PUBLICATIONS OF THE ASTRONOMICAL SOCIETY OF THE PACIFIC
99:191-222, March 1987

DAOPHOT: A COMPUTER PROGRAM FOR CROWDED-FIELD STELLAR PHOTOMETRY

PETER B. STETSON

Dominion Astrophysical Observatory, Herzberg Institute of Astrophysics
5071 West Saanich Road, Victoria, British Columbia V8X 4M6, Canada

Received 1986 October 13, revised 1986 December 5

ABSTRACT

The difficult art of stellar photometry in crowded fields is currently undergoing a surge of popularity, and a number of different computer programs for deriving photometric information from two-dimensional digital images are currently in use. This paper describes one such program, DAOPHOT, which was written and continues to be developed at the Dominion Astrophysical Observatory. Emphasis is placed on the various types of philosophical and technical complications which arise when accurate photometry is sought for blended stellar images, and on the mathematical algorithms with which DAOPHOT attempts to deal with these complications, rather than on details of the coding. Some ways in which DAOPHOT resembles or differs from other similar programs are mentioned, and a discussion is presented of known shortcomings of the current program as well as possibilities for future improvement.

Key words: data-handling techniques–photometry (general)

I. Introduction

Astronomers never seem to want to do anything easy. Currently, many observers are attempting the difficult task of obtaining precise photometry for stars in crowded fields, a problem which is now becoming tractable principally because in the past few years efficient, photometrically linear image-detectors have become available at a number of observatories. These detectors, principally video cameras, image photon-counting arrays, and charge-coupled devices (hereinafter CCDs; see, e.g., Boyle and Smith (1970) and Kristian and Blouke (1982)) permit the recording of two-dimensional data arrays which directly represent the brightness distribution in a small patch of night sky. Stars whose images are recorded in such digital frames can thereby be photometered: the number of photons detected within the image of a given star is determined in some well-defined fashion, while the number of photons contributed to the image of the star by diffuse emission from the night sky is estimated from another region within the same two-dimensional data array. Then the number of detected stellar photons can be converted to magnitudes and colors on a standard system by observations of stars with known photometric indices, just as is done for photon counts obtained with a photomultiplier.

The advantages of photometry from two-dimensional direct images have been recognized from the days when stars were first recorded on photographic plates: (1) When a number of program stars are contained within a small region of sky (typically many arc minutes on a side for photographic plates, a few arc minutes for electronic detectors), all can be recorded with a single exposure. This results in more efficient use of telescope time. In addition, relative photometry which is unaffected by temporal variations in the instrumental response or in the transmission of Earth's atmosphere may be obtained for stars within a given field. (2) Once an image has been recorded, stars can be measured without fear that they will move around relative to the photometer aperture. Thus, smaller measuring apertures can be used than are needed for photoelectric photometry, where the aperture must be large enough to contain image excursions due to seeing variations and telescope tracking errors. Smaller apertures allow smaller corrections for the diffuse light of the night sky, resulting in much more precise measurements for faint stars, as well as better separation of program stars from close neighbors. (3) The sky measurements are strictly simultaneous with the stellar observations, reducing photometric errors caused by short-term variability of the terrestrial night-sky brightness.

Video cameras, photon-counting arrays, and CCDs also have certain additional advantages over photographic plates. (1) Their higher quantum efficiency permits faint stars to be measured in much less time or with smaller telescopes. (2) Over a broad range of brightness, these

electronic detectors appear to be inherently linear in a photometric sense: the data numbers supplied to the astronomer by the analog-to-digital converter are directly proportional to the number of photons detected in each picture element ("pixel") of the detector. This means that no complicated conversion of data numbers to intensities, corresponding to the characteristic D-to-I curve of a photographic plate, need be found. When a star is well isolated in the frame, its instrumental magnitude can be derived very simply by synthetic aperture photometry. This involves simply summing the data numbers within some region containing the image of the star and estimating the surface brightness of the sky from a nearby, star-free region of the frame. Then

$$\text{magnitude} = \text{constant} - 2.5 \times \log[(\text{summed data numbers} - \text{sky})/(\text{exposure time})] ,$$

where the constant is the same from one frame to the next. These instrumental magnitudes can be converted to magnitudes and colors on a standard system by observations of stars with known photometric indices, just as if they had been obtained with a photomultiplier.

CCDs[1] also have certain disadvantages:

1. Current CCDs have much smaller photosensitive areas than photographic plates. This renders them unsuitable for projects requiring coverage of large areas while needing only moderate photometric accuracy—searches for large-amplitude variable stars in the Galaxy, for instance.

2. In spite of their small areas, CCDs produce enormous quantities of data. While the ability to observe many stars simultaneously collapses the observing time in proportion, as compared with photomultiplier-based photometry, the CCD image containing those stars may be represented by 150,000 to 640,000 data numbers[2] rather than the one or two data numbers per star which a photomultiplier would produce for the same observations. Because of the much larger volumes of data, CCD images require more sophisticated computer hardware and software for their acquisition and analysis than photomultiplier data. That is, (a) major hardware facilities are needed for displaying and processing the two-dimensional images, and (b) software must be specially developed for obtaining the desired scientific results from the raw data arrays.

3. CCD images require careful calibration to remove certain two-dimensional patterns imposed by the detector itself. Bias patterns manifest themselves primarily as systematic variations in the intensity zero point across the chip. Flat-field patterns, which are variations in the intensity *scale*, manifest themselves both globally across the chip and locally from one pixel to the next, and usually depend significantly on the color temperature of the illumination. In addition, position-dependent variations are often found in the dark counts due, for instance, to certain P-N junctions in the chip acting as light-emitting diodes, or to the on-chip amplifier's keeping one part of the detector constantly warmer than the rest (see, e.g., G. A. H. Walker *et al.* 1984). Finally, with some devices and in some photometric passbands strong spatial variations in the sky brightness are found, caused by interference within the detector of the light of night-sky emission lines. Due to the complexity and approximate nature of the preprocessing steps required to deal with these problems, CCD photometry is usually found to have poorer absolute accuracy than the best photomultiplier work, at least for bright, uncrowded stars.

4. CCDs in general have a smaller dynamic range than most photomultipliers, although it is greater than in many photographic emulsions. At the faint end the precision is limited by a finite readout noise (commonly equivalent to 10^1 to several $\times$ 10^1 photoelectrons) produced in the amplification of the photoelectric signal to levels detectable by macroscopic electronic circuits, and the accuracy may be limited by variations in the charge-transfer efficiency during readout for low exposure levels. At the bright end, a given detector element can collect only a finite number of photoelectrons before the accumulated charge affects the local quantum efficiency of the chip, producing a photometric nonlinearity; at still higher accumulated charge levels ($\sim 10^4$–10^6 electrons $\times$ pixel^{-1}, depending on the device) electrons are forced out of their potential wells and contaminate surrounding areas of the frame.

Given all these advantages and disadvantages, the overall gain in the efficiency of a CCD over a photomultiplier for measuring many stars in a small *uncrowded* field is substantial, but not as large as one would infer from the increased quantum efficiency and the possibility of recording dozens of stars at once. These gains are partially lost to the greatly increased difficulty of recording, transporting, and analyzing the vast quantities of data. The current limits to the accuracy of CCD photometry also deserve some special discussion. A. R. Walker (1984) has reported routinely obtaining subpercent broad-band photometry with the SAAO-UCL CCD camera: assuming that his phrase "mean deviation" refers to the mean absolute deviation of his observations from the best transformations, then the standard deviations of his magnitudes and colors generally lie in the range 0.004–0.005 mag.

[1]Because in recent years charge-coupled devices have become the instrument of choice for most electronic image detection at optical and near-infrared wavelengths, the remaining discussion will concentrate on them to the exclusion of other detector systems. For discussions of the properties of CCDs in astronomical applications, see, for example, G. A. H. Walker *et al.* (1984), Massey (1985), and Jacoby (1985).

[2]These figures are for the currently popular RCA and Texas Instruments CCD chips, whose active areas are roughly 300 $\times$ 500 and 800 $\times$ 800 pixels, respectively. New CCD chips being manufactured by Tektronix will contain 2048 $\times$ 2048 $>$ 4 $\times$ 10^6 pixels.

Work of this accuracy requires both extraordinary care and equipment which is of a quality not found in most CCD cameras. Here I will mention just two examples of the sort of limitations that are placed on *most* CCD photometry. First, although the fundamental standard stars in most photometric systems are so bright that even with one-meter class telescopes they would saturate the detector in a small fraction of a second, few CCD cameras have shutters permitting exposures of less than one second, whose durations and uniformity over the area of the chip are controlled to one part in a thousand. Most astronomers, therefore, are limited to secondary standard stars, whose photometric indices are themselves typically uncertain by a percent or more. Second, it is an easily observable fact that the form of the flat-field correction is dependent on the color of the illuminating light. Flat fields showing relative variations of order 1%–2% for a one-magnitude change in (say) the $(B-V)$ color of the illumination are not uncommon. This means that *the flat-field correction that should be applied to a program frame depends on the color of the stars one wishes to measure*. In principal, the image of a red star should be corrected using a different flat-field calibration than that for a blue star lying next to it. This limitation of CCD photometry may be due to the fact that, unlike conventional photomultiplier photometry, different parts of the detector are illuminated through different parts of the filter. If the wavelength-dependent throughput of the filter itself were to vary by a few percent across its face, this could account for the observed effect. Although in principle this difficulty can be resolved by sufficiently elaborate calibrations, in practice most researchers tend to satisfy themselves with standard errors of 1% or more.

The photometric difficulties become much worse when astronomers try to study the *crowded* fields which, after all, contain much of the most interesting stellar astrophysics which can be addressed by photometric techniques. Remote open clusters, globular clusters, Magellanic Cloud clusters, and resolved stars in external galaxies are usually so densely packed that the images of most stars overlap with the images of their nearest neighbors. When stars are blended it is obviously impossible simply to sum the CCD data-numbers corresponding to the image of each star and then to subtract an allowance for the diffuse sky emission. On the other hand, a rather different technique may be used for the measurement of stellar magnitudes in crowded fields. The inherent linearity of the detector is exploited *via* the assumption that the two-dimensional intensity profile of each star is known (or, at least, knowable). The positions and relative magnitudes of stars can then be accurately determined by using numerical fitting techniques to match the model profile to the observed light distribution. If nature has attempted to confuse us by blending the light of two or more stars, we retaliate by fitting a model in which two or more of the expected stellar profiles are superimposed: each model stellar profile is shifted in x and y and scaled in intensity, and one or more parameters describing the local distribution of diffuse sky light are varied, until a satisfactory fit of the overall model to the image data is achieved. Only the linearity of the detector and its freedom from adjacency effects make this technique possible.

A number of computer programs which exploit this capability of photometrically linear image detectors to perform stellar photometry in crowded fields have been developed in recent years, among which are those of Tody (1981), Penny and Dickens (1986), Buonanno and collaborators (1979, 1983), and Lupton and Gunn (1986). This paper presents a detailed description of another such computer program, DAOPHOT, which is now in use and undergoing additional development at the Dominion Astrophysical Observatory. DAOPHOT had its origins in the code POORMAN, which was developed by J. R. Mould and K. Shortridge of Caltech. A copy of this program was generously given by Mould to L. L. Stryker of the DAO, and she and E. W. Olszewski implemented it on the computer at DAO. Since a number of us at the DAO were interested in reducing more densely crowded fields than POORMAN was designed to handle, I replaced most of the existing algorithms in POORMAN with others of my own writing and added several new routines as well. At present, DAOPHOT retains the basics of the synthetic aperture photometry algorithm of POORMAN and Shortridge's utility routines for manipulating disk data-structure files; the other algorithms that will be described below were developed independently of POORMAN.

In Section II below I will give a general overview of the tasks that DAOPHOT is intended to perform. In Section III I discuss a few of the pivotal algorithms in greater detail. How the program is used is described in Section IV and some current performance characteristics are given. Finally, in Section V I discuss some of the known shortcomings of the code and sketch some possible future improvements.

II. Overview of the Problem

In the Introduction above I have mentioned both the basic problem and the essence of a solution: many astrophysically interesting populations of stars for which photometric measurements would be useful are so densely packed on the sky that any given star is likely to be blended with its neighbors. In many such cases useful photometry may be obtained by fitting a superposition of predicted stellar-light distributions, plus a model for the diffuse night-sky luminosity, to a two-dimensional digital image obtained by a photometrically linear detector. In practice, the computer software which will carry out this analysis should perform several functions:

1. Provision should be made for the automatic detec-

tion of those brightness enhancements in the digital picture which appear to represent the images of stars. When the number of stellar images in a data frame is small, this task is easily performed by a human eye and brain examining a pictorial representation of a digital image. However, currently-available detectors are easily capable of resolving many $\times 10^3$ stars in a single frame, and scores of such frames can be obtained in a single night. To identify and mark all the stars by hand and eye would be an atrocious waste of human effort. Furthermore, the CCD can record data over a much wider dynamic range of brightness than the human eye can perceive, all of which can be exploited by a computerized pattern-recognition algorithm. Thus, there is need for a reliable, automatic star-finding routine. For greatest usefulness the algorithm should have some ability to distinguish brightness enhancements which correspond to actual stars, however blended, from those caused by galaxies, cosmic ray events, cosmetic flaws in the detector, and random noise.

2. It should be possible to perform synthetic aperture photometry, preferably through a series of concentric apertures with a range of sizes. This is the easiest and potentially the most accurate method for obtaining instrumental photometry for bright, uncrowded stars, because it involves simple counting of detected photons. Thus, it avoids the computational effort of obtaining and fitting model stellar profiles, and it eliminates worry about the adequacy of the assumed model star image.

3. It must be possible to define a model stellar-light profile for each frame with the highest possible accuracy. If the brightness distribution in a star image can be specified with a maximum error not exceeding, say, 1% of its peak value, and if this star is blended with a companion 2.5 mag fainter, then the inaccuracy in our knowledge of the profile of the bright star may lead to an error of order 10% in our photometry of the faint companion. If the companion were 5 mag fainter than its primary, the error in its photometry could be $\sim$ 100%. Hence, model profiles which are accurate to a few percent or better everywhere would seem to be absolutely necessary for astrophysically useful stellar photometry in crowded fields. If the stellar profile is expected to depend significantly on position in the frame, on the magnitude or color of the star, or on any other variable, this must be allowed for in the definition of the model.

4. Obviously, once such a model profile has been derived, a procedure must exist which can determine the optimum spatial arrangement and intensity scaling of an ensemble of such profiles, together with a smooth, continuous brightness distribution representing the contribution of diffuse sources, to match the actual image data. To be effective, this procedure must utilize some fairly sophisticated and reliable artificial intelligence, such as the ability to recognize and disregard defective data and the power to decide when a star must be added to or deleted from the model to represent the data more realistically.

5. Once model profiles have been fitted to the stars in a digital image, it should be possible to subtract those positioned and scaled profiles from the data array, producing a two-dimensional map of the fitting residuals which may be displayed and examined by eye or subjected to further computer analysis. The human eye and, to a certain extent, suitable computer software can then identify failures of the model to fit the data: (a) images of stars which were in the frame but were not recognized by the star finder, perhaps due to having been blended with brighter companions; (b) stars whose images contained defective pixels; (c) images of galaxies which are almost, but not quite, stellar; and other, similar departures from a perfect fit may be recognized and flagged.

6. Given knowledge of the observed stellar profile, it is useful to be able to add realistic-looking artificial stars into the frame. Then the same algorithms which are used to find and photometer program stars can be used to find and photometer these artificial stars, whose derived positions and magnitudes can then be compared with their input values. This permits the astronomer to study the systematic and random properties of the analysis procedures within the context—crowding conditions, seeing, sky brightness, and so on—of the data frames actually being reduced. Such analyses can determine the discovery probabilities of stars as a function of magnitude (e.g., McClure *et al.* 1985, §IV.B), and the error distribution of the positional and photometric measurements—not merely the standard deviation of small, random measuring errors, but also the probability of large blunders, as when two blended stars are measured as if they were one (e.g., Smith *et al.* 1986, Fig. 5).

III. Specific Algorithms

In this section I present somewhat more detailed descriptions of the methods DAOPHOT uses to carry out the general tasks described above. The discussion will concentrate on general principles rather than on details of the coding; readers who are interested in more specific information are invited to request a copy of the source code for DAOPHOT, which is heavily documented. The reader should also note that a number of the routines described below are mutually independent. For instance, a completely different star-finding routine could be substituted for the one described here without affecting the other routines in the program.

A. *Finding Stars*

The task of finding stars involves:

1. detecting and locating small, *positive* brightness enhancements within an image;
2. distinguishing legitimate stellar images from
   a. random noise peaks in the data,

b. images of galaxies or other extended astronomical objects,
c. packets of electric charge generated by particles (cosmic rays, radioactive decay products) which impacted on the detector, and
d. cosmetic defects; and

3. recognizing when a seemingly extended object consists of two or more overlapping stellar images.

Even the human eye and brain are not capable of carrying out any of these tasks with absolute reliability, and the most modern computer software is no more trustworthy. However, algorithms may be devised which are fast enough, and which are right often enough, for valuable scientific results to be obtained.

In attempting to find stars in a two-dimensional image, the star-finding algorithm of DAOPHOT (called "FIND") exploits several pieces of a priori knowledge which the astronomer can provide: (1) the readout noise of the detector, in analog-to-digital converter units (ADU); (2) the number of photoelectrons corresponding to 1 ADU; (3) the observed brightness level, in ADU, above which the detector fails to operate linearly; (4) the approximate size (full width at half maximum) of an unresolved stellar image; and (5) the smallest brightness value which a pixel subject only to legitimate random noise may be expected to have. The first three of these quantities may generally be found in the users' manual for the CCD camera or may be determined independently by analysis of calibration frames; the fourth may be computed from the known image scale and seeing records or estimated by eye from the data frames themselves; and the fifth is readily computed, as will be discussed in Section IV below. Finally, the astronomer can also often provide (6) a catalog of known defective pixels, rows, and columns within the CCD chip.

