

SOME FACTORS AFFECTING THE ACCURACY OF STELLAR PHOTOMETRY WITH CCDs (And some ways of dealing with them)

Peter B. Stetson

*Dominion Astrophysical Observatory, Herzberg Institute of Astrophysics,
National Research Council
5071 West Saanich Road
Victoria, British Columbia V8X 4M6 / Canada*

ABSTRACT. There are many factors which make it difficult to relate instrumental CCD photometry to a fundamental standard system with an accuracy much better than about 1%. Here I will address only three of them: (1) infrared leaks in the filters; (2) the finite opening and closing times of mechanical shutters; and (3) changes in the air mass for long integrations. I will be approaching these subjects from the point of view of a visiting astronomer at someone else's observatory, who gets three or four nights of observing time a year, and has only the afternoons preceding those nights to perform whatever tests can be carried out while the equipment sits on the telescope.

1 Infrared leaks and flat fields

For normalizing CCD photometry of bright stars (those that are *not* very much fainter than the night-time sky, especially standards and stars observed with narrow-band filters), two types of flat-field calibration frames are particularly useful: dome flats, which are observations of the inside of the dome or of a white diffusing screen, usually illuminated with an incandescent lamp; and twilight-sky flats, which are observations of the bright sky just after sunset or just before sunrise. These frames are used to map out the spatial variations in the sensitivity of the detector, so that the derived apparent magnitude of a star can be made independent of where on the detector the star's image falls.

It is commonly found, however, that dome flats and twilight-sky flats differ by significant amounts, making it rather uncertain which – if either – of the two techniques is correct. For instance, in data from an observing run which Jim Rose and I had in 1986 on the University of Hawaii 2.2-m telescope with the IfA/Galileo 500² TI CCD, I found that the twilight flats in the *I*, *V*, *B*, and *U* filters had structure at the 7%, 9%, 11%, and 13% levels (root-mean-square), respectively. The dome flats in exactly the same filters, on the other hand, showed structure at the 5%, 6%, 7%, and 8% levels. The *U* dome flats much more closely resembled the *I* sky flats in general appearance than they did the *U* sky flats, and even the *B* and *V* dome flats appeared to show significant dilution by an *I*-like response pattern. This led me to suspect that the dome flats were being contaminated by very-long-wavelength light (longer-wavelength than *I*, probably), which is much stronger in the tungsten lamp than in the light of the twilight sky.

It seemed to me that I could see the same effect in dome- and twilight-flats which Bill Harris and I had obtained in 1984 and 1985 at Kitt Peak, albeit at a somewhat lower level.

Therefore I have carried out some crude numerical simulations based on published curves for “Mould” interference filters and RCA CCDs (which cover the principal peaks of the filter passbands, and the wavelength range 4000 – 9500 Å for the CCDs). Actual scans of the Kitt Peak “old Mould” filters out to 12.000 Å were kindly made for me by George Jacoby and Ed Carder, and I have taken rough curves for the reflectivity of aluminium and for the atmospheric opacity from Allen’s *Astrophysical Quantities* (3rd edition, The Athlone Press, London). Fig. 1 shows one example of such a simulation: the effective spectral response for a dome flat obtained through a “Mould” *B* filter. I start with the Planck curve for a 2300 K black body which, according to data kindly provided by Bill Schoening, is the temperature of the Kitt Peak 1-0.9m telescope’s flat-field lamp when operated at 7 volts – the voltage setting we used in 1984. This is multiplied by the throughput of the filter itself (shown in the second panel from the top) and by the response of the CCD, including the effects of two reflections off aluminium (which are unimportant). Note that the vertical axes in Fig. 1 represent the logarithm of the number of photons per unit wavelength interval.

