

Further Progress in CCD Photometry

Peter B. Stetson¹

¹*Dominion Astrophysical Observatory, Canada*

Abstract

I comment on some of the steps required to convert raw instrumental magnitudes, derived by profile-fitting or synthetic-aperture photometry from CCD images, to final calibrated photometry on a standard system. The status of the DAO program to obtain homogeneous *BV* photometry for star clusters and nearby galaxies will also be discussed.

1. Introduction

Photoelectric photometry is easy. If you want to measure some stars photoelectrically, take your list of coordinates and some finding charts to the dome. Point the telescope, identify the program star, center it in the little hole, and write the star's name down on the strip of paper that comes out of the printer. Integrate for a while through each filter, then move the star out of the hole and integrate on the sky. Reobserve that star a few more times if you want, or move on to the next star in your program. If you are fast you can observe a hundred stars or more in a night. Later, divide the star and sky counts by the integration times, subtract the skies from the star observations, and -2.5 times the logarithm of what's left is your instrumental magnitude; the difference between magnitudes in different filters is an instrumental color. You can do it with a piece of paper and a slide rule, although a pocket calculator is even better. Compare your instrumental magnitudes and colors for some stars observed at several airmasses and derive an extinction correction. Then take your extinction-corrected instrumental magnitudes and colors for some standard stars, compare these to the published magnitudes and colors, and derive transformation equations. Apply these to your program objects. Publish the results.

If you want to measure a lot of stars in a cluster you can do it photographically, which is even easier. At the telescope take several plates of your field. Once your plates are developed and dry, make a large-scale print of the best one. On this, next to the image of each star write a little number. Then take your plates into a darkened room, and place them one by one on the stage of an iris photometer. Consulting your finding chart, steer the machine to each star, center it up, turn a knob to make a needle go to zero, and write down the star's ID number and the number from the readout on the front of the machine onto a piece of paper. Later, plot up machine index *vs* published photometric magnitude for a sequence of stars that have been observed photoelectrically, sketch or compute a smooth curve through those points, and run the machine indices through this curve to provide magnitudes on the standard system. Publish the results.

Nowadays, we have the charge-coupled device (CCD). This remarkable chip combines the advantages of the photomultiplier and the photographic plate: it is linear and repeatable, so a photometric system can be accurately transferred from standard stars in one part of the sky to program stars elsewhere; it is an area detector, so many stars can be recorded with each integration. Observing with a CCD is not much more difficult than observing with a photomultiplier or plates: there may be a few more calibration observations to be made around sunset and sunrise, but on the other hand you don't get eyestrain trying to decide whether some faint puff of light is centered in the aperture, and you *certainly* don't have to slop around in chemicals next morning.

No, the price you pay for enjoying the benefits of the CCD is not in the observing, it's in the analysing. Not only is the analysis more complex with CCDs than with photomultipliers and plates, there are also immensely larger volumes of data. No longer are we talking about one measured datum per star per plate for a few hundred stars in a cluster, or a half dozen data per star for a hundred or so stars observed in a night with a photomultiplier. These days it's possible to obtain a hundred CCD frames in a night, each frame potentially containing of order 10^4 stars and of order 10^6 individual data-numbers. The process of converting these immense mounds of raw data into a publication comes in three basic stages, which may be called "preprocessing", "processing", and "postprocessing."

Preprocessing includes bias-subtraction and flat-fielding, and does not decrease the volume of data from an observing run; what it does do is to ensure that the data are on a consistent photometric zero-point and scale. Dozens of manuals and guidebooks and papers (including several at this meeting) have been written on how to preprocess your images. Perform these steps properly, and you now have images that contain accurate, linear representations of the brightness of objects in your program fields.

Processing — converting the rectified CCD intensity arrays into lists of positions and magnitude indices for the stars contained therein — is another matter. If your CCD frames contain few stars, it is often possible to simulate photomultiplier photometry by simply defining some geometric region as an "aperture": sum up the photons detected within an "aperture" centered on a star, sum up the photons detected within a star-free region of the same area elsewhere in the frame, and call the difference between these two numbers the signal from the star alone. Readout noise in the CCD and the flux contributed by the diffuse light of the nighttime sky make this method very imprecise for faint stars, and for crowded fields, as in star clusters, it may not be possible to define an aperture which is large enough to contain all of one star's flux while entirely excluding the flux of neighbors. For dealing with faint stars against noisy backgrounds and stars in crowded regions, model-profile fitting offers the best possible recovery of photometric information for stellar objects. A half-dozen or so profile-fitting packages written by astronomers originally for their own use have become fairly widely distributed in the field. Since the problem is complex, the various programs are quite individualistic, and adopt an interesting variety of

