calculated PSFs as would be seen under very well-sampled conditions while the bottom panel shows the same PSFs as they would appear when imaged by the SSI. The severe pixelization of the PSFs is apparent. The rows are for 7,000, 5,500, and 4,000 Å (top to bottom) and the five columns are (left to right) defocus values for the SSI camera in mm. From Howell & Merline (1991).

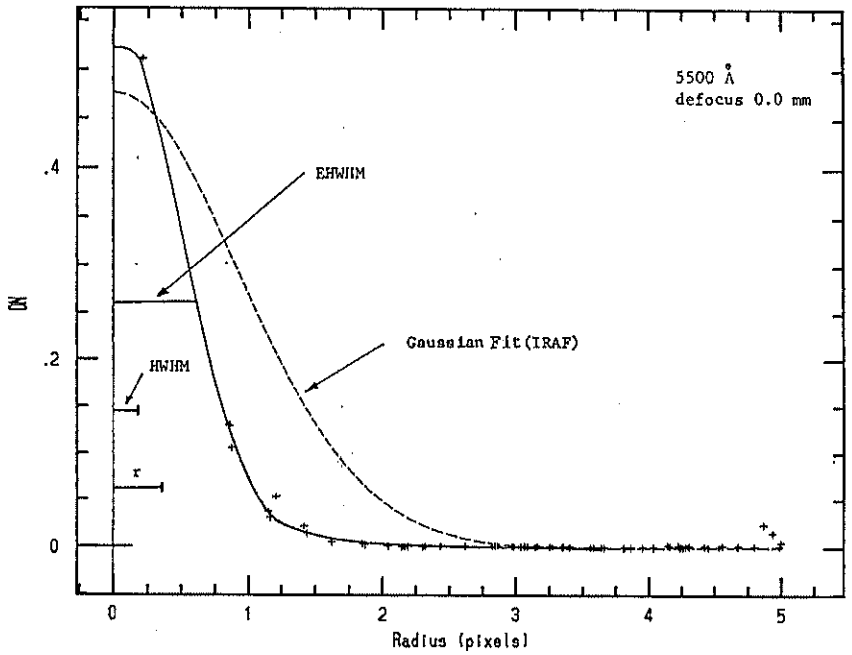

Fig. 7.3. Radial profile plot of a bright star imaged by the *Galileo* SSI. The plus signs are the actual CCD DN values (for $G = 380$ e$^-$/DN) and HWHM and r correspond to predicted values for an Airy disk imaged at 5,500 Å. Note that an approximation as a Gaussian profile is a poor representation of the actual PSF but the determined EHWHM for a single measurement is not far off. From Howell & Merline (1991).

value is assigned to it; second is the effect of digitization noise, which is large for the SSI. Figure 7.3 illustrates the first of these issues by presenting SSI data for a bright star. Because of the nature of the PSFs imaged with the SSI, one pixel (plus sign at $r = 0.25$) contains much more flux than any of the remaining ones. A standard Gaussian fit to these data (in this case made by *IRAF*) is seen to provide a complete misrepresentation of the image profile. Imagine the photometric error one would introduce by assumption of this type of profile and use of its shape as an indication of the total counts observed for this star. The effective FWHM (EFWHM) is defined as the apparent PSF width as determinable from a single undersampled image of a star. We see here that the EHWHM is 0.7 pixels, compared with the expected value (at 5,500 Å) of 0.55. The digitization effect present at the highest SSI gain setting leads to an uncertainty of ±379 electrons per DN. The above

effects combined lead to an overall relative photometric uncertainty of 5–10% and an absolute spectrophotometric uncertainty of 10–30% for SSI data. These are higher than the 2–5% uncertainties quoted for the WFPC2 camera and are directly in proportion to the greater undersampling and higher CCD gain values used in the *Galileo* SSI.

Further readings concerning the special conditions and circumstances of CCDs when used for space-based observations can be found with a quick search of the websites of the *Hubble Space Telescope* and other satellite and spacecraft observatories. Access to numerous internal technical, engineering, and calibration reports is given as well as literature articles containing applications of the findings to astrophysical objects.

7.2 Radiation Damage in CCDs

With the launch of the *Galileo* spacecraft and the *Hubble Space Telescope*, astronomical imagery with CCDs from outer space began. With this exciting new windows on the universe came many unexpected effects in the performance and output noise levels of the CCDs involved. The study of radiation damage in CCDs had occurred in a number of military projects, but the low light and low noise levels needed for astronomy required new laboratory work and the development of techniques to deal with or avoid radiation effects altogether.

The hostile conditions expected in outer space were not the only radiation source to be concerned about for CCDs. Satellites in low Earth orbit, such as the *Hubble Space Telescope*, pass through the South Atlantic Anomaly (SAA) periodically, receiving healthy doses of high energy protons. *Galileo's* radioisotope thermal electric generator (RTG) provides power for the spacecraft as well as a neutron dose that bathes the onboard CCD imager. These inherent radiation environments, along with the general space background of cosmic rays and high energy particles from such events as solar flares or planetary magnetic fields, cause both temporary and permanent damage to a CCD in addition to long-term degradation.

