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

BASIC PHOTOMETRY TECHNIQUES

G. S. Da Costa
Anglo-Australian Observatory, P.O. Box 296, Epping, NSW 2121, AUSTRALIA

ABSTRACT In this contribution basic techniques for producing photometry from CCD frames are described. The topics covered include methods for image center determination, routines for estimating the sky background and the techniques of both aperture and Point-Spread-Function fitting photometry. Procedures for transforming instrumental magnitudes to those on a standard system are also discussed.

1. INTRODUCTION

The advent of high quantum efficiency linear panoramic detectors, such as CCDs, has meant that many areas of Astronomy previously restricted only to those with access to large telescopes, are now also open to those with access to more modest facilities. This is particularly true in those areas of Astronomy that encompass "photometry" in all its forms. For example, in the past, "faint" broadband photometry meant measuring stars fainter than approximately 20th magnitude to calibrate photographic plates. With a single channel photoelectric photometer, this was a difficult task even with a large telescope (e.g. Sandage 1970). With a CCD however, photometry accurate to better than 5% at magnitudes fainter than V=20 is a relatively routine observation on 1m class telescopes (e.g. Sarajedini & Da Costa 1991) while with large telescopes at good sites, photometry accurate to better than 0.2 mag at levels fainter than V=24 is possible (e.g. Pritchet & van den Bergh 1988).

In this contribution the basic techniques for performing photometry on CCD frames are set out in some detail. The frames are assumed to have been "fully processed", i.e., they are assumed to be bias-subtracted, flattened, cosmic-ray cleaned, etc. It is also assumed that the scale of the CCD frame is such that there are at least two pixels per full width at half maximum of an unresolved image, i.e. the profiles of stars are not undersampled, though for many of the techniques described, this is not a fundamental requirement. The next section outlines the basic procedures of aperture photometry, while Point-Spread-Function fitting is discussed in the following section. These techniques are discussed in the context of photometry of stellar images but many of the techniques described, especially those for aperture photometry, can be equally well applied to the photometry of resolved objects such as galaxies. The analysis of CCD frames for surface brightness information however, is a different topic that will not be covered here. Davis *et al.* (1985), for example, can be consulted for information on the techniques involved in that application of CCD photometry.

The photometry techniques described in the first part of this contribution produce magnitudes that are on the "instrumental" system; i.e., that system which results from the combination of the bandpass of the filter and the response of the CCD used in the observations. However, in order to compare the observations with those of others or with the predictions of theory, the instrumental magnitudes must be transformed to a standard system. This process is discussed in the final section.

2. APERTURE PHOTMETRY

In principle, measuring the brightness of a star on a CCD frame should be easy. One merely needs to add up the counts in all the pixels that contain light from the star, estimate the contribution to these pixels from the sky background using nearby pixels, and then subtract the sky contribution to get the net signal from the star. Since the summation of the light from the star is usually done within a circular aperture, this technique is known as *aperture photometry*. However, in practice, things are not this simple. In fact the derivation of a magnitude from a stellar image on a CCD frame falls into three separate tasks: determining the center of the stellar image, estimating the sky background and then calculating the "total" amount of light from the star. Each of these steps is now discussed in more detail.

2.1 Image Center Determination

Accurate determination of the center of a stellar image in a pixel array is of fundamental importance in the field of astrometry, and so not surprisingly, astrometrists have developed a variety of methods to perform this task. However, only two, relatively simple techniques will be discussed here. The first is the *marginal sum method* (Auer & van Altena 1978, Stetson 1979).

In this method the first step is to extract from the CCD frame a subarray centered on an initial guess for the center of the star. The size of the subarray should be large enough so that it contains not only the star of interest but also enough pixels to allow an estimation of the sky background. This means that the subarray should be at least 5 times the full width at half maximum (hereafter FWHM) of the image in size. From this subarray the x and y marginal sums, $\rho(x_i)$ and $\rho(y_j)$ are formed by summing the pixel intensities down the columns and across the rows, respectively. That is, if I_{ij} is the intensity at the pixel ($x = x_i$, $y = y_j$), then the marginal sums are:

$$\rho(x_i) = \sum_j I_{ij} \quad \text{and} \quad \rho(y_j) = \sum_i I_{ij}$$

If a one-dimensional function, such as a Gaussian, is then fit to each of these marginal sums, an estimate of the image center in each coordinate will result. If the star of interest is isolated (and not too faint else the marginal sums are strongly affected by noise) then the center determined by this method will be perfectly adequate for placing the center of the measurement aperture. But if the star has nearby companions that fall within the subarray, then the derived center is likely to be invalid. This occurs because the four parameter (background level, center, width and height) fit is rather easily biased by the presence of nearby companions. Stetson (1979) describes a procedure for handling mild contamination but if the stars of interest are not isolated, then it is prudent to adopt more sophisticated centering algorithms than this simple technique.