approaches to each of the hurdles that must be overcome (see, *e.g.*, Stetson 1990a,b, 1991a, 1992a for previous discussions and references). The bottom line, however, is that each package reduces the two-dimensional data array representing an input CCD frame to an output list of presumed *bona fide* detections, their positions, and their apparent brightnesses on an instrumental system. At this point, my story today begins.

2. The many and various steps of postprocessing

To begin, let us assume each astronomical image obtained during our CCD run has been reduced to a list of stars: their positions within the coordinate frame of that image and their apparent instrumental magnitudes relative to some zero-point for that image. How do we get from here to the part about publishing the results? I illustrate the procedures required by describing a typical post-processing reduction path, using software with which I am personally familiar.

a. Aperture corrections

First, we note that there may not be a consistent photometric zero-point from frame to frame. While we expect that the CCD is inherently a consistent, stable, linear detector of photons, our measuring process may have mislaid some of that consistency. In profile-fitting photometry, the instrumental magnitude for a star comes from the height of a model point-spread function (“PSF”) for that frame scaled to the intensity values recorded inside the star’s image. Since the instrumental magnitudes for each frame are defined relative to a *different* PSF, it is difficult to be certain at the sub-percent level that the zero-points for all frames are identical. One solution to this problem has been discussed at length elsewhere (Howell 1989; Stetson 1990a,b), so I will summarize it here in a few words. Synthetic-aperture photometry is performed for several selected, bright, isolated stars in each CCD frame; if necessary, the profile-fitting results are used to subtract unwanted stars from the frame beforehand, to prevent their contaminating the aperture photometry. Either the selected stars are measured through an aperture large enough to contain a constant fraction of their total flux, independent of seeing and guiding differences from frame to frame, or they are measured through a series of apertures of increasing size, and these results are extrapolated on a frame-by-frame basis to an aperture large enough to contain a consistent fraction of their flux. We then rely on the stability of the CCD to provide a consistent relationship between the total number of photons detected from a star, and its true brightness as perceived through the atmosphere, telescope, and a filter, just as one does with a mechanical aperture and a photomultiplier. The integrated total instrumental magnitudes for several stars in each frame are stored in a so-called “total magnitude” file.

b. Cross-identification

Next, since we will eventually want to compare repeat measurements of a given star, whether to beat down noise, to measure variability, or to determine colors from observations through different filters, we must cross-identify re-observations of

the same star in different frames. The telescope may have been shifted between successive exposures, and we may want to combine data from different cameras or different telescopes, so we can not expect a given star to reappear at the same position in every image. While a human being armed with a finding chart *could* display each CCD frame and note the position of every star of interest, this could take forever. We want a computer program which, when presented with several lists of positions and magnitudes, can intercompare the lists, recognize prominent asterisms, cross-identify the corresponding stars, and estimate the geometric transformation equations that interrelate the coordinate systems of the various frames. The program must be able to cope with arbitrary translations, rotations, scale changes, and flips, and it must be effective when the overlap between the two lists is incomplete, as when some stars in one frame lie outside the area or beyond the magnitude limit of another.