In 1984 I had asked the Kitt Peak staff about the possibility of red leaks in the “Mould” filters, and was told that “the red leaks are negligible – less than 1% out to one micron.” You can see from Fig. 1 that this statement was completely accurate: in fact the red leak is less than about 0.2% out to one micron. What they and I had both forgotten, however, is that when dealing with light sources as cold as the flat-field lamp, the relative flux at *B* is also less than 1%. The very red color of the lamp has two effects on the spectral response of the system. First, the short-wavelength side of the filter’s principal passband is eroded by the rapid falloff in the lamp’s flux. Second, and much more important, a long red-leak region is raised up. True, the red leak never gets much above 2%, but on the other hand it may span as much as 5000 Å in wavelength. Numerical integration indicates that 10% of the light getting detected comes from wavelengths longer than the nominal cutoff of the filter at 5300 Å; this contaminating light has an effective wavelength near 10.000 Å – quite different from the 4500 Å light the filter is supposed to be letting through. Since the spatial character of the detector’s sensitivity is a strong function of wavelength (cf. data given above for the IFA/Galileo CCD), this means that our *B* dome flat will have superimposed on it a quite different pattern corresponding to the longer wavelengths.

I should point out one particular source of uncertainty in this simulation: the published sensitivity curves for the Kitt Peak RCA CCDs only extend to 9500 Å, and I have been unable to locate a quantitative description of how the detector response falls off longward of that. Apparently, the infrared sensitivity differs strongly from chip to chip, and is also a strong function of temperature. For this simulation I have assumed that the curve drops exponentially toward longer wavelengths with the last known slope, but the real situation with this CCD or with any other could be better or, of course, worse.

In the interval between our observing runs in 1984 and 1985, Kitt Peak introduced the use of “color-balance” filters to modify the spectral characteristics of the lamps employed for dome flats. These are commercial filters designed to let photographers use ordinary daylight-balanced color film in situations where the light comes from incandescent lamps. A tracing of one of these filters, once again kindly provided by George Jacoby, Ed Carder, and Bill Schoening, is shown as Fig. 2. The filter progressively attenuates longer-wavelength light throughout the visible region of the spectrum – just what is needed to make orange artificial light look as yellow as sunlight. However, longward of about 7000 Å – a region that most photographers don’t care about – the color-balance filter also has a red leak. In Fig. 3 I compare the spectral sensitivity of our 1984 dome flats (without color-balance filter;

denoted “u” for “unbalanced”) to that of our 1985 dome flats (with balance filter; denoted “b”). I also show here comparable simulations of the response of the system to the light of a Solar-type star (a 5700 K black body, including the effects of terrestrial atmospheric extinction; denoted “s” for “star”), and to the light of the twilight sky (modelled by starting with a 5700 K black body, attenuating this by the atmospheric absorption and *multiplying* it by the wavelength-dependence of the atmospheric scattering; denoted “t” for “twilight”).

Part of the improvement in going from the unbalanced dome lamp to the color-balanced one results from the fact that, with the balance filter in place, we could increase the voltage on the lamp from 7v to 10v – thus increasing its temperature from ~ 2300 K to ~ 2800 K – without getting exposure times that were too short. The higher temperature increases the overall ratio of *B*-to-*I* flux and reduces the erosion of the short-wavelength side of the filter’s principal passband. The rest of the improvement comes from the fact that, even with its own red leak, the color balance filter still reduces the throughput in the spectral region 6000 Å to 12.000 Å by a factor ranging from 4 to 20. Fig. 3 shows that the spectral response of the balanced *B* dome flat is a much better match to that of a typical star than is the unbalanced dome flat. Is it a better match to the star than the twilight-sky flat is? For wavelengths shortward of about 8800 Å the answer is “yes” – in spectral-energy distribution the color-balanced dome lamp matches starlight better than the twilight sky does. Longward of about 8800 Å the color-balanced dome lamp may still *look* like a better match, but remember: the plot is on a logarithmic vertical scale. In reality, 10^{-4} is more than nine times closer to 10^{-6} than it is to 10^{-3} . This means that if the CCD’s sensitivity continues to fall off exponentially (or slower) to a point *much* longward of 8800 Å, twilight-sky flats might be more reliable even than color-balanced dome flats. (This conclusion is based solely on the spectral properties of the incident light, and neglects the possibility that the dome flats might not illuminate the detector uniformly – another reason why twilight flats might be better).