More sophisticated centering methods make use of the fact that the star of interest is presumably near the center of the subarray, hence it makes sense to concentrate on those pixels nearest the center of the subarray. The simplest of these methods is called the *image centroiding method.* This method proceeds as follows. Suppose (x_0, y_0) is an initial guess at the image center. The first step is to compute the x and y marginal sums once more, but this time instead of using the whole subarray, the summation is confined to a box $2a \times 2a$ in size. The box size a is an adjustable parameter but best results are obtained if a is taken as approximately the FWHM of the image. The marginal sums are thus:

$$\rho(x_i) = \sum_{j=-a}^{a} I_{ij} \quad \text{and} \quad \rho(y_j) = \sum_{i=-a}^{a} I_{ij}$$

The next step is to compute the mean intensities, $\overline{X}$ and $\overline{Y}$, of each marginal:

$$\overline{X} = \frac{1}{2a+1} \sum_{i=-a}^{a} \rho(x_i) \quad \text{and} \quad \overline{Y} = \frac{1}{2a+1} \sum_{j=-a}^{a} \rho(y_j)$$

The final step is to compute the image centroid *using only those points that lie above the mean intensities.* These are presumably the points that have the most signal from the star. That is, a new estimate of the center of the stellar image is given by:

$$x_1 = \frac{\sum_{i=-a}^{a} (\rho(x_i) - \overline{X}) x_i}{\sum_{i=-a}^{a} (\rho(x_i) - \overline{X})}$$ where the summation is for those i such that $\rho(x_i) \geq \overline{X}$

and

$$y_1 = \frac{\sum_{j=-a}^{a} (\rho(y_j) - \overline{Y}) y_j}{\sum_{j=-a}^{a} (\rho(y_j) - \overline{Y})}$$ where the summation is for those j such that $\rho(y_j) \geq \overline{Y}$.

At this point it is important to check that the new estimate of the center (x_1, y_1) lies within one pixel of the initial guess (x_0, y_0). If this is not the case then it is necessary to repeat the process with (x_1, y_1) becoming the new initial guess (x_0, y_0). In other words, this centering method is an iterative process.

This image centroiding method is computionally quite efficient and provides accurate results provided the stars are neither very crowded (if there is another star in the $2a$ x $2a$ box, the results are likely to be biased) nor too faint (large errors in the image center can occur at low signal-to-noise). In either of these two cases, Point-Spread-Function techniques, which are discussed below, are to be preferred.

2.2 Background Determination

The aim of this section is to endeavor to determine what the signal would be in the aperture if the star of interest wasn't there! As noted by Stetson (1987) for example, the background signal is made up of a number of contributions. First, there are diffuse sources such as the terrestrial night sky emission, the zodiacal light, etc, as well as contributions from scattered light inside the CCD camera and dark signal from the CCD itself. There are also contributions from point sources - other stars and galaxies, whether detected or not, near the star of interest.

To determine this background, the usual procedure is to look at the signal in an annular region centered on the star. Use of an annulus, or indeed any symmetrically placed region, ensures that at least to first order any gradient present in the background should cancel out. Obviously, to avoid biasing the result, the inner radius of the annulus should be sufficiently far from the center of the star that its contribution to the signal in the background annulus is negligible. In practice, this means an inner radius that is at least a factor of several times the FWMH of the image. As regards the size of the annulus, ideally it should be sufficiently large that it contains many hundreds of pixels, thereby ensuring that the uncertainty in the background determined is small, at least in a statistical sense.

In the ideal case of a perfect detector and a completely isolated star, the histogram of pixel intensities for the sky annulus; i.e. a plot of number of pixels with intensity I against I for the pixels in the sky annulus, would have a Gaussian form. In this case the appropriate value to take for the sky background is the mean of the distribution. Further, in this ideal situation, the *mean* is the same as both the *median* and the *mode* of the distribution, with the mode defined as the most frequently occurring pixel intensity (i.e. the intensity at which the histogram peaks). In the real world however, contributions from the wings of bright stars, from faint stars and galaxies, from cosmic rays, etc, all act to add a positive skew to the histogram of pixel intensities. In such a situation the histogram is no longer symmetrical and the mean, median and the mode are no longer equal. Indeed, for a skew in this sense, the mean is most affected by the contaminated pixels and the mode least, with the median lying in between. So it is the mode of the distribution that provides the best estimate of the true background signal. The pixel intensity histogram is usually not actually calculated, instead the mode of the distribution is estimated from the formula (e.g. Kendall & Stuart 1977, p. 40):

$$\text{mode} = 3 \times \text{median} - 2 \times \text{mean}$$

provided that the median is less than the mean. If this is not the case then the contamination is probably small and the mean can be taken as the best estimate of the sky background. It is also usual that any strongly deviant pixel intensities are clipped iteratively before the final value is calculated. This scheme works as follows: first, compute the mean, the median, the mode (via the "3-2" formula) and the standard deviation for the intensities of the pixels in the sky annulus. Next, trim off from the (now sorted) distribution those pixels that lie further from the median than some predetermined multiple of the standard deviation. The multiple of the standard deviation used for the trimming can be a function of the number of pixels left in the distribution (see, for example, the comments in the sky determination code in Stetson's DAOPHOT program) but in most situations, numbers like 2.5 or 3 produce reasonable results. The third step is to recompute the mean, the median, the mode and the standard deviation once more, but this time using the trimmed distribution. If the value of the mode has stabilized then the process is complete; if not, continue iterating until it does.

2.3 Adding Up the Light

With the image center determined and an estimate of the background intensity found, the total signal from the star inside an aperture of radius R is now easily found. Expressed in magnitudes it is:

$$m = zpt - 2.5 \log I$$

where *zpt* is an arbitrary number used to produce reasonable output values for the magnitudes (typical values are 23.5 or 25.0) and

$$I = \sum I_{ij} - n_{pix}\, i_{sky}$$

In this expression n_{pix} is the number of pixels in the aperture, i_{sky} is the background sky value (per pixel) and the summation is carried out over those pixels whose distance from the center of the image is less than R.