See Murtagh 1992 for an extensive catalog of automatic pattern-matching algorithms which have been implemented, but let me describe mine briefly here. “DAO-MATCH” attempts to recognize triangles of stars, exploiting the fact that while their positions, orientations, and sizes may change, triangles won’t change *shape* as they are translated, rotated, and scaled. It considers brightest stars first, since these are the most consistently detectable objects in the field. For each target field, the astronomer selects one frame to serve as the “master” frame — normally this will be the deepest one, or among a set of partially overlapping images it may be the one closest to the center of the pattern. The program reads the star list from the master frame and sorts the stars by apparent instrumental magnitude; it then does the same for the next frame of that field. Taking the brightest three stars from each list, it encodes the shape of each triangle in two numbers: the ratio of the length of the second-longest side to the longest side, b/a , and of the shortest side to the longest side, c/a . These ratios are clearly independent of translations, rotations, and scale changes, and each triangle defines a unique point in two-dimensional $(b/a, c/a)$ -space. If the two points in $(b/a, c/a)$ -space coincide within a certain tolerance, the three stars may be provisionally cross-identified: the star opposite the longest side in frame 1 corresponds to the star opposite the longest side in frame 2, and so on. Transformations are estimated from the three presumed cross-identifications and are presented to the user for approval. If the scale and rotation factors look about right (*e.g.*, scale ≈ 1 , rotation ≈ 0 for frames taken with the same equipment), the user may accept these transformations and move on to the next frame. If, however, the first two triangles do not have the same shape, or if the user wants transformations based on more than three provisional cross-identifications, the program will take the fourth brightest star from each list and compute the three new triangles formed by this star with each pair of the previous three. If any of these have consistent shapes, new cross-identifications are made and the implied transformations are offered to the user. If these are still unacceptable, the fifth brightest star in each frame and six new triangles are taken under consideration. When the user is satisfied that the cross-identifications are correct and the geometric transformations are adequate, the coefficients of the transformations are written to a disk file (a “match file”). The user

then enters the name of the star list for the next image of that field, and the process is repeated. In my experience, this routine works reliably and without human help provided the degree of overlap between the input lists is greater than about 25%: if at least one-fourth of the stars which appear in list 1 also appear in list 2, and *vice versa*, the correct transformation constants will be found in a matter of seconds.

The number of possible triangles can be huge ($O(n^3)$), so it is not possible in general to cross-identify *every* star in the field in this way. Therefore another program has been written ("DAOMASTER"), which accepts the approximate transformation equations provided by DAOMATCH, reads in all the star lists for that field, and cross-matches *all* stars by spatial proximity: if after transformation to the coordinate system of the master frame, a star lies within a specified distance of a star in the master list, it is provisionally identified with that star; if it lies near *no* star in the master list it is added to the list as a possible new detection. As subsequent lists are considered, their stars may be identified with ones in the original master list or with stars added from other frames, or may themselves be added to the master list if no such identification is possible. After all possible cross-identifications have been made, the augmented master list is cleaned of "insignificant" detections, on the basis of user-supplied acceptance criteria. Then the current set of acceptable cross-identifications is used to refine the geometric transformation constants, and all available positional determinations for each star are averaged to provide a New, Improved! master list. The process is repeated with an increasingly stringent tolerance for positional agreement: at the beginning a rather lax agreement criterion is used because some allowance must be made for positional mis-match caused by errors in the provisional transformations provided by DAOMATCH; as the transformations are refined, a more strict criterion is used to minimize the number of false cross-identifications. Several passes of increasing strictness can be performed on several dozen star lists for a given field in a few minutes.

Once the user is satisfied that the geometric transformations are accurate and the master list is as complete as possible, DAOMASTER offers several types of output, such as: a copy of the complete master list with average positions and magnitudes on the system of the master frame (a "master file"); a file containing the new, more accurate transformation equations (a new "match file"); and a simple transfer table of cross-identifications (a "transfer file"): star 1 in input list 1 is the same as star 37 in list 2, it doesn't appear in list 3, it is star 196 in list 4, ... The output files can serve a wide variety of purposes. For instance, the master file contains an index of the frame-to-frame scatter in a star's magnitudes, which may be used to identify variables if all frames were taken through the same filter. The new match file also contains an estimate of the magnitude zero-point differences among the frames; taken together, the geometric transformations and photometric offsets allow the user to combine partially overlapping CCD frames in a single montaged image covering a larger area of sky (see, *e.g.*, Stetson & Harris 1988, Figs. 1–6.) The transfer file will be used later in the final photometric reductions.

c. Real-world identifications

At this point each star is still only a few entries in some lists. Although we can now identify each star of interest among all frames of a given field, we have yet to associate particular entries in the master list with real-world names, such as the stars' identifications in published standard-star lists or in previous photometric investigations of the cluster. This may be done with a routine called `FETCH`, which was once a FORTRAN program interacting with the DAO's VICOM image display; it has since been reincarnated as an IRAF script. The user displays the master frame or a montaged image of the target field. Then, having placed the cursor on a star, the user enters the star's real-world name; the routine then searches the master file for the star lying closest to the cursor's position. A "fetch file" is created which contains the stars' real-world names and the corresponding entry in the master list — in particular, the stars' (x, y) -coordinates in the system of the field's master frame. This user-interactive process needs to be done only once per field.

d. Preparing for the photometric transformations

We now have almost everything we need to derive the photometric transformations between our instrumental system and the standard system: we can identify a given star's real-world name with a particular (x, y) position in the master system for its field; either through the geometric transformations derived in §b above or using the transfer file from §c we can then find a named star in *any* frame of that field. Now we must provide a file containing the ancillary data for each exposure: the filter, the time of mid-exposure, the airmass, and the exposure time. This is not difficult. Now we are ready to collect the observed data for named stars from a night's worth of observing into a single file.