Similar comparisons for the “Mould” interference-filter *V*, *R*, and *I* passbands are shown in Figs. 4-6. We see that the advantage of using the color-balance filter quickly disappears as redder passbands are considered and, indeed, for *R* and *I* the color-balance filter does more harm than good. If we refer again to Fig. 2, we can see immediately why this happens: the color-balance filter suppresses the *R* and *I* flux even more than it does the red leak.

What are the consequences of using an incorrect flat-field frame? Since the flat-field correction merely scales one part of the image relative to another, and since the absolute throughput of the system will be determined by observations of standard stars made with the same device (and rectified by the same flat-field frame), using the wrong flat field should not introduce a net photometric error. What will happen is that the magnitude derived for a star will depend upon where that star fell on the detector. Thus, there will be additional scatter introduced into the photometry when stars from various places on the detector are measured. Of course, if the observer always places the standard stars at the center of the field, and then reduces program stars that come from all parts of the detector (such as in star-cluster observations), a systematic error can be introduced.

Here are my conclusions and recommendations: (1) *Know your filters and detector.* Don’t be satisfied with statements like, “The red leak is negligible,” or, “The detector dies beyond one micron.” Demand to know how negligible is negligible and how dead is dead. You’ll almost certainly find, as we have, that the people responsible for the instrument will be helpful and cooperative, but they are invariably busy and may only do extraordinary tests when they are specifically requested. (2) *Color-balance filters (for dome flats) may*

do more harm than good. Since 1986 Kitt Peak has introduced a new set of “Mould” filters with red leaks apparently less than 0.01% out to 12.000 Å. With these, the color-balance filters will probably do more good than harm, but with the previous filters the opposite may have been the case. In which situation are you? You probably won’t know until you’ve followed recommendation 1. (3) *Illuminate your dome flats with diffuse sunlight.* Nature has provided us with a 5700 K light bulb, with a spectral-energy distribution remarkably like that of a typical star. It has lots of flux for those troublesome U and narrow-band filters, too. Why not open the dome slit a few inches and let some of that free light illuminate your dome screen? (Turn the slit away from direct sunlight and check the screen for shadows before starting. Exposure levels will vary on days with scattered clouds, but this may not be a serious problem if enough frames are taken.) (4) *If you get a cloudy night, rejoice!* You can spend the time taking every sort of calibration frame you can think of: dome flats with balance filters, dome flats without balance filters, dome flats to low and high exposure levels – the works! If you gain a deeper understanding of your instrument, and can therefore reduce your systematic errors by 30%, that’s as good as a second clear night!

2 Shutter-timing errors

When you start a CCD integration, the instrument-control computer sends a signal to the camera head instructing it to open its mechanical shutter. Some time later the shutter begins to open, and at some still later time it is completely open. After the prearranged exposure time has elapsed, the computer sends another signal to the camera telling it to close the shutter; after some time the shutter actually begins to close, and eventually it is completely closed. These various delays are typically of order many milliseconds, and their net effect is to produce a difference, δ , between the presumed and actual exposure times. The shutter timing error, δ , may be either positive or negative. Since it is produced by events occurring just at the beginning and the end of the exposure, and since the computer’s own clock should be accurate to microseconds, at worst, δ should be nearly independent of total integration time. This shutter-timing error could become significant if sub-percent accuracy is desired for bright stars requiring exposure times of order one second. In principal, δ could be measured by the instrument scientist and then allowed for in the camera-control program, but since different shutters with different time-constants may be controlled by the same computer, this is often not done.

One way to check for shutter timing errors is to take dome flat-field exposures with different integration times, say, one and ten seconds. (Since we will use only the ratio of these frames, it does not matter whether the frame is absolutely uniformly illuminated, or whether there is a red leak.) In order to improve the signal-to-noise ratio of the short exposures and to avoid any non-linearity in the detector, I use the focussing sequence which allows me to open and close the shutter many times before reading out the chip. Thus, I make ten one-second exposures before reading out, and compare this frame with a single ten-second exposure. If many 10×1^s and 1×10^s exposures are taken alternately, any fluctuations in the dome lamp’s brightness can be averaged out. Repeat determinations on different days check the reliability of the results.