Ironically, as CCDs became better astronomical devices in terms of their low read noise and dark currents, they also became much more susceptible to damage by high energy radiation. The SAA, for example, provides about 2,000 protons per square centimeter per second with energy of 50–100 MeV, for each passage. *Galileo's* RTG was expected to produce 10^{10} neutrons per square centimeter at the location of the CCD over the expected six year mission lifetime. Passage through Jupiter's radiation

belts near the moon Io was predicted to provide a 2,500 rad dose of radiation to the CCD with each orbit. These levels of radiation do indeed cause damage to the CCD involved and methods of monitoring the changes that occur with time and the development of new manufacturing techniques aimed at radiation hardness were needed (McGrath, 1981).

The two major areas of concern in radiation damage to CCDs are (1.) high energy photon interactions, which result in fast electrons, which in turn cause simple, localized damage defects and the generation of numerous electron–hole pairs, and (2.) nuclear reactions caused by uncharged neutrons or high energy protons, which cause large area defects and are more likely to lead to partial or complete failure of a device (Janesick, Elliott, & Pool, 1988). The first of these radiation induced concerns is called an ionization effect and involves gamma rays or charged particles. The second, involving massive particles, is termed a bulk effect or displacement damage owing to its ability to displace silicon atoms from their lattice positions within the CCD.

Displacement damage can involve single silicon atoms or bulk damage involving clusters of atoms, all removed from their original lattice locations within the CCD. The vacancies remaining in the lattice structure create trapping locations, which in turn cause degraded or no CTE performance for one or more pixels in the array. As the result of lattice stresses, the trap locations become populated by one or more of the doping elements such as phosphorus. The presence of a phosphorus atom within the silicon lattice modifies the band gap energies locally and is thought to be the cause of the observed reduced CTE effects (Srour, Hartmann, & Kitazaki, 1986).

Repair of single lattice defects has been accomplished by heating the CCD (to room temperature or higher), while bulk defects are essentially impossible to repair. It has been noticed, however, that at low temperatures ($<-100°$C), the trapped charge accumulated at the defect location remains trapped and has little effect on the overall CTE. This temperature dependence has been shown to be proportional to $\exp(-E_T/kT)$, where E_T is the activation energy of the lattice traps (Janesick, Elliott, & Pool, 1988). Thus, one way to avoid lattice defects is to operate the CCD at temperatures as low as possible. Interestingly, high temperature operation ($>30°$C) allows trapped charge to be released very quickly, eliminating a deferred charge tail and providing good CTE. These various techniques involving temperature manipulation of a CCD system are hard to employ in space. The temporal behavior of the CCD involved can be unpredictable and may be different for each CCD, even those of the same type.

Ionization effects, caused by gamma rays or charged particles, cause a charge buildup in the CCD gate structures and can produce increases in the CCD dark current. Whereas a 150 keV electron is needed to cause an actual silicon atom displacement, only a few eV of energy deposited in the gate insulator is enough to change the potential and cause charge trapping. Even an intense UV flood (Schaeffer et al., 1990; Ditsler, 1990) with 2,500 Å photons can cause dark current increases or even render the CCD inoperable. The charge buildup causes new states to exist within the band gap of the silicon leading to easier generation of thermal electrons and thereby an increased dark current. The affected CCD pixels, that is, those which have been damaged by the ionizing radiation, will show increased dark current, while their neighbors will not. Histograms of the amount of dark current produced as a function of pixel signal level often show "spikes" of dark current at specific signal levels. This is taken to indicate a sort of quantized structure in the amount of damage that occurs per radiation site (Janesick, Elliott, & Pool, 1988).

Methods of protecting CCDs from radiation effects are varied. The type and amount of radiation expected and the scientific goals of the imager must be carefully weighed to produce the final compromise. For example, the CCD flown on the *Galileo* mission was initially tested for its ability to withstand gamma radiation similar to that expected in the Jovian magnetic fields, a test which it passed well. It was probably by shear luck that a test or two was also performed to understand its performance when exposed to neutrons. Increased dark current was prevalent and, without time and money to redesign the imager, operation at a temperature of −130°C was implemented compared with the original proposed value of −40°C. In the case of the *Hubble Space Telescope* CCDs, the increased dark current was not a large factor because of their initially colder operating temperature, but long-term degrading CTE and QE effects were noticed (Holtzman, 1990) and attributed to in-orbit radiation damage. A detailed report of the detectors in the *Hubble* telescope is contained in Brown (1993) and considerations on improving the ability of CCDs to counteract the effects of radiation are discussed in IEEE Nuclear Science Symposium (1988) and Bely, Burrows, & Illingworth (1989).

7.3 CCDs in the UV and EUV (300–3,000 Å) Spectral Range

Progress in the use of CCDs for UV and EUV observations has occurred on two main fronts. Coatings applied to the CCDs to down-convert high energy photons to visible light photons is one method. The other method

involves new manufacturing techniques that allow short wavelength photons to penetrate the silicon even given the very small absorption depths at these wavelengths.

Coating the CCD with a UV phosphor has been discussed previously in this book (Chapter 2) with regard to enhancement of the near-UV wavelength range. These same coatings can often increase the QE of a CCD to usable levels, down to wavelengths as short as 500 Å. Lumogen, a UV phosphor that absorbs photons of $\lambda < 4{,}200$ Å and reemits their energy near 5,200 Å, is a popular choice. This inexpensive material is thermally deposited onto the CCD surface while under vacuum into a layer about 6,000 Å thick. Use of Lumogen as a coating delivers a QE from 500–4,000 Å of around 15% (yielding about 3% with a solar blind filter in place) and actually increases the intrinsic QE from 5,000–8,000 Å, as it acts as an antireflection coating. Phosphor coatings can be deposited onto either front- or back-side illuminated CCDs (McLean, 1997c). Other coatings such as coronene and Metachrome II are also used for UV and EUV enhancement (Trauger, 1990; Geary et al., 1990; Schempp, 1990).

