INTRODUCTION TO TIME-SERIES PHOTOMETRY USING CHARGE-COUPLED DEVICES*

Steve B. Howell
Planetary Science Institute
2421 East 6th Street
Tucson, AZ 85719

Received: January 30, 1992

Abstract

This paper reviews in some detail the ideas of differential photometry using a CCD detector. It discusses the basic ideas of observational uses of a CCD for this kind of work and then gives details of the quantitative aspects. Signal-to-noise calculations for a point source, optimum extraction of the digital data via software, and correct statistical error assessment of the data are some of the topics covered. An observational example is presented.

"A pupil from whom nothing is ever demanded which he cannot do, never does all he can." -- John Stuart Mill

1. Introduction

Charge-Coupled Devices (CCDs) have, to repeat an oft-heard phrase, revolutionized astronomy in terms of allowing observers to "see" fainter, obtain higher signal-to-noise (S/N) data, and to measure more precise astrophysical values. But in my opinion, these are not the only areas where CCDs will make one of their biggest impacts in astronomy. I believe that these small pieces of semi-conductor material and their modest entourage of associated electronics and host computers will greatly increase the amount of good basic astronomical data collected. Small telescopes available at institutions from high schools to universities, amateur astronomers at all levels, and even the many small, usually unsupported, telescopes at major observatories, are now all possible sites of frontier research.

For example, imagine the surprises and information that would be contained in a program to observe a sample of variable quasars every week for two years, or the monitoring of any variable star every night for an extended period of time. These types of projects are happening now, mostly through the use of smaller dedicated telescopes and CCD systems. They will present unprecedented data sets that will aid both future observers and theorists alike.

I will start this review with a brief discussion of the basic ideas and uses of differential photometry. I will then get into some details of the usage of a CCD for this type of science, some examples to acquaint the reader with CCDs, the method of differential photometry, and ways to test the validity of the data obtained. I will next present some possible future directions and, finally, I have included, within the references, an additional list of relevant reading material.

---

* This paper is a condensed version of the original which appeared in *Astronomical CCD Observing and Reduction Techniques,* ed. S.B. Howell. It is used here by permission of the Astronomical Society of the Pacific.

2. What's the Use?

If you go back and examine the title of this paper, you will see that it contains some fancy words strung together, making it all sound pretty impressive. So what is it we are really talking about? We want to exploit CCDs for observing celestial sources, examine and measure the collected data, and then determine whether a particular source is variable in time and to what statistically significant level.

We will not be concerned very much here with absolute photometry or transforming our measured magnitudes to a standard photometric system. For time-series work, this step is not always necessary. But we will be concerned with the method by which many types of variable or suspected variable sources can be studied, ways of reducing the data in an optimum fashion, and statistical procedures we can use to determine the level of variability. During the past five years, stars, asteroids, comet nuclei, quasars, and active galaxies have all been studied using various forms of differential photometry.

The topic of interest to us here is the observation and study of temporal source variability using a CCD as the detector. There are great advantages to using a CCD compared with a photomultiplier tube (PMT); CCDs have higher quantum efficiencies than PMTs, they provide simultaneous background measurement, they eliminate the need for aperture-centering of the source during observation, and maybe most importantly, CCDs have less background (sky) noise per sensing element. PMTs do, however, still have some observational regimes in which they are a better choice than CCDs, mainly in the area of very high-speed sampling. This is likely to change in the near future as some groups are already working on high-speed CCD photometers (see e.g., Stover 1986; Abbott and Opal 1988).

The basic idea for time series photometry (Howell and Jacoby 1986) using a CCD as a detector is to sit on a source (or sources) for some extended time period during which the instrument samples the emitted flux. Commonly, these measurements are made through various filters to isolate certain wavelength regions of interest. White light photometry, i.e., using the instrumental characteristics to determine the effective bandpass instead of a filter, generally collects flux arriving from most of the visible part of the spectrum (CCDs, for example, are generally sensitive to incoming flux between about 3000 Angstroms and 1μ).

Differential photometry consists of comparing a source of interest to other sources present on the same frames. These comparison sources provide essentially a perfect method of dealing with sky, background, and instrumental noise, as they can be used to set limits on the noise levels, the instrument stability, and therefore the intrinsic accuracy of variability. In the simplest case, there are three objects of interest: the source which we want to test for variability, and two comparison sources. Let us call these, keeping with tradition, V (the source of interest or suspected variable), and C and K (two comparison sources). We also will limit our discussion here to point sources (stellar-like objects such as stars, QSOs, or active galactic nuclei), as the situation is much more difficult for two-dimensional objects (e.g., comets, fuzzy galaxies).

Figure 1 shows an example of time-series light curve data. Plotted are four light curves, three of which show common features (dips and wiggles), indicating that these (and of course all) sources on each frame were subject to similar variations due to transparency changes or clouds, unrelated to actual source variation. The source V quite obviously shows a different light curve pattern, giving us a hint that it may indeed be intrinsically variable with time. In fact, 25 candidate light curves were produced from this ≈2.5-hour sequence of CCD observations, actually allowing the discovery of the variable star from these data (see Figure 7).

3. The Basics

Howell and Jacoby (1986) looked into two methods of obtaining time-series photometric data with a CCD. One method, which has since not been further explored to the best of my knowledge, was to track the telescope at non-sidereal rates. This action produces trailed images of all the sources present, including any source of interest, and by picking the appropriate track rate, one could vary the time resolution acquired. This method has potential for obtaining high time-resolution data, but also has some weak points (see Howell and Jacoby for a discussion of some of its weak points).

The other method discussed by Howell and Jacoby, and the focus of this paper, was termed "repeated exposures". Simply put, the telescope is trained upon an object of interest and allowed to track it for extended periods of time. During this time, short exposures are made in a continuous fashion. Each of these frames therefore contains the object(s) of interest and (hopefully) two or more comparison sources. The entire data set is reduced in a more or less standard way (see Howell 1992). Two-dimensional aperture photometry is then completed for each source of interest on each frame, and a light curve is produced. When acquiring the CCD field to use for such a time-series data set, defects in the CCD, such bad pixels, hot pixels, and bad columns, should be avoided. Also, any object of interest should not be placed too near an edge, especially if no auto-guiding capability exists. This procedure ensures that small drifts in image placement due to guiding errors, telescope flexure, or atmospheric refraction won't cause the source to fall outside the CCD boundaries.

