

Stellar Photometry Software

KENNETH A. JANES

Boston University, Department of Astronomy, 725 Commonwealth Avenue, Boston, Massachusetts 02215
 Electronic mail: janes@buast5.bu.edu

J. N. HEASLEY

Institute for Astronomy, University of Hawaii, 2680 Woodlawn Drive, Honolulu, Hawaii 96822
 Electronic mail: heasley@hoku.ifa.hawaii.edu

Received 1992 October 16; accepted 1993 February 25

ABSTRACT. We describe here the Stellar Photometry Software (SPS) that we have developed at Boston University and the University of Hawaii. SPS combines in a single program procedures for locating stars, computing a mean stellar point-spread function (PSF), and performing aperture and/or multiple PSF-fitting photometry, along with related bookkeeping functions. The software can be run either interactively or in batch mode on computers using the UNIX operating system. The performance of SPS is compared to that of the photometry programs DOPHOT and IRAF/DAOPHOT using both real and simulated CCD observations. A direct comparison of the instrumental magnitudes shows that all three programs produce comparable results.

1. INTRODUCTION

The growing flood of digital astronomical images has made it imperative that suitable software capable of exploiting the full potential of this data be developed. The basic processing of digital (CCD) data for removal of sensitivity variations and other electronic effects can now be effected in a systematic and uniform fashion using software packages such as IRAF. However, a variety of approaches and philosophies to doing stellar photometry are possible. In the case of simple, uncrowded images, straightforward procedures simulating a fixed focal-plane aperture are quite effective. These are limited only by the requirements that there be no neighboring stars within the aperture and that it be possible to model the background sky in some suitable fashion. For more complex star fields, however, it is necessary to model the detailed instrumental response to a point (stellar) source to determine the relative magnitudes of blended or nearly blended stellar images. This point-spread function (PSF) modeling forms the basis of most programs to compute stellar magnitudes from digital images. These include CAPELLA (Llebaria et al. 1989), DAO-PHOT (Stetson 1987), DOPHOT (Mateo and Schecter 1989), HAOPHOT (Gilliland 1990), INVENTORY (Kruszewski 1989), ROMAFOT (Buonanno et al. 1983; Buonanno and Iannicola 1989), STARMAN (Penny and Dickens 1986), and others.

In spite of the differences among these programs, each of them works by fitting a model of the PSF to individual stellar images, and comparisons of these programs (see, e.g., Ortolani and Murtagh 1989) seem to indicate that the results are not dramatically different. Indeed, given the diversity of approaches adopted by the various program developers, the apparent good agreement between these software packages is encouraging to the user. What differences there are in the photometric results certainly relate to differences in the sophistication of the numerical and statistical techniques employed, especially in how they treat

multiple PSF fitting to groups of stars in very crowded fields. Inevitably there are trade-offs, although the results do not appear to differ greatly. However, the photometry programs vary considerably in the ease of use and in the computational resources required.

We have been engaged for some time in the development of a stellar photometry program at Boston University and the University of Hawaii. We began this project before the other programs were generally available, and we chose to continue developing our program with the goal of incorporating what we see as the best features of some of the other programs with ideas we have developed. Our aim has been to develop a system that makes minimal computational demands, runs quickly, and is easy to use, but does not sacrifice statistical rigor or photometric precision. Our program has reached a relatively fixed state, and we have a number of projects nearing completion that make use of the software. Further, since we have received requests from colleagues to use our software, it is appropriate to describe the system, which we have called Stellar Photometry Software (hereafter SPS).

2. DESCRIPTION OF THE ALGORITHMS

The SPS combines in a single program the functions of locating star images, computing a mean PSF, and performing aperture and/or multiple PSF-fitting photometry, along with various bookkeeping and diagnostic functions. It can be run in interactive sessions or in a batch mode, and its flexible menu-driven operation in interactive mode makes the program easy to learn and use. The program is written almost entirely in FORTRAN with a small number of compatible C language subroutines used for special tasks. At the present time, the program runs under the UNIX operating system (including variants such as Ultrix). The complete reduction of a relatively uncrowded image containing a few hundred stars can be done in an interactive session of a few minutes at a workstation, while a

complex and crowded field containing several thousand stars may require as much as an hour of work. The total cpu time required per frame is rarely more than a few minutes on a Sun or DEC workstation.

2.1 Making the PSF

We have chosen to represent the PSF by the sum of a circular Gaussian and a table of remainders, that is, the mean deviations from the Gaussian. Starting from an initial list of stars, whose instrumental magnitudes have been determined from aperture photometry (see Sec. 2.4), a Gaussian function is fit to the stellar data. It is important to note that the PSF may or may not actually be Gaussian in shape. The assumption is only that it is approximately Gaussian and that it is approximately circular in profile; in practice, rather large deviations from a Gaussian shape or circular profile can be tolerated. As noted by Stetson (1987), if the basic Gaussian shape is subtracted from a normal stellar profile, the remaining components of the stellar profile will generally be relatively small, with shallow slopes. Therefore, interpolation in such a table of remainders can be done quickly and accurately. Our approach differs somewhat from that of Stetson in that he approximates the stellar profile imaged onto the detector by a Gaussian and then integrates this profile over the detector pixel elements (see, e.g., his Eq. 6) before subtracting from the measured stellar profile to obtain the table of remainders. We fit the Gaussian directly to stellar images, which in practice accomplishes the goal of removing the steep gradients from the PSF and leaving a table of slowly varying residuals.

The remainders table is derived from the residuals left after fitting the preliminary Gaussian function, appropriately scaled and shifted, to the PSF stars. Each entry in the table is the mean of the residuals corresponding to that position in the PSF, weighted by the brightnesses of the individual stars, and scaled to the same units as the Gaussian. To facilitate the calculation of the PSF, the remainders table is evaluated at tenth pixel intervals in x and y . The stars are, of course, randomly located with respect to pixel boundaries, so the remainders table is built up by doing a bilinear interpolation in the arrays of residuals at each of the tenth pixel increments.