Modern manufacturing processes have again come to the rescue through the development of techniques that allow the CCD itself to have more sensitivity to short wavelength photons. We have mentioned before that the gate structures of a CCD can absorb short wavelength photons before they enter the actual pixel itself, thus reducing or eliminating the collection probability at these wavelengths. Solutions to this problem consisted of using a back-side illuminated, thinned device or making CCDs with transparent gates. Both of these techniques can be further improved and employed for detection of UV and shorter wavelength photons.

During the thinning process, the back side of a CCD forms an oxide layer, which can include a surface layer of incomplete bonds that are positively charged. With the very short absorption depths for UV and shorter wavelength photons, any photoelectrons produced within the silicon are more likely to be attracted to the positive bonds where recombination occurs, and not to the pixel potential well for collection, storage, and output. Various techniques have been developed to reduce the number of positive incomplete bond sites that exist in a thinned CCD (Janesick et al., 1985; Janesick et al., 1989; Baily et al., 1990). Additionally, methods that allow removal of the oxide layer produced in the thinning process have been developed and consist of precision etching of the oxide layer under controlled conditions (Bonanna, 1995).

Recent detailed work, aimed at understanding the characteristics and performance of CCDs in the wavelength range of 300–4,000 Å, has been

performed. The laboratory setup used and various types of short wavelength enhanced CCDs produced and tested are described in Bonanna (1995), which concludes that QE values of 10–60% can be achieved over the wavelength range of 300–2,500 Å. These same CCDs also have up to 80% QE at 6,000 Å. Phosphor coatings and manufacturing improvements are about equal in their ability to enhance UV and EUV performance; however, both types of improvement appear to show a decrease in their QE with time, once the CCD is cooled and put under vacuum. Contamination by outgassing within the vacuum and subsequent freezing of the contaminants onto the CCD surface are thought to be the most likely cause of the reduced QE.

To overcome the low QE of CCDs at short wavelengths, even after the above enhancements have been performed, some high-energy applications employ standard unmodified CCDs as detectors, preceded in the optical path by a device such as a microchannel plate (MCP). MCPs can produce up to about 500,000 electrons per incident high energy photon, and this electron cloud strikes a phosphor coated photocathode producing visible light photons that are collected and imaged by the CCD (Eccles, Sim, & Tritton, 1983). MCPs operate at high voltages (a few keV) and are inherently solar blind as they require high energy photons for activation. Final QE values of up to 20% are possible with a well-constructed device. This increased QE is the largest advantage of instruments that use intensified CCDs, while poor spatial resolution, phosphor decay effects, and smaller dynamic range (compared with a normal CCD) are the major disadvantages (McLean, 1997c; Longair, 1997).

7.4 CCDs in the X-Ray (<500 Å) Spectral Range

Figure 7.1 provided hints that CCDs may also be useful detectors for the X-ray region of the spectrum, as the absorption depth within silicon rises shortward of about 1,000 Å. Figure 7.4 shows us a similar result, only this time we express it as the quantum efficiency of the CCD as a function of photon energy or wavelength. We note that within the X-ray region, back-side thinned CCDs are extremely efficient detectors, approaching a quantum efficiency of 100% at times. Currently planned X-ray telescopes such as *XMM* and *AXAF* are using CCDs as their detectors (Longair, 1997).

X-ray detection by CCDs works in a slightly different manner than detection of optical photons. An incident optical photon creates a photoelectron within the silicon lattice, which moves from the valance to

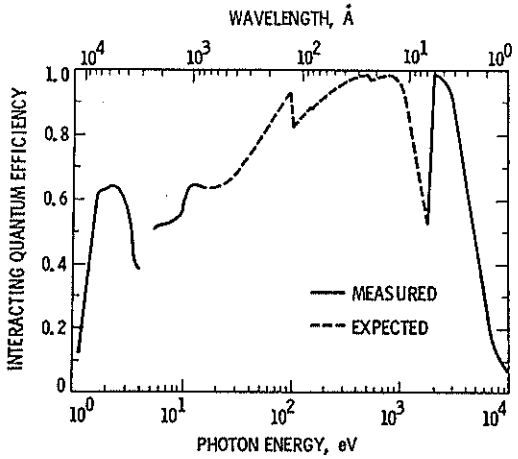

Fig. 7.4. Quantum efficiency for a typical thinned, back-side illuminated CCD from the X-ray to the optical spectral regions. From Janesick et al. (1988).

the conduction band and is then held there (in a pixel) by an applied potential. The absorption of an X-ray photon by silicon ejects a free photoelectron of energy $E - b$, where $E = h\nu$ and b is the binding energy of the electron to the silicon atom, typically 1,780 eV. As this free electron moves through the silicon lattice, it produces a trail of electron–hole (e–h) pairs, each requiring an average of 3.65 eV of energy to be produced. If all the energy of the free electron went into e–h pair production, the energy of the incident X-ray could be precisely knowable simply by counting the ADUs produced within the CCD pixels. However, a small undetermined amount of the free electron's energy goes into various phononic states of the silicon lattice, thereby causing some uncertainty in the value of the incident photon's energy. The level of this uncertainty, the "Fano" factor,† is so small that to obtain Fano-noise limited CCD performance, the CCD read noise must be less than about 2 electrons (Janesick et al., 1988).