For small apertures, or equivalently for large pixels, it is necessary to be concerned with the finite size of the detector pixels since, for example in the case where there are only a dozen or so pixels contributing to the intensity sum, the inclusion or exclusion of individual pixels can substantially affect the resulting magnitude. Correcting for this effect is called allowing for *partial pixels* and there are a number of different schemes that can be employed. Such schemes include the simple expedient of assuming each pixel near the aperture boundary is made up of 4 sub-pixels, each with corresponding intensity 1/4 I_{ij} and including only those sub-pixels within the aperture in the intensity sum, or the following scheme which is that employed in the aperture photometry section of the DAOPHOT program (Stetson 1987). In this scheme, if r_{ij} is the distance of pixel (x_i, y_j) from the center of the star, then if $r_{ij} < R - 0.5$ include all I_{ij} (pixel completely in aperture), or if $r_{ij} > R + 0.5$ exclude all I_{ij} (pixel completely outside). For those pixels on the boundary, i.e. those with $R - 0.5 < r_{ij} < R + 0.5$, include a fraction $[(R + 0.5) - r_{ij}]$ of I_{ij} in the intensity sum. This corresponds to linearly weighting the pixel intensity as it changes from being completely inside the aperture to completely outside it. In any case, regardless of whether a partial pixel scheme is employed or not, it is important to remember that a circular aperture is in practice being approximated by an irregular polygon.

2.4 Choice of Aperture

The final question to be considered is what to choose for R, the size of the aperture with which to measure the brightness of the star. The seemingly obvious answer to this question is "the radius that contains all the light from the star" but in fact the answer is not this simple. For example, the profile of stellar image is much larger than one would naively think (see, for example, King 1971) so that an aperture that actually includes "all" the light would have to be *very* large. But for faint stars such a large aperture is definitely not desirable, instead, as is now demonstrated, the choice for the "optimum" aperture with which to measure a particular star depends on its brightness.

That this is the case can be seen from a consideration of the errors involved in aperture photometry (for simplicity only photon statistics and readout noise are considered here; for a more detailed discussion see Newberry 1991). Assuming that an aperture of radius *R* contains a total intensity I_{st} from a star, and that the CCD system gain is D electrons per ADU, then the number of photoelectrons generated is just $I_{st}D$. Photons follow Poisson statistics so that the standard deviation associated with signal $I_{st}D$ photons is just $(I_{st}D)^{1/2}$ photons or $(I_{st}/D)^{1/2}$ in ADU. The magnitude error is then:

$$m \pm \delta m_1 = zpt - 2.5 \log (I_{st} \pm \sigma(I_{st})) \quad \text{or}$$

$$\delta m_1 = 1.09 \frac{\sigma(I_{st})}{I_{st}} = \frac{1.09}{I_{st}D} (I_{st}D + r^2 n_{pix})^{1/2}$$

if we also include the readout noise r (in electrons). As before n_{pix} is the number of pixels included in the aperture. This expression demonstrates that, provided readout noise does not dominate, the error in the magnitude decreases with increasing aperture size, since I_{st} increases with increasing aperture radius. However, this is not true indefinitely because once the aperture includes "almost all the light", δm_1 remains constant or actually increases due to the readout noise contribution of the extra pixels.

Now there is a second source of noise that arises as follows: the sky background i_{sky} is presumably well determined (at least in a statistical sense) since it is derived from a large number of pixels. But the actual sky signal in the aperture is subject to Poisson statistics in the same way as the object intensity. It is straightforward to show that the error in the stellar magnitude from this source is just:

$$\delta m_2 = \frac{1.09}{I_{st}} \left(\frac{n_{pix} i_{sky}}{D}\right)^{1/2}$$

Now n_{pix} is proportional to the square of the aperture radius, so it is evident that δm_2 increases with increasing aperture size and this error is clearly *minimized for small apertures*. The combined error is then:

$$\delta m = \frac{1.09}{I_{st}D} (I_{st}D + n_{pix}(r^2 + i_{sky}D))^{1/2}$$

Inspection of this equation indicates that there are two limiting cases:

a) "bright" stars for which $I_{st}D >> n_{pix}(r^2 + i_{sky}D)$. In this situation large apertures that contain "all" the light are acceptable.

b) "faint" stars for which $I_{st}D$ does not dominate $n_{pix}(r^2 + i_{sky}D)$. In this situation, small apertures, *which will not contain all the light*, are to be preferred.

In this context, it is also worth noting that I_{st} and i_{sky} are both proportional to the integration time while r^2 is fixed. Consequently, a single long integration is preferable to a sum of shorter ones (cosmic ray frequency willing) since it minimizes the effect of readout noise. Further, when $i_{sky}D > r^2$, the photometry is said to be "*sky noise dominated*" but "*readout noise dominated*" when the reverse is true. With broadband filters and modern CCDs only rarely are data readout noise limited, but this is often not the case for narrow band imaging.

This difference in the choice of optimum aperture size depending on the brightness of the star, leads naturally to the concept of aperture corrections.