The value of the PSF at a given (x,y) position relative to its center is just the sum of the Gaussian term plus the entry from the remainders table, obtained by a bilinear interpolation in the table. The program includes various diagnostics and checks that can be performed on the PSF, such as removing a star from the list or editing out bad pixels. Figure 1 shows, for a typical PSF array, a cut across the center of the remainders array, together with its associated Gaussian. In the figure the remainders have been scaled up by a factor of 10 for visual clarity. The problem of setting the zero point of the photometry still must be resolved by suitable application of aperture photometry to the stars used to create the PSF. Because of the way we use the aperture photometry magnitudes to scale the individual stars that contribute to the average PSF, our




FIG. 1—Cut across the center of an average PSF. The dashed line is the Gaussian component subtracted from the actual mean stellar PSF, while the solid line shows the corresponding remainders. The latter have been scaled up a factor of $10\times$ to make their variations more visible. The PSF for this example is for the CCD image shown in Fig. 2.

final PSF produces instrumental magnitudes that have the same zero point as the aperture photometry, so no separate “aperture” corrections need to be applied to the final photometry.

The generation of a “clean” PSF is the most difficult and time-consuming task in performing photometry in crowded fields. It is often necessary to resort to an iterative process, using successive approximations to the final PSF to remove contaminating neighbors to the bright stars being used to create the PSF (see, e.g., the discussion by Stetson 1989 and the examples in his Fig. 7). This process is labor intensive for the user. In SPS, once a preliminary PSF has been defined, either by identifying stars manually or allowing the program to select them automatically, a “first pass” is made through the data to find stars and perform multiple PSF fitting. SPS allows the development of successive approximations to the final PSF with minimal user intervention by designating some stars in the primary star list as PSF stars, either manually or automatically, by setting limits on the acceptable estimated magnitude errors from the first pass photometry. In generating an improved PSF, the program will use the designated stars for the PSF, after removing the effects of all neighboring stars from the PSF stars. If the image is uncrowded, the entire process can be automatic. In moderately crowded situations, it may be desirable to select the preliminary list of PSF stars manually. Only in the most crowded situations is it necessary to do a completely interactive calculation of the PSF. In such cases it may also be necessary to repeat the procedure several times in order to weed out any stars with companions too close to be found by the automatic star-finding algorithms.

2.2 Star Finding

Stellar coordinates can be entered into the program in several ways: by reading from an existing file, by entering coordinates manually from the terminal, by selecting stars from the image displayed on the workstation, or by an

automatic finding routine. The normal procedure is the latter, especially when there are hundreds or thousands of star images in the frame.

The automatic star-finding procedure works by computing the cross correlation between the image and the PSF, centered successively at each position in the image. In effect, a map of the cross-correlation coefficient is produced, in which values near one represent high probability locations of point sources (stars). The user can adjust the size of the correlation radius used to compute the cross-correlation map; we find, as a rough guide, that it is best to keep the correlation radius somewhat over half the FWHM of the PSF. Stars are assumed to be located at local maxima above a user-selected threshold value of the correlation coefficient. If the user selects to run the star-finding procedure before a PSF has been constructed, say to find stars for aperture photometry, the program uses a Gaussian approximation to the PSF with a FWHM specified by the user. In practice, using the actual PSF or the Gaussian approximation gives very similar results in locating stars, with the differences arising only for some of the stars at the faint magnitude limit of a particular frame.

The procedure is highly efficient at finding stars in all but the most crowded fields, but nonstellar image features, such as cosmic rays or galaxies, generally will have a lower correlation coefficient and will not be selected. We do not attempt to perform any further discrimination on the objects detected, e.g., examining the spatial moments of an object to determine if it is elongated. The running time for the star-finding procedure is only weakly dependent on the number of stars found, but does of course increase in direct proportion to the total number of pixels in the image. It takes approximately 70 s to search an image 1024×1024 pixels in size (on a Sun IPX workstation).

The star-finding procedure in SPS operates on the “residual” image array, i.e., a copy of the image from which all stars that have been found and successfully fit have been subtracted. (On program initialization this array holds a copy of the original frame.) Thus, once a “pass” has been made through the data, one can rerun the star-finding procedure to search for stars previously missed because they were crowded by a brighter neighbor.

2.3 PSF Fitting

Given an estimate for the position of a star and a model of the PSF, the stellar magnitude is found by a nonlinear least-squares fit of the scaled and shifted PSF to the observed stellar image. The full nonlinear fit may require the determination of four parameters for each star: the stellar magnitude, x and y positions, and sky level. However, in some circumstances the complexity of the fit may be reduced depending upon whether the star’s position or the background level are specified in advance. If the stellar position is known in advance, then the problem reduces to a linear least-squares fit for the magnitude and/or background. The usual procedure is to fit to the magnitude and position, assuming that the sky has been separately mea-

sured (see Sec. 2.5). The algorithm used in the nonlinear least-squares fitting is the Marquardt gradient-expansion procedure (see Bevington 1969).