There are two very interesting characteristics of using CCDs for X-ray detection. One, just alluded to above, is that X-ray imaging with a CCD provides not only an image but also the ability to produce a spectrum from the image. Each incident X-ray photon collected produces a measureable number of e–h pairs, thus yielding a method by which one

† The term Fano factor is due to U. Fano who, in 1947, formulated a description of the uncertainty in the energy of ion pairs produced in a gas by ionizing radiation.

can backtrack and obtain the incident photon energy (Longair, 1997). X-ray detection can also be used as an excellent diagnostic tool of the CTE of a CCD. Using, for example, a ^{55}Fe source (which produces X-ray photons due to the 5.9 keV Kα Fe line), each absorbed X-ray photon will deposit 1,620 electrons into each pixel. Acquisition and examination of a ^{55}Fe CCD image can help discover problems with the device (Janesick et al., 1988; Janesick et al., 1987b). CCD pixel defects or problems with the associated electronics manifest themselves as pixels in the array with less than 1,620 collected electrons, smeared charge trails showing poor CTE, and other visible diagnostic events.

The use of CCDs as detectors for X-ray imaging and spectroscopy is still in an early phase. This chapter has provided a very brief introduction to the use of CCDs for detection and study of photons with energies higher then those in the optical wavelength region. The literature is just beginning to grow with articles related to high energy uses of CCDs in astronomical observations. For example, the laboratory work involved in the realization of the CCD detectors for the *AXAF* and *XMM* satellite missions has involved the development of new technologies and new analysis procedures. The interested reader is referred to the recent works by McLean, Longair, and others listed in Appendix A.

A
CCD Reading List

This appendix provides a reading list covering the aspects of CCD development, research, and astronomical usage. There are so many articles, books, and journal papers covering the innumerable aspects of information on CCDs that the material presented in a book this size or any size can only cover a small fraction of the details of such work. Even the list presented here does not cover all apsects of interest concerning the use of CCDs in astronomy, but it does provide a very good starting point. The growth of information on CCDs has risen sharply over the past ten years and will, no doubt, continue to do so. Thus, the student of CCD science must constantly try to keep up with the latest developments both in astronomy and within the field of opto-electronics, both areas where progesss is being made.

This appendix lists the reading material by broad catagory, but keep in mind that many works contain information on specific charateristics of certain types of CCD as well as information as to their use and problems. Scientific journal articles are not listed here since they are referenced within the papers presented, are provided throughout the text, and are easily searched for using the journal indexing systems provided in printed and web based forms.

Numerous nontechnical articles related to CCDs have appeared in such magazines as *Astronomy*, *Sky & Telescope*, *Scientific American*, and *CCD Astronomy*.† The reader is cautioned that in some of these descriptions of CCDs, the details are not always completely correct nor have the methodologies always been rigourously tested. However, some of the "popular" articles contain a wealth of information that would be

† *CCD Astronomy* is no longer published but has been incorporated into *Sky & Telescope*.

unlikely to appear elsewhere. We refer the reader to the websites of these magazines for a list of relevant articles.

Every major observatory is a storehouse of knowledge on CCDs. They have websites filled with useful information, large numbers of engineering and technical reports covering aspects of instrument design and construction, manuals for the observatory instruments available, and often newsletters that contain very worthwhile information on their particular CCDs. Many laboratories working on CCDs also have websites that are extremely valuable (see Appendix B).

A.1 General CCD References

Listed below are some general references containing details on CCDs. They are listed by year of publication, and the end of the list contains yearly or periodic references.

Eccles, M. J., Sim, M. E., & Tritton, K. P., 1983, *Low Light Level Detectors in Astronomy*, Cambridge University Press.

Borucki, W. & Young, A., eds., 1984, *Proceedings of the Workshop on Improvements to Photometry*, NASA Conference Publication 2350.

Dereniak, E. & Crowe, D., 1984, *Optical Radiation Detectors*, Wiley.

Hearnshaw, J. B. & Cottrell, P. L., eds., 1986, *Instrumentation and Research Programmes for Small Telescopes*, D. Reidel.

Walker, G., 1987, *Astronomical Observations*, Cambridge University Press.

Borucki, W., ed., 1988, *Second Workshop on Improvements to Photometry*, NASA Conference Publication 10015.

Robinson, L. B., ed., 1988, *Instrumentation for Ground-Based Optical Astronomy*, Springer-Verlag.

Jacoby, G. H., ed., 1990, *CCDs in Astronomy*, ASP Conference Series Vol. 8.

Philip, A. G. D., Hayes, D., & Adelman, A., eds., 1990, *CCDs in Astronomy II: New Methods and Applications of CCD Technology*, L. Davis Press.

Howell, S. B., ed., 1992, *Astronomical CCD Observing and Reduction Techniques*, ASP Conference Series Vol. 23.

Butler, C. J. & Elliot, I., eds., 1993, *Stellar Photometry – Current Techniques and Future Developments*, IAU Colloquium 136, Cambridge University Press.