Given these pieces of information, FIND attempts to locate stars by going through the picture pixel by pixel, at each location asking the question, "If there is a star centered in this pixel, how bright is it?" For each pixel, a numerical answer to this question is estimated by fitting an analytic Gaussian profile to the brightness values in a surrounding subarray of pixels. If by chance there should happen to be a star image centered in that pixel, the fit will be good and the central height of the fitted Gaussian profile will be proportional to the brightness of the star. If there is *not* a star image centered in the pixel—if the pixel lies in an empty region of sky or in the *wings* of a star image—then the central height of the best-fitting Gaussian will be near zero or even negative. Thus, a location in the image where the central height of the best-fitting model Gaussian profile achieves a large, positive value probably lies near the center of a star image.

In more specific detail, the computations may be described as follows. Let the observed brightness (in ADU) in the (i,j)-th pixel be represented by $D_{i,j}$, and let G represent the unit-height, circular, bivariate Gaussian function

$$G(\Delta i, \Delta j; \sigma) = e^{-(\Delta i^2+\Delta j^2)/2\sigma^2} .$$

Then the central brightness of the hypothetical stellar profile which best fits the pixels around the point i_0,j_0 in the picture, H_{i_0,j_0}, and an estimate of the local background brightness, b, are given by the set of equations

$$D_{i,j} \doteq H_{i_0,j_0} G(i-i_0, j-j_0; \sigma) + b, \qquad (i,j) \text{ near } (i_0,j_0) ,$$

where the symbol "$\doteq$" represents a least-squares fit to the data for some set of pixels (i,j) in a defined region around and including (i_0,j_0). The numerical value of H_{i_0,j_0} is then given by simple linear least squares:

$$H_{i_0,j_0} = \frac{\Sigma(GD) - (\Sigma G)(\Sigma D)/n}{\Sigma(G^2) - (\Sigma G)^2/n} , \qquad (1)$$

where n is the number of pixels used in the fit. The point (i_0,j_0) is allowed to vary over the pixels in the input image, and the values of H thus computed are stored in an array nearly as large as the original picture. (For reasons which will be given below, a margin equal in width to the radius of the fitting region for H is left around the perimeter of the image.) Local maxima in the derived array H are then identified and regarded as candidate astronomical objects. Equation (1) is the arithmetic equivalent of convolving the original picture with a truncated Gaussian which has been lowered to give it a zero integral: the reader can easily show that equation (1) is equivalent to

$$H_{i_0,j_0} = \sum_{i,j} (W_{i-i_0,j-j_0} D_{i,j}) ,$$

where

$$W_{\Delta i,\Delta j} = \frac{G(\Delta i, \Delta j; \sigma) - (\Sigma G)/n}{\Sigma(G^2) - (\Sigma G)^2/n} .$$

Note that $\Sigma W \equiv 0$.

Figure 1 is a graphic representation of the conversion of the original image data D to the derived data array H, in the analogous one-dimensional case. From left to right in Figure 1, the simulated objects are (a) a single, isolated star with critical sampling (σ = 1 pixel), (b) a blended pair of stars, (c) a barely-resolved galaxy, (d) a cosmic-ray hit, and (e) a defective, low-valued pixel. These objects are superimposed on a sloping sky background. Figure 1(a) represents the appearance of these simulated objects in image data, D, while Figure 1(b) displays the result of convolving D with the truncated, lowered Gaussian function, shown as connected x's, to obtain H. Searching for local maxima in H rather than in D has several significant advantages:

1. The search is optimized for objects the size of stellar images. First, the estimated FWHM of the stellar profile, which the astronomer provides to the program, immediately implies the value of σ to be used for the fitting function G: $\sigma = FWHM/2.355$. Second, the radius of the

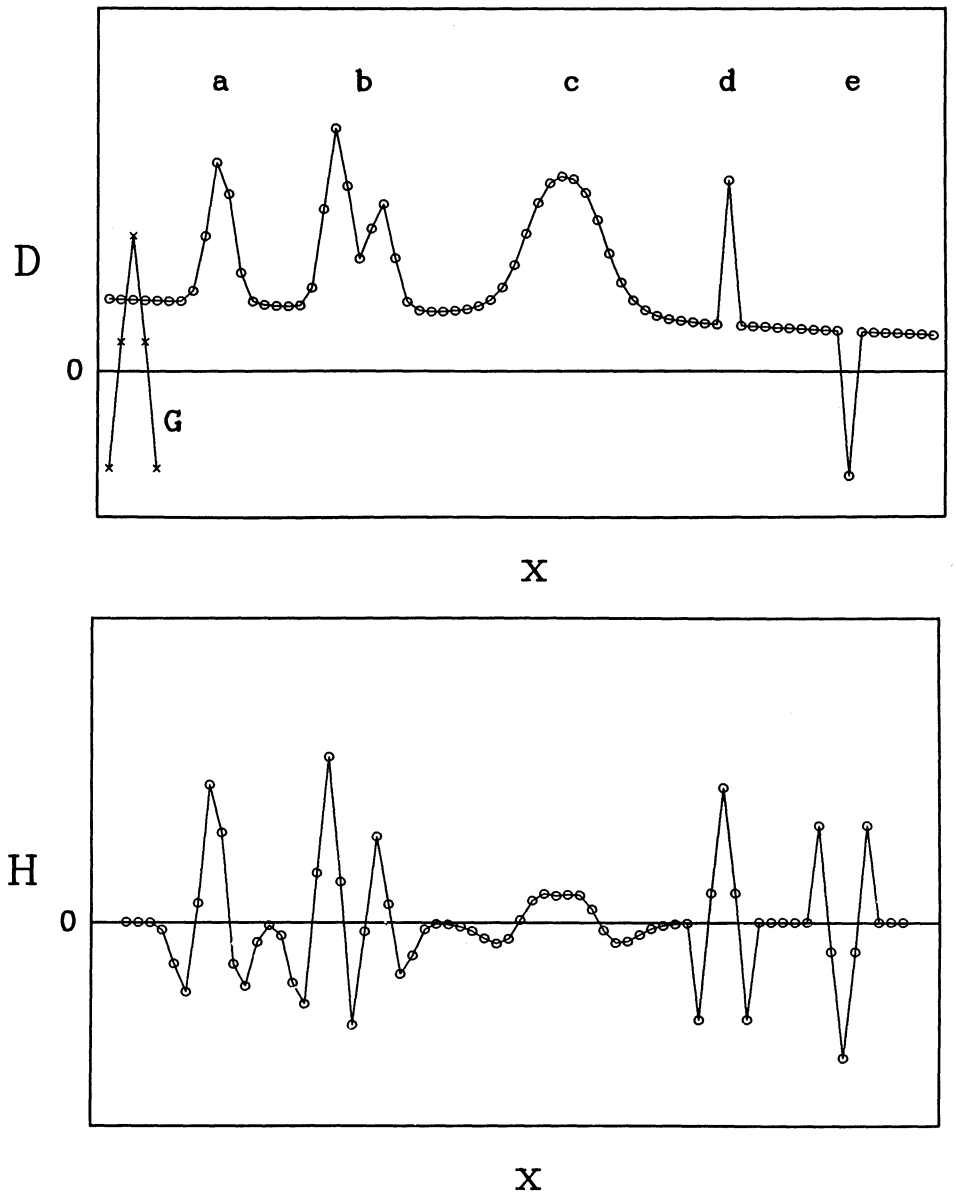

FIG. 1–(a) A synthetic one-dimensional brightness distribution $D(x)$ to illustrate the operation of the FIND algorithm of DAOPHOT. The brightness enhancements represent (a) a star image, (b) a blended pair of star images, (c) a barely-resolved galaxy image, (d) a cosmic-ray event, and (e) a low-valued bad pixel. The connected x's identified with a "G" represent the lowered Gaussian function with which the brightness data will be convolved. (b) The result of convolving $D(x)$ with the lowered Gaussian function to produce $H(x)$. Items to note: (1) The broader galaxy profile is suppressed by the convolution, in comparison with the stellar profiles. (2) The background value of H is zero, even though in D the background was nonzero and sloping. (3) The blended double is better separated in H than in D. (4) In this example, the stellar images are depicted as critically sampled, FWHM ~ 2.4 pixels, and are hardly distinguishable from the cosmic ray in H; the cosmic ray and the maxima in H found on either side of the low-valued pixel will have to be rejected on the basis of their *sharp* indices.

region around (i_0,j_0) which is included in the sums is also determined from the FWHM specified. Because of the nature of least-squares fits, objects with characteristic radii $>> \sigma$, such as diffuse nebulae, or those with radii $<< \sigma$, such as image defects, will tend to have smaller values of H than a legitimate star-like image (radius ~ σ) with the same central value of D. Thus, searching through H rather than D introduces a beneficial bias against significant, positive, but nonstellar brightness enhancements. (For the same reason, this is *not* a good algorithm to use when a complete inventory of the stars *and galaxies* in a digital image is desired.)

2. The local sky brightness, b, does not appear in the calculation of H—the quantity H represents the estimated central brightness of the star *above the local sky*. The user specifies not an absolute brightness threshold

above which candidate objects are found but rather a brightness enhancement. Similarly, because the Gaussian function being fitted to the data is symmetric, any slope in the sky cancels in the computations for H, as may be seen in Figure 1. Thus, if the diffuse sky background changes across the frame, such as for stellar objects embedded in large, partially-resolved galaxies or in regions of diffuse nebulosity, the brightness criterion for detection will not, to first order, be a function of the local sky brightness. (However, increased noise may enhance the probability for faint stars to be rejected on the basis of secondary criteria, such as the sharpness and roundness indices discussed below.)

3. Peaks in H tend to be narrower than peaks in D, which helps to distinguish blended binaries (cf. Fig. 1).

4. The significance level of a detection is readily estimated. According to the theory of least squares (e.g., Hamilton 1964, Ch. 4), if s_1 represents the random error of a brightness measurement in a single pixel in a star-free region of the frame, then the random scatter in H when no stellar image is present will be given by

$$s_H^2 = \frac{s_1^2}{\Sigma(G^2) - (\Sigma G)^2/n} \quad . \tag{2}$$

The value of s_1 may be computed from first principles

$$\begin{aligned} s_1^2 = &\text{ (readout noise in ADU)}^2 \\ &+ \text{(approximate sky level in ADU)/} \\ &\text{(number of photons per ADU)} \quad , \end{aligned} \tag{3}$$

or it may be measured directly in a star-free region of the frame. Now, given s_1 and $\Sigma(G^2) - (\Sigma G)^2/n$ (which the program computes once the expected FWHM of a stellar image has been specified), s_H may be calculated, and it is then trivial to specify the minimum value of H, $H_{\min}$, which will represent a five-sigma detection, say. Note that if the sky background *does* vary strongly across the frame, then using a minimum value of H for the detection of objects will produce a variation in the *significance* of discoveries as a function of background, but not a direct variation in the limiting magnitude. In principle, $H_{\min}$ could have been made a function of b so that the significance limit would be independent of the local background, but then the magnitude limit would vary.

5. Bad pixels, such as known cosmetic defects which the user has masked out, or brightness values which are found during reduction to be above the known saturation level of the chip, or more than 5 s_1 (say) below the average sky level, are omitted from the fits for H. While the values for H found near these omitted pixels will be slightly noisier than average because of the missing data, the omission of these pixels will cause far less contamination of the principal star-finding criterion than their inclusion would. (It is this ability to omit bad pixels from the sums in eq. (1) which makes it more accurate to describe this algorithm as a series of least-squares fits rather than as a convolution, even though the arithmetic would otherwise be the same.)

Having computed an array H comparable in size to the original data frame D, the routine runs through H looking for the positions of local maxima, that is, for values of (i_0,j_0) such that $H_{i_0,j_0} \geq H_{\min}$, and $H_{i_0,j_0} \geq H_{i,j}$ for all pixels (i,j) less than some distance from (i_0,j_0) (this distance, again, is determined from the estimated stellar FWHM supplied by the user). This produces a list of positive, apparently significant brightness enhancements which is biased against known defective pixels, major cosmic-ray events, and faint galaxies; some nonstellar objects and less-obvious image flaws remain to be excluded.

In designing further selection criteria, it was kept in mind that this algorithm was intended to work in crowded fields; consequently, criteria stressing properties of the image *cores* rather than the shape of image *wings* were devised, in the hope that they would be less confused by the possible presence of neighboring stars. Furthermore, an effort was made to minimize both the number and the complexity of the pattern-recognition criteria so that they would retain validity for the faintest detections possible: criteria involving the loss of too many degrees of freedom or resting on subtle details of image profile require high signal-to-noise ratios in the stellar image for meaningful definition. Also, for speed, criteria which could be computed from simple weighted sums of raw image data were preferred.

A dominant class of nonstellar detections, including most cosmic-ray events and many cosmetic flaws, may be recognized by their being narrower than a seeing-broadened stellar image. In addition, narrow local maxima in H tend to be found on all sides of a defective pixel with a brightness value less than the average sky background, as illustrated in Figure 1. Some other detections, such as nonnucleated galaxies, appear broader than stars. Therefore, a sharpness criterion was devised which compares H_{i_0,j_0} (the height of the best-fitting Gaussian function) to the height of a two-dimensional delta-function, d, defined by taking the observed intensity difference between the central pixel of the presumed star image and the mean of the remaining pixels used in the fits of equation (1):

$$\begin{aligned} d_{i_0,j_0} \equiv D_{i_0,j_0} / \langle D_{i,j} \rangle \quad , \quad (i,j) \text{ near } (i_0,j_0) \text{ but } \neq (i_0,j_0) \quad , \\ sharp \equiv d_{i_0,j_0}/H_{i_0,j_0} \quad . \end{aligned} \tag{4}$$

For a very narrow profile, such as that caused by a cosmic-ray event, all of the intensity will be contained in the central pixel, while the computed central height of the Gaussian function which best fits this intensity spike and the area of the frame around it will be somewhat less, hence $sharp \gtrsim 1$. Conversely, for the spurious detections found around low-valued pixels, d should take on a relatively small value ($0 \lesssim sharp << 1$). Some faint galaxies may also have small values of d and, hence, of $sharp$. For

moderately-peaked objects, such as star images, which are presumed to have a constant profile shape, *sharp* should scatter about a value significantly less than unity and greater than zero. Detections whose sharpness indices fall significantly above or below this value are easily identified and excluded (Fig. 2(a)).

Another class of false detections arises in the charge overflow columns and rows from grossly overexposed objects in the frame. These are readily recognized because they are much more elongated in x or y than star images. Therefore, a roundness criterion is computed, which compares the peakedness of the enhancement in the x-direction with that in the y-direction: the height of the best-fitting one-dimensional Gaussian function of x,h_x, is compared to the height of the best-fitting one-dimensional Gaussian function of y,h_y. These fits are effected by a formula involving sums of the image-brightness values identical in form to equation (1)—the only difference is that one-dimensional Gaussian functions $g_x(\Delta i) = e^{-\Delta i^2/2\sigma^2}$, $g_y(\Delta j) = e^{-\Delta j^2/2\sigma^2}$ are substituted for G. If

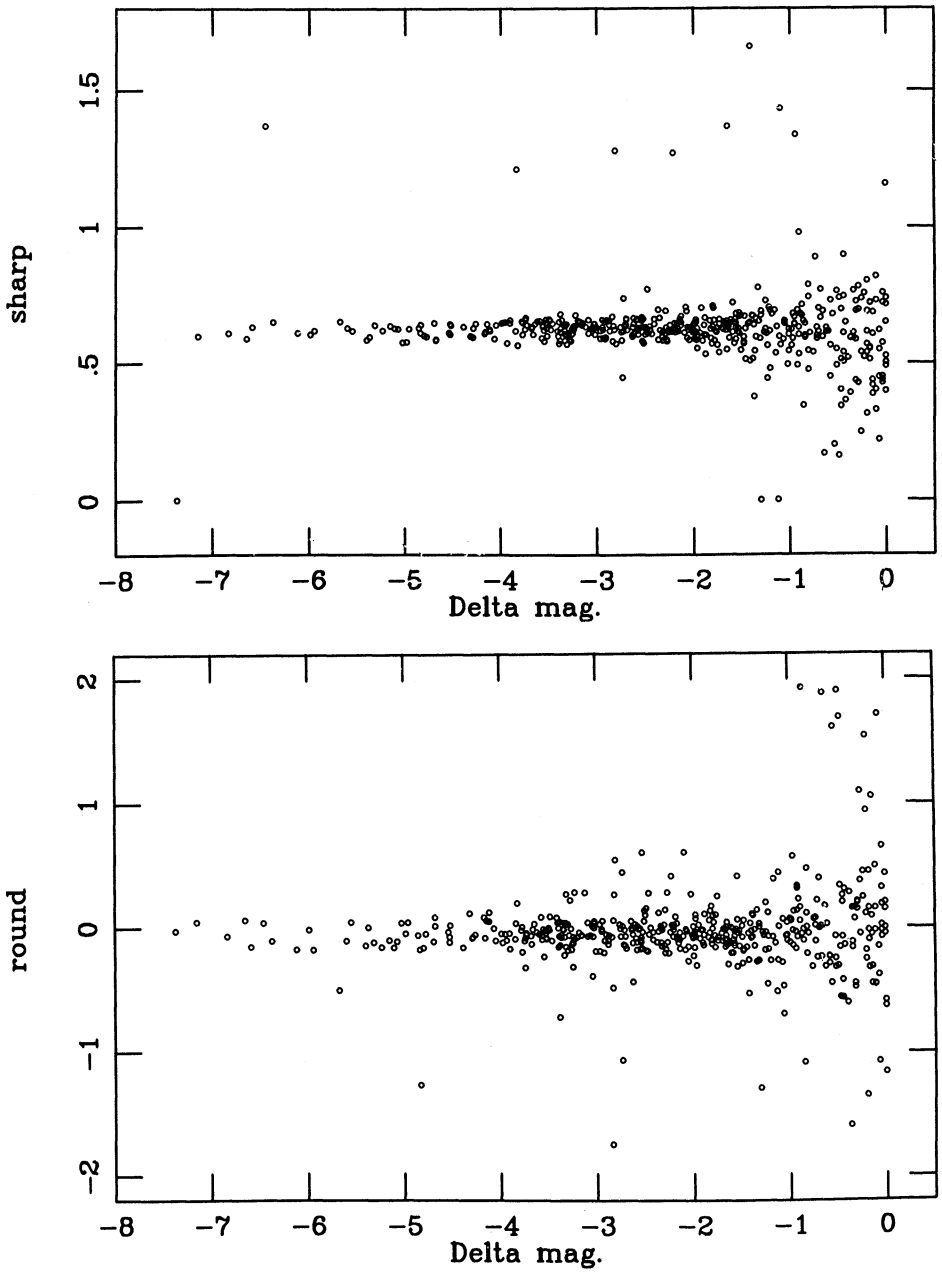

FIG. 2–(a) The distribution of *sharp* as a function of magnitude for candidate detections in a typical critically-sampled image. The brightness is measured in magnitudes relative to the detection threshold: Δ mag $= -2.5 \log(H_{i_0,j_0}/H_{\min})$. (b) The same as (a), except showing *round* as a function of brightness. The usual cutoffs, $0.2 \leq$ *sharp* ≤ 1.0, $-1.0 \leq$ *round* ≤ 1.0 would be adequate in this case, but more stringent acceptance limits might be even better.