The program solves for only one star at a time, even in the most crowded fields. If there are other star images crowding the one being fit, they are subtracted from the image (using the current best estimate of their positions and magnitudes) before doing the fit for the star in question. When the new solution for a star gives a significantly different estimate of its magnitude, then all the surrounding stars are refit using the new values of position and magnitude for the primary star. In turn, if any of these are significantly changed, a new solution for the primary star is made. This procedure almost always converges in 2–5 iterations and permits an enormous increase in speed over programs where all contiguous stars in a “group” are fit simultaneously. The overall approach is that of a relaxation method, and in general the derived magnitude of a star will be determined by its own brightness and that of its closest neighbors. On any pass through the star list, we proceed in order of decreasing brightness (based upon the results of a previous fitting cycle or from aperture photometry). This in effect strengthens the relaxation approach in that the brightest stars, which will dominate their own magnitude determination, are fit first and any contributions they make to their fainter neighbors are estimated early and reasonably accurately in the iterative process. There is in principle no limit to the size of a “group” that affects the magnitude of an individual star—all stars in the current star list can contribute to the solution of any other star. In practice, however, we limit the maximum size of a group to 100 stars, with a group being defined in such a fashion as to expand spatially away from the star being measured. Thus, even if the group is full, multiple passes through the relaxation cycles will insure that the influence of more distant stars will diffuse through the entire star list. The least-squares solution for a star’s position and magnitude also generates an estimate of the standard error of the values of the fitting parameters, in particular the magnitude.

In performing the fit, the individual pixels are weighted according to the estimated variance of the data, as computed according to the model

$$\sigma_{\text{pix}}^2 = \sigma_{\text{sky}}^2 + \sigma_{\text{rn}}^2 + F \times 1.0 / (e^- \text{ per ADU}) + \sigma_{\text{flat}}^2 \times F^2 + \sigma_{\text{psf}}^2 * (F - \text{sky})^2 + \sum_k \sigma_{F_k}^2$$

(see Gilliland and Brown 1988; Stetson 1987). The quantity F is the flux in a given pixel measured in the instrumental units. In addition to the variances of the sky, read noise, flat field, and PSF, the variance of the data includes as the final term the uncertainty in the flux contributed by all neighboring stars. This sigma is computed from the magnitude error of the k th contaminating star, in flux units, multiplied by its PSF value at the current pixel location for the contaminating star.

2.4 Aperture Photometry

Just as in single-channel photoelectric photometry, the zero point of the instrumental magnitude of one star or frame is tied to other stars and frames only if the effective measuring aperture is fixed. The actual stellar profile extends beyond the size of any practical aperture, regardless of the seeing or the telescope image scale (King 1971). The aperture must be large enough so that small irregularities in image centering, fluctuations in seeing, or field variations of the stellar profile do not seriously affect the flux within the chosen radius; however, the aperture must not be so large that the sky background unnecessarily dominates the measurement uncertainty.

The SPS program includes an aperture photometry routine. Aperture photometry is used to normalize the PSF, and it can also provide quick and accurate photometry in moderately uncrowded fields. The user specifies a reference aperture (in pixels), chosen according to the conditions and instrumental characteristics. The routine takes the preliminary position of the star to calculate a more precise position, by doing a cross correlation between the marginal distributions in x and y and their reversed distributions. Given the location of the star, a synthetic aperture is created, the distance of each pixel from the center of the star image is computed, and the successive sums in stellar flux are computed out to each radius. The sums are sorted according to increasing distance from the center of the stellar image, and the value of the sum corresponding to the desired radius is found by fitting a quadratic least-squares function to the sums on either side of the desired radius. The error of this quantity is found by accumulating the (squared) sums of the individual pixel errors, as computed from the previous equation. The sky level is found as described in Sec. 2.5.

If all stars on the frame are to be measured with aperture photometry, then the radial profiles of stars are accumulated into a mean stellar profile, evaluated at integral pixel positions. The mean profile is compared with the individual profiles, and whenever a stellar profile deviates significantly from the mean profile, it is removed and the mean profile is recomputed. The mean profile (and its estimated error) can be used to improve the aperture photometry (see, e.g., Howell 1989). Thus for a faint star, the sky noise will dominate the error in all but the central few pixels, whereas a bright star can be measured reliably at much larger radii. If the selected reference aperture is sufficiently large, then the fainter stars and possibly crowded star images will not be measured well, so at the option of the user, the individual aperture photometry magnitudes can be corrected by the mean profile. The fainter, more crowded stars, which are more accurately measured with small apertures can be corrected to the larger reference aperture using the mean profile. In SPS the aperture photometry routine can automatically fit the truncated profile of an individual star to the mean profile, thereby providing a normalization to the full aperture radius.

The use of the mean profile to correct aperture photometry requires the assumption that the shapes of star images

do not depend on position in the image. It is not necessary to assume that the images are well focused, or that they are even round. In the case of large or trailed images, the magnitude corrections for small apertures will be large, because a larger proportion of the light will fall outside the aperture. However, the larger the aperture correction is, the larger the uncertainty in the correction, so poor images will lead to poor photometry, as would be expected. It is also the case that if the calculated PSF is a good model of the light distribution in a stellar image, then PSF fitting will always lead to better photometry of faint stars than aperture photometry, because the former solution is constrained by the *a priori* knowledge of the image shape and the latter is not.

2.5 Finding the Sky

The sky can be determined in three ways: it can be specified in advance, it can be included as a constant level in the least-squares solution, or it can be determined separately from the magnitude determination as the mode of the distribution of pixel values around the star. The latter is the method that we normally employ and that will be emphasized here.

The question of the relative merits of using the mean, the median, or the mode of pixel values as the best representation of the sky background has been discussed many times (see, e.g., Stetson 1987; Eaton 1989). After studying those discussions, we have decided to adopt the mode as the most appropriate in crowded field situations. The sky is estimated from those pixels lying outside the reference radius for aperture photometry and inside some outer radius specified by the user. A typical outer radius is of the order of 20–25 pixels. Assuming an inner radius of seven pixels, then there will 1000 or more pixels in the distribution. In a best-case situation, with no nearby bright stars contaminating the distribution, one would expect a Poisson distribution of pixel values. For a typical case, where the sky value is of order 100 electrons, with a corresponding expected $\sigma \sim 10$, the expected σ_{sky} , found from 1000 pixels would be roughly one-thirtieth ($1000^{1/2}$) of that value, or about 0.3 electrons. The practical uncertainty will, of course, be substantially larger than this value.