Kilkeny, D., Lastovica, E., & Menzies, J. W., eds., 1993, *Precision Photometry*, South African Astronomical Observatory Publication.

Sze, S. M., 1994, *Semiconductor Sensors*, Wiley.

Philip, A. G. D., Janes, K., & Upgren, A., eds., 1995, *New Developments in Array Technology and Applications*, IAU Symposium No. 167, Kluwer.

Reike, G., 1996, *The Detection of Light from Ultraviolet to the Submillimeter*, Cambridge University Press.

McLean, I. S., 1997, *Electronic Imaging in Astronomy*, Wiley.

Longair, M., 1997, *High Energy Astrophysics*, Vols. 1 & 2, Cambridge University Press.

Holst, G. C., 1998, *CCD Arrays, Cameras, and Displays*, SPIE Press.

Beletic, J. W. & Amico, P., eds., 1998, *Optical Detectors for Astronomy*, Kluwer.

Martinez, P. & Klotz, A., 1998, *A Practical Guide to CCD Astronomy*, Cambridge University Press.

Modern Technology and Its Influence on Astronomy, Chapters 20 and 21, J. V. Wall & A. Boksenberg, eds., Cambridge University Press.

Photonics Handbook. Yearly editions available from Laurin Publishing Company Inc., Pittsfield, MA.

Optical Engineering. The primary journal of the International Society for Optical Engineers, has numerous special editions devoted to CCDs.

Proceedings of SPIE. The International Society for Optical Engineering. Numerous volumes over the past twenty years or so. Of special interest are the volumes on "Instrumentation in Astronomy," and "Advanced Technology Optical Telescopes."

B

CCD Manufacturers and CCD Website Information

Table B.1 lists present day CCD manufacturers and this appendix lists internet sites of interest to CCD users. In addition to the companies that make CCDs, a number of other internet locations provide detailed information of interest. These are laboratories that study, characterize, and modify CCDs for astronomical use, manufacturers of CCDs and complete systems, and software houses that specialize in CCD codes. The list below is not exhaustive but does provide a good jumping off point. A search of the internet using "CCD" will provide at least an afternoon of reading material and numerous other sites to visit.

Solid State Devices Group at the University of Leicester
http://www.star.le.ac.uk/solidstate/ssdg.html
University of Arizona, Steward Observatory, CCD Lab
http://sauron.as.arizona.edu/ccdlab/
CCD Section of the Electronic Universe Web Project
http://zebu.uoregon.edu/ccd.html
San Diego State University CCD Lab
http://mintaka.sdsu.edu/ccdlab/LabMain.html
University of Hawaii CCD Lab
http://ccd.ifa.hawaii.edu/
MIT CCD X-Ray Group
http://space.mit.edu/ACIS/Welcome.html
European Southern Observatory
http://www.hq.eso.org/odteam/generalccd/generalccd.html
National Optical Astronomy Observatories (KPNO & CTIO)
http://www.noao.edu/

Hubble Space Telescope
http://www.stsci.edu/
Royal Observatory
http://www.roe.ac.uk/
South African Astronomical Observatory
http://www.saao.ac.za/
Anglo-Australian Observatory CCD Imaging manual
http://www.aao.gov.au/local/www/cgt/ccdimguide/manula.html
Joint Astronomy Center in Hawaii
http://www.jach.hawaii.edu
UCO/Lick Detector lab
http://gardiner.ucolick.org/ccdev/ccddist.html
CFHT CCD Selection
http://www.cfht.hawaii.edu/manuals/detectors/CFHT_CCD.html
In Situ Tests for CCDs at the Telescope
http://www.cfht.hawaii.edu/tmca/top/top.html
Photometrics Corp.
http://www.photomet.com/
Scientific Imaging Technologies (SITe): CCD Tutorial
http://www.site-inc.com/tutorial.htm
Royal Greenwich Observatory CCD Page
http://www.ast.cam.ac.uk/apo/docs/ccds.html
CCDs at the La Palma Observatories
http://www.ing.iac.es (see information for users)
Apogee Instruments Inc.
http://www.apogee-ccd.com/
PixelVision Inc.
http://www.pv-inc.com/
Princton Instruments
http://www.prinst.com/
SpectraSource Instruments
http://www.optics.org/spectrasource/
Observatory Techniques Magazine (amateur)
http://home.midco.net/otm/OTM.HTML
Osservatorio Astrofisico di Catania, Istituto di Astronomia
http://www.ct.astro.it/
Osservatorio Astrofisico di Catania, Istituto di Astronomia: CCD Lab
http://www.ct.astro.it/ccdlab/index.html