the candidate detection is really a charge-overflow column, $D_{i,j}$ will be a strongly-peaked function of x and a nearly flat function of y ($h_x >> 0$ and $h_y \approx 0$), and conversely for detections in charge-overflow rows. Star images, on the other hand, are strongly peaked functions of both x and y, so that for round, well-guided images $h_x \approx h_y$. Thus, a roundness criterion

$$round \equiv 2\left(\frac{h_y - h_x}{h_y + h_x}\right) \tag{5}$$

readily distinguishes stars ($round \approx 0$) from bad rows and columns ($round \approx \pm 2$) (Fig. 2(b)). This procedure does not select against objects which are elongated in a direction inclined to the rows and columns of the data arrays, such as edge-on spiral galaxies.

In the determination of both *sharp* and *round*, known or recognizably bad pixels may be omitted from the sums. This broadens somewhat the distribution of index values for legitimate stars, but still allows reasonable classification in the majority of cases. The roundness criterion can only be meaningfully computed, however, when there are a significant number of valid pixels on all four sides of the candidate detection. For this reason *round* cannot be computed for pixels lying near the edge of the frame and, therefore, FIND does not bother to compute H in a band a few pixels wide around the perimeter of the image, as mentioned above.

Nearly always, adequate limits on the sharpness and roundness criteria for real stars are $0.2 < sharp < 1.0$, and $-1.0 < round < +1.0$, but if desired in individual cases, more appropriate limits can easily be determined by examining plots like those in Figure 2. Once a given brightness enhancement has passed muster as a probable stellar image, its brightness relative to the detection limit is crudely estimated: $\Delta\text{mag} = -2.5 \log(H_{i_0,j_0}/H_{\min})$.

B. *Aperture Photometry*

The aperture photometry routine ("PHOT") which is currently in DAOPHOT is little changed from the one in POORMAN, which was kindly donated to the DAO by J. R. Mould of Caltech. The aperture routine of POORMAN in turn rests on algorithms originally developed for the Mountain Photometry code of Kitt Peak National Observatory, so they will not be discussed in detail here. However, the way in which the sky brightness is estimated will be discussed in some detail, even though this is not an original algorithm developed at the DAO, because it is important for the credibility of the photometry obtained by the rest of the program.

Just as for photometry obtained with a photomultiplier, synthetic aperture photometry rests on integrating the total light received by the detector from some solid angle of sky which includes a star of interest. The typical amount of light received per unit solid angle *exclusive* of that particular star is independently estimated from a nearby region, and the difference between those two numbers gives an estimate of the total amount of light received from the star alone. In photomultiplier photometry, the sky brightness is measured by placing the aperture in some apparently star-free region of sky, identified either by eye at the telescope or from examination of a deep photographic reproduction of the star field. Similarly, in interactive synthetic-aperture photometry the astronomer can select a star-free region of the data frame by eye and use the mean brightness measured there as an estimate of the number of sky photons per unit solid angle. However, in *automatic* CCD photometry, where the astronomer cannot afford the time to make these decisions for each star, some other estimator of the local sky brightness must be employed.

At this point it is useful to reflect in more philosophical detail on just what is meant by "sky brightness." At any given position in the data frame, the data numbers supposedly representing detected photons come from any of several sources: (1) detected stars, (2) undetected stars, (3) localized image defects and cosmic-ray events, and (4) diffuse sources, including but not necessarily limited to (a) the terrestrial night sky, (b) diffuse astronomical sources (e.g., zodiacal light, emission nebulae, underlying galaxies), (c) scattered light in the camera, and (d) dark signal. The distinction between (1) and (2), detected and undetected stars, is unimportant here but will be of interest below in the discussion of profile-fitting photometry.

For aperture photometry, the problem is this: having identified a significant local-brightness maximum in a digital image and presuming this to represent a star, having centered a synthetic aperture at that position, and having summed the data numbers corresponding to pixels lying within the aperture (for precision, partial weight is given to pixels lying partly within the aperture), how many data numbers would have been detected in the aperture if the particular star in question were not there? This quantity is to be subtracted from the observed sum to correct for all the sources of signal other than the program star.

A usable answer is obtained as follows. Consider a region of the frame which is near the position of the program star but far enough away that a negligible fraction of the star's own light falls within it. It should be symmetrically disposed about the star, so that any gradient in the distribution of the other sources of signal (a falloff in the surface density of stars, or spatial variations in diffuse nebulosity, for instance) will cancel at least to first order. For the greatest possible precision, the comparison region should contain many more pixels than the stellar aperture, so the random scatter in the sky region will not tend to add significantly to the unavoidable uncertainty of the star measurement. A circular annulus with an inner radius several times the stellar FWHM is a convenient comparison region meeting these specifica-

tions. Now, given the set of brightness values found in the comparison region, the *mode* of that set is a good estimator of the sky brightness per pixel in that part of the frame. This number, representing the most-frequently observed brightness value in the pixels of the comparison region, is to be multiplied by the area of the stellar aperture, in $(\text{pixels})^2$, and subtracted from the summed data numbers within the aperture.

In this practical method of dealing with the problem, we use the modal value of the brightness measurements in the comparison region not because it is a robust estimator of the mean sky brightness, as is sometimes assumed, but rather because it is itself a maximum-likelihood estimator—it represents the most probable value of the brightness of a randomly chosen pixel in that part of the image, including *all* sources of signal other than the star in question. As a maximum-likelihood estimator the mode is guaranteed to lie near the center of the shortest possible confidence interval of a given confidence level. In other words, the probability of a large error in the predicted value of the non-program-star signal within the aperture is minimized when the mode is taken as the estimator. This technique differs in an important philosophical sense from certain other sky-finding algorithms, such as the asymmetric clipping algorithm of Ratnatunga and Newell (1984), which are designed to answer a different question altogether: what would the diffuse sky brightness be in some specified region of the frame if *none* of the contamination sources of types (1), (2), and (3) (see above) were present? In the case of synthetic-aperture photometry, we know that these additional sources *are* present and that they are as likely to appear within the stellar aperture as anywhere else in that part of the frame. Therefore, an estimate of the sky brightness *allowing for*, not *corrected for*, the presence of these contaminants is required.

C. *The Point-Spread Function*

In the case of the *Hubble Space Telescope*, whose image-forming properties are well mapped a priori and whose guiding will be so good that its effect on image profiles may be either imperceptible or derivable from the fine-guidance error signals, the point-spread function (PSF)[3] may be known independently of the data frames one is trying to reduce. For ground-based imagery the vagaries of seeing and guiding are such that a different model PSF must be adapted to the observed profiles of bright stars in each data frame. Two approaches to the derivation of model image profiles from observed ones spring immediately to mind.

1. The analytic approach employs some mathematical function $I \equiv I(x - x_0, y - y_0)$ to describe intensity as a function of two-dimensional distance from the centroid, (x_0, y_0), of a star; parameters in this function are adjusted to give the best possible representation of actual image profiles in a given frame. The function must be sufficiently flexible to map both the complex instantaneous seeing profile (see King 1971, for example) and the effects of erratic image motion and telescope tracking to an accuracy of the order of a few percent of the peak intensity. The analytic PSF has the major advantage that it can be integrated numerically over the area of each pixel in a stellar image, so the adverse effects of finite pixel size in an undersampled image are minimized. The disadvantages are (a) these numerical integrals can be very time consuming, and (b) imperfectly formed images may require either an inconveniently large number of parameters for their description or else some compromise in the maximum error that can be tolerated. Analytic point-spread functions were developed for the reduction of linearized photographic images (e.g., Penny 1976, 1979; Buonanno *et al.* 1979) and have more recently been applied to data from image photon-counting systems and CCDs (e.g., Peterson *et al.* 1978; Buonanno *et al.* 1983; Penny and Dickens 1986). They have also been used for astrometric measurements (Monet and Dahn 1983).

2. The empirical approach uses bivariate interpolation to estimate intensity values at fractional-pixel positions within an observed stellar profile. The images of several bright stars can be interpolated to a common grid and summed to increase the signal-to-noise ratio, and the resulting PSF can then be numerically interpolated and scaled to match the observed data for each program star. This approach has been adopted in the computer programs RICHFLD (Tody 1981; Stryker 1981; and references therein) and VISTA (Lauer 1983). The empirical point-spread function has the advantages (a) that interpolations can usually be computed much more quickly than two-dimensional numerical integrals of arbitrary analytic functions, and (b) that it avoids entirely the question of whether enough free parameters have been included in the model profile to represent the true PSF adequately. It has the striking disadvantage that it is poorly suited to undersampled data, because the interpolation formulae make simplifying assumptions about the shape of the true point-spread function between the sample points. For instance, cubic-spline interpolation (e.g., RICHFLD) assumes that the fourth and higher derivatives of the true PSF are negligible; sinc interpolation (e.g., VISTA) assumes that all spatial frequencies above the Nyquist frequency have amplitudes effectively zero. It turns out that neither of these assumptions is adequate when the data are undersampled. When the full width at half maximum of the PSF is less than two to three pixels, the specific brightness of the stellar profile can change by an order of magnitude from one pixel to the next; under these circumstances none of the textbook formulae can guarantee interpolations correct to a few percent. The breakdown of

[3]The *point-spread function* is the two-dimensional brightness distribution produced in the detector by the image of an unresolved source, such as a star.

the method is most severe near the center of the profile where the amplitudes of the neglected higher-order variations are greatest; unfortunately, this is precisely where most of the photometric information resides.

The error of interpolation could be compounded if the user attempts to improve the signal-to-noise ratio of the PSF by combining the observed profiles of several bright stars in the frame. When different stars are superimposed, they must be interpolated to a common grid—for instance one precisely centered on each stellar centroid—before they are added together. Otherwise, the summed PSF would be broader than the individual star images because one star may be centered on a pixel, while the centroid of another may lie near the corner of a pixel. Subsequently, when the PSF is fitted to a program star the mean model profile must be interpolated again to match the sample grid of the program star. This need to interpolate twice, once in the original data for each PSF star to align their centroids, and again in the mean PSF to match the sampling for the program star, could double the interpolation error inherent in the empirical approach. However, the error of the second interpolation can be reduced by storing the empirical PSF as a more finely-spaced grid of points than the original data; for example, since the error of cubic-spline interpolation in one coordinate goes as $\Delta x^4 \frac{\partial^4(\mathrm{PSF})}{\partial x^4}$, the second interpolation can be made arbitrarily precise by choosing suitably small values of Δx and Δy for the tabular PSF. Of course, this could produce an empirical PSF requiring an array comparable in size to the original image. Furthermore, this trick cannot be used to eliminate the error of the first interpolation, since in this step the sample spacing is imposed by the detector.

DAOPHOT uses a model for the point-spread function which attempts to exploit the best aspects of both the analytic and empirical methods. The "PSF" routine first fits an analytic bivariate Gaussian function to the central regions of a bright star as follows (King (1971) has shown that the core of a seeing-broadened image is nearly Gaussian). For each pixel, the two-dimensional integral of the Gaussian function over the area of that pixel is computed, and the centroid, one-sigma widths in x and y, and the height of the Gaussian function are determined from the actual intensity data by standard nonlinear least squares, in which

$$D_{i,j} - sky \doteq H \int_{i-\frac{1}{2}}^{i+\frac{1}{2}} \int_{j-\frac{1}{2}}^{j+\frac{1}{2}} G(x-x_0,\, y-y_0;\, \sigma_x, \sigma_y)\, dx\, dy \quad , \qquad (i,j) \text{ near } (i_0,j_0) \tag{6}$$

is solved for H, x_0, y_0, σ_x, and σ_y. This Gaussian serves as a first-order approximation to the actual stellar profile. The observed residuals, $\Delta_{i,j}$, of the actual data from the best-fitting Gaussian are computed at pixel locations (i,j) around the star and are stored in an array. Then, a look-up table, $C(x-x_0, y-y_0)$, representing corrections from the Gaussian function to the actual, observed stellar profile is computed by bivariate cubic interpolation within this array of residuals; the look-up table is centered on the centroid of the star, and values of C are tabulated at half-pixel intervals in $x-x_0$ and $y-y_0$.

If the user wishes to improve the signal-to-noise ratio of the point-spread function by adding another star, then the current PSF is shifted and scaled to the new star by the same fitting routine as will be used subsequently for the photometric profile fits, in order to ensure that the centroid of the new star is matched as precisely as possible to that of the current PSF. Then the scaled Gaussian function is subtracted from the data for this new star to produce another array of residuals. These residuals are again interpolated to half-integral values of $x-x_0$ and $y-y_0$ and are added into the look-up table of profile corrections, C. This procedure may be repeated with as many stars as desired. Later, when this PSF is evaluated to compare with the brightness observed in a pixel some nonintegral distance from the centroid of a program star, the program first integrates the Gaussian approximation over the area of the pixel (eq. (6)). Then, it computes a correction from the integral Gaussian function to the model PSF by bivariate cubic interpolation in the look-up table, $C(\Delta x, \Delta y)$. The sum of these two numbers represents the computed value of the point-spread function for that pixel.

Thus, this method of defining a point-spread function is empirical—it defines the PSF on the basis of a two-dimensional look-up table containing brightness values actually observed within the profiles of bright stars, rather than solely in terms of the parameters of a best-fitting analytic function. The intermediate step of defining a Gaussian first approximation and then making empirical corrections from it to the observed profile is performed in order to improve the accuracy of the interpolation: rather than approximating the stellar profile by a formula whose fourth and higher derivatives are everywhere zero, we employ a formula whose fourth and higher derivatives are precisely those of the integrated Gaussian function most closely resembling the core of the stellar image. Nearly all of the brightness variation from one pixel to the next in a stellar profile is represented by the integral Gaussian function which, because it is analytic, can be interpolated with arbitrary precision. The residual non-Gaussian structure in the image profile, which is represented by the look-up table, has much smaller amplitude, and hence the cubic interpolation is subject to smaller absolute errors.

This may also be demonstrated using a numerical-analysis approach. As mentioned above, the error of a cubic interpolation goes as the fourth derivative of the function being approximated. A function of two indepen-

dent variables has five fourth derivatives, but using $\frac{\partial^4}{\partial x^4}$ to represent all of these, then while the error of simple cubic interpolation goes as the fourth derivative

$$\epsilon(\text{simple interpolation}) \sim \frac{\partial^4(\text{PSF})}{\partial x^4} ,$$

for the hybrid interpolation scheme,

$$\epsilon(\text{hybrid interpolation}) \sim \left[\frac{\partial^4(\text{PSF})}{\partial x^4} - \frac{\partial^4(\text{Gaussian})}{\partial x^4}\right] \equiv \frac{\partial^4(\text{PSF}-\text{Gaussian})}{\partial x^4} . \quad (7)$$

Experiments with critically-sampled data show that the Gaussian first approximation is generally accurate to the 10% level or better everywhere. If $|(\text{PSF}-\text{Gaussian})| \lesssim 0.1(\text{PSF})$, then $\epsilon(\text{hybrid}) \lesssim 0.1\epsilon(\text{simple})$, and so using a hybrid model point-spread function is clearly of significant value. This method of evaluating the PSF is nearly as fast as simple interpolation, because the integral of the bivariate Gaussian is the product of two one-dimensional integrals, which can be evaluated very quickly. It also retains the robustness against undersampling enjoyed by the analytic PSF because the Gaussian first approximation is explicitly integrated over the area of each pixel, while the look-up table of empirical corrections, $C(\Delta x, \Delta y)$, itself represents observed values of

$$\iint_{\text{pixel}} (\text{PSF}-\text{Gaussian}) dx\, dy .$$

Finally, it may be mentioned that this method could be made still more accurate, provided some particular function were known a priori to yield an even better first approximation to the stellar profile than the Gaussian function does (cf. eq. (7)). This would probably involve some sacrifice in speed, however, since then the two-dimensional integrals of the analytic approximation over the area of each pixel would not in general reduce to the product of two one-dimensional integrals.