In some instances it is advantageous to include the sky directly in the least-squares fits, especially if there is a significant gradient in the sky across the field that makes it difficult to find the peak of the histogram of sky values at the star's position, or in the most crowded images, such as near the cores of globular clusters. In the latter case, the background of a moderately bright star in a dense field of stars is highly variable on a scale of a few pixels. The brighter stars can be fit including a local estimate of the sky within the fitting radius (typically set equal to about twice the FWHM when running in this mode). After the bright star is subtracted from the image, somewhat fainter stars can then be located and fit, albeit with considerably lower precision. On a subsequent pass through the star list, when the fainter stars are subtracted first, the sky computation for the brighter stars can be determined much more accu-

rately. This in turn leads to an improvement in the photometry of the fainter stars. In this way, stars can be successively “peeled away” from the image, permitting photometry for at least the brighter stars that may be of adequate precision for some purposes, even right through to the middle of a globular cluster.

3. OPERATION OF SPS

We have attempted to keep the program simple to use and to understand, as well as fast. The current version operates under the UNIX operating system and variations; duplicate copies are maintained at Boston University and the University of Hawaii, with the former running on DEC workstations and the latter on Sun workstations. The program can be run with or without image display capability. Our current implementation of the program uses the SAO-image display utility that runs under the X11 window environment (VanHilst 1991) to provide image display and cursor readback from the display.

3.1 User Interface and Program Operation

SPS is designed to be a flexible, interactive photometry tool for deriving instrumental magnitudes. The program is normally run in interactive mode, in which the user is presented with a series of menus to operate the program. The user will normally interact with a setup menu, where one specifies the image to process and adjusts parameters that must be set to perform the various photometric tasks, and the main or primary menu, through which specific tasks relating to the photometry are executed. On-line help is available to the interactive user within each of these menus.

The various files used by the program can be specified in the setup menu, with a standard naming convention being used as a default. Ordinarily the user specifies an image name, which will be used to create the names of the other files used by the program. These include a PSF file, a basic output data file, a similar file containing just stars used in making the PSF, and a file containing standard star measurements. A final (optional) file is a map of bad pixels. The image format is assumed to be disk images of the FITS standard files.

To model the noise structure of an image, it is necessary to include information about the CCD characteristics at the setup stage—the gain of the CCD in electrons per analog-to-digital unit (ADU), the estimated uncertainty in the sky background, the estimated error introduced by the flat-field processing, the estimated error in the sky background, and the CCD read noise. In addition, the saturation level of the CCD can be specified. Any individual pixels above the given saturation level will be ignored by all calculations in the program.

Various parameters for performing photometry are specified in the setup menu. The PSF diameter, the effective fitting radius (which may be smaller than the full PSF size), the aperture photometry radius, the outer sky radius, and the estimated FWHM of stellar images may be entered in the setup menu. While the program provides defaults for

each of these parameters, there is no guarantee that these will provide useful measurements for a particular CCD frame, and individual frames must be examined prior to attempting photometry with SPS to determine whether the defaults are suitable. The estimated FWHM is used only in the star-finding procedure when no PSF is available. In that situation, a Gaussian PSF with that FWHM is used for the cross correlation.

The primary SPS menu allows the user to perform specific tasks appropriate to doing stellar photometry on digital images, e.g., loading the image, generating the PSF, finding stars in the frame, and performing aperture or PSF-fitting photometry. Common bookkeeping and data manipulation functions are also provided in this menu, e.g., loading or saving a coordinate/photometry list, specifying stars to add or delete from the star list, and eliminating duplicate stars. This menu also provides interfacing to the external image display routines and cursor interface. Finally, an extensive set of “fake star” routines has been included as an option under the primary menu, allowing the user to generate artificial images or add artificial stars to an image being analyzed to test for photometry accuracy and completeness.

The primary menu configuration allows the user to execute the commands in any sequence that makes logical sense (e.g., you cannot perform photometry on a frame until it has been loaded into memory). In addition, one can toggle between the primary and setup menus as needed to adjust the parameters set in the latter and then resume the actual photometric measurements.

While the SPS software was designed primarily for interactive use, it is possible to perform batch or automated photometry with the software. As is common in most modern operating systems, SPS can run in batch (or background) mode by trivially accepting terminal input redirected from a file. An “autophot” mode is available from the primary menu (or a command line option) which allows a frame to be reduced fully automatically using the setup parameters in the startup file (or the defaults if such a file is not found). The program will find the stars in the frame, construct an initial estimate of the PSF, perform a first pass fit to all the stars and subtract these from the frame. A second pass of star finding is then performed on the residual frame and then all the stars are fit again. A new PSF is then constructed from all the available photometric information and a final set of PSF fitting. In practice, this fully automatic reduction works well only for relatively uncrowded images. The major limitation of the fully automated mode is optimal selection of suitable stars for generating the PSF in very crowded fields. An intermediate option is for the user to work interactively to identify PSF stars manually and then do the rest of the processing automatically. A final check of the residual images can be done at the end, and any stars missed or other problems can be fixed quickly in interactive mode.

3.2 Performance—Speed, Storage Requirements

The complete reduction of an image can ordinarily be accomplished in minutes. In a typical case, a moderately




FIG. 2—CCD image of the central region of the open cluster NGC 103. Specifics of the image are summarized in Table 1. This image is one of the three CCD frames used in comparing the instrumental magnitudes from SPS, DOPHOT, and IRAF/DAOPHOT and was provided by Dr. Randy Phelps. North is up and east to the right.

crowded image of an open cluster (Be 39) covering 512×512 pixels and containing about 1000 star images was reduced in 30 min of “wall clock” time, with care taken to generate the best possible PSF and then to remove every detectable star from the image. An extremely crowded globular-cluster image containing several thousand stars may take an hour or longer to reduce.