Define R as the mean brightness of a 10×1^s exposure divided by the mean brightness of a 1×10^s exposure. Then

$$R = \frac{10 \cdot (1 + \delta)}{10 + \delta} \quad \Rightarrow \quad \delta = \frac{10 \cdot (R - 1)}{10 - R}$$

During 1988 June, I determined the shutter errors of the RCA5 detector on CTIO's 4m telescope. On two consecutive days I got values of $\delta = 12 \pm 1$ msec and 9 ± 1 msec. The agreement isn't perfect, but it's adequate to correct my shortest exposure times to better than 0.1%.

However, there is no reason to stop here. In fact, these mechanical shutters tend to open in one place first and to close in the same place last: with iris shutters, for instance, the center of the field tends to get longer exposures than the corners. By taking the ratios of the $10 \times 1''$ and $1 \times 10''$ images, rather than of their mean intensities, we can create a two-dimensional map $\delta(x, y)$:

$$\delta(x, y) = \frac{10 \cdot (R(x, y) - 1)}{10 - R(x, y)}$$

An example of such a map derived from the same 1988 June CTIO data is shown in Fig. 7, where the contour levels represent loci of constant δ , with the values indicated in milliseconds. One way to use this map would be to enter it with the known (x, y) coordinates of some star, and read out from it a shutter correction for that position. There is a still cleverer way to use it, however. Say we have some frame which has a nominal exposure time of two seconds. Take our two-dimensional map of δ and add a constant value of 2; now we have a map of $(2 + \delta)$. Divide our program image by this map and multiply it by 2. Now we have a flat, uniform exposure time of exactly two seconds everywhere in the frame. This correction can easily be applied when the raw data are flat-fielded, and can be ignored thenceforth.

3 Mean airmass for long integrations

In many cases, especially for data taken at Cassegrain foci, integration times can be very long – up to an hour or more. The program field is moving across the sky this whole time, so what should be used as the effective air mass to correct the observation for atmospheric extinction? The simplest way out would be simply to use the air mass at mid-observation; a slightly more sophisticated approach would be to average the initial and final air masses. But, as Jeff Pier has pointed out (1982, *A.J.* 87, 1515), neither of these is satisfactory. Consider the case of an exposure which passes through the meridian at mid-observation: the air mass at mid-observation is the *minimum* air mass for the entire observation, and the average of the initial and final air masses is the *maximum* air mass for the observation – clearly neither of these represents a sensible average. What we really want is the mean air mass that correctly predicts the total number of photons accumulated by the detector during the course of the exposure. In other words, recognizing that the instantaneous apparent magnitude of a star, v , is related to its extra-atmospheric magnitude, V , and to the instantaneous value of the air mass, X , by

$$v = V + K \cdot X,$$

we want to determine the mean air mass \bar{X} such that

$$(t_1 - t_0) \cdot 10^{-0.4K\bar{X}} = \int_{t_0}^{t_1} 10^{-0.4KX(t)} dt$$

This is a very unpleasant equation, because it indicates that the effective mean air mass of the observation depends upon the extinction coefficient, K , which is an unknown. In

fact, it is in order to determine K that we want to compute the air mass in the first place. Furthermore, the effective extinction coefficient depends somewhat on the color of the star being studied, $K \approx k_0 + k_1 (B - V)$, which means that in principle we may have to compute different mean air masses for stars of different colors recorded in the same frame!