Many CCD systems are available today in which these types of observations can be made, from national observatories to amateur telescopes. The biggest problem with most of these systems is in handling the sheer volume of data that can be generated. For example, a typical light curve obtained on one night of observing might consist of 6 hours of 3-minute integrations on a source, resulting in $\approx$100 CCD frames. For a CCD of dimension 512 x 512 pixels (average-to-small by today's standards), this equals, at $\approx 1/2$ megabyte per image, a total of 50 megabytes of data. For a week-long run, this becomes $\approx$300 megabytes. The numbers go up rather quickly and if this amount of data is to be manipulated and reduced quickly and efficiently, some thought must be given to the storage, retrieval, and reduction processes to be used.

One of the easiest procedures to use to avoid large amounts of data is to use only a piece of the CCD. Many software systems (such as those in use at Lowell or Kitt Peak National Observatory) allow for windowing the CCD to some rectangular size of dimension smaller than full frame. For many applications with typical images scales, ($\approx$0.3-0.8 arcsec/pixel), a CCD format area of 256 x 256 pixels is usually quite sufficient. Another easy solution is to use chip binning. This allows the user to sum together pixels with the CCD electronics before they are stored as digital numbers. These "superpixels" cause a slight loss in spatial resolution (while increasing the S/N for a given integration time), and accomplish a reduction in the stored CCD frame size. Both of these answers to reducing the final stored CCD size are to be used with caution, however. The image scale of a given camera combination (i.e., the mapping or arc seconds on the sky to each CCD pixel) should be considered. Using too small a window might not allow a sufficient number of comparison sources to be present, and binning the pixels may undersample each point source such that photometric accuracy and source-centering capability are sacrificed (see e.g., Stone 1989; Stetson 1987; Buonanno *et al.* 1989; Merline and Howell 1991).

For systems in which autoguiding is available, it should certainly be used. It will allow much faster data reduction since each source of interest is located at or near the same pixels in every frame. Datasets obtained without being guided, even for good tracking telescopes, often present reduction software problems and may need to have each frame reduced separately, greatly increasing the user interaction time required.

Another consideration is the deadtime associated with CCD readout. Typical astronomical CCD readout times are from 20-45 seconds, during which time the stored charge in each pixel is converted into a digital number (sometimes called a data number [DN] or an analog-to-digital unit [ADU]) and stored usually on a magnetic disk. Software overhead due to other items such as image preparation and storage time must also be considered. The effective sampling time for a given CCD integration is therefore the sum of the integration time + readout time + any other associated overhead time needed by the computer software or hardware. Typical numbers might be that for a 60-second integration, one would need to add 5 seconds for image and CCD preparation, 30 seconds for readout, and another 5 seconds for storage of the image. This would make the effective sampling time equal to 100 seconds. All of these numbers, except the integration time, are usually dependent on the CCD format size used and are independent of the integration time.

Even though the deadtimes for most CCD systems seem exceedingly long as compared to PMTs, the actual efficiency of a CCD for photometric observations is much higher. Remember that with a CCD one measures the variable, some comparisons, and the sky simultaneously, so the practice of offsetting to comparison stars and centering them in the aperture, and then offsetting to blank sky positions, is eliminated. The final two-dimensional digital image also allows for better sky determination and software measuring apertures that can be fit after the fact and are much smaller than those used with PMTs. A typical PMT observation might use a 10- or 20-arcsec aperture compared to a software-selectable, few-arcsec aperture one can use with CCD data. This illustrates a big advantage of CCDs in terms of noise per sensing element, i.e., 300+ square arcsec of sky for a PMT aperture imaged onto the photocathode compared to ≈0.2-1.0 square arcsec of sky per CCD pixel. Programs requiring time resolutions shorter than 10-15 seconds will not be served well by most currently available CCD systems. More sophisticated CCD controllers and software, as well as creative thinking about data handling (e.g., real-time reductions or sub-area pixel readout), will likely eliminate most of the sampling time constraints with CCDs in the near future.

A problem that has always plagued photometrists is non-photometric weather. Using differential techniques with CCD, observations can progress even when non-photometric conditions prevail (see Figure 1 and Howell and Jacoby 1986). During these non-photometric times, a loss of S/N or time resolution occurs due to the presence of thin clouds and/or transparency variations, but that is heavily offset by the large increase in the total number of usable nights. While observing on these nights however, absolute photometry cannot be performed. But, since one measures all the sources in the CCD frame simultaneously, an image taken on a photometric night, at some previous or later time, of the same field in the same filter, allows each differential measurement to be converted into an absolute value. For differential measurements, air mass effects, color terms, and differential refraction are essentially eliminated and do not have to be treated. Observations of a few photometric standards on a given night, and without correction for color terms and air mass, will allow magnitudes accurate to a few tenths to be obtained for the program objects.

4. The Details

Before we start into the realm of differential photometry itself, let us first describe some basic ideas we will need later on. These are: the signal-to-noise ratio (S/N) and what it means for point sources imaged on CCDs, how we can extract the data optimally, the use of growth curves, and some comments on background determination.

4.1 S/N Ratio

Photons behave statistically according to the rules of Poisson (or photon) statistics. Many statistics books and basic astronomical photometry and instrument texts contain information on Poisson statistics (see e.g., Henden and Kaitchuck 1982 or Warner 1988). Thus, for a perfect photon counting device (noiseless or essentially so), the S/N (equal to $1/\sigma$) is given by

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_*}} = \sqrt{N_*} \tag{1}$$

where N_* are the total counts in the source.

However, just as we eventually learned after our freshman physics classes, no system is really perfect. Therefore, we must augment our perfect equation with some other terms that bring reality into play.

For a CCD, the standard equation for the S/N of a measurement of a point source is given by

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{pix}(N_S + N_D + N_R^2)}} \tag{2}$$

where each of the terms is described in Table 1.

Table 1. Terms used in equations 1-3

N_*	=	total counts collected in source (electrons/integration)
n_{pix}	=	total number of pixels used in measuring aperture
N_S	=	sky level (electrons/pixel/integration)
N_D	=	dark level (electrons/pixel/integration)
N_R	=	read noise (rms electrons/pixel/read)
n_B	=	total number of pixels used in background determination
G	=	gain of CCD (electrons/ADU)
G_o	=	gain used when CCD parameters (such as read noise) were determined (usually = G)
σ_f^2	=	variance of digitization noise

Equation (2) appears in many papers, CCD instrument manuals, and on blackboards in instrumentation classes. But where does it come from? There are many Society of Photo-Optical Instrumentation Engineers (SPIE) publications that contain treatments of its derivation (see e.g., the *SPIE Proceedings,* volumes 264, 290, 331, 445, 501, 627), (see also Mortara and Fowler 1981; Howell 1992).