Table B.1. *CCD detector manufacturers*

Company	Company Location	CCD Name	CCD Types	Internet Web Address	Comments
Dalsa Inc.	Ontario, Canada	Various	Large formats; small pixels	http://www.dalsa.com/	Designs CCD systems; fabricates CCDs via foundaries
Eastman Kodak Co.	Rochester, NY, USA	Kodak	Large formats; small pixels	http://www.kodak.com/	Wide range of CCD types; interline and frame transfer devices
EG&G Reticon	Sunnyvale, CA, USA	EG&G	Various	—	Developed large format arrays with Lick Observatory
English Electric Valve	Essex, England	EEV	Scientific; X-ray	http://eco-web.com/register/00281.html	High performance scientific devices; close ties to astronomy; foundry runs
Loral/Fairchild Imaging Sensors	Newport Beach, CA, USA	Fairchild/Ford	Large formats	—	Close ties with astronomy; foundry runs; Fairchild + Ford Aeronutronics
MIT/Lincoln Labs	Lexington, MA, USA	Lincoln Labs	High speed read out; low noise	http://www.ll.mit.edu/	Federally funded R&D facility, not publically available
Orbit Semiconductor	Sunnyvale, CA, USA	Orbit	Very large formats	http://www.eedesign.com/VendorGuide/Orbit.html	CMOS and CCD foundaries; produce CCD foundary runs
Philips	Eindhoven, Netherlands	Philips	Unique CCDs, new products	http://207.87.19.21/	CCDs produced for internal use
David Sarnoff Research Center	Princton, NJ, USA	Various	Various	http://www.sarnoff.com/	Former RCA Labs; contract work
Scientific Imaging Technologies	Beaverton, OR, USA	SITe	Various	http://www.site-inc.com	Former Tektronix imaging group; foundry service
Texas Instruments	Dallas, TX, USA	TI	Various	http://www.ti.com/sc/docs/msp/ccd.htm	No longer produces scientific grade CCDs, industrial and consumer varieties
Thomson	Cedex, France	Thomson	Various; large formats	http://www.esi.es/About_esi/Members/thomson.html	Scientific & commercial CCDs; close ties with ESO

Table B.1. *(Continued)*

C

Some Basics of Image Displays

Most computer screens and image displays in use are 8-bit devices. This means that the displays can represent data projected on them with $2^8 = 256$ different greyscale levels or data values of resolution. These greyscale levels can represent numeric values from 0 to 255 and it is common to only have about 200 levels actually available to the image display for representing data values with the remaining 50 or so values reserved for graphical overlays, annotation, etc. If displaying in color (actually pseudo-color), then one also has available about 200 separate colors, each with a possible grey value of 0–255, or the famous "16 million possible colors" listed in many computer ads.

On the display, the color black is represented by a value of zero (or in color by a value of zero for each of the three color guns, red (R), green (G), and blue (B)). White has R = G = B = 255, and various grey levels are produced by a combination of R = G = B = N, where N is a value from 0 to 255. Colors are made by having R $\neq$ G $\neq$ B or any combination thereof in which all three color guns are not operated at the same intensity. A pure color, say blue, is made with R = G = 0 and B = 255 and so on. You may have noticed that color printers have three (or four) colors of ink in them. They contain cyan, blue, and magenta (and black) inks, which are used in combination to form all the output colors. This difference (cyan etc. vs. RGB) in the choice of colors is simply because display screens mix light whereas printers mix ink to form specific colors.

Terms one hears but rarely uses in astronomy are hue, saturation, and brightness. Hue means the color of the image, saturation is the relative strength of a certain color (fully saturated = 1), and brightness is the total intensity of a color where black = 0. When you change colors (RGB) you are really changing the hue, saturation, and brightness of

the image display. These three terms are fully explored in Gonzalez & Woods (1993) as well as in almost any text introducing image processing techniques.

Almost all CCD data obtained today have a dynamic range of much greater then 8 bits. Thus, in order to display the CCD image, some form of scaling must be performed to allow the image to be shown on a display with only 8 bits. A common technique (often performed by the software without user intervention) is called linear scaling. This type of scaling divides the entire true data range into say 200 equal bins, where each bin of data is represented by 1 of the 0–200 available greyscale levels. In a pseudo-color representation, the 200 levels are further divided into three bins, one for each color. For example, if an image has real data values in the range from 0 to 100,000 ADUs, linear scaling will place the real data values between 0 and 500 ADU into the first scaled bin and will display them as a 0 on the screen. If your image is such that all the interesting astronomical information have real values of 0 to 2000, this linear scaling scheme will represent all the real image information for the values of 0 to 2000 within only 4 display bins, those having values of 0–4.

To avoid such poor scaling and loss of visual information, two alternatives generally exist: Either one uses a linear scaling but within a specific data window or one uses a different type of scaling altogether. The first option is accomplished by having the software again perform a linear scaling but this time using its 200 output display levels to scale image data values only within the data window of say 0 to 2000. Different scaling options allow for nonlinear modes such as log scaling, exponential scaling, histogram equalization, and many others. These are easily explored in any of the numerous image processing software packages used today.

A method commonly used to aid the eye when viewing a displayed image is that of interactive greyscale manipulation. You probably know this as changing the image stretch or contrast and perform it via movement of a mouse or trackball while displaying an image. The actual change that is occurring is a modification of the relation between the input data values to those output to the display screen. The software mechanism that controls this is called a look-up table or LUT. Some sample LUTs are shown in Figure C.1, where we see the relation of input to output data value. For greyscale images, all three color LUTs are moved in parallel, while in pseudo-color mode, each color LUT can be individually controlled. Changes in the slope and intercept of the LUT