2.5 Aperture Corrections

The subject of aperture corrections is discussed in some detail in Howell (1989) and in Stetson (1990). As a result, the subject need be covered only briefly here. The fundamental principle that underlies the concept is that, with a linear detector such as a CCD and in the absence of distortion introduced by the camera optics, the *Point-Spread-Function*, which can be defined as the two dimensional intensity distribution produced at the detector by the image of an unresolved point source such as a star, *is the same for all stars regardless of their brightness and position on the frame*. That is, the images of bright and faint stars differ only by a scaling ratio. Consequently, *the fraction of the total in an aperture of a given size is the same regardless of the brightness of the star*. Clearly then, the best way to proceed is to measure faint stars through a small aperture, thereby minimizing the errors, and to measure bright stars through both small and large apertures. The difference between the small and large aperture measures for the bright stars, called the *aperture correction*, can then be subtracted (smaller magnitudes are brighter!) from the small aperture measures for the faint stars to produce an effective large aperture measurement.

This leaves one remaining question to be addressed: how large should the "large" measurement aperture be? The answer to this question is that the aperture has to be large enough that measurements made with it are independent of changes in the seeing, telescope focus, etc, that can occur during a night's observing (see, for example, Massey *et al.* 1989). This means that the aperture should have a radius as large as 4 to 5 times the FWHM of the images. In practice an appropriate approach is to measure the star through a number of apertures and plot a *growth curve*, i.e. a plot of the measured magnitude as a function of aperture size. It is then relatively straightforward to define the asymptotic limit of such a curve, which can then be taken as the "total" (or at least seeing independent) magnitude of the star. For this procedure to work, the star has to be isolated but for moderately crowded stars, a "total" magnitude can be derived by fitting the inner part of the growth curve to that for an isolated star with similar FWHM. This procedure is discussed in detail in Stetson (1990). In Figure 1 (normalized) growth curves for stars observed in different seeing conditions are shown. These curves illustrate the points discussed above - if seeing independent magnitudes accurate to the 0.01 mag level are desired, then relatively large apertures are needed.

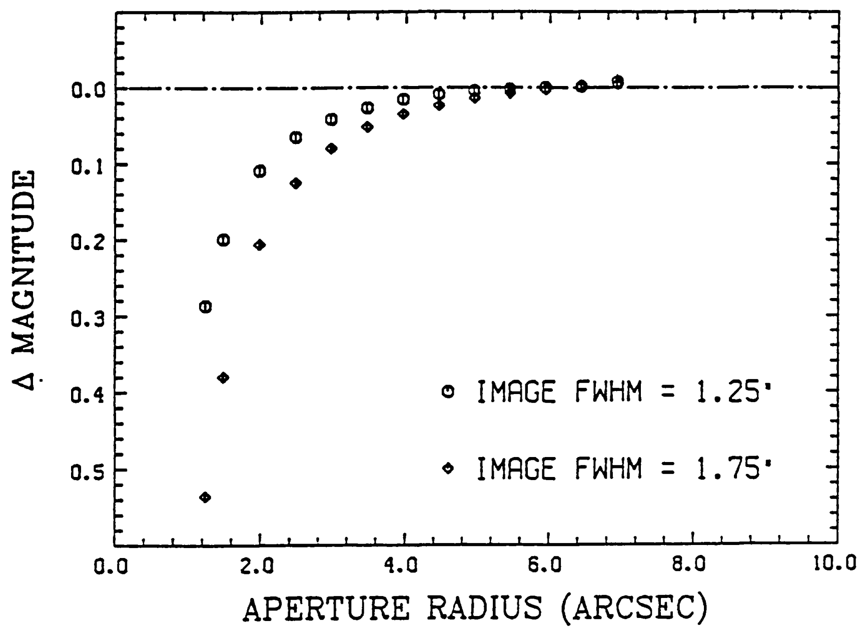

Fig. 1. Growth curves for stellar images in different seeing conditions, normalized to the same asymptotic magnitude. Δ magnitude is the difference between the aperture magnitude and the asymptotic value indicated by the dot-dash line.

3. POINT SPREAD FUNCTION PHOTOMETRY

Aperture photometry as discussed in the previous section is adequate if the stars of interest are neither too faint nor too crowded. However, it goes without saying that much astronomy of interest is contained in situations where these criteria are not met. In these cases, Point-Spread-Function (PSF) photometry is a more appropriate technique. The basic principle of the technique is that all images on a CCD frame have, again baring distortions introduced by the camera optics, the same form and thus differ from one another only by a scaling ratio. Then fitting a suitably defined PSF to a series of images will give relative magnitudes via the equation:

$$m = zpt - 2.5 \log (\text{scaling ratio})$$

in which *zpt* is the magnitude assigned to the PSF.

Point-Spread-Function fitting programs can be grouped into two distinct classes. The first, typified by Stetson's famous DAOPHOT program (Stetson 1987), uses the data itself to define an *empirical* PSF, while the second uses the data to fix the parameters defining a *model* PSF. This latter approach is used in, for example, Penny's STARMAN program (Penny & Dickens 1986) and in the DoPHOT code (Mateo & Schecter 1989). There are advantages to both approaches. The first makes the least assumptions about the shape of the PSF but is only defined on the pixel grid. This can cause difficulties when interpolating the PSF to a different center, especially if the images are nearly critically sampled (i.e.