Equation (2) is the typical equation used by most observers to calculate the S/N of a point source observation or to estimate integration times needed before going to the telescope. There are, however, other subtle terms that are not included in it which can become quite important under some conditions. Howell and Merline (1992) derive the S/N equation in gory detail and discuss the implications of each term. Their final result is a revised CCD equation of the form

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{pix}\left(1 + \frac{n_{pix}}{n_B}\right)(N_S + N_D + N_R^2 + G^2\sigma_f^2)}} \tag{3}$$

This equation has two terms in it not present in equation (2). The term

$$\left(1 + \frac{n_{pix}}{n_B}\right)$$

accounts for the error associated with the background (sky) determination which is dependent on the number of background pixels used. The term $(G_0\text{-}G)\sigma_f^2$ deals with the digitization error in the output data. The digitization noise is a consequence of the conversion from analog to digital output. Basically, any fractional part of the signal "left over" in each pixel, after the conversion to an integral number of ADU, is simply chopped or lost. For example, a CCD gain of 10 electrons/ADU will yield an output value of 1 ADU for 10 electrons, 2 ADU for 20 electrons, but will also give a value of 1 ADU for 11 to 19 electrons. This process may lose up to almost one full ADU per pixel in any integration. In practice the problem is complex and the average value of the lost signal depends on the type of A/D converter used and on the system gain. This problem is discussed by Opal (1988), who also provides a good discussion of A/D converters useful for digital photometry. Digitization noise is treated in detail in Howell and Merline (1992). This latter term is especially important when one is using a CCD system with a large gain, i.e., when the noise floor is not well sampled (e.g., a system with read noise = 7 electrons and a gain of 100 electrons/ADU). See Howell and Merline (1992) for an account of more than you probably ever wanted to know about the *S/N* equation.

4.2 Poisson Statistics and Equations 1 - 3

As in all approximations to reality, they are good enough up to a certain point, after which deviations occur. In what regimes can we use equation (1) and when do we need to invoke the full-blown treatment given by equation (2) or (3)? This is in principle easy to see from the equations themselves. In equation (2), we see that if, $N_* >> n_{pix}(N_S + N_B + N_R^2)$, the equation is approximately given by $\sqrt{N_*}$ or the (noiseless) Poisson result. This will be our working definition of an observation of a bright source.

Figure 2 shows a plot of the percentage error which would occur if equation (1) were used to calculate the *S/N* instead of equation (2) (see Howell 1989). For a specific range in the observations (i.e., bright sources), $S/N \approx \sqrt{N_*}$. But it very quickly deviates from this approximation as the additional noise term in the denominator of equation (2) begins to become significant (i.e., faint or low *S/N* sources).

4.3 Optimum Data Extraction

We now look at the *S/N* of a point source as a function of the software aperture applied to it during the extraction of photometric data. Figure 4 shows a point source of moderate *S/N* imaged onto a CCD. Close examination of this image reveals that the wings of the point source cover much more area than does the core and most of this area is noisy. In fact, if we plot the calculated *S/N* for a point source vs. radius of the extraction aperture, we get the results shown in Figure 5 (see Howell 1989).

There are three important features to glean from this figure. The *S/N* of a point source reaches a maximum at fairly small radii, approximately given by, but not in general equal to, the half-width at half maximum (HWHM) of the point spread function (PSF). (Note: In Howell [1989] a typographical error led to this being listed as the full width at half maximum [FWHM].) From this maximum, the *S/N* decreases in both directions. Secondly, the maximum *S/N* is not necessarily at the same radius for all point sources, even when the sources are measured from the same CCD frame. Lastly, the decrease away from maximum *S/N* is quite rapid in all cases.

We can see that to obtain the maximum *S/N* for a given point source observation, a software aperture equal to the maximum *S/N* for each particular source can lead to a large photometric improvement (Howell 1990a). Too small an aperture does not allow for enough signal and too large an aperture includes many noisy pixels. In practice, an aperture suited to the faintest point sources of interest, while being slightly too small for brighter sources, works quite well.

4.4 Growth Curves

The method of data extraction described in the Optimum Data Extraction section above does not, however, include all the light collected by the CCD for a source of interest. So if optimum extraction techniques are to be used, some correction scheme (sometimes called an aperture correction) must be applied in order to determine the true total flux. Summing over the areal distribution of light in steps of increasing radii, a growth curve for a given point source can be produced (DaCosta *et al.* 1982; Howell 1989, 1990b; Stetson 1990). Figure 6 gives some examples of typical CCD point source growth curves. Howell (1990b) and Stetson (1990) discuss these in detail and give examples of their uses.

4.5 Background Determination

The standard method of two-dimensional aperture photometry (see e.g., Adams *et al.* 1980) consists of placing a software aperture of some radius around a source of interest and summing the flux within it. Next, an annulus of inner radius larger than the source radius and outer radius determined by other considerations, such as field crowding or wanting a large enough sample of the background, is summed over, and an estimate of the background is found (see Howell and Merline 1992 for a discussion of errors due to background determination). This estimate might be a simple average, the mode, the median, or some other statistical construct. The single number that results from the background determination is then multiplied by the total number of pixels in the source aperture and subtracted from that sum. One is thus left with an estimate of the source flux.

For bright sources, where N_* is large compared to the rest of the noise terms, this background estimate and subtraction process is a very small source of error, and almost any method used to get the background value will suffice. Unfortunately, for most of our observations background determination is the critical part of the entire analysis and is thus a very important step in performing good two-dimensional digital photometry. When the background is determined from any CCD image it is not likely that usage of any two methods (such as the mean and the median, for example) will yield the same estimate for the background or "sky" value. Therefore, by using one method over another, different answers are produced. Which is correct?

Even a difference in the sky value determined by two methods of as small as 1 ADU may cause a difference of $n_{pix} * 1\ ADU \approx 80$ ADU for a source. For a CCD of gain 10 electrons/ADU, this gives 800 electrons = 800 photons uncertainty in N_*, possibly a substantial fraction of N_* itself. This topic is discussed in Howell (1989), who concludes that for low S/N observations, any sky determination scheme may lead to problems and suggests growth curves as a way out of the dilemma.