It is easiest for me to describe the process by following the series of steps taken by my computer program, "COLLECT." The user begins by specifying the name of the file with the ancillary data (filter, airmass, exposure time, etc.) for all the CCD frames of the night. Then, frame by frame, the user enters the name of the file with the large-aperture photometry for selected stars in that frame, which was obtained as described in §a above. It reads in both these data and the profile-fitting magnitudes for *all* stars the same frame from the output of the profile-fitting routine. It identifies the stars with aperture photometry among the profile-fitting results by positional coincidence, and uses the comparison of these data to determine the additive correction required to place the profile-fitting magnitudes onto the system of the aperture magnitudes; this allows a consistent comparison with data from other exposures. Next the user enters the name of the fetch and match files for the field. Named star by named star, the program takes the star's master-frame position from the fetch file, uses the geometric transformations to determine the star's position in *this* frame, and searches for the star among the input lists of aperture and profile-fitting photometry. If the named star has both aperture and profile-fitting results, the program considers the standard error of each: the large-aperture magnitude may be uncertain because of read noise and sky noise in the large aperture, or because it has been extrapolated from a much smaller aperture by means of a poorly defined growth curve; the profile-

fitting magnitude may be uncertain because the correction from the profile-fitting zero-point to the large-aperture zero-point may be badly determined for that frame. Whichever of the two total-magnitude estimates is less uncertain is written to the output file, which I call an “observation file,” along with the star’s real-world name, and the filter, time, airmass, and exposure time of the observation. If no concentric-aperture photometry is available for that particular star, then it is the profile-fitting magnitude corrected to the system of the large-aperture photometry that is written out. The process of building up the observation file proceeds as fast as the user can type a file name followed by several carriage returns; a night’s data for stars with real-world names can be assembled in around a half hour.

e. Deriving the transformations

I have described my transformation procedure elsewhere (*e.g.*, Stetson & Harris 1988, Stetson 1992b), so here I will provide only a brief summary. The observed magnitudes on the (large-aperture) instrumental system are fitted by robust least-squares to transformation equations consisting of terms involving the standard-system magnitudes and colors of the standard stars, and the airmasses and times of observation of the frames. For instance, for *BV* photometry the equations may be of the form

$$\begin{aligned}v_{ij} &= V_j + F_1(V_j, (B - V)_j, X_i, T_i) \\ b_{ij} &= B_j + F_2(V_j, (B - V)_j, X_i, T_i)\end{aligned}$$

for a measurement of standard star *j* in frame *i*, where maybe

$$F_1 = A_0 + A_1 \cdot (B - V) + A_2 \cdot (B - V)^2 + A_3 \cdot X + A_4 \cdot (B - V) \cdot X + A_5 \cdot T \cdot X + \dots$$

I want to stress that the *observed, instrumental magnitudes* are fitted by functions of the *standard photometric indices* and the time and airmass of the observation — no attempt is made to derive or employ instrumental colors. Not only is this proper least squares, since the errors of the fit are dominated by uncertainty in the observed instrumental magnitudes, but it is also more convenient: unlike the case with photoelectric photometers, where it is easy to cycle quickly through the filters (*e.g.*, *VVBBBBVV*), so each observation can be made symmetric about its midpoint and both filters therefore have the same effective airmass, it is difficult to define consistent instrumental colors with a CCD. The airmasses of the *B* and *V* exposures of a given frame pair will necessarily be somewhat different, and the difference between the two will change depending upon the declination, hour angle, and duration of each exposure. The *standard-system* magnitude and colors are constant attributes of a (non-variable) star, and the transformation equations can be consistently expressed in terms of them.