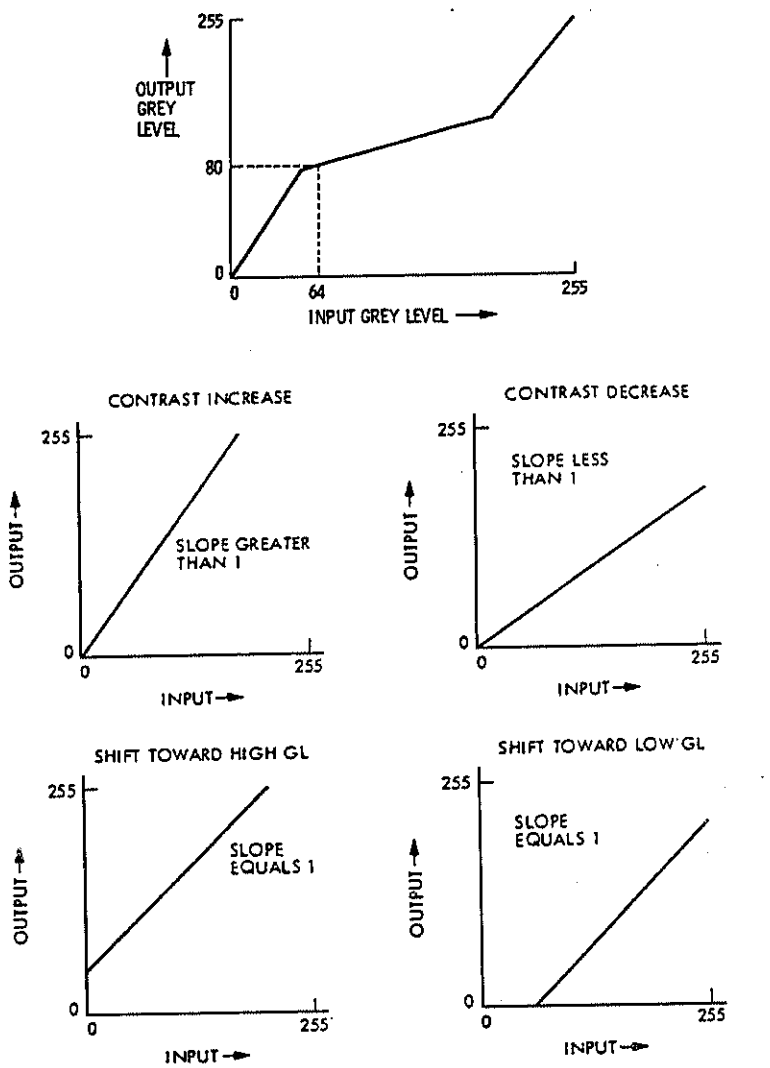

Fig. C.1. The top panel shows an example LUT, which converts input pixel values or grey levels to output grey levels on the video display. In this particular case, an input pixel value of 64 will be displayed with a grey level value of 80. Note that the relationship between input and output is not a single-valued linear transformation. The bottom panel shows various linear stretch operations. If the mapping passes through the origin and has a slope >1, the effect will be to increase the output contrast. A slope of <1 decreases the contrast. If the linear mapping intersects the y axis (i.e., the output value axis) the input values will be systematically shifted to higher output grey levels. Intersecting the horizontal axis causes a net shift to lower output grey levels. From Gonzalez & Woods (1993).

transfer function control changes in the image brightness, contrast, and color. Terms such as "linear stretch" simply refer to a LUT using a linear transformation function.

In the above discussion, we have discussed pseudo-color, but what about true color? True color attempts to display imagery in a manner that is similar to how our eye might perceive a scene. In the visible bandpass, this amounts to trying to match the RGB response of our eye by displaying three separately scaled 8-bit images. Thus, true color displays must have a way to display three separate 8-bit images at once. This is done using a 24-bit display consisting of three individual 8-bit image planes. For true color representation of data from outside the range of the human eye, the user must approximate what the scene might look like in each of the three "colors." The user or display software creates some method of scaling the nonvisible data to the RGB values within the display device itself.

Further information on general image processing techniques can be found in Gonzalez & Woods (1993), while additional details on image displays can be found in Hanisch (1992). User manuals for astronomical image processing software packages such as IRAF and MIDAS often have numerous examples and routines useful for dealing with CCD imagery.

References

Adams, M., Christian, C., Mould, J., Stryker, L., & Tody, D., 1980, *Stellar Magnitudes from Digital Pictures*, Kitt Peak National Observatory publication.

Amelio, G., Tompsett, M., & Smith, G., 1970, *Bell Systems Technical Journal*, **49**, 593.

Baily, P., et al., 1990, *Proc. SPIE*, **1344**, 356.

Baum, W. A., Thomsen, B., & Kreidl, T. J., 1981, in *Solid State Imagers for Astronomy*, eds. J. C. Geary & D. W. Latham, *Proc. SPIE*, **290**, 24.

Belton, M., et al., 1992, *Space Sci. Rev.*, **60**, 413.

Bely, P.-Y., Burrows, C., & Illingworth, G. (eds.), 1989, *The Next Generation Space Telescope*, Space Telescope Science Institute publication.

Bland, J. & Tully, R., 1989, *Astron. J.*, **98**, 723.

Blouke, M., Yang, F., Heidtmann, D., & Janesick, J., 1988 in *Instrumentation for Ground-Based Optical Astronomy*, ed. L. B. Robinson, Springer-Verlag, p. 462.

Bonanna, G., 1995, in *New Developments in Array Technology and Applications*, eds. A. G. D. Philip, K. A. Janes, & A. R. Upgren, Kluwer, p. 39.

Born, M. & Wolf, E., 1959, *Principles of Optics*, MacMillan, Chap. VIII.