The current configuration of the program will handle images as large as 1024×1024 pixels. Two versions of the image are used by the program, so the program itself requires well in excess of 8 MBytes of memory. As the size of modern CCD detectors continue to grow, the memory demands on the processing software do as well. However, at some point it becomes impractical to keep increasing the space required to keep the entire image in memory. Digital images with larger than the maximum allocated storage can be processed by defining the origin of a 1024×1024 pixel “window” of the original frame. Thus, it is possible to process very large frames with multiple passes of the SPS software over the same frame. The photometry lists of each of the frames give the actual x and y positions relative to the original image, not the window origin, so it is a straightforward matter to relate stellar positions between the output files for a given large frame.

4. COMPARISON WITH OTHER PHOTOMETRY SOFTWARE

We turn now to the comparison of our SPS program with two widely distributed photometry packages, namely, DOPHOT (Mateo and Schecter 1989) and the IRAF implementation of DAOPHOT (hereafter IRAF/DAOPHOT; Stetson et al. 1990). Our tests used version 1.0 of DOPHOT, which was kindly supplied by Dr. Mario Mateo and Dr. Paul Schecter. We used the version of IRAF/DAOPHOT that has been distributed with version 2.10 of IRAF. This implemen-




FIG. 3—CCD image of the sparse globular cluster E3. This image was kindly provided by Dr. James Hesser and Dr. Robert McClure. It has been distributed as a test image for verifying the operation of Dr. Peter Stetson’s DAOPHOT software. North is up and east is to the left.

tation of the DAOPHOT algorithms in IRAF is based upon Stetson’s original code (Davis 1992) and does not include some of the improvements he has incorporated in DAOPHOT II (Stetson 1991).

For test frames we used three actual CCD images and a simulated CCD frame generated by the computer. The first real CCD image was that of the sparse globular cluster E3 obtained by Dr. James Hesser and Dr. Robert McClure, the second was an image of the Large Magellanic Cloud cluster NGC 2041 obtained by Dr. Mario Mateo, and the final frame covers the central region of the open cluster NGC 103 and was obtained by Dr. Randy Phelps. The NGC 103, E3, and NGC 2041 frames are shown in Figs. 2, 3, and 4, respectively. All three fields were imaged in the V band. The E3 and NGC 2041 frames are or have been distributed as “test images” with the (original) DAOPHOT and DOPHOT software, respectively. These three images are very similar in their PSF characteristics, with each having the FWHM of the PSF of \sim three pixels. They are quite different, however in their crowding characteristics, with the E3 image being rather empty, and NGC 103 showing crowding only near the frame limit, while the NGC 2041 frame has a dense cluster core with many bright stars. As the latter frame is our example of an extremely crowded field, we have restricted our PSF fitting to a 200×200 pixel region centered on the cluster itself. We have, however,




FIG. 4—CCD image of the large magellanic cloud cluster NGC 2041. This image was kindly provided by Dr. Mario Mateo. It has been distributed by him as a test image for verifying the operation of the DOPHOT software. North is up and east is to the right.

TABLE 1
Properties of Test Images

Image	E3	NGC 2041	NGC 103
Telescope	CTIO 4 m	CTIO 4 m	NOAO 0.9 m
Observer	Hesser & McClure	Mateo	Phelps
Date	10 Feb 83	28 Oct 87	22 Nov 90
Exposure (sec)	150	250	90
CCD	RCA	RCA	Tektronix
Rows \times Columns	284 \times 492	312 \times 508	800 \times 800 ^a
Gain (e^- /ADU)	9.77	7.4	9.6
Read noise (e^-)	82	45.1	10 ^b
FWHM (pixels)	2.9	2.9	2.8
Sky level (ADU)	461	50	30

^aCenter of a 1024 \times 1024 frame.

^bAdopted.

used stars in the less crowded parts of the frame to generate the PSF to be used in the photometric reductions. (The remainder of the NGC 2041 frame is no more difficult for the photometry programs to process than our other sample images.)

Table 1 summarizes the specifics of the three CCD images used for our comparison. Two minor restrictions were imposed on our comparison tests with DOPHOT, in that we wanted to use the software as distributed rather than attempting to modify it, thereby possibly introducing errors in the coding. Only the central 800 \times 800 pixel region of the original NGC 103 frame was used in the comparison so as to not overflow the storage dimensions of DOPHOT. Further, DOPHOT required that we convert our NGC 103 images from IRAF into a simple 16-bit unscaled FITS image with `bzero=0` and `b scale=1`. This will result in some slight loss of significance in the data but should not greatly affect the comparisons with SPS or IRAF/DAOPHOT, which use the original 32 bit FITS versions of the data.