The equations above are suitable for photometric nights, where the throughput of the system depends only upon the color of the star and the airmass of the observation (with perhaps an allowance for a secular time-variation of the atmospheric

transparency — the $X \cdot T$ terms). On cirrusy nights, the transparency may fluctuate unpredictably from frame to frame. We nevertheless expect that the effective exposure for different stars *in the same frame* will be the same, and we may define analogous transformation equations: *e.g.*,

$$F'_1 = A_{0,i} + A_1 \cdot (B - V) + A_2 \cdot (B - V)^2 + A_4 \cdot (B - V) \cdot X + \dots$$

Here, $A_{0,i}$ is not a constant applying to all frames for the night, but is rather an instantaneous photometric zero-point which applies to the i -th frame alone. The other constants A_k have precisely the same meaning as in the clear-sky reductions. Provided at least some of the frames taken during the night each contain several photometric standards with a range of colors, the color-transformation terms can be computed; once they have been established, any frame which contains even one photometric standard can have its individual zero-point determined.

My program to perform these solutions (“CCDSTD”) requires a file containing a standard-star library, where each star is identified by its real-world name accompanied by its published photometry on the standard system, a short “transformation file” expressing the form of the desired transformation equations in a simple code, and the observation file created in §d above. It then computes the least-squares transformation equations, whether in the clear-sky or cloudy-sky mode. The encoded transformation equations and the derived values of the transformation coefficients are written out to a “calibration file.”

f. Applying the calibrations

Once the photometric calibrations for a night, or an observing run, or a series of runs have been derived, they can be applied to program stars in an inverse sense, to infer their values of V and $B - V$ (in the current example) from their observed values of v and b . This is done by Newtonian approximation:

$$\begin{aligned} V_j &= v_{ij} - F_1^{(j)}(V_j, (B - V)_j, X_i, T_i) \\ B_j &= b_{ij} - F_2^{(j)}(V_j, (B - V)_j, X_i, T_i) \end{aligned}$$

This is another reason why the transformation equations are expressed in terms of *standard-system* magnitudes and colors: although we do not yet know these values for our program stars, we can utilize *all* the v and b observations we ever made of the star — from different frames, different nights, different runs, different telescopes, whatever — to determine them. Each observation has a corresponding set of calibration constants A so, having assembled all the observations, the program merely postulates a crude guess at the standard-system values of V and $B - V$ for star j , and computes new, least-squares estimates of V and B from the observed residuals $\langle v - F_1 \rangle$ and $\langle b - F_2 \rangle$. The new values are fed back into the right sides of the equations and the process is iterated to convergence, which happens very quickly. By the end

of the process, each observation has been calibrated with the *best possible* estimate of the star's true magnitude and colors. In contrast, had we tried to define the transformation equations in terms of the *instrumental* colors (e.g., $V = v + f_1(v, b-v, X, T)$) then one night's observation of the star would have been calibrated using that night's value of $b-v$ (which might be ill-determined), another night's observation would be calibrated using a different, comparably poor, value of $b-v$, and data from a night on which the star was accidentally observed in only one filter would have to be discarded. Not so with the present scheme.

There are two situations to which this method may be applied: (1) to stars with real-world names which appear in the observation file created in §d above — these could be standard stars whose magnitudes and colors you want to redetermine from your own observations, or stars with names of your own devising without prior photometry; or (2) stars in some cluster or other field, most of which have no names yet, but which exist as a set of cross-references obtained by DAOMASTER (§b). In essence, these situations may be thought of as photomultiplier-like and photographic-like reduction modes. The approaches taken in the two cases differ.

In the first situation the user simply runs a program named "CCDAVE" and types in the names of all the observation files to be included in the analysis. The program reads in the observed data and the corresponding transformation constants from the appropriate calibration files. It then assembles all the available data for each star — identified by real-world name — and performs the inverse transformation by Newtonian approximation as described above. The program produces a file with the final, averaged photometry and its standard errors, as well as an index of the goodness-of-agreement which can be used to help spot variable stars.

The second situation is slightly more complex. A program named "FINAL" accepts the transfer file generated by DAOMASTER, which contains the names of the profile-fitting photometry files for all the CCD frames of the field, and the transfer table of cross-references for all observations of stars in the final master list. It then accepts the final averaged photometry for named stars produced by CCDAVE above. Next the user specifies the name of the fetch file for the field; from this the program learns the real-world names of local standards in the field, and recovers their standard-system magnitudes and colors from the averaged-photometry file. It also uses these stars' positions as recorded in the fetch file plus the cross-references provided by the transfer file to recover each star's instrumental, profile-fitting magnitude from each CCD frame. Finally, it reads in the photometric transformation constants for each frame from the calibration file for the night on which the frame was obtained. Now we are ready to fly.