D. *Profile-Fitting Photometry*

1. The Basics

Now that a list of probable stars has been obtained from the "FIND" routine (§III.A above), crude apparent-magnitude estimates and sky determinations have been made for all these stars using "PHOT" (§III.B), and a model point-spread function for the frame has been derived with "PSF" (§III.C), photometry may be obtained for all the stars in the frame by means of least-squares profile fits. Before we can proceed to a description of that process, however, one more problem must be addressed. Given ordinary computer resources, it is not currently convenient to reduce all stars in a frame simultaneously, even though this is conceptually the most straightforward and accurate method when stars are blended together. This approach would require the repeated solution of systems of simultaneous equations involving at least $3N$ unknowns, where N is the number of stars—for each star two centroid coordinates and one brightness scale factor must be determined—and requiring $O(N^3)$ arithmetic operations. Solving for additional parameters to describe the brightness distribution of the sky could involve a large number of extra unknowns, especially for a field containing emission nebulosities or an underlying galaxy. However, if the problem can be broken up into ν independent smaller problems, each containing of order N/ν stars, the effort of reduction would go from $O(N^3)$ to $O(\nu(N/\nu)^3) = O(N^3/\nu^2)$: for example, reducing 3000 stars as 100 groups of 30 stars each produces a 10,000-fold decrease in the effort required. Moreover, the patch of sky around a small grouping of stars would require fewer parameters for its accurate description than would the frame taken as a whole. Therefore, an automatic routine to divide the star list for a given frame into optimum subgroups is required.

DAOPHOT includes a routine named "GROUP" which performs this function. Assume that some critical separation, in units of pixels, can be defined such that two stars separated by less than this distance should be reduced together. In practical terms, such a separation may be defined as follows: let the *image radius* be the distance from the centroid of the brightest program star to a point where its profile disappears into the background noise; let the *fitting radius* be defined as the radius of a region such that only pixels within one fitting radius of the centroid of a star will actually be included in the least-squares normal equations for the determination of the position and brightness of that star. That is, for each step of the iterative profile fits, the model PSF will be shifted and scaled to the current estimate of each star's centroid and brightness, and will be subtracted from the original image data within one *image radius* of the star's center; the residuals of this subtraction only within a region one *fitting radius* from the star's center will be considered in the least-squares computation of the incremental corrections to that star's fitting parameters. The fitting radius is generally taken to be somewhat smaller than the image radius: while in a critically-sampled image the profile of a bright star may be traced out to a distance of ten or more pixels from the center, in fact nearly all of the star's light and, hence, nearly all of the photometric information (although a somewhat smaller fraction of the astrometric information; see King (1983) for a discussion of this distinction) will be contained in a region only a few pixels in radius. Furthermore, using the smallest possible fitting radius greatly reduces the sensitivity of the fits to problems such as the presence of stars not included in the star list, and errors in the sky value. The *critical separation* alluded to above is readily defined as the sum of the image radius and the fitting radius. When two stars are separated by more than this distance, the profile of neither

extends into the fitting region around the centroid of the other. Two stars which are separated by less than this distance should be reduced together.

The actual separation of a star list into autonomous groupings is conceptually simple. Beginning with the first star, the remainder of the list is searched for stars lying within one critical separation. If none are found, the star in question is thereby shown to be isolated and may be fitted by a single model profile. If other stars are found to lie within the critical distance of this star, they are assigned to its group. When the end of the star list has been reached, the stars remaining unassigned are searched for any lying within the critical distance of the second star in the group; if such are found they are likewise assigned to that group. This process continues until *no* star remains among the unassigned stars which is within the critical distance of *any* star in the group. When this condition has been met the group is completely self-contained and independent of the remainder of the data frame. The first star remaining in the unassigned list then becomes the nucleus for the formation of a new group, and this process continues until the entire original star list has been assigned to mutually-independent groups consisting of one or more stars. These groups may then be reduced individually.

Of course, in badly crowded frames some groups formed by this procedure may contain hundreds of stars—still too many for practical simultaneous reduction. Such groups may simply be rejected from further consideration as too crowded for satisfactory photometry. Alternatively, the stars in the overly large groups may be segregated from the rest and run again through the grouping procedure with a smaller numerical value for the critical separation. When the photometry for these stars is derived it will be inferior in quality to that derived for stars assigned to small groups in the first place.

The actual fit of the model profiles to the stars in a group is performed by a routine named "NSTAR," which employs the iterated, linearized least-squares method described by Brown (1955; see Eichhorn and Clary 1974). Starting guesses at the centroids and instrumental magnitudes of all the stars in the group are required for this linearized scheme and are provided by the star-finding and aperture-photometry routines above. Assuming these guesses to be correct, the PSF is first shifted and scaled to the position and brightness of each star, and each profile, out to one *image radius* from its centroid, is subtracted from the original image data. This results in an array of residuals containing the sky brightness, random noise, and systematic errors due to inaccuracies in the estimates of the stellar parameters. Then, knowing the first derivatives of each model profile with respect to the x-centroid, y-centroid, and brightness of the star it represents (these are easily obtained from the PSF), first-order corrections to the stellar parameters are computed by least-squares solutions of the following system of equations:

$$\Delta_{i,j} \doteq sky + \sum_{k=1}^{N} \left[\Delta x_{0,k} \left(\frac{\partial P_k}{\partial x_{0,k}} \right)_{i,j} + \Delta y_{0,k} \left(\frac{\partial P_k}{\partial y_{0,k}} \right)_{i,j} + \Delta h_k \left(\frac{\partial P_k}{\partial h_k} \right)_{i,j} \right] , \tag{8}$$

where

$\Delta_{i,j}$ = residual in the (i,j)-th pixel after subtraction of the preliminary profiles,

P_k = model PSF, shifted and scaled to estimated position and brightness of k-th star,

$x_{0,k}$ = estimated x-centroid of the k-th star,

$y_{0,k}$ = estimated y-centroid of the k-th star,

h_k = estimated brightness scale factor of the k-th star,

N = number of stars in the group, and

(i,j) = the set of pixels lying less than one fitting radius from any star in the group.

Given the known residuals, Δ, from the previous iteration, the known derivatives of the current model PSF's, and an estimate for the sky, these equations are solved for the corrections to the stellar parameters, $\Delta x_{0,k}$, $\Delta y_{0,k}$, and Δh_k, which are then added to the current estimates $x_{0,k}$, $y_{0,k}$, and h_k, to create new parameter estimates for the next iteration. The solution of equation (8) is repeated in this way until the corrections become negligible. After the positions and magnitudes of all the stars in every group have been computed in this fashion, another routine, "SUBSTAR," may be used to subtract the shifted, scaled model stellar profiles from the original image, producing a two-dimensional representation of the fitting residuals which may be examined by eye or subjected to further computer analysis.

2. The Subtleties

Several points which were glossed over in the preceding discussion of the profile fits will be dealt with here.

a. Modeling the Sky. The contribution of the diffuse sky illumination to the profile fits was represented in equation (8) by the italic word "*sky.*" In most published applications of the profile-fitting method a model for the sky brightness distribution, typically consisting of a polynomial function of x and y, is fitted to the image data at the same time as the model stellar profiles (e.g., Penny 1976, 1979; Penny and Dickens 1986; Buonanno *et al.* 1979, 1983). This extension of the problem is conceptually straightforward: for instance, determining a constant sky background simultaneously with the stellar parameters for a group of N stars merely expands the problem from one involving $3N$ unknowns to one with $3N + 1$ un-

knowns. Similarly, fitting a least-squares plane to the sky background, i.e., $sky = ax + by + c$, makes it a problem of order $3N+3$.

My own experiments with the simultaneous determination of sky parameters have been less than satisfactory, however. This appears to be because fitting the sky brightness at the same time as the stellar profiles requires that the fitting radius (defined above) be quite large, so as to sample adequately the region where each stellar profile levels off into the background. This increases the likelihood that the fits will be disturbed by problems such as faint stars present in the frame, but undiscovered and not included in the star list; errors in the wings of the model PSF; and slightly defective pixels. To illustrate this point, consider a typical case where the stellar FWHM is of order two or three pixels. If the fitting radius is taken to be, say, six pixels in order to obtain an adequate sample of the background, then the profile fit for a single star will be based on about $\pi 6^2 \approx 100$ pixels. The sky determination will be dominated by the pixels near the perimeter of this region, consisting of roughly $2\pi 6 \sim 40$ pixels. Contrast this with an alternate solution where, instead of solving for a new estimate of the sky brightness as a part of the profile fit, the sky value derived during the aperture photometry (§III.B) is substituted directly into equation (8) and a fitting radius of two pixels is adopted for the determination of the three stellar parameters alone. If the sky annulus employed during the aperture photometry had inner and outer radii 12 and 25 pixels, this sky measurement would be based on roughly 1500 pixels, as compared with 40 in the previous case. The new fitting region for the star, now consisting of approximately $\pi 2^2 \approx 12$ pixels, contains virtually all of the stellar photometric signal that was contained in the larger fitting region, but it has a smaller likelihood, by a factor of 9, of containing undetected stars or bad pixels. Therefore, the second method, adopting as correct the sky values obtained during the aperture photometry, appears to offer two advantages for better photometry: (1) the average sky value can be determined much better because it is based on many more pixels, and (2) the fitting region for the star is less likely to be contaminated because it contains fewer—and only the best—pixels. It has the disadvantage that it requires the sky brightness to have a constant gradient over a region of order 25 pixels in radius, as opposed to a region 6 pixels in radius for the example given above. (Because the sky annulus is symmetric about its center, the expectation value of the mode is unchanged in the presence of a brightness gradient across the region, although, if the gradient is large compared to the random noise in a pixel, the mode becomes a less efficient statistic than it would be in the absence of such a gradient.) This method also has the severe philosophical disadvantage that the mode-finder used for aperture photometry measures the wrong kind of sky. In Section III.B above it was remarked that the mode estimates the contamination arising from (1) detected stars, (2) undetected stars, (3) discrete image defects, and (4) diffuse sources, all of which must be considered in correcting the aperture measurement of the program star. In the case of multiple-profile-fitting photometry, however, only contaminants (2), (3), and (4) should be represented in *"sky,"* because (1)—detected stars—will be accounted for by the explicit profile fits. That, for at least some applications, these admittedly serious drawbacks are of less consequence than the more subtle advantages of the method can be demonstrated by examining the intrinsic widths of the principal sequences of star clusters which have been reduced in both ways.

In Tables I and II are summarized photometric results from B,V frame pairs for a field in the globular cluster M 12 = NGC 6218 = C1644-018, and for a field in 47 Tucanae = NGC 104 = C0021-723. This latter field also includes a large sample of background stars in the Small Magellanic Cloud. These four frames[4] have been reduced in several different ways: (1) adopting the "sky" values found during the aperture photometry as correct, and using equation (8) to determine the stellar parameters alone, with a fitting radius of 2.0 pixels (this is the method and the fitting radius which I routinely use for data such as these); (2) solving for a flat (constant) sky-brightness distribution as a part of the multiple-profile fit for each stellar group, employing fitting radii of 2.0, 4.0, and 6.0 pixels; and (3) solving for a sloping (planar) brightness distribution as part of the profile fits for each group, again with fitting regions 2.0, 4.0, and 6.0 pixels in radius. Tables I(a) and II(a) give for each of the regions indicated in Figures 3 and 4: the mean shift in color of the main sequence derived by each reduction method from the mean of all reduction methods ($\langle\Delta\rangle$); the observed width of the cluster sequence in instrumental $(b-v)$ color for each of the reduction methods ($\langle|\Delta|\rangle$); and the number of stars that successfully converged to a fit (n). The mean main sequence was defined by the parabola which best fits the data for all stars within some maximum distance from the main sequence, $|\Delta(b-v)|<\Delta_{\rm lim}$, for each magnitude range. The values of $\Delta_{\rm lim}$ used, which were chosen to be approximately $3\langle|\Delta|\rangle$, are given in the tables. The mean absolute residual is used to quantify the width of the main sequence because it is less sensitive to the inclusion or exclusion of stars near $\Delta_{\rm lim}$ than the traditional root-mean-square statistic would be. To a reasonable approximation, however, the standard deviation of data distributed according to a normal error distribution is about 25% larger than the mean absolute deviation. In addition, for this

[4]All four frames were obtained at the prime focus of the CTIO 4-m telescope, the M 12 frames by J. E. Hesser and D. A. VandenBerg in April 1985, the 47 Tuc frames by J. E. Hesser and W. E. Harris in October 1983.

Table I

(a) Photometric Scatter on the Main Sequence of M12

Sky-fitting method	$<\Delta>$ (mag)	$<\lvert\Delta\rvert>$ (mag)	n	$<\Delta>$ (mag)	$<\lvert\Delta\rvert>$ (mag)	n
	Region I $\lvert\Delta\rvert < 0.07$			Region II $\lvert\Delta\rvert < 0.09$		
none, r=2	−0.0032	0.0151	121	+0.0008	0.0228	343
flat, r=2	−0.0074	0.0186	122	−0.0006	0.0277	339
flat, r=4	+0.0024	0.0208	116	−0.0033	0.0282	330
flat, r=6	+0.0072	0.0228	114	+0.0018	0.0298	324
plane, r=2	−0.0056	0.0196	118	+0.0012	0.0283	341
plane, r=4	+0.0034	0.0200	115	−0.0057	0.0270	333
plane, r=6	+0.0036	0.0234	94	+0.0040	0.0283	259
	Region III $\lvert\Delta\rvert < 0.15$			Region IV $\lvert\Delta\rvert < 0.27$		
none, r=2	+0.0011	0.0425	176	+0.0042	0.0845	217
flat, r=2	+0.0039	0.0551	176	−0.0035	0.0866	218
flat, r=4	−0.0049	0.0450	173	−0.0012	0.0843	211
flat, r=6	−0.0051	0.0485	168	+0.0016	0.0880	212
plane, r=2	+0.0041	0.0506	178	−0.0053	0.0939	208
plane, r=4	−0.0040	0.0444	172	−0.0048	0.0862	219
plane, r=6	−0.0001	0.0481	131	+0.0014	0.0838	152

(b) Overall comparisons

Sky-fitting method	$\left\langle \frac{<\lvert\Delta\rvert>}{<\lvert\Delta\rvert>_{none}} \right\rangle$
none, r=2	1.00
flat, r=2	1.19
flat, r=4	1.17
flat, r=6	1.25
plane, r=2	1.21
plane, r=4	1.14
plane, r=6	1.23

Table II

(a) Photometric Scatter on the Principal Sequences of 47 Tucanae and the Small Magellanic Cloud

Sky-fitting method	$<\Delta>$ (mag)	$<\lvert\Delta\rvert>$ (mag)	n	$<\Delta>$ (mag)	$<\lvert\Delta\rvert>$ (mag)	n
	Region I $\lvert\Delta\rvert < 0.05$			Region II $\lvert\Delta\rvert < 0.07$		
none, r=2	−0.0004	0.0137	49	+0.0021	0.0174	55
flat, r=2	+0.0016	0.0136	47	+0.0087	0.0223	51
flat, r=4	−0.0025	0.0138	47	−0.0039	0.0191	53
flat, r=6	−0.0031	0.0137	48	−0.0065	0.0202	57
plane, r=2	+0.0067	0.0165	45	+0.0154	0.0238	50
plane, r=4	−0.0025	0.0134	46	−0.0057	0.0206	52
plane, r=6	−0.0021	0.0132	47	−0.0041	0.0205	56
	Region III $\lvert\Delta\rvert < 0.17$			Region IV (SMC) $\lvert\Delta\rvert < 0.27$		
none, r=2	+0.0005	0.0478	84	−0.0021	0.0778	220
flat, r=2	+0.0145	0.0584	71	+0.0009	0.0980	198
flat, r=4	−0.0024	0.0523	79	−0.0136	0.0878	206
flat, r=6	−0.0108	0.0557	79	−0.0100	0.0823	201
plane, r=2	+0.0095	0.0611	75	+0.0085	0.1012	194
plane, r=4	−0.0003	0.0594	81	−0.0045	0.0794	196
plane, r=6	−0.0078	0.0533	80	−0.0035	0.0779	207

(b) Overall comparisons

Sky-fitting method	$\left\langle \frac{<\lvert\Delta\rvert>}{<\lvert\Delta\rvert>_{none}} \right\rangle$	CPU time (hr)
none, r=2	1.00	12.5
flat, r=2	1.19	16.1
flat, r=4	1.08	16.9
flat, r=6	1.10	20.8
plane, r=2	1.29	19.3
plane, r=4	1.11	17.4
plane, r=6	1.06	21.9

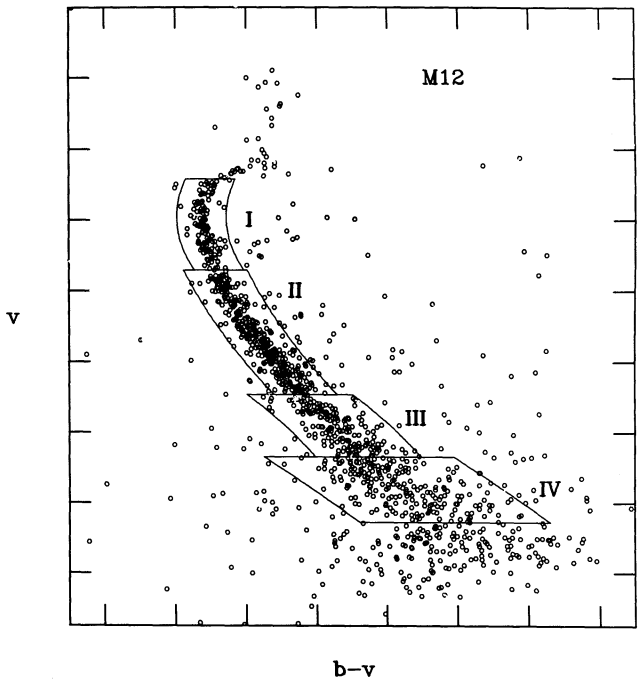

FIG. 3–Instrumental color-magnitude diagram from a frame pair of a field in the globular cluster M 12. Tick marks are placed at one-magnitude intervals in apparent v magnitude and at 0.2-magnitude intervals in $(b-v)$ color. The numbered boxes enclose the regions of the CMD which were analyzed to determine the thickness of the main sequence for the various reduction methods.