While a detailed comparison of the photometry from SPS with that of the other two programs can be very illuminating, it is only relative in the sense that we do not know what the “sky truth” of the observations really is. To attempt to illuminate how well each of the programs deals with reproducing a known “reality,” we generated a simulated CCD frame with parameters appropriate to the NGC 103 field. This simulated frame was created with a program that was completely independent of any of the photometry programs so that none of them would have an unfair “advantage” in the tests (e.g., because the simulation code assumed some specific form for the PSF). Because the PSF is a critical piece of information for these simulations, we adopted an analytical form similar to that of a real star. The specific shape used was Gaussian in the stellar core

$$P(r) = e^{-(r/w)^2}, \quad r \leq r_m,$$

with an extended wing of the form

$$P(r) = \frac{C}{[1 + (r/R)^2]^\beta}, \quad r \geq r_m,$$

where the constant C is chosen to make the two representations (but not their derivatives) continuous at the mesh radius r_m . All the radii are measured in units of pixels from the center of the star. We fixed $r_m=2$, $R=1.0$, and $\beta=1.5$




FIG. 5—(a) Comparison of PSF-fitting photometry with SPS, with the input magnitudes used to generate the simulated CCD frame described in the text. The ordinate is the difference in magnitude between the SPS photometry and the input magnitudes after adjusting for the mean magnitude difference between the two data sets. The abscissa is the input magnitude scale. (b) Same as (a), but comparing the instrumental magnitudes from DOPHOT with the input magnitudes of the synthetic frame. (c) Same as (a) but for the instrumental magnitudes from IRAF/DAOPHOT.

for the simulation. Stellar positions and relative magnitudes for the simulation were taken from an SPS fit to the actual NGC 103 frame. Stars were added into a frame with a constant sky level of 10 ADU. The read noise and CCD gain for the NGC 103 frame from Table 1 were adopted. The noise model for the simulation allowed at each pixel for CCD read noise plus Poisson “photon” noise from the star and sky.

Both the real and simulated CCD frames were processed in the same way. For the comparison software we followed the procedures described in the DOPHOT *User's*

TABLE 2
Comparisons with the Simulated Frame

	SPS	DoPHOT	IRAF/DAOPHOT
Number of stars in the input list	2462	2462	2462
Number of stars found	1623	3701	1762
Number matched to input list	1537	1452	1650
RMS magnitude difference	0.101	0.117	0.133
RMS difference in x position	0.091	0.098	0.093
RMS difference in y position	0.091	0.115	0.103

Manual (distributed with the photometry code) and the *IRAF User's Guide for Stellar CCD Photometry* (Massey and Davis 1990). As SPS and the other programs require similar inputs, e.g., the PSF size and fitting radii, aperture photometry, and sky radii, we attempted to keep these parameters the same for each program. Exact duplication of parameters between the various codes is not possible because of how they implement their photometry algorithms, thus without extensively modifying SPS or the other programs it was difficult (or impossible) to get an *exact* comparison between these programs for an individual test frame. However, our tests should produce the results a "typical" user of any of these photometry should obtain. In SPS and IRAF/DAOPHOT, we used identical inner and outer radii for computing the sky prior to the nonlinear fitting. In DoPHOT, the sky "box" was set to be large enough that the wings of the stellar profile had decayed sufficiently that "real" sky is reached inside this region, and the actual sky determination was done as part of the nonlinear fitting procedure (see the discussion by Mateo and Schechter 1989). The selection of PSF stars also differed between the three programs. For IRAF/DAOPHOT, we selected approximately 20 stars manually using the image display and constructed the PSF following the guidelines in the IRAF photometry user's guide. DoPHOT automatically updates its PSF information as it successively processes to fainter stars within the frame (see again the discussion by Mateo and Schechter 1989). With SPS we constructed initial PSFs by allowing the program to select a large number of candidate stars, and this list was subsequently refined by deleting those PSF stars that showed large differences between the aperture photometry and PSF-fitting magnitudes.

The SPS photometry for the simulated CCD frame is shown in Fig. 5(a). The magnitude differences have been adjusted for a zero-point offset and are plotted versus the known input magnitudes. To determine the zero-point offset, we calculated the mean magnitude difference for those stars brighter than instrumental magnitude 18, and having $\sigma_{\text{mag}} < 0.02$ mag.

One of the most extensive comparisons of digital photometry software took place at the first ESO/ST-ECF Data Analysis Workshop, wherein the participants reduced a common set of data frames. The comparisons were, for the most part, qualitative in that only the shapes and the general distribution of stars in instrumental color-magnitude diagrams were contrasted (e.g., Ortolani and Murtagh 1989). Here we want to compare directly the instrumental magnitudes from SPS, DoPHOT, and IRAF/DAOPHOT for the test images. Figures 5(b) and 5(c) illustrate the com-

parison of both DoPHOT and IRAF/DAOPHOT instrumental magnitudes with the actual input for the simulated image. The PSF-fitting results for all three codes appear to faithfully reproduce the input magnitudes of the simulation frame.

Because we have a specific list of input stars of known magnitudes and positions that form the simulated image, it is possible to make a quantitative summary of the ability of each of the three codes to recover the input values. To the extent that the simulated image is a reasonable representation of a real CCD image, this summary should be indicative of the performance of the three codes in actual use. We compared each of the three output files separately with the original input data by matching stars in the output lists to the corresponding stars in the input list by their coinci-




FIG. 6—(a) Comparison of the SPS instrumental magnitudes for the NGC 103 test frame with those found by DoPHOT. (b) Same as (a) for IRAF/DAOPHOT. (c) Comparison of the DoPHOT instrumental magnitudes with those from IRAF/DAOPHOT for the NGC 103 test frame.






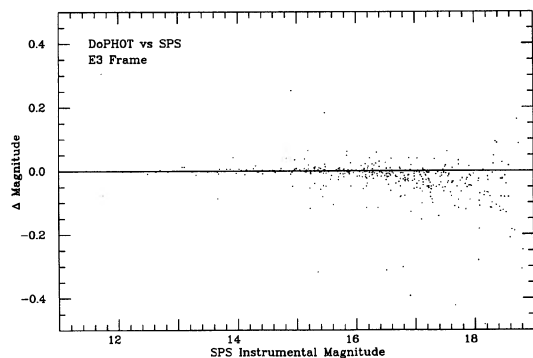


FIG. 7—(a) Comparison of the SPS instrumental magnitudes for the E3 test frame with those found by DoPHOT. (b) Same as (a) for IRAF/DAOPHOT. (c) Comparison of the DoPHOT instrumental magnitudes with those from IRAF/DAOPHOT for the E3 test frame.