5. Getting a Look at your Data

Before we begin with the details associated with differential time series photometry, let us look at some subtle elements concerning point sources imaged onto CCDs. Merline and Howell (1991) have developed a model for point sources imaged onto a CCD. We make use of that model here to show the reader three-dimensional profiles of point sources from high through low S/N. This will illustrate some of the reasons necessary for having a good handle on understanding the noise terms and how to calculate the errors correctly in CCD data (see Howell and Merline

1992 for more examples).

Figure 6 shows three models for stars of 17th, 19th, and 21st magnitude, imaged onto a TI CCD with a one-meter telescope. Table 2 gives complete details of the models.

Table 2. Model Parameters for Figure 6

Sky magnitude (mag/sq. arcsec)	21.5
FWHM (arcsec)	1.5
Telescope diameter (meters)	0.9
Filter (λ_c:$\Delta\lambda$)	5500A/980A
CCD Quantum Eff (%)	70
Pixel size (microns)	15
n_{pix} (pixels)	69
Read noise (e^-/pixel/read) $\equiv N_R$	8
Sky value (ADU/pixel) $\equiv N_S$	70
Dark current (e^-/pixel/read) $\equiv N_D$	0
Gain (e^-/ADU)	4.15
Integration time (seconds)	120

	Figure		
	6a	6b	6c
Magnitude	17	19	21
Total counts (ADU) $\equiv N_*$	12,701	1,992	307
S/N (eq. 2)	185	42	6.5

The Axum[TM] graphics package (from TriMetrix, Inc.) was used on a PC/AT to make the plots in Figure 6. The 3-dimensional plots are made as pixel histograms rather than the common spline surface fits. We prefer to use pixel histograms rather than the common spline fits which tend to smooth out the real features in three-dimensional data, masking the detailed nature of the image.

For the 17th magnitude star, we see that it can be classed as marginally bright as $N_* > n_{pix}(N_S + N_D + N_R^2)$. Thus, the background is a somewhat unimportant contribution in terms of noise. For the 21st magnitude object, just the opposite is true. Dealing with the noisy background would be the biggest uncertainty in performing accurate photometry on this object.

From Figure 6c, one can easily see how a single number does not represent the background (sky) value very well. It is this large uncertainty we want to avoid by not using the large area contained in the stellar wings. Keep these profiles in mind when you are working with CCD images, as this is what your data really look like to whatever software you use. Routines like Gaussian fits, sky estimators, and two-dimensional photometry packages should all be used with the understanding that knowing your noise characteristics is as important as the signal in which you are interested.

6. Differential Photometry

Differential photometry using a CCD has been used now for some years with much success. It allows highly precise photometric measures to be obtained even for faint sources and with small telescopes.

Differential photometry is a very powerful technique which allows observers to use many traditionally unusable nights, attain very precise measures, and does not eliminate the possibility of obtaining absolute magnitudes for the source of interest. I will now discuss the method of differential photometry (see Howell *et al.* 1988) in its simplest form, including an appropriate test to determine if true variability was observed.

For each CCD frame we can measure three sources (*V*, *C*, and *K*) and form the differences (in magnitude) *V-C* and *C-K* or the ratios (in flux) *V/C* and *C/K*. The *C-K* light curve data are used for three purposes:

1) to check for variability or trends in the comparisons themselves,
2) as a measure of the intrinsic accuracy provided by the instrumental setup used, and
3) as a reference to determine if *V* is indeed variable and at what level.

Since this differencing technique essentially eliminates any drifts that occur between all the sources simultaneously (including clouds, air mass effects, color terms, and differential refraction), these types of observations can be made even during non-photometric weather (i.e., clouds) with the only real loss being in terms of total flux collected during passing clouds for a given integration time (see Figures 1 and 7).

Differential photometry has been around for years and is not a technique new to CCDs (see Young *et al.* 1991). When this technique was used with PMTs, one could integrate on each of the sources independently, choosing different integration times such that the output *S/N* for each source was the same. Therefore, the output light curves could be directly compared. When using a CCD however, all the sources present on a given frame are of the same integration time, which means that in general, none of their associated *S/N* values will be equal. Thus a direct comparison of each source light curve is not possible (see Howell 1990a). What is needed is a scale factor for each source which provides a "correction" for all the individual *S/N* values. This gives all of them equal footing and this correction must be applied before comparing *V-C* with *C-K*.

We make the following assumptions about any variability seen in the sources:

$$\sigma^2_{V\text{-}C} = \sigma^2_{V\text{-}C}(\mathrm{VAR}) + \sigma^2_{V\text{-}C}(\mathrm{INST}) \tag{4}$$

and

$$\sigma^2_{C\text{-}K} = \sigma^2_{C\text{-}K}(\mathrm{VAR}) + \sigma^2_{C\text{-}K}(\mathrm{INST}) = \sigma^2_{C\text{-}K}(\mathrm{INST}). \tag{5}$$

Each dataset has two likely components of variability: one due to instrumental variations (a generic term for all noise, like photon statistics, read noise, etc., not including any actual source variability) and, secondly, any actual source variation itself. For *C-K*, we assume that $\sigma^2_{C\text{-}K}(\mathrm{VAR}) \equiv 0$, and it is up to the observer to verify that the comparison stars are non-variable.

Our goal is to compare $\sigma^2_{V\text{-}C}$ with $\sigma^2_{C\text{-}K}$ in such a way as to determine if *V* is, in fact, variable. This is accomplished by the use of a scaling factor, Γ^2, in the following way:

$$\sigma^2_{V\text{-}C}(\mathrm{INST}) = \Gamma^2 \sigma^2_{C\text{-}K}(\mathrm{INST}) \tag{6}$$

The expression for Γ^2 contains terms which relate the *S/N* of each differential light curve. In this way, variations seen in the suspected variable data can be compared to what is expected merely due to statistical noise (i.e., $\sigma^2_{V\text{-}C}(\mathrm{INST})$). The application of equation (6) can be done for each CCD frame separately or applied globally to an entire dataset. Figure 8 gives an example of this. In general, if one's data does not contain any large excursions (such as moonrise during a time-series observation, for example), then equation (6) globally applied is usually sufficient.