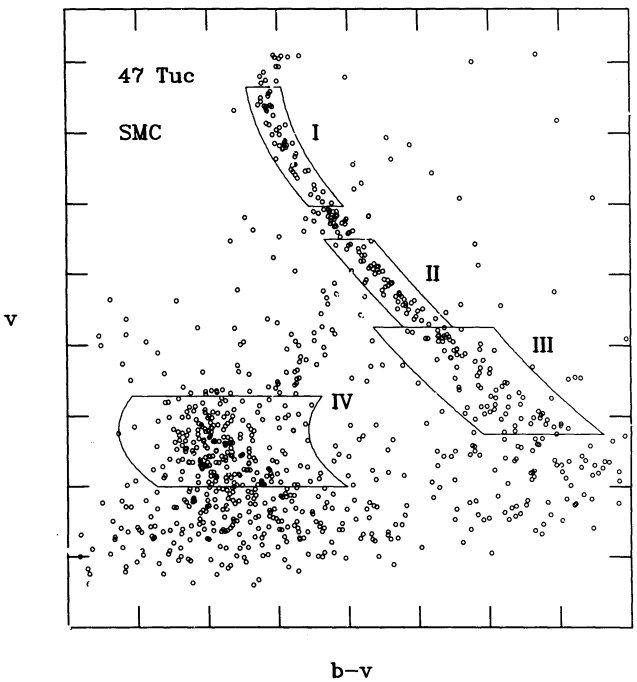

FIG. 4–Instrumental color-magnitude diagram for a frame pair of a field in the globular cluster 47 Tuc, including a region in the Small Magellanic Cloud. Otherwise the same as in Figure 3. The gap between the uppermost two boxes on the main sequence of 47 Tuc was left in order to omit the SMC field's red horizontal-branch clump, which is visible in the space between the boxes.

filter set it is known that the standard $(B-V)$ color is of order 1.2 to 1.25 times larger than the instrumental $(b-v)$ color, so taking these two corrections together we may estimate $\sigma_{B-V} \sim 1.5\langle|\Delta(b-v)|\rangle$. (Note that M 12 is reddened by $E_{B-V} \sim 0.2$ mag, so some of the scatter in the main sequence may be due to variable reddening. The scatter in 47 Tuc is probably all observational.)

Tables I and II show that there is no particular dependence of the overall position of the main sequence on the reduction method used, so any *systematic* difference in the photometry obtained by the various methods is probably negligible. It does seem, though, that the photometric *scatter* about the main sequence tends to be larger for those reductions where the sky was determined explicitly from the profile fits than for the reductions where the modal brightness values from the annular comparison region were imposed on the solutions. This is emphasized by Tables I(b) and II(b), where the overall ratios of the main-sequence widths for the sky-fitting to the non-sky-fitting reductions are tabulated. Furthermore, it seems that when the sky is determined during the profile fits, fewer stars converge successfully. This is probably due to the unavoidable correlations between the stellar brightness and the sky brightness, and between the stellar centroid and the sky gradient (for the sloping-sky solutions), which appear in the normal equations of the least-squares problem—the algorithm has a difficult time distinguishing between a faint star and a slight change in the sky brightness when both are allowed to be completely free unknowns. Finally, I have included in Table II(b) the total reduction time required for one pass of the frame pair for 47 Tuc + SMC through NSTAR, using the VAX 11/780 computer at the DAO for each of the reduction methods; in all, approximately 1500 stars were reduced in the visual frame, 2200 in the blue frame. For the sky-fitting solutions the reduction time is significantly increased for three reasons: (1) the least-squares problem for each stellar group is increased by whatever number of free parameters is used to describe the sky; (2) an increase in the fitting radius means that more pixels must be processed in the fit for each star; and (3) the correlations between sky parameters and stellar parameters make the least-squares problem less stable, requiring more iterations on average per stellar group.

Thus it is clear that, for these frames at least, solving explicitly for the local diffuse sky brightness produces consistently more photometric scatter at the cost of greater effort than simply adopting the modal value found in a surrounding annulus of pixels, even though the latter method is fundamentally wrong. Of course, as mentioned above, this conclusion cannot be expected to remain valid in situations where the diffuse brightness distribution changes slope significantly over a distance comparable to the radius of the sky annulus. Otherwise explicit model fits or other sophisticated corrections for the sky's brightness distribution are indispensible. Although NSTAR routinely imposes the modal sky value on the profile fits, the source code includes explicit instructions for modifications to permit the sky value to be determined as a free parameter.

b. Weighting. The noise properties of the CCD as a function of brightness level are knowable a priori and can be used to assign weights to the individual pixels for the least-squares solution of equation (8). To first approximation the random scatter in a brightness measurement follows from the known readout noise and the number of photons per digital brightness unit (substitute "$D_{i,j}$" for "*approximate sky level*" in eq. (3)). To this may be added an allowance for the difficulty of accurately calibrating the frame, e.g., the impossibility of deriving a flat-field frame which is correct for stars of all color temperatures simultaneously and for all times during the course of a night's work. However, even after reasonable values have been assigned to these sources of scatter, the brightness residuals near the centers of bright stars are always found to be larger than expected. This is the result of the failure of the approximations and interpolations used in the definition of the PSF: for example, variations of the true point-spread function from place to place within the frame and any remaining fourth- and higher-order errors in the profile interpolations. Experience shows that, for critically sampled data, fitting errors usually persist at about the level of 2%–3% (rms error per pixel) of the peak value of the profile. Therefore, the scheme of assigning weights to the individual pixels in a profile fit must allow for this effect.

To estimate quantitatively how these residual profile errors should depend on the degree of under- or oversampling, reconsider equation (7), which represents the lowest-order terms neglected in the approximation of the empirical PSF. Let $s_{\rm int}$ represent the root-mean-square interpolation error, P the true PSF, and G the Gaussian first approximation,

$$s_{\rm int} = \left\langle \left(\Delta x^4 \frac{\partial^4 (P-G)}{\partial x^4} \right)^2 \right\rangle^{1/2} .$$

Recognizing that $\frac{\partial^4 G}{\partial x^4}$ scales as $\frac{1}{\sigma^4}$, it follows that

$$s_{\rm int} = \alpha \left(\frac{\Delta x}{\sigma} \right)^4 \qquad (9)$$

in some average sense at least, where x is a generic spatial coordinate, and σ is the standard width parameter of the Gaussian approximation. The quantity α represents the degree to which the actual point-spread function typically differs from the Gaussian function, averaged over the profile: $\alpha \sim \left\langle \left(\frac{P-G}{G} \right)^2 \right\rangle^{1/2}$. The quantitative value of α is obtained by trial and error, requiring that the standard error per unit weight averaged over many stars come out

to be unity. As mentioned above, from critically-sampled images ($\sigma \sim 1$ pixel, FWHM ~ 2.4 pixels) obtained with a number of telescopes, $\alpha \sim 0.02$ to 0.03. The $1/\sigma^4$ dependence in equation (9) indicates that the fitting errors would be reduced to $\sim 1\%$ pixel^{-1} for a FWHM ~ 3 pixels, and increased to $\sim 15\%$ pixel^{-1} for a FWHM ~ 1.5 pixels.

Taking into account all these potential sources of error in the profile fits, NSTAR assigns to an arbitrary pixel (i,j) a weight given by

$$\begin{aligned} weight &= s^{-2} \\ s^2 &= (\text{readout noise})^2 + \left(\frac{D_{i,j}}{\text{photons per ADU}}\right) \\ &\quad + (0.0075 D_{i,j})^2 + \left(0.027 \frac{D_{i,j} - sky}{\sigma_x^2 \sigma_y^2}\right)^2 . \end{aligned} \quad (10)$$

The quantities σ_x and σ_y are the standard deviations of the Gaussian first approximation to the PSF, $\sigma_x^2\sigma_y^2$ representing a compromise value of σ^4. The number 0.0075 represents the quantitative estimate of the average flat-field error (the value of this constant, basically an educated guess for the typical scatter in a typical astronomer's typical data, is not critical since this is seldom a dominant contributor to the error estimate), and 0.027 represents the estimated average interpolation error, α. The pixel weight defined by equation (10) will be subject to additional modifications to be discussed below.

c. Assisting Convergence. One of the plagues of iterative linearized least squares is the reluctance of occasional solutions to converge, which can be a severe handicap when thousands of reductions are to be performed noninteractively. Discontinuities in the statement of the problem represent a principal cause of this defect. For instance, I have stated above that only pixels lying within one *fitting radius* of the current estimate of some star's centroid shall be included in the solution for a group. This will sometimes lead to an oscillating solution: occasionally there will be a pixel lying about one fitting radius away from some star's center such that when this pixel is included in the star's fit, the centroid coordinates are adjusted so as to place this pixel outside the fitting radius; when the pixel is omitted from the fit the centroid is shifted so as to bring it back within one fitting radius, and so on. This could go on indefinitely. One solution to this problem is to weight the pixels according to their distance from the stellar centroid, with pixels in the center of the fitting circle receiving unit weight, and the weight falling continuously to zero at the fitting radius. In this way, the pixel just described need not be accepted or rejected absolutely; instead, an intermediate situation can be approached asymptotically, where the pixel achieves partial acceptance near the fitting radius. It is not certain that such a compromise result is more accurate than either of the two oscillating states, but at least it does guarantee convergence to *some* reasonable intermediate answer, allowing the computer to proceed to the next set of stars. For this purpose, DAOPHOT multiplies the pixel weight derived in the previous section by a radial function of the form

$$weight \propto \frac{5}{5 + \frac{1}{R^2/r^2 - 1}} ,$$

where r is the distance of the pixel from the star's currently-estimated centroid, and R is the fitting radius. This function is illustrated in Figure 5.

Another problem is the tendency for occasional solutions to run away. The least-squares algorithm of Brown and of Eichhorn and Clary represents a linear approximation to a nonlinear problem. It uses a multivariate version of Newton's method for finding the zero-crossing of a function to derive those values of the fitting parameters for which the first derivative of χ^2 with respect to each of the parameters is zero simultaneously, thus defining the desired minimum in χ^2. The first approximations to the stellar positions and magnitudes provided by the star-finding and aperture-photometry algorithms are generally quite good, and for the most part the solutions converge quickly to the desired solution. Occasionally, however, an inconvenient correlation in the normal equations will cause a large change in the current values of two or more of the parameter estimates, which largely compensate each other in the computation of χ^2, but which nevertheless throw the new parameter estimates out of

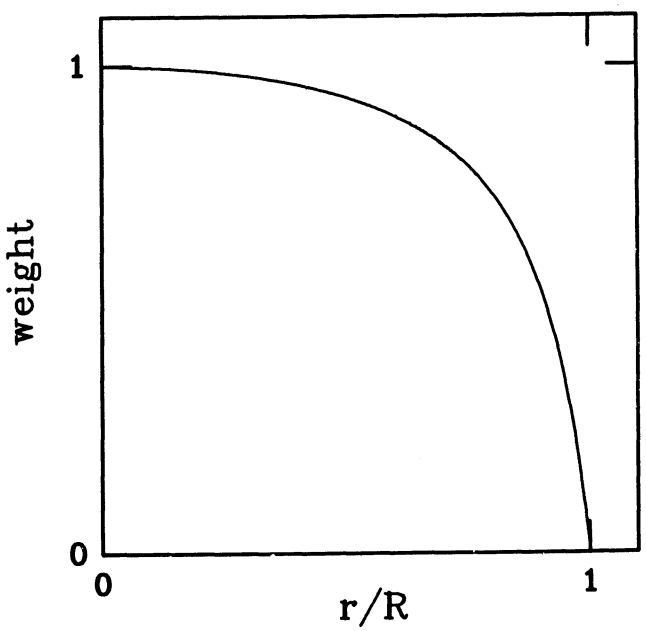

FIG. 5–The radial-weighting function applied to the image brightness data in the fitting region around each star image. The abscissa is the radial distance of the pixel from the star's centroid, in units of the radius R of the circular fitting region; the ordinate is the weight.

the capture radius for the true minimum. One way of preventing this is to accept the fact that the star-finder and the aperture photometry algorithm do produce reasonably good initial estimates of the stellar parameters, and to limit the size of the corrections that any one iteration may apply. Given some parameter etimate p_i used as input for the i-th iteration of a solution, the correction Δp_i derived from that iteration, and a value P representing the magnitude of the maximum permitted correction, NSTAR obtains the corrected parameter estimate p_{i+1} from

$$p_{i+1} = p_i + \left(\frac{\Delta p_i}{1 + \frac{|\Delta p_i|}{P}} \right) .$$

The quantity in brackets has limits Δp_i as $|\Delta p_i|/P \rightarrow 0$ and $\pm P$ as $|\Delta p_i|/P \rightarrow \infty$. The quantity P acts as a clamp—when the computed parameter correction is small it is applied virtually unchanged, but when the computed correction is unreasonably large, the correction actually applied is no larger than P. NSTAR uses clamps set initially at 0.4 pixels for the x- and y-centroid position and at 2 mag for the stellar brightness. In addition, NSTAR keeps track of the sign of each parameter correction, and when the sign of the correction to a given parameter is seen to change from one iteration to the next, the magnitude of the clamp is cut in half. Thus, as long as the solution continues to move in a given direction, it is allowed to go virtually as far as it wants, although the size of the individual steps is restricted. At the same time oscillating solutions are strongly damped: instead of being able to switch repeatedly between two different states, each of which points to the other, the algorithm is eventually forced to consider intermediate, compromise solutions. The final convergence is tested on the basis of the computed parameter corrections Δp_i rather than on the restricted corrections that are actually applied, so there is no danger that the algorithm will confuse a situation where the clamps have merely become very tight with a true convergence.

d. Resisting Bad Data. In virtually every astronomical image there are invalid elements in the data array. Some of these result from defects in the detector, which may be mapped ahead of time and which may therefore be flagged "to be ignored" for the reductions. In other cases the causes of the errors are unpredictable, such as impacts of cosmic rays onto the detector. Again, in some cases these are so obvious that they may be automatically recognized and ignored (see discussion of the FIND routine, §III.A above) such as, for instance, a single saturated pixel in an otherwise level sky. A problem arises when a datum is only *slightly* erroneous, and even more so when it is involved in the profile of a star, where data values are expected to change rapidly from one pixel to the next even under the best of circumstances. These pixels can sometimes be recognized *a posteriori*. That is, after all the profile-fits have been performed and after the suitably shifted and scaled point-spread functions have been subtracted from the image, visual examination of the resulting frame can identify individual large residuals surrounded by smaller residuals of the opposite sign. The bad pixels can then be flagged "to be ignored" and the profile-fits repeated. This has at least two disadvantages: (1) it requires an additional pass through the data, and (2) it is inherently discontinuous: a datum is either accepted or rejected—possibly producing a major change in the results—with no room for compromise in cases where a residual seems a bit large without being completely unreasonable.

NSTAR's profile fits are made resilient against bad data by a scheme of dynamically reweighting pixels on the basis of their residuals from previous iterations: after the first few iterations of the least-squares fit, the weight of each pixel (§b and §c above) is multiplied by a factor which is nearly unity for small-to-moderate residuals and approaches zero for very large ones. The precise formula used is

$$weight \propto \left[1 + \left(\frac{|\Delta|}{a \times s} \right)^b \right]^{-1} ,$$

where Δ is the intensity residual of the pixel from the profile fit, s is the standard error of the pixel (§III.B.2.b above), and a and b are parameters which modify the shape of the falloff: for $b > 1$ a pixel with a residual small compared to a standard deviations has its weight not materially changed, one with a residual exactly equal to $a \times s$ is given precisely half its normal weight, and for large residuals the weight falls off roughly as Δ^{-b}. As long as good pixels outnumber the bad ones in any given star's image, after the first few iterations the good pixels should have smaller residuals than the bad ones. When the reweighting is then turned on, the good pixels receive higher weights than the bad ones, and as the solution progresses it will attempt to match the good pixels ever more closely, while continuing to reduce the influence of the bad ones still further.