First, FINAL redetermines the photometric zero-point of each frame from the local named stars with standard-system photometry generated from CCDAVE above: for each frame small correction to A_0 is estimated from

$$\hat{A}_0 = \langle v - V - F_1 \rangle$$

The reason for this step is precision. There are any number of reasons why any given

CCD frame may depart systematically from the optimum photometric calibration for its night by a few hundredths or a few thousandths of a magnitude: uncertainty in the growth-curve correction, a contrail drifting past the telescope during the exposure, a touch of mist on the telescope mirror, . . . Such frame-to-frame jitter in the photometric zero-point can introduce extra scatter in a cluster's color-magnitude diagram, since not all stars will necessarily be measured in all frames. By recorrecting all frames to a single consistent zero-point defined by a specific sequence of well-measured local standards, one does not place the average cluster photometry on the standard system with any greater accuracy, but one does produce sharper sequences in the color-magnitude diagram.

The rest is easy. The set of cross-references contained in the transfer table allows FINAL to assemble all observations of each star in the master list for the program field, and the photometric calibrations — complete with the revised zero-points — are used to correct these observations to the standard system by Newtonian approximation, as before.

Now you may publish the results.

3. Current status of the DAO cluster-photometry program

Since 1983, a number of us at the DAO have been collecting observations in the *BV* photometric system for open clusters, globular clusters, and nearby galaxies. Contributing observers include Roger Bell, Mike Bolte, Bill Harris, Jim Hesser, Bob McClure, Linda Stryker, Nick Suntzeff, Don Vandenberg, and me. Many of these observations have been reduced — mostly by means rather more crude than described above — and published. But in bits of time scattered between other projects, I have been trying to rereduce these data using the latest software and consistent reduction techniques. As of this moment, early August 1992, data from 16 observing runs have been reduced. These comprise a total of 52 nights (33 photometric, 19 partly cloudy) on five telescopes (CTIO 4m and 0.9m, KPNO 4m and 0.9m, CFHT 3.6m), from which 2,556 useable CCD frames of astronomical objects have resulted.

The central pillar of our standard system is the work of Landolt (1973, 1992), but additional standards have been taken from a number of sources: Graham 1982 (E regions); W. E. Harris, unpublished (M11 cluster); Christian *et al.* 1985 and L. E. Davis, unpublished (M92 and NGC 7006); Heasley and Christian 1986 (M92); Baade and Swope 1961 and Stetson 1979 (Draco); Hawarden 1970 and Anthony-Twarog *et al.* 1979 (Mel 66); Stetson 1981 (NGC 1851); Tift 1963 (field near NGC 121); and Graham 1981 (field near NGC 300) bring the total number of "external" standards which we have used to date to 255, most of them observed by us many times. Through the intermediary of our own observations we have derived zero-point corrections to place the results of these latter studies on the system of Landolt's work — in the mean, at least — after the fashion of Stetson and Harris 1988 and Stetson 1991b.

From our own observations we are also able to define a number of new secondary "standards", which serve two important purposes. First, since they have been re-

observed a number of times on different nights and with different telescopes, adding well-observed program stars to the standard list strengthens the internal homogeneity of our various reduced data sets. Second, the secondary standards in program fields will be the “local standards” used to refine each field’s photometric zero-points in the second mode of applying the transformations as described above. In addition to the 255 primary standards, we have so far been able to define 2,325 secondary standards in a total of 95 CCD-sized fields scattered over the sky. The secondary standards all satisfy the criteria $N(V) \geq 4$, $N(B) \geq 4$, $\sigma(V) \leq 0.02$ mag, $\sigma(B - V) \leq 0.03$ mag, and $\chi < 2$, where χ represents the ratio of the observed, external, frame-to-frame scatter in the results for the star to the scatter expected from the photon statistics, read noise, quality of the profile fits, and quality of the nightly transformation solutions; a large χ value could indicate intrinsic variability or crowding/seeing problems, either of which would disqualify a star as a useful standard.

The work continues, and may even see publication some day.