However, not all is lost. For mathematical simplicity, let's convert this to powers of e rather than powers of 10. Also, we recognize that \bar{X} is a constant, even if it is unknown, so we can bring it to the right side of the equation and into the integral:

$$1 = \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} e^{-0.921 K (X(t) - \bar{X})} dt.$$

If we approximate this integral using Simpson's 1/3-rule, we obtain

$$1 = \frac{1}{6} \left[e^{-0.921 K (X_0 - \bar{X})} + 4e^{-0.921 K (X_{1/2} - \bar{X})} + e^{-0.921 K (X_1 - \bar{X})} \right],$$

where X_0 , $X_{1/2}$, and X_1 are the instantaneous air masses at the beginning, the middle, and the end of the observation, respectively. Simpson's rule is very good; the error term is $\left[\frac{(t_1 - t_0)^4}{2880} \cdot \frac{\partial^4 X}{\partial t^4} \right]$. Now, taking a first-order Taylor expansion of this we get

$$1 = \frac{1}{6} \left\{ 1 - 0.921 K (X_0 - \bar{X}) + 4 \left[1 - 0.921 K (X_{1/2} - \bar{X}) \right] + 1 - 0.921 K (X_1 - \bar{X}) + 0 \left(0.4 K^2 \overline{(X - \bar{X})^2} \right) \right\}$$

Note that $\frac{1+4+1}{6} = 1$, which cancels the 1 on the left side, and then we can divide through by 0.921 K , ultimately arriving at

$$\bar{X} = \frac{X_0 + 4X_{1/2} + X_1}{6} + 0 \left(0.08 K \overline{(X - \bar{X})^2} \right)$$

The first part of this equation, up to the plus sign, is our adopted approximation for the mean air mass of the exposure. It is very conveniently *independent* of the extinction coefficient, K . The error term at the far right is pleasantly small: for reasonable values of $K \sim 0.25$ magnitude and $X_{max} - X_{min} < 0.25$, the error is of order 10^{-4} .