Figure 6 illustrates how the fits respond to simulated bad data when this weighting scheme is applied. For this test, on a typical frame three isolated stars possessing instrumental magnitudes near 12.0, 14.5, and 17.0 were selected; for comparison, in this frame stars became saturated in their centers near an instrumental magnitude of 11.5, and the 5-sigma detection limit was around 18.5. For each of the stars a single-profile fit was performed in the normal fashion, with a fitting radius of 2.0 pixels. Then the brightness value in the central pixel of each star image was corrupted by adding a certain number of ADU, and the profile fits were performed anew. Figure 6 shows how, in each case, the derived magnitude depends on the size of the artificial defect for $a = 2.5$ and various values of b. In the uppermost panel, no continuous reweighting

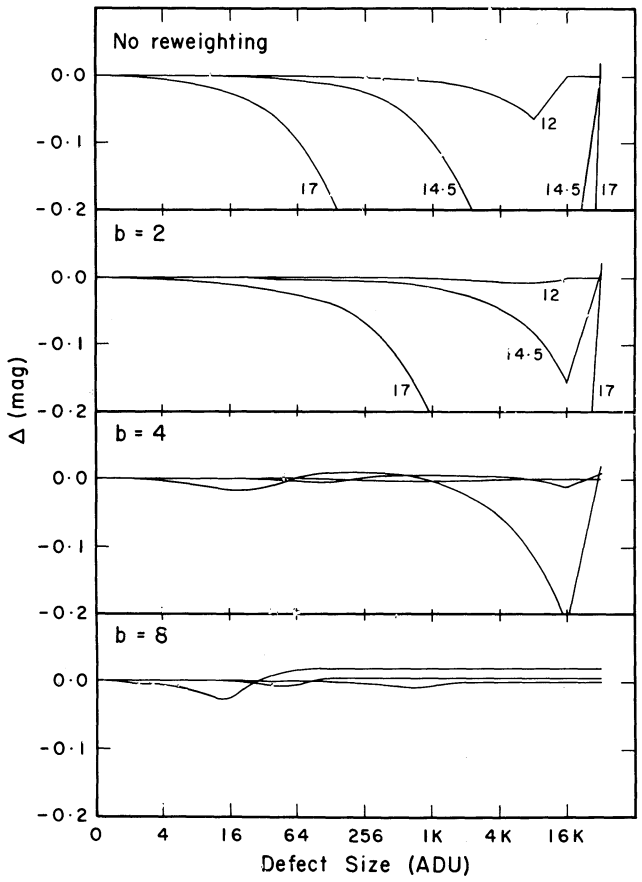

FIG. 6–The responses of the computed magnitudes of three stars to artificial image defects, as described in the text. In each panel, the difference between each star's magnitude as derived from the defective data and the magnitude as derived from the original data is plotted as a function of the size of the imposed defect. The curves corresponding to stars with original instrumental magnitudes near 12, 14.5, and 17 are labeled in the uppermost two panels. The different panels illustrate the effect of different degrees of bad pixel rejection as quantified by the parameter b (see text). As would be expected, in each case the fainter the star, the greater the sensitivity to a defect of a given amplitude. However, increasing b clearly reduces the sensitivity of the magnitude determinations to the defective data.

was employed so that, as the number of ADU in the central pixel increased, the computed etimate of each star's magnitude grew brighter until the point was reached where the central pixel exceeded the "maximum valid data value" which had been specified for the frame. At this point the derived magnitude of the star jumped discontinuously to the value representing the best fit to the other pixels in the image, exclusive of the central one. As would be expected, the fainter the star, the more sensitive is its derived magnitude to the number of ADU by which the central datum was falsified. For positive values of b, the sensitivity of the stars' magnitudes to the size of the spurious datum is reduced, and they approach the "without-the-central-pixel" magnitude estimates in a more continuous fashion. For $b = 8$, the algorithm is very effective at paying progressively less attention to the central pixel as its defectiveness becomes more and more obvious. In this case the magnitude of the faintest star departs by no more than 0.024 mag from its original value, no matter how large a defect is imposed, and the magnitudes of the brighter stars wander by even less.

Unfortunately, when the bad pixels outnumber the good ones in some star's image this scheme will not consistently produce the right answer. On the other hand, neither would any other simple, automatic scheme that I am aware of. The estimate of the standard error which is returned along with each star's derived magnitude is based on an honest evaluation of the scatter of the individual pixels' residuals from the derived profile fit, even though the weights of extreme residuals are fudged in performing the fits themselves. Therefore, cases where a number of badly-defective pixels are included in the fit may be recognized by the large standard error returned by the program.

e. Rejecting Stars. Obviously, an accurate inventory of the star images in a given frame is highly desirable. Not only is such an inventory useful in itself for luminosity-function studies of stellar populations, but is is also a practical necessity for accurate photometry: in order to obtain a precise measure of the intrinsic brightness of a given star, the existence and the approximate brightnesses of any detectable companions must also be known. Conversely, there is also a more subtle danger in the possibility of perceiving stellar images where none are in fact present. When a random noise peak within the profile of a bright star is erroneously taken to represent a faint companion, with the result that the superposition of two model point-spread functions is fitted to the data, the brightness of the real star will be underestimated by whatever fraction of the luminosity is assigned to the fictitious companion. Therefore, some automatic mechanism must be put in place which will, with some degree of reliability and consistency, recognize when an entry in the star list probably does *not* represent a real star, and will then eliminate the entry.

It is obviously not possible to recognize every physical or optical double in a given set of stars. Since there will always be some composite-light binaries with angular separations smaller than the resolution of the optical system, principal sequences in color-magnitude and color-color diagrams will never be infinitely thin. Therefore, it is pointless to insist that a star-rejection algorithm be absolutely *reliable* in distinguishing all optical or physical doubles among a given set of stars. It is much more important that it be *consistent* in deciding whether to treat a given image as a single or a double: it is when a star is reduced as a single in one frame and as a double in another that ridiculous derived colors are most likely to occur. Richer and Fahlman (1986, private communication), Crotts (1986), and I have had success in using

computerized routines external to DAOPHOT which intercompare star lists for different frames of a given field to arrive at a master list of "real" detections, which is then imposed uniformly on all the frames. Such a procedure is probably the best available method to ensure the consistent reduction of different data frames. Still, there is need for an automatic star-rejection algorithm to work within NSTAR which utilizes only the information available in a single image.

As the linearized least-squares solution proceeds, at each iteration the standard error of each star's brightness is derived from the corresponding diagonal element of the inverse normal matrix. Following each iteration after the fourth the least-certain detection in a given stellar group is tested for reality: it is regarded as spurious and is removed from the star list (a) if after the fourth iteration the star's brightness is less than its standard error, (b) if after the eighth iteration its brightness is less than 1.5 times the standard error, (c) if after the twelfth iteration its brightness is less than twice the standard error, or (d) if the solution for the group seems to converge before the twelfth iteration and the brightness of the least-certain star is less than twice its standard error. No more than one detection at a time is eliminated in this way from any stellar group. In addition, any detection whose centroid is less than one full width at half maximum from the centroid of a brighter star is subjected to the same tests, even though it may not be the least-certain star in the group. Any detection whose centroid is less than 0.4 times the FWHM from the centroid of a brighter star is rejected out of hand. When an entry is deleted from the star list in any of these ways, the iteration counter is backed up by one and the iterative fits continue from that point. This allows the new model two full iterations to adjust itself to the data before other apparent detections are subjected to reality tests.

IV. Use of the Program

Preparing a raw CCD image for eventual photometry involves several steps which must be performed before the data are ready for analysis with DAOPHOT. First, the actual electrical signal corresponding to no detected photons, which may drift with time, is generally measured for an individual exposure by reading out more pixels than the detector actually contains. This produces a set of unexposed "pixels" from which the current value of the baseline DC voltage in the amplifier output may be determined; this value is then subtracted from the data in the rest of the frame. Fixed-pattern spatial variations in the photometric zero-point of the output image are removed by subtracting suitably-averaged frames of zero exposure time. For some devices it is also necessary to subtract long-exposure frames taken with the camera shutter closed, to remove a spatially-varying dark signal. Then, to ensure a uniform photometric scale over the entire image, the data frame is divided by a flat-field exposure which has been obtained with the camera system exposed to uniform illumination of the appropriate color temperature. Next, in some cases it is necessary to remove interference-fringe patterns produced by strong night-sky emission lines, which appear as rippling spatial variations in the sky brightness. A map of the fringe pattern is created by combining long-exposure images of regions of dark night sky which contain as few stars as possible; the fringe map is then appropriately scaled in intensity and subtracted from the program frame. Finally, it is useful to mask out the data corresponding to known defective pixels, so that during subsequent processing the computer will not confuse them with legitimate astronomical objects.

To be sure, any of these preprocessing steps may be more complex than outlined here. For instance, the flat-field exposure referred to above may be a judicious combination of out-of-focus images of the interior of a diffusely-illuminated dome with exposures of the twilight and/or nighttime sky, in order to achieve simultaneously the uniform illumination, correct color-temperature, and high signal-to-noise ratio required for proper flat fielding. Nevertheless, the raw data frames must be subjected to a preprocessing procedure like that outlined above before they are ready for photometric analysis.

Given a suitably prepared digital image, the user should know several numerical quantities before proceeding to the use of DAOPHOT, as mentioned in Section III.A. Among these are (1) the readout noise per pixel in the calibrated digital image, (2) the number of photons corresponding to one digital brightness unit, (3) the maximum brightness level at which the detector operates linearly, (4) the full width at half maximum of a star image in the data frame, and (5) the approximate diffuse sky brightness in the data frame. Items (1), (2), and (3) will be constant for a given detector (although the numerical values will change if multiple individual exposures are added or averaged for the photometric analysis). Item (4) will change somewhat from frame to frame depending on the seeing and the accuracy of the guiding, but even a crude estimate is adequate for the purposes to which it will be applied. Item (5) will obviously change greatly with exposure time, filter bandpass, phase and altitude of the moon, and many other factors, but a routine exists in DAOPHOT which estimates the average sky brightness from a few thousand pixels distributed uniformly across the data frame, employing the mode finder mentioned in Section III.B.

Once these quantities are known, from (1), (2), and (5) it is possible to compute the random background noise in the image (eqs. (2) and (3)). From this number the user then obtains the brightness enhancement corresponding to a stellar detection of a given significance level ($H_{\min}$, §III.A), and a minimum legitimate data value, which

may be set at some number of standard deviations below the average sky brightness—DAOPHOT will presume that any pixel whose datum is below this value is defective, and will ignore it. The user is then ready to run FIND, which requires as inputs the estimated FWHM of a stellar image, the minimum and maximum legitimate data values, and the minimum significant brightness enhancement. The program will then produce a list of detected objects each with its (x,y) coordinates, its estimated brightness in magnitudes above the detection limit, and its sharpness and roundness indices. If the user feels that too many legitimate stars are being rejected, or that too many image defects are being accepted as astronomical objects on the basis of the imposed limits on brightness, sharpness, and roundness, these limits may be adjusted and the star finder reinvoked.

Having arrived at a satisfactory initial star list, the user then performs synthetic aperture photometry on the detected objects with the PHOT routine. The desired radii, in units of pixels, of from one to twelve concentric stellar apertures and of the inner and outer edges of the sky annulus must be specified. At this point, the program also requires the numerical values of items (1) and (2) above: the readout noise and the number of photons per digital brightness unit. These will be used to compute the standard errors of the derived aperture magnitudes. All of these numerical values required by the aperture photometry routine are stored in tabular form and may easily be recalled for convenience when many data frames obtained with the same equipment are to be reduced. The program will then go through the star list generated by the FIND routine, computing for each star a local sky brightness and a magnitude in each of the specified stellar apertures.

Once a user has acquired a little experience with reducing data from a given detector system, neither the star-finding nor the aperture photometry needs to be performed interactively, nor do they require the routine use of an image-display system. Batch jobs can be set up which will produce star lists and aperture magnitudes for all data frames obtained in a given observing run without direct astronomer supervision. The aperture photometry produced in these two steps is usually adequate for uncrowded fields, such as exposures of bright photometric standard stars.

If the field is crowded, however, the user must now obtain the empirical point-spread function for each data frame. This is an interactive procedure which requires the use of an image-display system. On this display system the user examines the images and the environments of the brightest unsaturated stars in the frame (a sorting routine exists within DAOPHOT which permits the user to order the stars in any of the data files produced by the program according to apparent magnitude), and chooses those which appear to be relatively isolated. If some of the bright stars in the frame have a few faint companions, these may also be used in the determination of the frame's PSF if the astronomer uses an interactive procedure to remove the faint neighbors. To demonstrate the effectiveness of this process, twelve stars from a typical image[5] are illustrated in Figure 7, both before and after the neighbor-removing procedure outlined below; the images of two of these stars as they evolve through the procedure are also illustrated in Figures 8 and 9.

The neighbor-removing procedure goes as follows. Having chosen suitable stars, the user supplies the sequential identification number of the first of them to the PSF routine in DAOPHOT, which computes the hybrid, empirical model profile for the star, out to a user-specified image radius. If the star is blended with fainter companions, they will appear as bumps in the initial point-spread function. The user continues to specify the ID numbers of other desired PSF stars, whose observed profiles are first fitted and then added to the empirical PSF, which is thus built up star by star. At the same time, the routine writes the coordinates of all the PSF stars and of all their known neighbors into a disk file. When the twelve PSF stars in our sample frame have been added together in this way, the point-spread function illustrated in the middle left panels of Figures 8 and 9 results. Obviously, because some of the stars were blended with nearby neighbors, this model PSF is not an accurate representation of the true point-spread function in the frame. Nevertheless, using this first estimate of the PSF, multiple-profile fits to the groups consisting of the PSF stars and their known neighbors can be performed. When the fitted profiles have been subtracted, residuals like those shown in the bottom left panels of Figures 8 and 9 are produced. Clearly, because of the inadequacy of the model PSF, the fits are poor and the residual frames are messy. However, if at this point only the neighbors of the PSF stars are subtracted and the PSF stars themselves are allowed to remain in the frame, then images like those at top center in Figures 8 and 9 result. Because the model PSF has been scaled down in intensity to match the profiles of the faint companions, the errors in the model PSF are likewise reduced in importance: even though the neighbors have been subtracted imperfectly, still in the new image they contaminate the profiles of the bright stars less than they did in the original image.

The neighbors of all twelve PSF stars having been thus subtracted (approximately), they may be added together again to produce the new model point-spread function

[5]This is the average of sixteen 250-second visual exposures of a field in the globular cluster NGC 3201 obtained by J. E. Hesser and D. A. VandenBerg using an RCA CCD camera at the prime focus of the CTIO 4-m telescope. The full width at half maximum of the stellar profile is approximately 1.75 pixels = 1.″0.

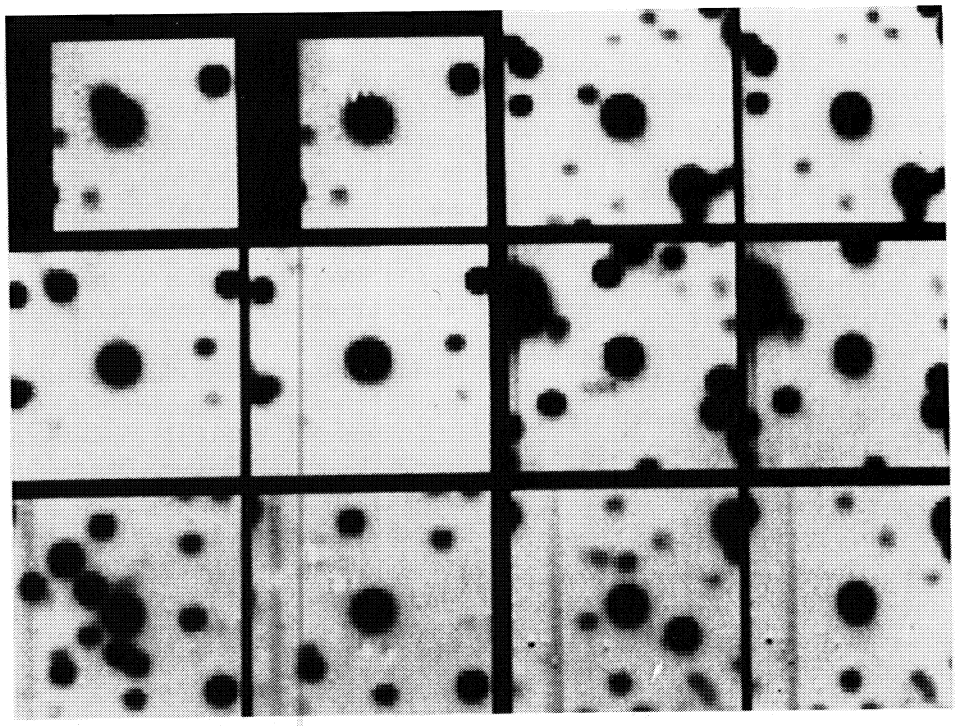

FIG. 7(a)–Six of the twelve bright stars chosen from a typical crowded frame for the determination of the empirical point-spread function. Each star is shown twice, in the left panel of each pair as it appears in the original image, and in the right panel as it appears after the close neighbors have been removed by the procedure described in the text.

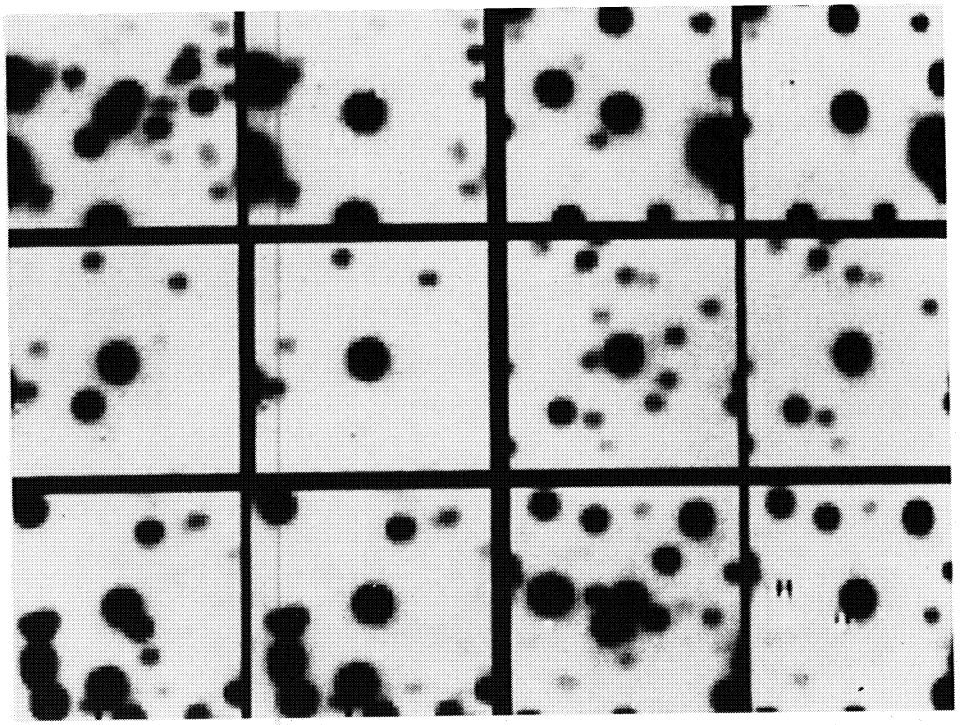

FIG. 7(b)–The remaining six of the twelve bright stars chosen from a typical crowded frame for the determination of the empirical point-spread function. Each star is shown twice, in the left panel of each pair as it appears in the original image, and in the right panel as it appears after the close neighbors have been removed by the procedure described in the text.