FIG. 8—(a) Comparison of the SPS instrumental magnitudes for the NGC 2041 test frame with those found by DoPHOT. (b) Same as (a) for IRAF/DAOPHOT. (c) Comparison of the DoPHOT instrumental magnitudes with those from IRAF/DAOPHOT for the NGC 2041 test frame.

dence in position; the closest match found within a one pixel diameter circle was assumed to be the correct one. We counted the numbers of stars matched along with the rms differences between the input and output magnitudes, the x positions and the y positions, as shown in Table 2.

As the table shows, none of the three programs found all stars in the input list, because the “exposure” of the simulated frame was such that many of the fainter stars were too faint to be found. All three programs missed nearly all of the faintest 750 or so stars in the input list. Conversely, all three codes made spurious detections, and in fact, most of the stars found by DoPHOT do not really exist. The rms magnitude and position differences are similar in all three programs.

In Figs. 6–8 we plot magnitude differences of SPS and

DoPHOT or IRAF/DAOPHOT for the three actual CCD frames. Again, after adjusting for zero-point differences, the agreement between the different programs is gratifying. Although the differences among the three programs are small overall and consistent with the measurement errors in general, nevertheless, slight nonlinearities can be discerned in the comparisons among the three programs. In particular, as shown in Fig. 7, there is a systematic difference between the IRAF/DAOPHOT and SPS photometry of E3 such that the fainter stars differ from the brighter ones by several hundredths of a magnitude. This slight curvature can be explained entirely by differences in the way the background sky level is computed. On the scale of the SPS magnitudes in this frame, a change in the sky level of one ADU out of 465 will change the photometry of an 18th

magnitude star by approximately 0.05 mag. In fact, the SPS sky values are slightly higher than the IRAF/DAOPHOT sky values in the central part of the image, where the star density is highest. This difference (about 0.6 ADU) is enough to explain completely the slight curvature seen in Fig. 7. The sky values computed for this frame by DOPHOT are intermediate between the IRAF/DAOPHOT and SPS values, but have a considerably larger scatter.

As can be seen in Table 1, the E3 frame has a much higher sky level than the other two images. Since the random fluctuations in sky values are at least as large as the square root of the data numbers (and probably considerably more), the sky uncertainty in this frame must be substantially larger than that in the other frames, based on Poisson statistics alone. Therefore, unless the sky were to be evaluated in exactly the same way in all three programs, slight systematic and random variations in computed sky level are inevitable. From the present data, there is no way to tell which, if any of the programs, has the "correct" sky values, or even what should be the proper definition of the background level. This situation illustrates a fundamental problem in digital stellar photometry. In a region crowded with stars, the background, which includes scattered light from many individual stars, as well as contributions from a large number of stars, undetectable individually, can become indeterminate at a sufficiently high level to have a significant, systematic effect on the photometry. In the other two frames, where the sky level is much lower and presumably better defined, the three programs are in much better agreement.

In each of the comparisons shown in Figs. 5–8, there are several relatively bright stars with unusually large residuals. For example, in the comparisons of the E3 frame, there are 38 stars with (SPS) instrumental magnitudes brighter than 17, and that have magnitude differences greater than 0.05 in Figs. 7(a), 7(b), or 7(c). In some of these cases, one of the three codes will differ from the other two, or in about a third of the cases, all three differ substantially from one another. There is no obvious common factor among these stars. The most pervasive condition is that almost all of the stars with unusually large residuals are in crowded fields with several companions contributing to the solution to a star. Since the weighting of pixels is treated slightly differently in each of the three programs, it should be expected that the resulting magnitudes should be somewhat different in the more complex solutions.

Not surprisingly, the largest scatter in these comparisons is found in Fig. 8. The core of NGC 2041 represents the most "severe" crowding case we have considered, and we would not expect any of these three programs to do very well in such extremely crowded regions. All three programs do reasonably well, and the resulting magnitudes, while they should be treated with some caution, are probably suitable for many applications.

5. SUMMARY

Although the three stellar photometry programs discussed in this paper were developed independently and use

somewhat different algorithms, nevertheless they all yield rather similar results. Under different situations each gives slightly different answers, but overall, the comparisons among the three give no indication about which, if any, of the programs is superior. The decision to use one or the other will more likely be determined by familiarity and convenience rather than by potential scientific gain.

The preceding evaluation serves to demonstrate that the analysis of CCD imagery has reached a certain maturity. In fact, it is possible to demonstrate that the current generation of photometry software generally leads to results that are about as good as can be expected theoretically. Each of the three programs makes an estimate of the precision of the calculated magnitude for a star, based on the known sources of error. These estimates indicate the best one can do with a given image, and it is unlikely that they are overestimates, in the mean at least, but they are simply estimates, based on the statistical model. In the case of the simulated NGC 103 image, where the "truth" is known, we can compare the estimated errors for each star in the image with the actual error, for each of the three programs. This comparison can be made by looking at the normalized error

$$\sigma_N = \frac{(|\epsilon_{\text{true}}| - \sigma_{\text{est}})}{(|\epsilon_{\text{true}}| + \sigma_{\text{est}})},$$

where ϵ_{true} is the difference between the measured instrumental magnitude and the actual (known) magnitude, and σ_{est} is the magnitude error estimate generated in the photometry program. For a perfect algorithm, σ_N should be scattered randomly about zero, with a rms value of 0.5. We find that $\sigma_N = -0.07 \pm 0.42$, -0.15 ± 0.43 , and -0.16 ± 0.43 for stars measured by SPS, DOPHOT, and IRAF/DAOPHOT, respectively. Thus, for ordinary, relatively crowded images such as the NGC 103 simulation, the existing programs do about as well as can be expected theoretically.