FWHM ~ 2 pixels). For this reason most empirical PSF codes store the PSF as the sum of an analytic function (e.g. a Gaussian) and a look-up table of residuals from the analytic fit. This reduces significantly the interpolation uncertainties. The second approach has the advantage that the necessary integrations and interpolations are more easily carried out, but of course has the disadvantage that the chosen model profile may not be a good representation of the actual images. The analytic functions used in these codes to model the profile of a stellar image are usually either a Gaussian G(r), or a modified Lorentzian L(r), or a Moffat function M(r), or occasionally the sum of a Gaussian and a Lorentzian. In their simplest form these functions are (see Stetson *et al.*, 1990, for a more complete discussion):

$$\text{Gaussian:} \quad G(r) \propto e^{-r^2/2\alpha^2}$$

$$\text{Modified Lorentzian:} \quad L(r) \propto \frac{1}{1+(r^2/\alpha^2)^\beta}$$

$$\text{Moffat:} \quad M(r) \propto \frac{1}{(1+r^2/\alpha^2)^\beta}$$

Both types of code work by fitting the PSF to data for an image inside a fixed aperture, the so-called *fitting radius*, and maximizing some "goodness-of-fit" criterion. The process is easily extended to fitting many images simultaneously with the fit to N stars requiring the determination of $3N$ parameters - the center (x_c, y_c) and the scaling ratio for each image. Obviously this ability to fit many images simultaneously allows photometry in much more crowded situations than is the case for aperture photometry.

One final point that must not be overlooked is the requirement to bring the relative magnitudes produced by the PSF-fitting code onto a proper instrumental system. This is achieved by determining total magnitudes via aperture photometry for the brightest least crowded stars on the frame. The (mean) difference between the magnitude from PSF-fitting and the total magnitude for these stars then defines the correction to be applied to all the PSF-fitting derived magnitudes to place them on the instrumental system. In this procedure it is often necessary to first subtract out neighboring stars from the vicinity of those used to define the PSF-to-total correction. Fortunately, the very nature of the PSF-fitting process makes this easy to do. However, it is usually the case that this PSF-to-total magnitude correction is the largest source of systematic uncertainty in the whole photometry process.

4. TRANSFORMATION TO A STANDARD SYSTEM

While there are some applications of CCD photometry, such as time series analysis of variable stars, that yield results from the instrumental magnitudes themselves, in most situations it is necessary to transform the instrumental magnitudes onto a standard system. Standard systems can be classified as *broadband* (filter widths ~ 1000 Å), *intermediate-band* (filter widths ~ few hundred Å) or *narrow-band* (filter widths ~ few tens of Å). Examples of the first type are the widely used UBVRI system or the Washington CMT_1T_2 system, while the Stromgren uvby, the Thuan-

Gunn uvgriz, and the DDO system are all examples of intermediate-band systems. Narrow-band filters are usually used for emission line (e.g. Hα or OIII) imaging and flux determination. In all cases the standard systems are defined by sets of *standard stars*, stars whose magnitudes and colors in the standard system are known.

In principle the transformation of the observed instrumental magnitudes to a standard system should be straightforward - one merely needs observations of a few standard stars which can then be used to define a transformation to the standard system. In practice however, the transformation process contains a number of subtleties which must be kept in mind if precise answers are required. Indeed it is a good idea to remember the following principle. Instrumental magnitudes from a CCD frame can easily have internal precisions that are better than 0.01 mag, but this fine internal precision will be wasted if the transformation to the standard system introduces uncertainties at the 0.05 mag level! The moral of this section then is, if you are interested in producing photometry on a standard system (e.g. to compare your photometry with other data or with theoretical predictions) then *observing standards is just as important as observing program objects*. It is a waste of telescope time to compromise otherwise excellent data by failing to acquire sufficient standard star observations!

For the same reasons it is important that the CCD/filter combination used, which defines the instrument response, provide a good approximation to the bandpasses of the standard system. While non-standard bandpasses (e.g. wider or narrower filters, different central wavelengths, etc) may appear to produce satisfactory transformations, the standard stars used to derive the transformations are almost inevitably solar abundance dwarfs. But the spectral energy distributions of the program objects, e.g. metal-poor giants, may be sufficiently different from that of the standards that when combined with the non-standard bandpassess, serious (and most likely undetected!) systematic differences from true standard system magnitudes will result.

4.1 An Example: Transformation to the BVRI System

The (U)BVRI system is the most commonly used broadband photometric system in Astronomy. The UBV system originated with Johnson in the 1950s (Johnson & Morgan 1951, Johnson 1955) and is based on a set of glass filters and a 1P21 photomultiplier. The system is inherently not a well defined system; for example the long wavelength cutoff of the V bandpass was set by the falloff in the response of the 1P21 photocathode and not by the filter. This creates an immediate difficulty in reproducing the standard V bandpass for a red sensitive detector. Similarly, the Cousins R, I system (Cousins 1976,1978), which incidentally is not the same as the less well setup Johnson R, I system, is also not completely filter defined since the red cutoff of the I bandpass is set by the response falloff of a cooled GaAs photocathode.