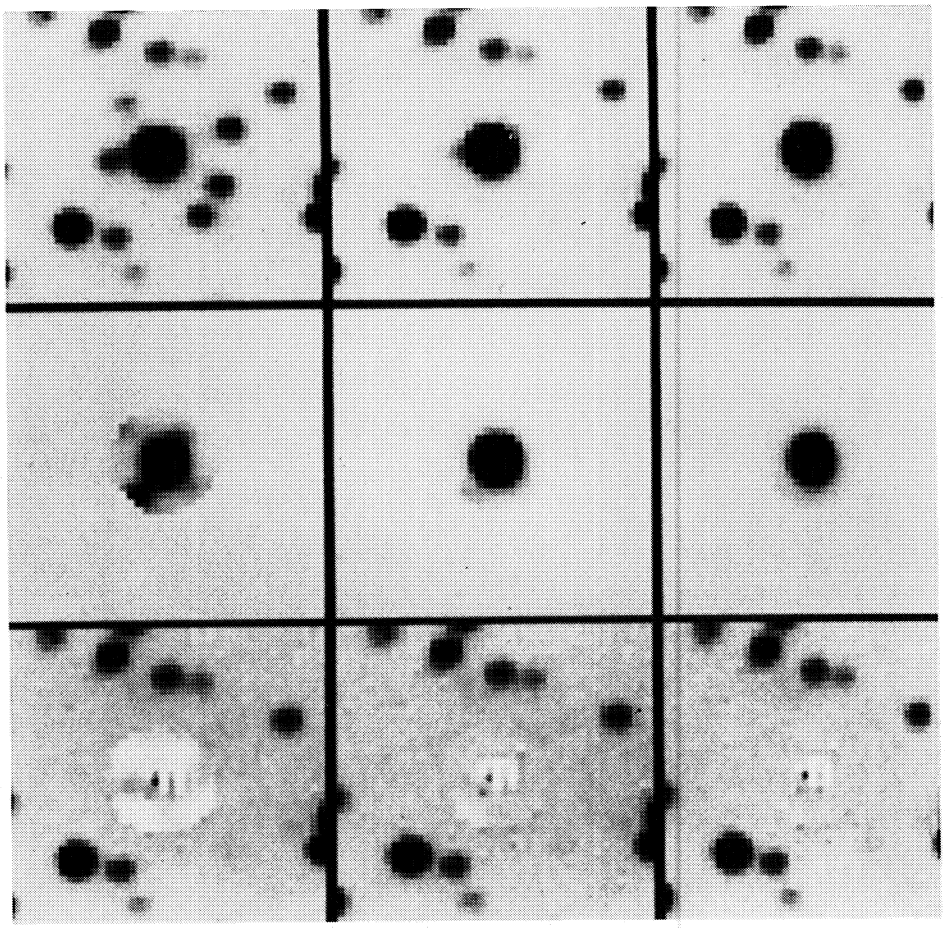

FIG. 8–Snapshots of one of the stars in Figure 7 as it evolved through the neighbor-elimination procedure described in the text. Time increases most rapidly from top to bottom, more slowly from left to right. (top left) The appearance of the star in the original frame. (middle left) The first model point-spread function obtained from the twelve stars in Figure 7. (bottom left) The residual frame resulting from fitting the PSF illustrated in the panel above to the PSF star and its neighbors, and then subtracting them all from the image. (top middle) The modified frame created by subtracting from the original frame only the neighbors of the PSF star, leaving the PSF star itself in the image. (middle middle) Second-generation point-spread function generated by adding together the twelve PSF stars, after the first PSF had been used to subtract the neighbors. (bottom middle) Residual frame resulting from fitting and subtracting the second generation PSF. (top right) Frame after subtraction of neighbors with the second-generation PSF. (middle right) Third-generation point-spread function. (bottom right) Residual frame resulting from application of the third-generation PSF.

illustrated in the center of Figures 8 and 9. When this second-generation profile is once again fitted to the PSF stars and their neighbors in the original frame, and they are once again subtracted, the fits are found to be considerably better (Figs. 8 and 9, bottom center). By once more subtracting the neighbors and producing a third-generation PSF from the twelve bright stars, it is found that it is now possible to fit and subtract the groups of stars with almost no perceptible evidence of systematic errors caused by the bright stars' neighbors (Figs. 8 and 9, right-most column). To be sure, there are still noticeable defects near the centers of the regions where bright stars have been subtracted, but these are random and unavoidable residuals due to Poisson noise and the difficulty of interpolating the PSF to better than a few percent. The largest single residual in the bottom right panel of Figure 9, which is at the center of the bright neighbor above and to the left of the PSF star, represents 5% of the peak brightness of the PSF star. The typical residual is considerably smaller than this, and by the time twelve stars have been added together the mean point-spread function is defined to ~ 1% (root mean square), in spite of the fact that many of the stars used to define the function had neighbors in the original frame.

Now that the model stellar profile has been defined, the star list produced by the aperture-photometry routine, which contains estimated centroids, instrumental magnitudes, and local sky values for all stars identified by the star-finding routine, is split into autonomous groupings by the GROUP algorithm (§III.D). For this routine the astronomer needs to specify only the numerical value of the critical separation, in pixels. For very crowded

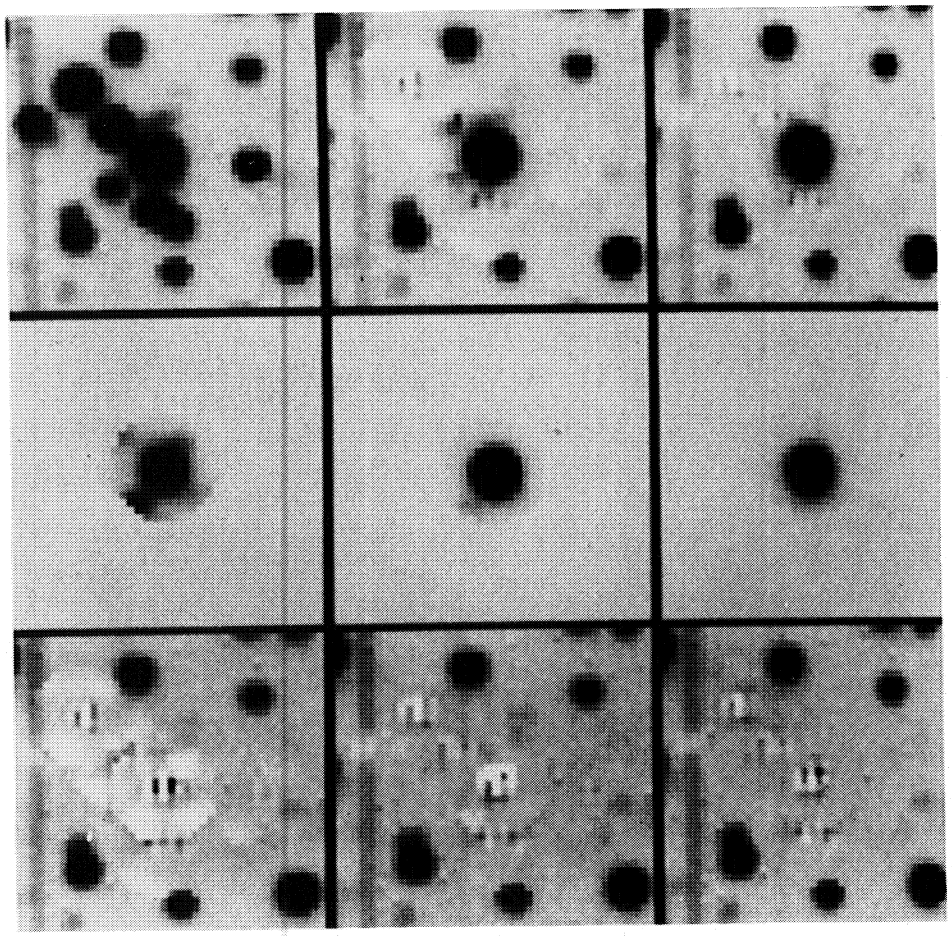

FIG. 9–Same as Figure 8, for a different PSF star.

fields, one or more groups may be formed which contain more stars than can conveniently be reduced simultaneously (at present, the maximum group size the NSTAR routine is allowed to attempt is 60 stars). Therefore, DAOPHOT contains a SELECT command which permits the user to extract from the file created by the GROUP command only those stellar groups within a certain range of sizes. Those groups which are small enough to reduce are put in a file by themselves, while the groups which are too large may either be discarded or may be run again through the GROUP routine with a smaller numerical value for the critical separation.

The names of files containing the stellar groups and the model point-spread function may now be specified as inputs to the multiple-profile-fitting routine, NSTAR, which will produce an output file containing the improved stellar centroids and magnitude determinations. In the standard version of the program, the individual sky values which were determined for all stars in the course of the aperture photometry are unchanged during the profile fits, but modifications are indicated in the code which allow the sky value for a given stellar group to be adjusted as well.

The photometric results produced by the profile-fitting routine may then be used in conjunction with the model PSF to subtract the known stars from the data frame, perhaps revealing hitherto unrecognized features. Figure 10 shows a magnified region of a data frame both before and after subtraction of the known stars. Figure 10(b) shows a number of faint stars which were missed by the star-finding algorithm because they were blended with much brighter companions. They are prominent now that the brighter stars have been removed. This new frame may be resubmitted to FIND which will then identify the previously-hidden stars with high reliability. The new set of stellar coordinates is run through the aperture photometry routine (employing the original data frame) to obtain homogeneous sky estimates for the new stars, and is then appended to the original star list to obtain a nearly complete inventory of the stellar objects in the image. After this augmented star list has been run through the grouping, fitting, and subtracting routines, the resulting residual frame may be inspected by eye to identify any stellar images which the computer was unable to recognize on

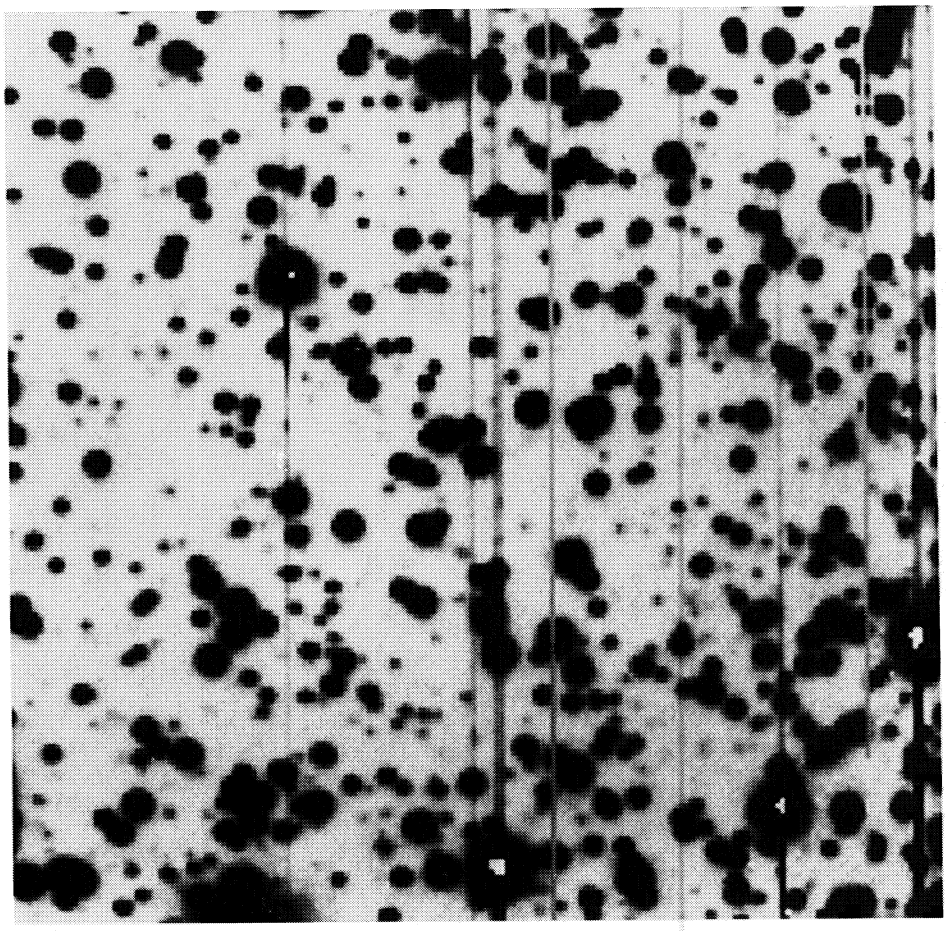

FIG. 10(a)–Section of a typical CCD image of a field in the outskirts of a globular cluster.

the two previous passes. If desired, these may be run through the aperture photometry routine and appended to the previous star list, which is then regrouped and rereduced with NSTAR. Our experience at the DAO is that two-thirds to three-quarters of the stars in a typical, moderately-crowded image are found in the first automatic star-finding pass; most of the remainder are found when the first subtracted image is run through FIND, leaving perhaps 1%–3% to be identified by hand and eye if absolute completeness is desired.

This find-fit-subtract procedure is rather expensive of computer time but is extremely economical of human time and produces a star list which is as complete and homogeneous as the computer and the human eye together can make it. The most astronomer-intensive operation in the whole reduction process is the derivation of the model PSF, which typically requires between one and three hours of direct interaction for a rather crowded field, such as the one illustrated above. Add to this several tens of minutes for identifying the 1%–3% of stars that FIND missed, plus overhead to set up the batch reductions, and the total cost in astronomer time is of order two to four hours for a frame containing perhaps several thousand stars. The rest of the work, which may well involve tens of computer hours, is done wholly by the machine.

Another routine in DAOPHOT performs the inverse of star subtraction. ADDSTAR adds the model stellar profile into a copy of the image, after shifting and scaling it to randomly-generated positions and instrumental magnitudes. When an image of a starfield has been thoroughly analyzed as described above, a number of new data frames may be generated by adding various different sets of artificial stars to this image. If these new frames are subjected to the same reduction process as the original, the derived results for the artificial stars may be compared with their input parameters to build up statistical information on the completeness of the star-finding and the photometric-error distribution as functions of magnitude. It is important that small-number statistics be combated by creating several different artificial frames, each with a different random arrangement of a relatively small number of artificial stars. To add very many artificial stars into a single image would significantly alter the mean crowding conditions and incorrectly simulate the completeness

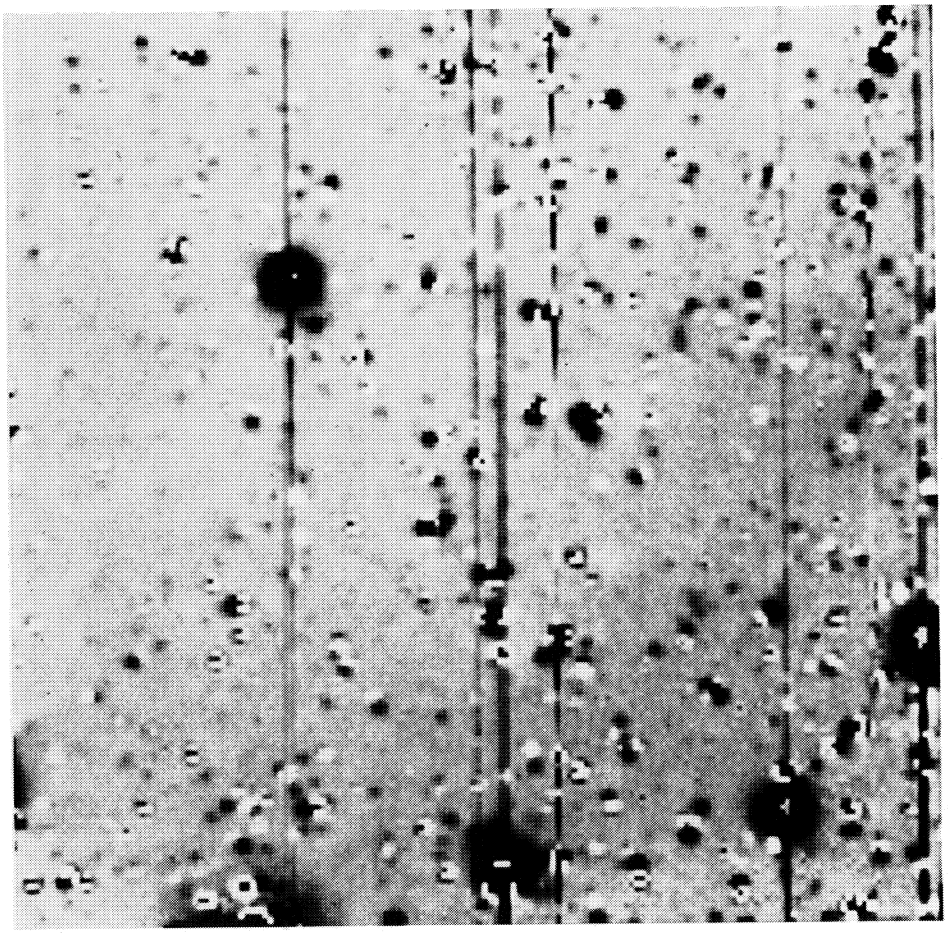

FIG. 10(b)–The same region as in Figure 10(a), as it appears after application of the profile-fitting and -subtraction routines to the initial star list for the region. Many secondary components of blended doubles have been revealed, and most of these can be detected and cataloged by running this derived image through the FIND routine.

and photometric accuracy of the original data.