That is not to say that the art of stellar photometry cannot be improved. In addition to improvements in such things as the ease and simplicity of use, there are several significant unresolved computational problems:

(1) As CCD detectors continue to grow in size, the data *management* problem is becoming severe. Detectors with 1000 or 2000 pixels on a side are now in common use. It is common to think in terms of doing photometry of tens of thousands of stars in a single field, for example, near a globular cluster, or in fields located in the nearby dwarf companion galaxies to our own. When mosaics of adjacent fields are made, data management becomes the critical issue. In our experience, high-precision CCD photometry is accomplished from detailed analysis of multiple frames of the same field (usually with slightly different pointing); thus keeping track of large numbers of stars from different frames, taken through different filters, can become as difficult as performing the actual photometry. Further, the level of photometric processing increases significantly if one is using the artificial star routines to investigate photometric errors and completeness, as each frame should be

reanalyzed in the same way as the original measurements to avoid biasing the results.

(2) The larger spatial coverage of the newest generations of CCDs poses a much more serious problem: It is no longer possible to assume that the PSF is constant across a frame. Apart from focal plane variations of the telescope, variations may result from slight misalignment of the CCD sensor in its dewar or from an intrinsic curvature of the device's surface. Variable PSFs can also arise with active optics and fast-guiding systems (see, e.g., McClure et al. 1991). Regardless of the cause of the PSF variation across the detector, this complication must be addressed in order to obtain the most precise photometric information out of the data. This problem has been discussed by Stetson (1991), but to date no satisfactory solution to the variable PSF problem exists. At the present time, this is probably the limiting factor in doing photometry on large-format digital arrays.

The development of the SPS program and various related photometric projects have been made possible in part through support from the National Science Foundation grants number AST-8818360 and AST-8915444 to KAJ, NASA grant NAGW-2679 to Michael Mendillo of Boston University, and by support from Boston University and the University of Hawaii's Institute of Astronomy. KAJ is especially grateful to the staff of the IfA for its hospitality on numerous visits. Randy Phelps and Kent Montgomery have contributed many valuable suggestions and located many annoying computer bugs. We would like to thank James Hesser and Robert McClure for permission to use the image of E3, and Mario Mateo and Paul Schechter for allowing us to use the image of NGC 2041. Finally, we would like to thank Lindsey Davis for her helpful comments and suggestions on the use of the IRAF/DAOPHOT routines.

REFERENCES

- Bevington, P. R. 1969, *Data Reduction and Error Analysis for the Physical Sciences* (New York, McGraw-Hill)
- Buonanno, R. and Iannicola, G. 1989, *PASP*, 101, 294
- Buonanno, R., Buscema, G., Corsi, C. E., Ferraro, I., and Iannicola, G. 1983, *A&A*, 126, 278
- Davis, L. 1992, private communication
- Eaton, N. 1989, in *ESO Conference and Workshop Proceedings 31, 1st ESO/ST-ECF Data Analysis Workshop*, ed. P. J. Grosbol, F. Murtagh, and R. H. Harris (Garching, ESO), p. 93
- Gilliland, R. L., and Brown, T. M. 1988, *PASP*, 100, 754
- Gilliland, R. L., 1990, in *ASP Conf. Ser. 8, CCDs in Astronomy*, ed. G. H. Jacoby (San Francisco, ASP), p. 281
- Howell, S. B. 1989, *PASP*, 101, 616
- King, I. 1971, *PASP*, 83, 199
- Kruszewski, A. 1989, in *ESO Conference and Workshop Proceedings 31, 1st ESO/ST-ECF Data Analysis Workshop*, ed. P. J. Grosbol, F. Murtagh, and R. H. Harris (Garching, ESO), p. 29
- Liebaria, A., Perichaud, L., Leporati, L., and Debray, B. 1989, in *ESO Conference and Workshop Proceedings 31, 1st ESO/ST-ECF Data Analysis Workshop*, ed. P. J. Grosbol, F. Murtagh, and R. H. Harris (Garching, ESO), p. 85
- Massey, P., and Davis, L. E. 1990, *A User's Guide to Stellar CCD Photometry with IRAF* (Tucson, NOAO)
- Mateo, M., and Schechter, P. L. 1989, in *ESO Conference and Workshop Proceedings 31, 1st ESO/ST-ECF Data Analysis Workshop*, ed. P. J. Grosbol, F. Murtagh, and R. H. Harris (Garching, ESO), p. 69
- McClure, R. D., Arnaud, J., Fletcher, J. M., Nieto, J., and Racine, R. 1991, *PASP*, 103, 570
- Ortolani, S. and Murtagh, F. 1989, in *ESO Conference and Workshop Proceedings 31, 1st ESO/ST-ECF Data Analysis Workshop*, ed. P. J. Grosbol, F. Murtagh, and R. H. Harris (Garching, ESO), p. 11
- Penny, A. J., and Dickens, R. J. 1986, *MNRAS*, 220, 845
- Stetson, P. B. 1987, *PASP*, 99, 191
- Stetson, P. B. 1991, in *ESO Conference and Workshop Proceedings 38, 3rd ESO/ST-ECF Data Analysis Workshop*, ed. by P. J. Grosbol, and R. H. Warmels (Garching, ESO), p. 187
- Stetson, P. B., Davis, L. E., and Crabtree, D. R. 1990, in *ASP Conf. Ser. 8, CCDs in Astronomy*, ed. by G. H. Jacoby (San Francisco, ASP), p. 289
- VanHilst, M. 1991, *User Manual for SAOimage* (Cambridge, Smithsonian Astrophysical Observatory)