Since Γ^2 is a scaling factor relating two variances, its value should be kept as close to 1.0 as possible. This will occur if all three sources are of equal magnitude. Rarely does this happen, however, and the best compromise is having *V* or *K* as close as possible in magnitude, with *C* slightly brighter. In reality, one must use whatever sources are available on a given frame and, as Γ^2 gets further from unity, a less stringent test of variability will occur. Once the Γ^2 value is calculated, application of equation (6) will show if the actual measured $\sigma^2_{V\text{-}C}$ is as expected (i.e., $\sigma^2_{V\text{-}C}$(INST)), or if $\sigma^2_{V\text{-}C}$ appears to be greater, thereby possibly indicating real source variability. An F-test can then be applied to check for variability in the source of interest, *V*, at a predetermined confidence level. This test for variability, at, say, the 95% confidence level, does not determine what kind of variable signal the data show (i.e., sinusoidal), merely if the source *V* is indeed variable (see Howell *et al.* 1988 for details and Howell 1990b).

7. The Future

CCDs are becoming very commonplace in today's well equipped observatory. They will undoubtedly increase the sheer amount of data collected by orders of magnitude compared to previous years of observation. New technologies of dealing with all these data and how and what amount to keep are large issues that still need to be decided. In the visible part of the spectrum, the quantum efficiencies of CCDs are currently pretty high and read noise per pixel per read has dropped to a very low level. Therefore, I don't believe that further advances in these parameters will be the big areas of improvement in the future years of CCD research. The areas of tremendous progress will be:

1) the collection of previously unthinkable data sets in terms of precision and temporal coverage, and
2) the way in which we deal with these data after collection.

New software techniques, real-time data reduction, and particularly, new ideas of how to reduce the data collected are currently advancing and will provide a new revelation to astronomy. So get involved, come on, we are all counting on you!

8. Acknowledgements

I would like to thank the National Science Foundation for providing partial support for this work, Lowell Observatory for generous allocations of observing time, and all of the following for their efforts and support during my research into CCD techniques both by collaborations with me and through their published and unpublished work and personal discussions: Jason Cardelli, George Jacoby, Jim Janesick, Tobias Kreidl, Sig Kutter, Keith Mason, Bill Merline, Ken Mitchell, Mike Shara, Peter Stetson, Paula Szkody, R. Mark Wagner, and Archie Warnock.

References

Abbott, T. M. C., and Opal, C. B. 1988, in *Instrumentation for Ground-Based Optical Astronomy: Present and Future,* ed. L. B. Robinson, Springer Verlag, New York, p. 380.

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

Allen, C. W. 1973, *Astrophysical Quantities,* Third Edition, Athlone, London.

Buonanno, R., and Iannicola, G. 1989, *Publ. Astron. Soc. Pacific,* 101, 294.

DaCosta, G. S., Ortolani, S., and Mould, J. 1982, *Astrophys. J.,* 257, 633.

Henden, A. A., and Kaitchuck, R. H. 1982, *Astronomical Photometry,* VanNostrand Reinhold, New York.
Howell, S. B., and Jacoby, G. 1986, *Publ. Astron. Soc. Pacific,* **98**, 802.
Howell, S. B., Mitchell, K. J., and Warnock III, A. 1988, *Astron. J.,* **95**, 247.
Howell, S. B. 1989, *Publ. Astron. Soc. Pacific,* **101**, 616.
Howell, S. B. 1990a, in *CCDs in Astronomy II: New Methods and Applications of CCD Technology,* eds. A. G. D. Philip, D. S. Hayes, and S. J. Adelman, L. Davis Press, Schenectady, NY, p. 133.
Howell, S. B. 1990b, in *CCDs in Astronomy,* ed. G. Jacoby, *Publ. Astron. Soc. Pacific Conference Series,* **8**, 312.
Howell, S. B. 1992, "Astronomical CCD Observing and Reduction Techniques", *Publ. Astron. Soc. Pacific Conference Series,* **23**.
Howell, S. B., and Merline, W. H. 1992, *Photonic Astronomy,* in press.
Merline, W. H., and Howell, S. B. 1991, *Astronomical Data Analysis Software and Systems,* eds. D. M. Worrall, C. Biemesderfer, J. Barnes, *Publ. Astron. Soc. Pacific Conference Series,* **25**, 316.
Mortara, L., and Fowler, A. 1981, in "Solid State Imagers for Astronomy", *Proc. Soc. Photo-Optical Instrumentation Eng.,* **290**, 28.
Opal, C. B. 1988, in *Second Workshop on Improvements to Photometry,* ed. W. J. Borucki, NASA Conference Pub. 10015, p. 179.
Stetson, P. B. 1987, *Publ. Astron. Soc. Pacific,* **99**, 191.
Stetson, P. B. 1990, *Publ. Astron. Soc. Pacific,* **102**, 932.
Stone, R. C. 1989, *Astron. J.,* **97**, 1227.
Stover, R. 1986, *Proc. Soc. Photo-Optical Instrumentation Eng.,* **627**, 195.
Warner, B. 1988, *High Speed Astronomical Photometry,* Cambridge University Press, NY.

Reading List

Bevington, P. R. 1969, *Data Reduction and Error Analysis for the Physical Sciences,* McGraw-Hill, NY.
Bottaccini, M. R. 1975, *Instruments and Measurement,* Charles E. Merrill Publishing Co., Columbus, OH.
Diego, F. 1985, *Publ. Astron. Soc. Pacific,* **97**, 1209.
Djorgovski, S. 1984, in *Proceedings of the Workshop on Improvements to Photometry,* ed. W. J. Borucki and A. Young, NASA Conference Pub. 2350, p. 152.
Holtzman, J. A. 1990, *Publ. Astron. Soc. Pacific,* **102**, 806.
Jacoby, G. H., ed. 1990, "CCDs in Astronomy", *Publ. Astron. Soc. Pacific Conference Series,* **8**.
Janesick, J. R., Hynecek, J., and Blouke, M. M. 1981, in "Solid State Imagers for Astronomy", *Proc. Soc. Photo-Optical Instrumentation Eng.,* **290**, 165.
Jayroe, R. R. 1984, *NASA Technical Paper 2396.*
Jayroe, R. R., and Schroeder, D. J. 1984, *NASA Technical Paper 2397.*
King, I. R. 1971, *Publ. Astron. Soc. Pacific,* **83**, 199.
King, I. R. 1983, *Publ. Astron. Soc. Pacific,* **95**, 163.
Lucy, L. B. 1974, *Astron. J.,* **79**, 745.
MacKay, C. 1986, *Ann. Rev. Astron. Astrophys.,* **24**, 255.
McMillan, R. S., and Stoll, C. P. 1982, in "Instrumentation in Astronomy IV", *Proc. Soc. Photo-Optical Instrumentation Eng.,* **331**, 104.
Nather, R. E. 1972, *Publ. Astron. Soc. Pacific,* **84**, 149.
Philip, A. G. D., Hayes, D. S., and Adelman, S. J., eds. 1990, *CCDs in Astronomy II: New Methods and Applications of CCD Technology,* L. Davis Press, Schenectady, NY.

Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T. 1986, *Numerical Recipes, The Art of Scientific Computing,* Cambridge University Press, Cambridge.
Walker, A. 1990, *Astron. Soc. Pacific Conference Series,* ed. G. H. Jacoby, 8, 319.
Walker, A. 1984, *Mon. Not. Roy. Astron. Soc.,* 209, 83.
Young, A. T. 1974, *Astrophys. J.,* 189, 587.
Young, A. T. *et al.* 1991, *Publ. Astron. Soc. Pacific,* 103, 221.

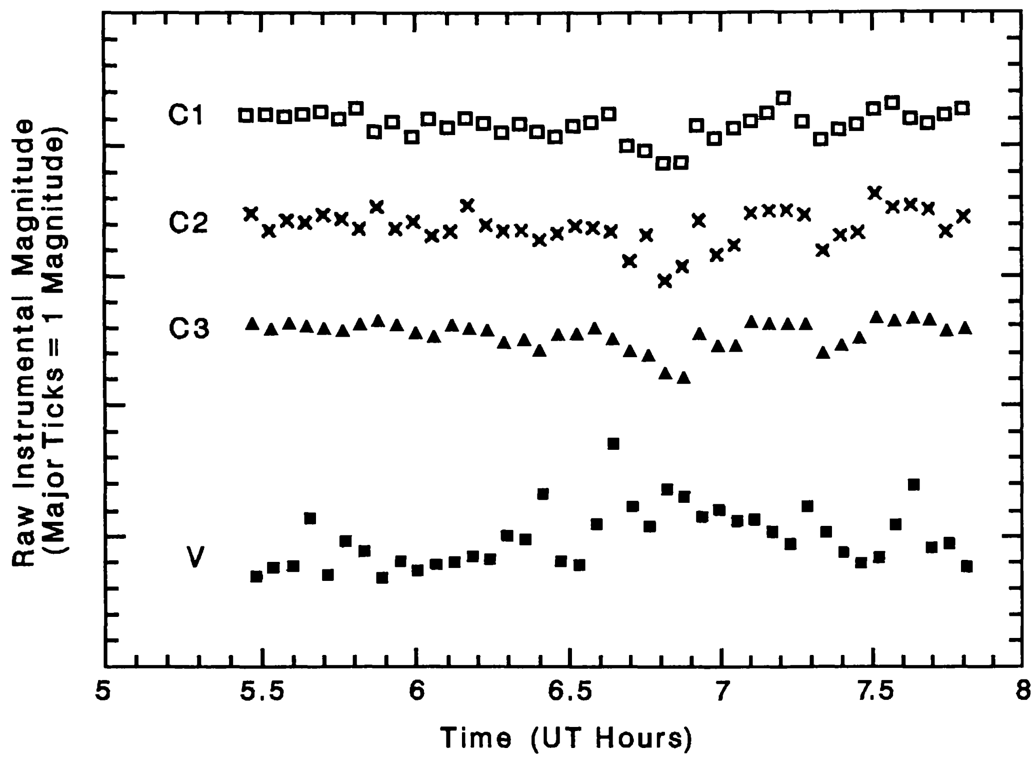

Figure 1. Four raw instrumental light curves vs. time. The real magnitudes of all four objects are approximately equal (V ≈ 19), and they have been plotted here with offsets applied. Notice how the three objects, labeled with C's, track each other, showing passing clouds or transparency changes.

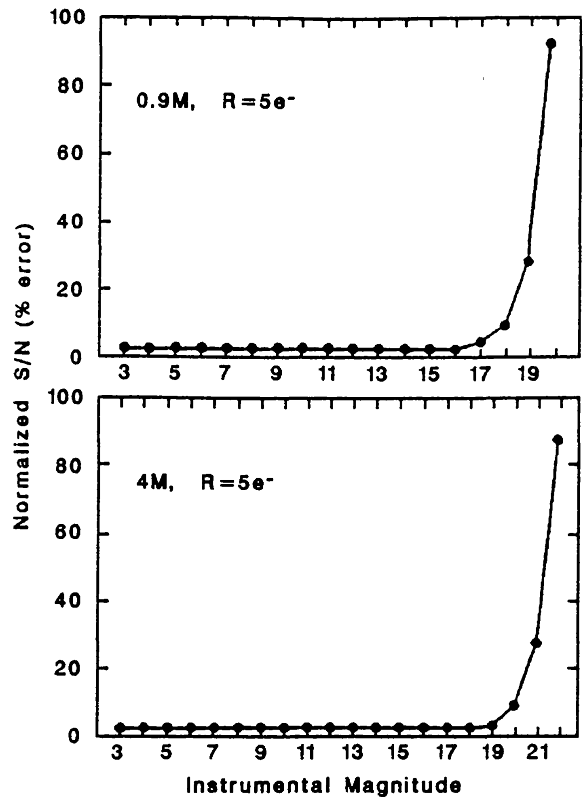

Figure 2. Two examples, for a 0.9- and a 4-meter telescope, showing the regimes over which simple Poisson statistics work and then how quickly they fail. These data are for 300 second exposures. For example, use of equation (1) for a 19th magnitude point source on a 0.9m telescope would yield ≈30% error as compared to using equation (2). R = read noise value.