Figure 1




Figure 2

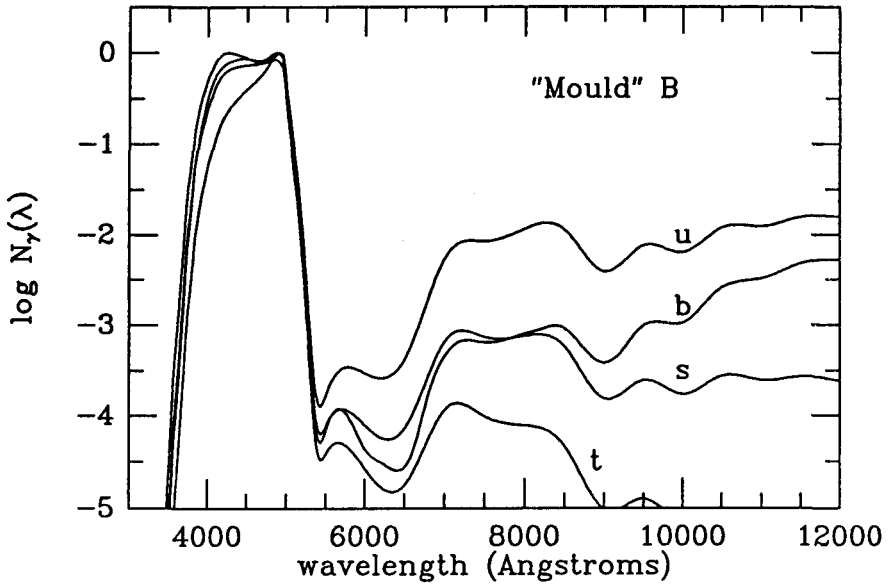


Figure 3

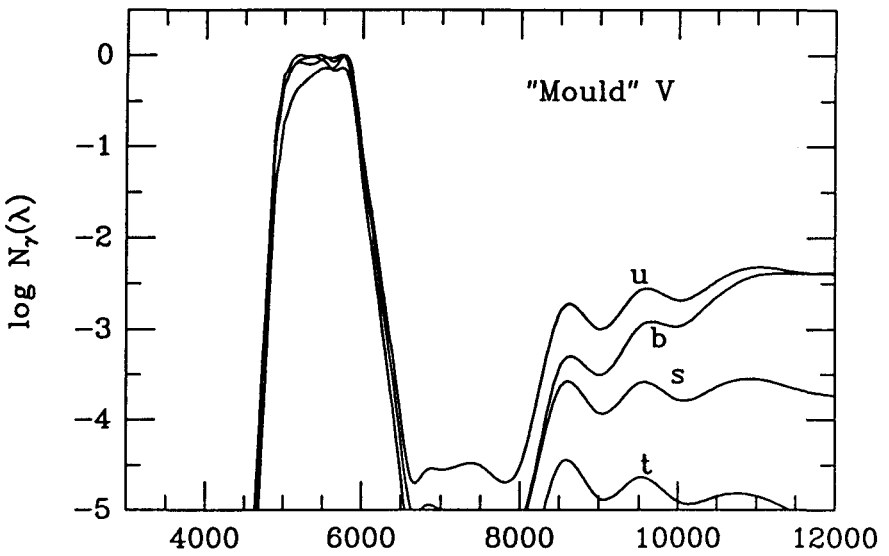


Figure 4

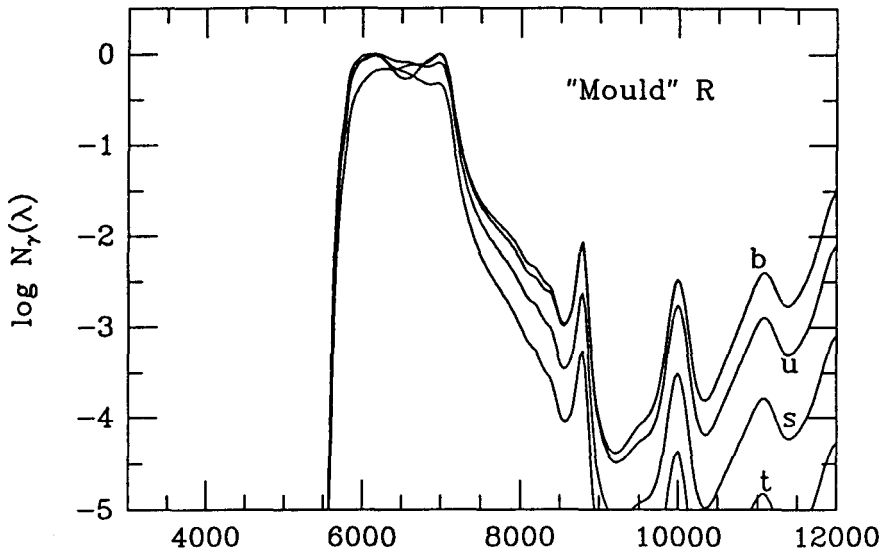


Figure 5

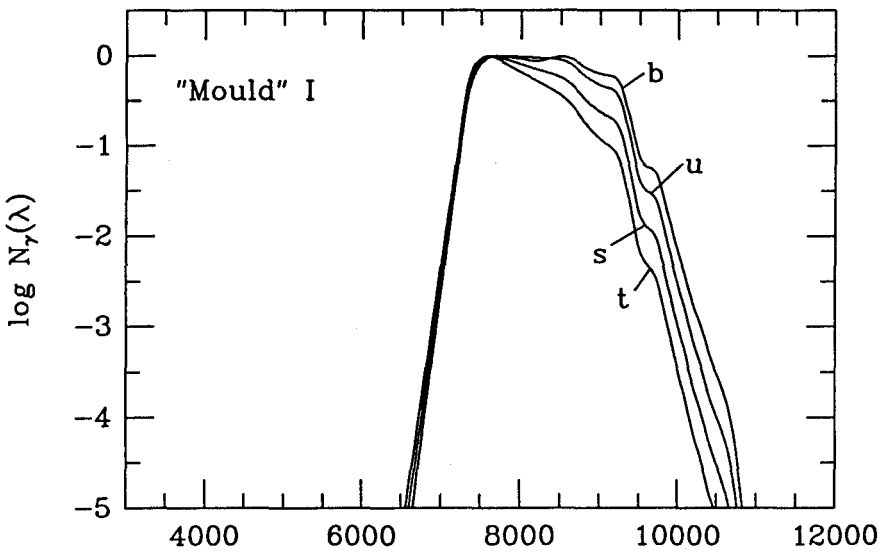


Figure 6




Figure 7