The actual bandpasses of the BVRI system are tabulated in Bessell (1990) and are illustrated graphically in Figure 2. It is important to remember that because this is a broadband system, the effective wavelength of each bandpass depends sensitively on the spectral energy distribution of the object under study. For example, Bessell (1990) gives the following information:

	λ_{eff} (Å)			
Filter	B	V	R	I
A0 V	4363	5448	6407	7982
K0 III	4520	5524	6535	8028

The largest changes in effective wavelength occur for the B band, because of the large difference in spectral energy distribution at these wavelengths, and for the R band because the long red tail in the R band response.

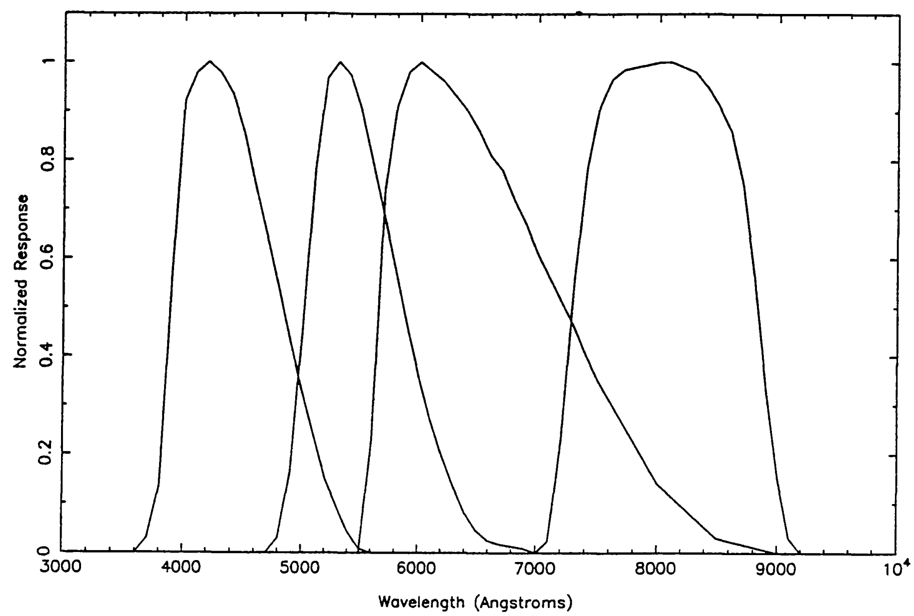

Fig. 2. The bandpasses of the BVRI system from Bessell (1990).

4.2 Practicalities

Assuming that the prescriptions of Bessell (1990) or others have been followed in the choice of filters, what then are the practical requirements for determining a good transformation from the instrumental to the standard system? The first and most obvious one is that it be a photometric night! It is hardly worth saying that there is no point in observing standard stars if sky conditions are not cloud-free. This is particularly relevant if the observing program is "all sky photometry", i.e. the standard stars and the program objects are scattered all over the sky. The second requirement is that the standard stars observed be those with accurate standard system magnitudes. This means in practice choosing stars from the lists such as those of Landolt (1973, 1983a,b), Graham (1982) and Menzies *et al.* (1989, 1991). While this frequently means that at most only a few standards can be observed per CCD frame, this should not be considered a disadvantage! Instead it means that the instrumental magnitude can be well determined and that the standard magnitude itself is likely to have been accurately measured. Indeed, it is perhaps worth emphasizing that for those situations where accurate photometry on a standard system is required, there is little to be gained from observing so-called "standards" in relatively crowded fields such as those in globular clusters. Rarely are the standard magnitudes of stars in these fields known to be accurately on the standard system, or to be precise at the 0.01 mag level. Thirdly, regarding such questions as the number of standards to be observed each night, how frequently they should be observed, etc, the simplest answer is to let common sense be your guide. Two standards observed at evening twilight are not likely to be enough to calibrate an entire night's observing! Reasonable requirements are the following: (a) Observe perhaps 15 to 20 standard star fields throughout the night (not just at the beginning and end!) with some reasonable fraction observed more than once. This will guarantee enough observations to check for things like changes in extinction during the night. (b) Observe standards with a wide range of colors, preferably encompassing the likely color range of the program objects. This will ensure that the color dependence of the transformations is well determined. (c) Make sure that the range of zenith distances (or equivalently the range of airmasses) over which the standards are observed exceeds the range in the program object observations. This will minimize the possibility of errors resulting from uncertainty in the atmospheric extinction.

If these (or similar) requirements are fulfilled, then a transformation to the standard system can be derived (see Massey *et al.* 1989 or Da Costa 1990 for a description) by writing equations (using V as an example, similar equations can be written for B, R and I) of the following form:

$$V_{std} = v_{inst} + a_0 + a_1(B-V)_{std} + a_2X + a_3X(B-V)_{std} + a_4UT + \ldots \quad \text{(other terms as required).}$$

Here v_{inst} is the instrumental magnitude (corrected for exposure time), X is the airmass for the observation (for zenith distances $z_d < 60$ degrees, $X = \sec(z_d)$) and UT is the time of the observation. V_{std} and $(B-V)_{std}$ are the standard system magnitude and color of the star. Then with a series of standard star observations it is possible to use a code (e.g. Stetson's CCDCAL program or the GAUSSFIT package (Jefferys *et al.* 1988)) to evaluate the coefficients a_i and their significance. The significances attached to the determined coefficients should not be ignored: with these types of programs there is always the danger of overdetermining the problem through having too many coefficients and too few standard star observations to constrain them. Use only the minimum number of coefficients required for a satisfactory fit! The terms are:

a_0 the *zeropoint* - a measure of the sensitivity of the CCD/telescope system in this bandpass;

a_1 the *color-term* - a measure of how well the instrumental system approximates the standard system. In most circumstances a_1 should be less than 0.1;

a_2 the *first order extinction coefficient* - this depends on the atmospheric transmission and therefore varies from site to site and often from night to night.

a_3 the *second order extinction coefficient* - this term is necessary, particularly for B, because extinction varies rapidly across the bandpass so that combined with the change in effective wavelength with spectral type, a variation in effective extinction with spectral type results.