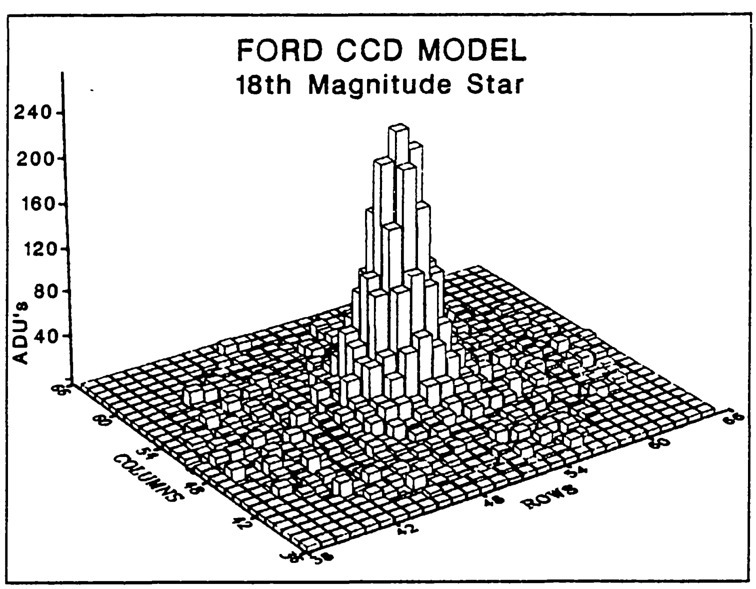

Figure 3. An 18th magnitude point source imaged at fairly high S/N, on a Ford Aerospace CCD. Note how the core is well defined and fairly Gaussian in shape. The large area contained in the wings is quite noisy due to both read noise and sky noise.

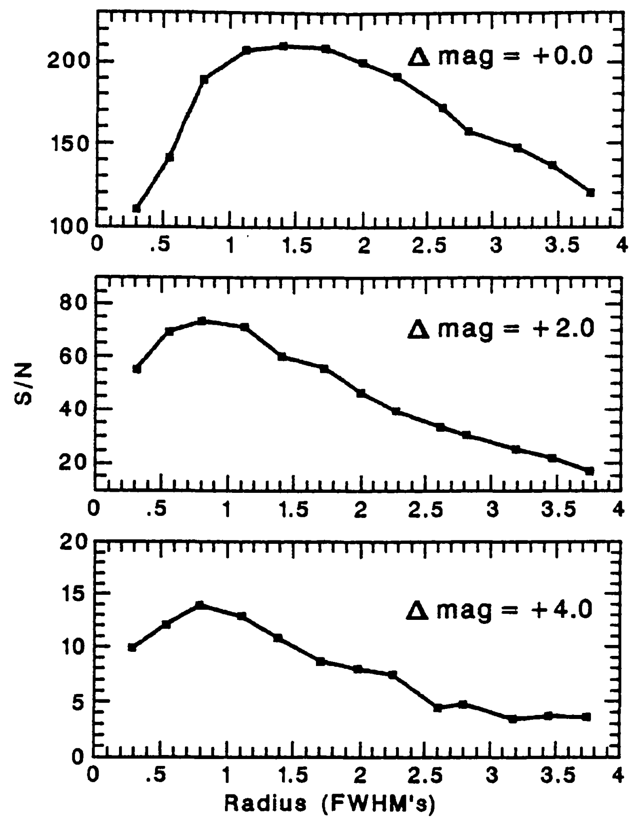

Figure 4. Plots of the S/N calculated from equation (3) for three point sources. The middle curve is for a source 2 magnitudes fainter than the top source, while the bottom source is 4 magnitudes fainter. The optimum extraction aperture for the two fainter sources is near 0.75 FWHM.

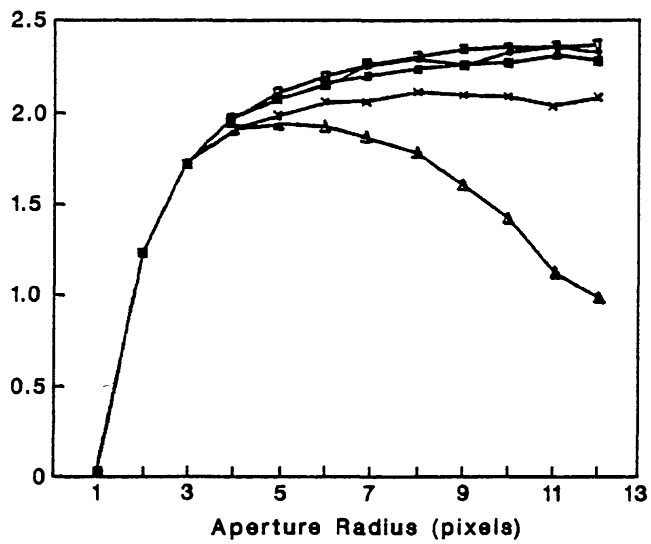

Figure 5. Growth curves for five stars from a single CCD frame. This example shows how well-sampled stars define the growth curve for a given frame, and how fainter stars deviate from this mean curve when they are extracted at large (usual) radii. The mean curve can be used to "correct" the low S/N sources back to their real values.