References:

- Anthony-Twarog, B. J., Twarog, B. A. & McClure, R.D. 1979, *ApJ*, 233, 188
Baade, W. & Swope, H. 1961, *AJ*, 66, 300
Christian, C. A., Adams, M., Barnes, J. V., Butcher, H., Hayes, D. S., Mould, J. R., and Siegel, M. 1985, *PASP*, 97, 363
Graham, J. A. 1981, *PASP*, 93, 29
Graham, J. A. 1982, *PASP*, 94, 244
Hawarden, T. G. 1970, *MNRAS*, 174, 471
Heasley, J. N. & Christian, C. A. 1986, *ApJ*, 307, 738
Howell, S. B. 1989, *PASP*, 101, 616
Landolt, A. U. 1973, *AJ*, 78, 959
Landolt, A. U. 1992, *AJ*, 104, 340
Murtagh, F. 1992, *PASP*, 104, 301
Stetson, P. B. 1979, *AJ*, 84, 1149
Stetson, P. B. 1981, *AJ*, 86, 687
Stetson, P. B. 1990a, in *CCDs in Astronomy II: New Methods and Applications of CCD Technology*, eds. A. G. D. Philip, D. S. Hayes & S. J. Adelman, L. Davis Press, Schenectady, p. 71
Stetson, P. B. 1990b, 102, 932
Stetson, P. B. 1991a, in *3rd ESO/ST-ECF Data Analysis Workshop*, eds. P. J. Grosbøl & R. H. Warmels, European Southern Observatory, Garching, p. 187
Stetson, P. B. 1991b, *AJ*, 102, 589
Stetson, P. B. 1992a, in *Astronomical Data Analysis Software and Systems I*, eds. D. M. Worrall, C. Biemesderfer & J. Barnes, *Astr. Soc. of Pacific Conf. Series*, 25, p. 297
Stetson, P. B. 1992b, *J. Roy. Astron. Soc. Can.*, 86, 71
Stetson, P. B. & Harris, W. E. 1988, *AJ*, 96, 909
Tift, W. G. 1963, *MNRAS*, 125, 199

Discussion

C. Morossi: *How do you take into account shutter problems in your data reduction procedure?*

Stetson: See the *Highlights in Astronomy* article referred to by Alistair Walker. The problem can be largely circumvented by avoiding the shortest exposure times – less than, say, 10 seconds.

M.J. Stift: *In profile-fitting one has to be aware that the shapes of the profiles can depend on colour when using broadband photometry. Already, at airmass 1.2, asymmetric profiles (positive or negative skewness) are possible. Taking the maximum intensity of a profile can lead to systematic errors in magnitude.*

Stetson: I believe that that would be a multiply differential effect. To first order, the profile fits conserve profile volume, so for isolated stars it is a second-order effect, and for faint stars inside the profile of brighter stars it would be a first-order problem. The model profile is derived from several stars in each field, and in some sense it is appropriate to a star of average colour; I think only a star of extreme colour could be seriously affected. Finally, the problem would affect standards just as it does program stars, so it would come out in the colour terms of the calibration.

R.M. Genet: *How much human time is involved in the reduction? Can it be totally automated?*

Stetson: I spend of the order of 10-15 minutes of my time per CCD frame. It is totally automated now, if you don't want to do intermediate visual checks.

A. Walker: *Are programs such as DAOGRON etc. available for the general user in much the same way as DAOPHOT is?*

Stetson: I've never refused a copy to anyone who asked me nicely.

T. von Hippel: *Do you have plans to add astrometry to your package?*

Stetson: People at the U.S. Naval Observatory are extracting astrometric information from DAOPHOT reductions. The process of correcting for scale, rotation, and other transformation terms requires another post-processing package analogous to the one I've described for photometry.

A.J. Penny: *Do the programmes take into account the change in atmosphere extinction due to the change in the zenith distance across a frame?*

Stetson: No. For small fields this is negligible if observations are made at reasonable air masses.

S.B. Howell: *I'd just like to point out that timing errors in the shutter are a big problem when using brighter standards such as the original Landot stars. Getting many sets of*

fainter standards is important.

A. T. Young: There must be angular effects in CCDs, which should show up in defocussed images. The response depends on angle as well as position. A pixel in an out-of-focus image receives light from only a small part of the pupil. But in flat fields, it sees the whole pupil. These effects, at least, should vary according to the Fresnel reflection coefficients in back-illuminated CCDs, and should be larger due to shadowing by electrodes in front-side illumination.

Stetson: I remain to be convinced that these effects would cause photons to be systematically lost or gained at the 0.5% level in slightly defocussed images.