Other terms such as time or color-squared dependences are not usually required but the residuals from the fit should always be investigated to confirm the lack of correlation with time, airmass, color, magnitude, etc. In particular, a color-squared dependence would only occur if the instrumental system is far from the standard system.

It is reasonable to expect that the color-terms should remain constant during a single run and perhaps from run to run if the same CCD and filters are used. The zeropoints should also be stable within a run since there is no reason to expect the overall sensitivity of the CCD/telescope system to change on a nightly basis. On the other hand it is quite likely that the extinction will change from night to night as atmospheric conditions vary, emphasizing the need for observations of standard stars over a range of airmasses each night, in order to determine the extinction correctly. Note that this approach of an assumed stable instrument and nightly varying extinction is opposite that often adopted for photoelectric photometry reductions, where mean extinction coefficients for a site are used and the zeropoint is allowed to fluctuate nightly. Given the inherent stability of CCDs the former approach seems more reasonable. Figure 3 shows typical transformation equation residuals as a function of time and V-I color for a set of standard stars observed with an RCA CCD during a photometric night on the 0.9m telescope at CTIO.

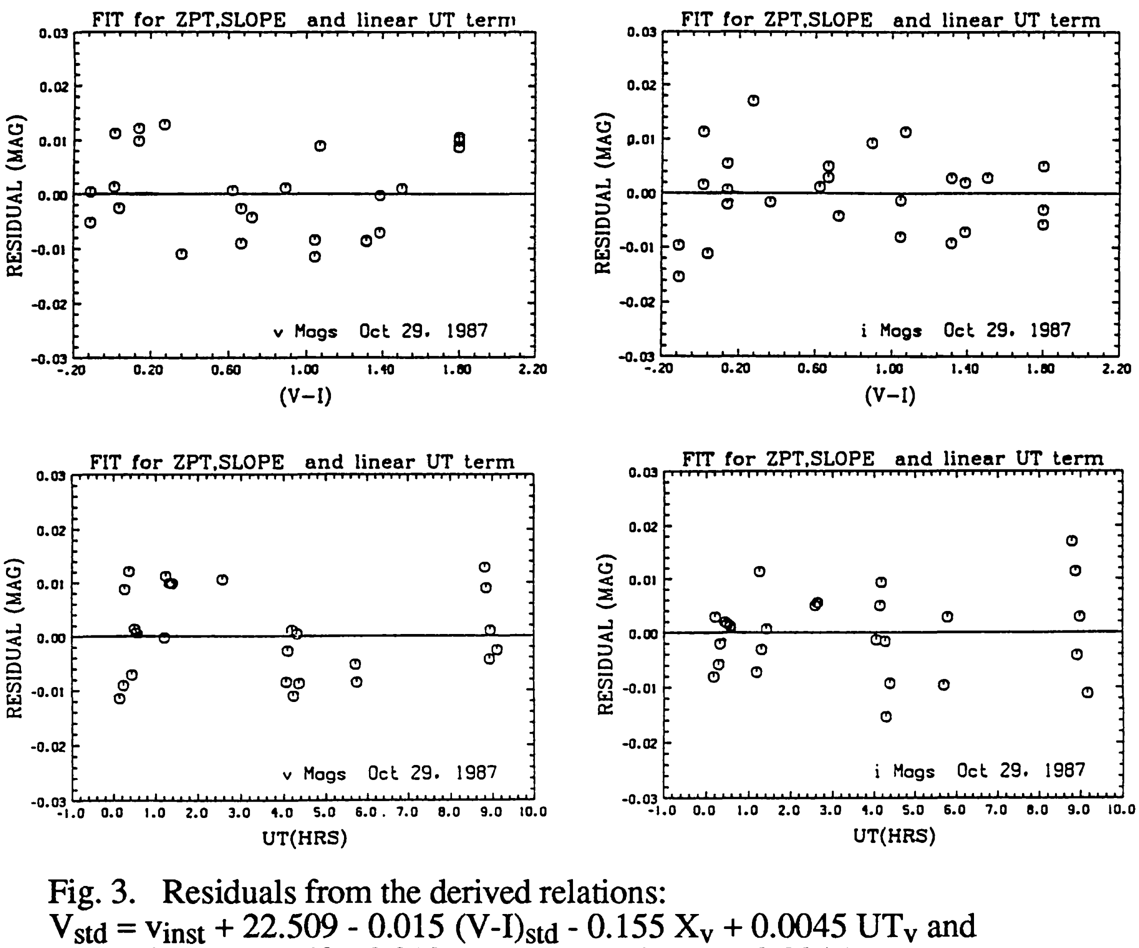

Fig. 3. Residuals from the derived relations:

$V_{std} = v_{inst} + 22.509 - 0.015\ (V\text{-}I)_{std} - 0.155\ X_v + 0.0045\ UT_v$ and

$I_{std} = i_{inst} + 21.248 - 0.018\ (V\text{-}I)_{std} - 0.063\ X_i + 0.0045\ UT_i$

are plotted as a function of standard V-I color and UT (expressed in decimal hours) for the standard star observations on a photometric night.