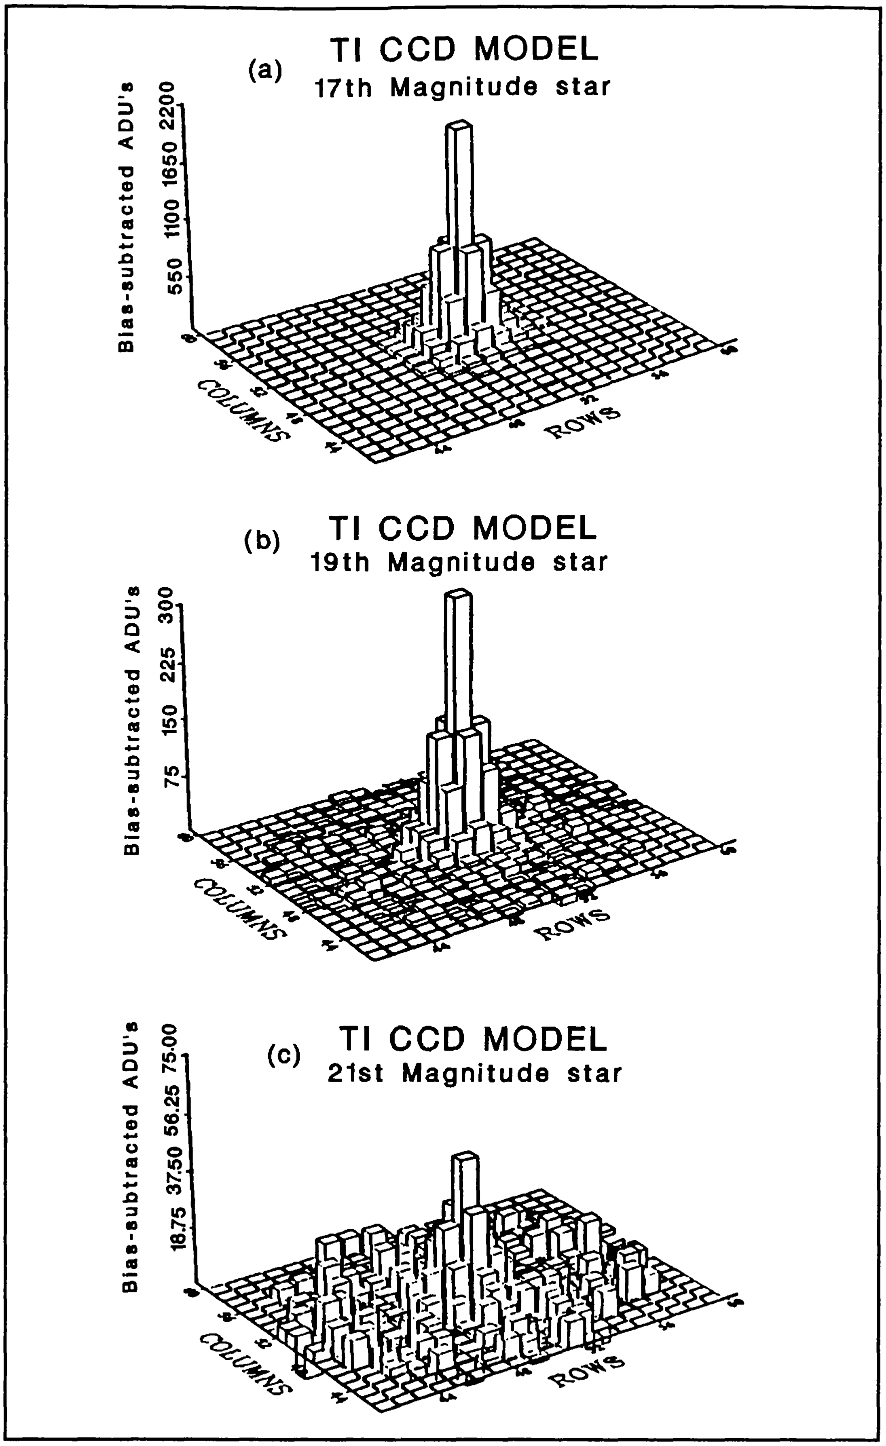

Figure 6. Models of a) 17th, b) 19th, and c) 21st magnitude stars are shown. They have S/N values of 185, 42, and 6.5, respectively. Note how the profiles become less well defined with respect to the relative increase in the noise level.

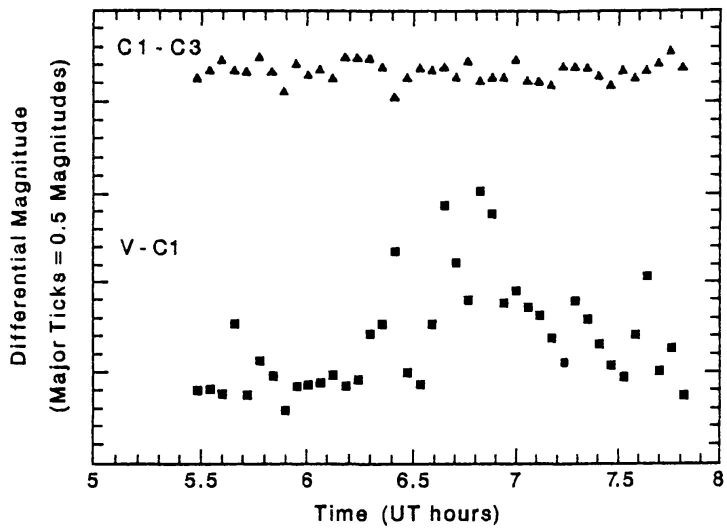

Figure 7. Differential light curves produced from the data in Figure 1. The C1-C3 light curve has a 1σ error of 0.053 magnitude, equal to its predicted value (see equation (10) in Howell *et al.* 1988). The V-C1 light curve shows that *V* is indeed variable and the F-test discussed in the text shows it to be variable at >99.5% confidence level.

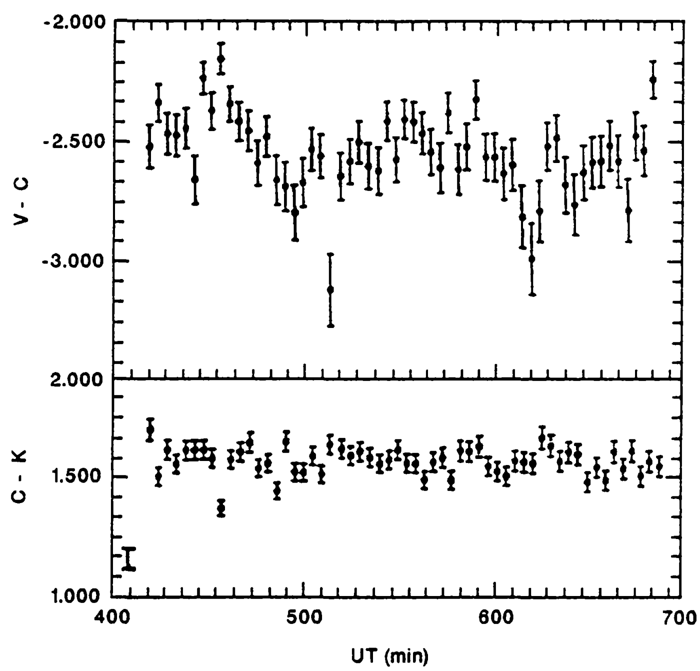

Figure 8. A near-infrared light curve for the star DM Dra. It was measured when at a mean magnitude of V = 21.7. The top panel shows the variations in the star itself, while the bottom panel shows two comparisons referenced against each other as a control. Equation (6) has been applied to these data on a frame-by-frame basis, thus each measurement is assigned its own error. The single error bar at the lower left is the global value calculated for *C-K* from equation (6). Note that the global error matches well the individual *C-K* scatter and how the error in the *V-C* data increases when the star grows fainter.