Boulade, O., Vigroux, L., Charlot, X., Borgeaud, P., Carton, P., de Kat, J., Rousse, J., & Mellier, Y., 1998, in *Optical Astronomical Instrumentation*, ed. S. D'Odorico, *Proc. SPIE*, **3355**, 614.

Bowen, I. S., 1960a, in *Astronomical Techniques*, ed. W. A. Hiltner, University of Chicago Press, Chap. 2.

Bowen, I. S., 1960b, in *Telescopes*, eds. G. Kuiper & B. Middlehurst, University of Chicago Press, p. 1.

Boyle, W. & Smith, G., 1970, *Bell Systems Technical Journal*, **49**, 587.

Brar, A., 1984, *AURA Engineering Technical Report No. 76.*

Broadfoot, A. L. & Kendall, K., 1968, *Journal of Geophys Res. Space Physics*, **73**, 426.

Brown, R. (ed.), 1993, *The Future of Space Imaging*, Space Telescope Sciene Institute publication, Chap 8.

Buonanno, R. & Iannicola, G., 1989, *Publ. Astron. Soc. Pac.*, **101**, 294.

Burke, B., Reich, R., Savoye, E., & Tonry, J., 1994, *IEEE Trans. Electron Devices*, **41**, 2482.

Chiu, L.-T., 1977, *Astron. J.*, **82**, 842.

Schempp, W. V., 1990, in *CCDs in Astronomy*, ASP Conference Series Vol. 8, ed. G. H. Jacoby, p. 111.
Schmidt, M. & Gunn, J., 1986, *Astrophys. J.*, **310**, 518.
Schmidt, M., Schneider, D., & Gunn, J., 1986, *Astrophys. J.*, **306**, 411.
Smith, M., 1990a, in *CCDs in Astronomy II: New Methods and Applications of CCD Technology*, eds. A. G. D. Philip, D. Hayes, & S. Adelman, L. Davis Press, p. 31.
Smith, R. M., 1990b, in *CCDs in Astronomy*, ASP Conference Series Vol. 8, ed. G. H. Jacoby, p. 153.
Srour, J., Hartmann, R., & Kitazaki, K., 1986, *IEEE Transactions on Nuclear Science*, NS-33, No. 6.
Sterken, C., 1995, in *New Developments in Array Technology and Applications*, eds. A. G. D. Philip, K. A. Janes, & A. R. Upgren, Kluwer, p. 131.
Stetson, P., 1998, *Publ. Astron. Soc. Pac.*, **110**, 1448.
Stetson, P. B., 1987, *Publ. Astron. Soc. Pac.*, **99**, 191.
Stetson, P. B., 1992, *Publ. Astron. Soc. Pac.*, **102**, 932.
Stetson, P. B., Davis, L. E., & Crabtree, D. R., 1990, in *CCDs in Astronomy*, ASP Conference Series Vol. 8, ed. G. H. Jacoby, p. 289.
Stone, R. C., 1989, *Astron. J.*, **97**, 1227.
Stover, R. J., Brown, W., Gilmore, D., & Wei, M., 1995, in *New Developments in Array Technology and Applications*, eds. A. G. D. Philip, K. A. Janes, & A. R. Upgren, Kluwer, p. 19.
Timothy, J. G., 1988, in *Instrumentation for Ground-Based Optical Astronomy*, ed. L. B. Robinson, p. 516.
Tobin, W., 1993, in *Stellar Photometry – Current Techniques and Future Developments*, eds. C. J. Butler & I. Elliott, IAU Colloquium 136, Cambridge University Press, p. 304.
Tonry, J., Burke, B., & Schechter, P., 1997, *Publ. Astron. Soc. Pac.*, **109**, 1154.
Trauger, J. T., 1990, in *CCDs in Astronomy*, ASP Conference Series Vol. 8, ed. G. H. Jacoby, p. 217.
Tüg, H., White, N. M., & Lockwood, G. W., 1977, *Astron. Astrophys.*, **61**, 679.
Tyson, J., 1990, in *CCDs in Astronomy*, ASP Conference Series Vol. 8, ed. G. H. Jacoby, p. 1.
Tyson, J. A. & Seitzer, P., 1988, *Astrophys. J.*, **335**, 552.
Vogt, S. & Penrod, G. D., 1988, in *Instrumentation for Ground-Based Optical Astronomy*, ed. L. B. Robinson, Springer-Verlag, p. 68.
Wagner, R. M., 1992, in *Astronomical CCD Observing and Reduction Techniques*, ASP Conference Series Vol. 23, ed. S. Howell, p. 160.
Walker, A., 1990, in *CCDs in Astronomy*, ASP Conference Series Vol. 8, ed. G. H. Jacoby, p. 319.
Walker, G., 1987, *Astronomical Observations*, Cambridge University Press.
Whitmore, B. & Heyer, I., 1998, *WFPC2 Instrument Report 97-08*, Space Telescope Science Institute Instrument Science Reports.
Wright, J. & Mackay, C., 1981, in *Solid State Imagers for Astronomy*, eds. J. Geary & D. Latham, SPIE.
Young, A. T., 1974, in *Methods of Experimental Physics*, Vol. 12, Part A, ed. N. Carleton, Academic Press, Chap 3.
Zaritsky, D., Shectman, S., & Bredthauer, G., 1996, *Publ. Astron. Soc. Pac.*, **108**, 104.

Index