Note that for this particular night, the transformation equations required a linear UT term, presumably as a result of the mean extinction changing uniformly during the night.

4.3 An Example

To illustrate that the procedures described in this contribution can actually produce accurate results, it concludes with a comparison of two completely independent sets of CCD photometry of stars in the globular cluster Pal 12. The first set is that of Stetson *et al.* (1989). Their data consist of B and V observations mostly obtained at the CHFT, using 2 different RCA CCDs, but with some additional data from the CTIO 4m taken with a third RCA CCD. The observations were calibrated primarily with standards from Landolt (1973) but also with standard stars from a field in M92 (Stetson & Harris 1988), from Landolt (1983a) and from Graham (1982). The frames were analyzed using the DAOPHOT PSF-fitting code. The second set of data is that from Da Costa & Armandroff (1990). Here V and I observations were obtained at CTIO using 2 different RCA CCDs on the 0.9m telescope. Standards from Landolt (1983a,b) and from Graham (1982) were used to calibrate the data and the Pal 12 stars were measured via the aperture photometry techniques described above (i.e. measurement with a small aperture followed by application of an aperture correction).

There are 17 stars, with $14.5 < V < 18.0$, in common between the two studies. For these stars, the mean difference in the V magnitudes is:

$$V(\text{DaC \& A}) - V(\text{PBS } \textit{et al.}) = -0.003 \text{ mag}$$

and the standard deviation of the differences is:

$$\sigma[V(\text{DaC \& A}) - V(\text{PBS } \textit{et al.})] = 0.008 \text{ mag}$$

The small size of these numbers was extremely gratifying to both groups!

Other examples of comparisons between CCD and conventional photoelectric photometry are given in Da Costa & Armandroff (1990) and in Walker (1990).

ACKNOWLEDGEMENTS

The author is grateful to Russell Cannon for his comments on the manuscript and to Steve Howell for the travel funding that enabled participation in the CCD School.

REFERENCES

Auer, L.H., and van Altena, W.F. 1978, AJ, 83, 531
Bessell, M.S. 1990, PASP, 102, 1181
Cousins, A.W.J. 1976, Mem. RAS, 81, 25
Cousins, A.W.J. 1978, Mon.Not.Astr.Soc.S.Afr., 37, 8
Da Costa, G.S. 1990, in CCDs in Astronomy, ASP Conf. Series, 8, 289
Da Costa, G.S., and Armandroff, T.E. 1990, AJ, 100, 162
Davis, L.E., Cawson, M., Davies, R.L., and Illingworth, G. 1985, AJ, 90, 169
Graham, J.A. 1982, PASP, 94, 244
Howell, S.B. 1989, PASP, 101, 616

King, I.R. 1971, PASP, 83, 199
Jefferys, W.H., Fitzpatrick, M.J., and McArthur, B.E. 1988, Celestial Mechanics, 41, 39
Johnson, H.L. 1955, Ann. d'Ap., 18, 292
Johnson, H.L., and Morgan, W.W. 1951, 114, 522
Kendall, M. and Stuart, K. 1977, The Advanced Theory of Statistics, Vol. 1, Distribution Theory, 4th edition, (Charles Griffin & Co., London)
Landolt, A.U. 1973, AJ, 78, 959
Landolt, A.U. 1983a, AJ, 88, 439
Landolt, A.U. 1983b, AJ, 88, 853
Massey, P., Garmany, C.D., Silkey, M., and DeGioia-Eastwood, K. 1989, AJ, 97, 107
Mateo, M., and Schecter, P. 1989, in Proceedings of the First ESO ST-ECF Data Analysis Workshop, edited by P.J. Grosbøl, F. Murtagh and R. M. Warmels (ESO, Garching)
Menzies, J.W., Barfield, R.M., Cousins, A.W.J., and Laing, J.D. 1989, S.Afr.Astr.Obs. Circ., 13, 1
Menzies, J.W., Marang, F., Laing, J.D., Coulson, I.M., and Engelbrecht, C.A. 1991, MNRAS, 248, 642
Newberry, M.V. 1991, PASP, 103, 122
Penny, A.J., and Dickens, R.J. 1986, MNRAS, 220, 845
Pritchet, C.J., and van den Bergh, S. 1988, ApJ, 331, 135
Sandage, A. 1970, ApJ, 162, 841
Sarajedini, A., and Da Costa, G.S. 1991, AJ, 102, 628
Stetson, P.B. 1979, AJ, 84, 1056
Stetson, P.B. 1987, PASP, 99, 191
Stetson, P.B. 1990, PASP, 102, 932
Stetson, P.B., and Harris, W.E. 1988, AJ, 96, 909
Stetson, P.B., VandenBerg, D.A., Bolte, M., Hesser, J.E., and Smith, G.H. 1989, AJ, 97, 1360
Stetson, P.B., Davis, L.E., and Crabtree, D.R. 1990, in CCDs in Astronomy, ASP Conf. Series, 8, 289
Walker, A.R. 1990, in CCDs in Astronomy, ASP Conf. Series, 8, 319