One further caveat must be kept in mind when artificial stars are analyzed, at least with the version of the program that is currently available to the general astronomical community. Unless the astronomer goes to the effort of obtaining completely independent model PSFs from different sets of stars in the image, the model profile which is used to insert stars will be the same as will subsequently be used to fit them. The next generation of DAOPHOT, currently undergoing tests at the DAO, does include simulated Poisson shot noise in the profiles of artificial stars, but for many applications this refinement of the procedure is largely of academic interest. For relatively faint stars, even for those that are isolated, the readout noise and the Poisson scatter in the sky background will dominate over errors in the shape of the PSF. Thus, adding a faint artificial profile on top of the background noise already present in the real data frame is quite adequate to simulate the image of a faint star. Similarly, in crowded regions the effects of the crowding itself will dominate over the artificially good match between the model profile and the fictitious star image. Finally, for bright blended stars, *systematic* profile errors due to variations in telescope focus across the chip, or to variations in the charge-transfer properties as functions of spatial position or exposure level, may dominate over the pixel-by-pixel Poisson noise in the stellar profile. Accounting for these problems will require a considerable advance in the sophistication of model point-spread functions. Thus, for now at least, artificial-star analyses are largely impractical for bright, blended stars. The observational errors in the absolute photometry of such stars are best evaluated by frame-to-frame, detector-to-detector, and telescope-to-telescope comparisons of calibrated results.

V. Potential Improvements

Crowded-field stellar photometry is such a complex task and the circumstances under which it is to be performed are so diverse, that definitive, universally-applicable algorithms will never be developed. The simultaneous fitting of multiple, overlapping, model profiles to spatially-resolved images, however, seems intuitively to

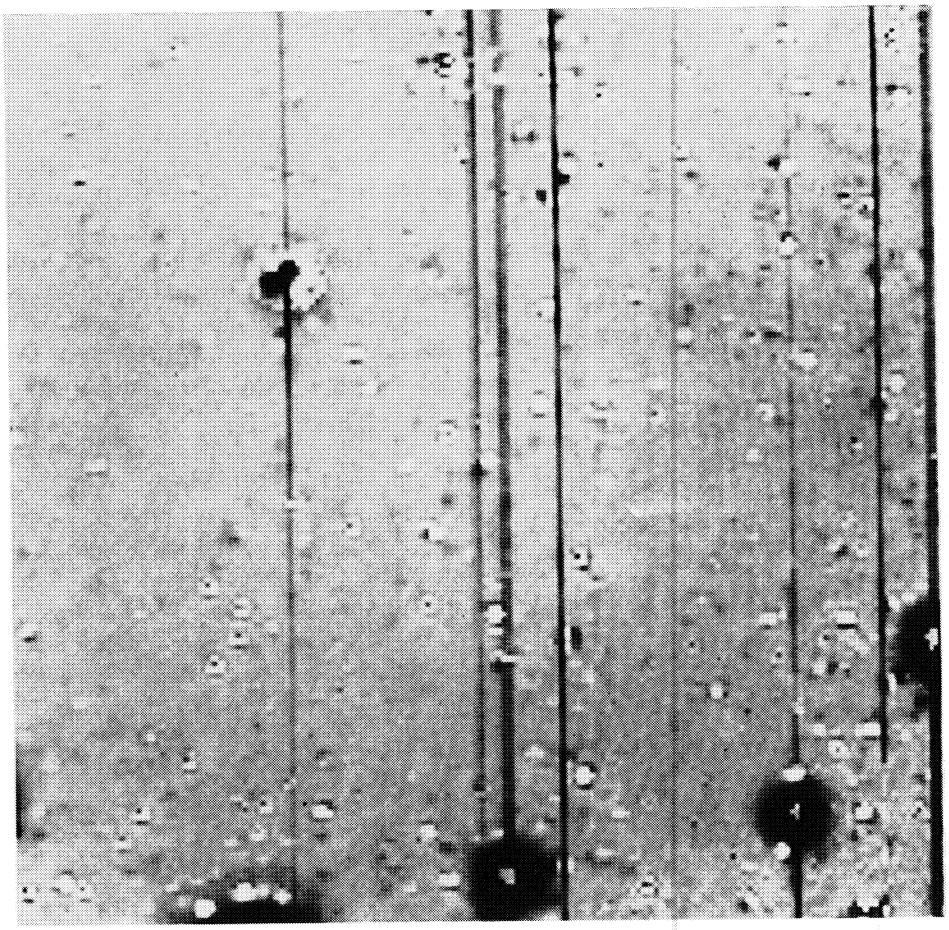

FIG. 10(c)–The same region as in Figure 10(a), after two passes through the find-fit-subtract loop.

be the best available approach, so new developments of this general technique will probably continue for many years.

The reader who has made it through the preceding material will already have noticed several places where the algorithms in DAOPHOT could be improved by pushing techniques like those outlined here to one more level of sophistication. For one example, the star-recognition criteria employed by FIND could easily be made more complex, e.g., by generalizing the *round* index to identify elongated objects that are inclined with respect to the rows and columns of the data array. As another example, the Gaussian first approximation to the point-spread function might be replaced by a more elaborate formulation, such as the convolution of a power law with a tilted, elliptical Gaussian. For the most part, the potential gains ensuing from such modifications probably would not repay the additional computational effort.

Investigators who already have experience with crowded-field stellar photometry will be aware that the algorithms described above deal with some problems only inadequately, and with others not at all. One of these problems is the fact that as magnitude limits become fainter and fainter, especially with the launch of the *Hubble Space Telescope*, the stars we are interested in studying will be blended not with other stars, but with galaxies. Galaxies differ from stars in that no two of them have precisely the same brightness profile, and, hence, no simple analog to the point-spread function can be used to deconvolve blended galaxy images. "Standard galaxy profiles"—Hubble laws, exponential disks, and the like—will permit the derivation of results that will be adequate for many purposes, but will probably never allow crowded-field photometry of a precision to rival sparse-field results. However, more complete solutions to some of the other outstanding problems with DAOPHOT can be envisioned, and it is useful to try to anticipate what some of them may be.

A. *Evaluating the Sky Brightness*

As mentioned in Section III, evaluating the diffuse component of the observed brightness at a given position in a crowded star field is extremely difficult. The complexities concealed within the word "*sky*" in equation (8)

probably represent one of the two greatest weaknesses in DAOPHOT. Even the meaning of the phrase "diffuse sky brightness" cannot be perfectly defined because different things are meant under different circumstances. For instance, as discussed above, for aperture photometry or for the fitting of a single, isolated model stellar profile, the question to be asked is, "What brightness value would we expect to observe in a particular set of pixels (containing a program star) if *this star* did not exist?" For multiple-profile fits the question becomes, "What brightness value would we expect to observe if *none of the stars that we know about* existed?" As discussed in Section III.B, a probabilistic answer to the former question can be derived from the histogram of observed brightness values in a region of the frame closely resembling that containing the star. Under some circumstances and to some reasonable degree of accuracy the same number can be used as an adequate answer to the second question, as discussed in Section III.D.2.a.

To a higher degree of accuracy, the answer to the latter question could be improved iteratively by fitting the model stellar profiles and subtracting the known stars from the frame, computing new estimates of the diffuse sky brightness from this new frame, reperforming the profile fits on the original frame using the new sky estimates, and so on. Thus, we are attempting to eliminate the known stars the better to estimate the sky brightness that remains. This process converges rapidly because although the light from a given star may affect many pixels, most of that light falls on only a few pixels. When we wish to determine the brightness of a star we examine those pixels where most of its photons were detected, so that an error in the initial guess of the sky brightness will produce the smallest possible error in the estimate of the star's apparent magnitude. Then when the star is subtracted from the image, in that vast majority of the pixels which are inhabited by the wings of the stellar profile, the amount of starlight to be subtracted is small and the error in the amount to be subtracted resulting from the inaccuracy in the star's initial brightness estimate is much smaller still. Thus, the new sky estimate will be wrong by much less than the original one, and in an iteration or two the error will be negligible. This method of improving the sky estimate is possible with the current level of software development—experiments on badly-crowded fields near the centers of globular clusters are currently under way at the DAO—but obviously it requires multiple reductions of each data frame to achieve the best results. It would be preferable to obtain a better sky estimate from the very beginning.

A number of sophisticated algorithms have been developed to provide accurate diffuse-brightness estimates for a region of sky in spite of the presence of contamination by stars and image defects; these are mostly used to obtain surface photometry of large galaxies. For the most part, these algorithms are not well suited to determining the sky brightness for aperture photometry because they are designed to measure the diffuse brightness that would be characteristic of a region if *no* discrete sources of contamination were present, as discussed in detail in Section III.B above. Conversely, for the same reason it might well be profitable to apply one of these methods, such as the asymmetric clipping algorithm of Ratnatunga and Newell (1984), to the determination of the sky brightness for multiple-profile fits in moderately crowded regions. However, it is possible to imagine even more sophisticated methods which will make use of all the available information to yield still better sky estimates.

One characteristic of the Ratnatunga and Newell method which limits its applicability to this problem is that it identifies and then simply ignores data values which appear to be contaminated. Recognizing that in astronomical images the contamination occurs in the form of predominantly *positive* brightness perturbations, the algorithm iteratively trims off the upper tail of the brightness histogram and then reflects the lower tail of the histogram about the resulting mean sky brightness to test the validity of the most recent trim. The method soon converges to a reasonable and symmetric corrected brightness histogram. Meanwhile, no quantitative use has been made of the shape of the bright tail of the histogram, which may in fact represent most of the pixels in the sample region; the data there have been identified as contaminated and have been excluded from further consideration. But, given the fact that we have at our disposal a model profile for star images in the data frame, we should be able to combine this knowledge with empirical observations of the stellar luminosity function in the field to model directly and quantitatively the entire brightness histogram observed in some region of a data frame. The diffuse sky brightness, corrected for the effects of contamination, would fall naturally from such an analysis.

Expressed in more practical terms, the upper end of the brightness histogram plus the known stellar-brightness profile could tell how to correct the lower end of the histogram for the wings of the bright stars; the observed or presumed stellar luminosity function plus the model profile could tell how to account for the presence of faint stars. Similar correction factors for the luminosity function and average morphological properties of background galaxies in the field and for the frequency and energy-spectrum of cosmic-ray hits could also be applied. Thus, a quantitative correction from the observed mode, median, or mean of the overall brightness histogram to the best estimate of the diffuse component of the sky brightness could be determined directly. Given suitable tabular or analytic forms for the stellar profile and for the various luminosity functions, there is no reason why this algorithm should require any more execution time than either the Kitt Peak mode-finder or the Ratnatunga/

Newell clipping algorithm.

B. *A More Sophisticated Grouping Algorithm*

As discussed in Section III.D.1 above, tremendous gains in reduction speed can be realized by subdividing the star list into the smallest possible self-contained groups for multiple-profile analysis. That section presented one method for producing groups which, while being conceptually simple, ensured that the photometric computations for no star belonging to one group could be altered by the existence of a star which had been assigned to a different group. A straightforward generalization of this technique allows a more realistic evaluation of the potential of one star to overlap with another, and so permits the breaking up of the total star list into even smaller subunits. This is accomplished by allowing the "critical separation," defined above, to be a function of apparent magnitude.

Let us reconsider the meaning of the "image radius." In Section III.D.1 this was defined as the distance from the centroid of the brightest program star to the point at which its profile disappears into the background noise; this quantity is easily measured from an enlarged image of a bright star on an image display. Now let us modify this definition. When a given pair of stars is being considered to determine whether they belong in the same group, let the image radius be momentarily defined as the radius at which the *brighter of these two stars* disappears into the background noise. If there is no pixel which is simultaneously less than one fitting radius from the centroid of the fainter star and less than one image radius from the centroid of the brighter, then the brighter star cannot affect the photometry of the fainter and the two stars need not be reduced together.

The calculations for this decision are made as follows. First, identify the fainter of the two stars (on the basis of the approximate magnitudes derived by aperture photometry, a preliminary run of the profile-fitting photometry, or the brightness values observed in a few pixels near the estimated centroid). Find the point one fitting radius from the centroid of the fainter star and lying on the line segment connecting the centers of the two stars. Evaluate the scaled model point-spread function for the brighter star at this point. If this brightness value is greater than some user-specified constant, ρ, times the standard error of the brightness value in a sky pixel (eq. (3) above, §III.A), then the stars overlap significantly and should be reduced together.

In this approach the user specifies not a physical radius of a bright star-image, in pixels, but rather the degree of image overlap that can be tolerated, in units of the random noise in the frame. The computer then determines the critical radius corresponding to this degree of overlap, automatically and separately for each pair of stars considered. Depending on the crowding conditions in the data frame or in subregions of a data frame, ρ can typically be assigned values of order 0.1 to 10. In the former case the errors due to crowding are guaranteed to be much smaller than those due to simple readout noise, Poisson statistics, and profile errors; in the latter case, the user may be assured that crowding errors will often dominate. Thus, the value of ρ which the user finds empirically to be needed for dividing a star list up into groups small enough for practical reductions itself provides a quantitative index of the seriousness of the crowding in the frame or subregion.

A star-grouping routine operating according to these principles has been in use at DAO since autumn 1985, but has not yet been generally released.

C. *Variable Point-Spread Functions*

It is unfortunately true that, because of optical aberrations or charge-transfer imperfections in the detector, some CCD cameras display point-spread functions which vary systematically across the image. If the problem is severe enough—more than a few percent—then it is capable of invalidating the crowded-field fitting techniques described in this paper.

An actual, observed PSF is the result of an instrumental profile, which may vary in some systematic, continuous, and repeatable fashion with position in the frame, convolved with a seeing and guiding profile, which may to good approximation be considered constant across a given frame (if the exposure time is long enough for the statistical properties of the seeing to become well established) but variable from one frame to the next. The resulting profile is then convolved with a square step-function and sampled at one-pixel intervals—the mathematical equivalent of integrating the function over finite pixels. This statement of the problem suggests that variable PSFs may be dealt with in the Fourier domain.

If the instrumental profile is a repeatable function of position, though perhaps also a function of telescope focus, camera temperature, or other independently-determinable variables, then a library of tabular or analytic instrumental profiles could be built up from laboratory tests or observations made on nights of excellent seeing. For a given program frame the broadening function due to seeing and guiding would be determined by comparing the Fourier transforms of actual stellar profiles in the frame with those of the corresponding library profiles. To predict the profile of an arbitrary program star, the transform of the broadening function determined in this way would be multiplied by the transform of the appropriate instrumental profile, further multiplied by the transform of a step function with a suitable phase shift, and inverse-transformed to yield the anticipated brightness profile as it would appear when integrated over discrete pixels. This would be a time-consuming procedure, and if the instrumental profile as a function of temperature and of

telescope focus, altitude, and azimuth (for instance) is not sufficiently repeatable, it may not be practical. For the *Hubble Space Telescope* and particularly for the wide-field camera, which at some wavelengths is seriously undersampled, this technique may be worth pursuing, since the camera is expected to be physically stable, and the instrumental profile will have been carefully mapped at high spatial resolution before launch.

Meanwhile, an interim solution much closer to the formulation described in Section III.C may be used to deal with modest variability of the image profile. This involves taking the first Taylor expansion of the PSF with respect to the x- and y-coordinates of the star's centroid in the image frame. If (X_k, Y_k) are the spatial coordinates of the k-th PSF star in the coordinates of the data frame, and for convenience letting X_0 and Y_0 represent the brightness-weighted centroid of the distribution of PSF stars in the frame, $X_0 = \frac{\Sigma_k h_k X_k}{\Sigma_k h_k}$ and $Y_0 = \frac{\Sigma_k h_k Y_k}{\Sigma_k h_k}$, then instead of evaluating one look-up table of corrections from the best-fitting Gaussian first approximation to the actual stellar profile, we compute three:

$$C(\Delta x, \Delta y) = \frac{1}{\Sigma_k h_k} \sum_k h_k (\text{observed profile} - \text{Gaussian})_{\Delta x, \Delta y},$$

$$C'_X(\Delta x, \Delta y) = \frac{1}{\Sigma_k h_k} \sum_k h_k (X_k - X_0)(\text{observed profile} - \text{Gaussian})_{\Delta x, \Delta y},$$

and

$$C'_Y(\Delta x, \Delta y) = \frac{1}{\Sigma_k h_k} \sum_k h_k (Y_k - Y_0)(\text{observed profile} - \text{Gaussian})_{\Delta x, \Delta y}.$$

As before, $(\text{observed profile} - \text{Gaussian})_{\Delta x, \Delta y}$ is calculated by interpolation in a grid centered precisely on each PSF star's centroid. The three arrays C, C'_X, and C'_Y are sufficient to define a point-spread function which varies linearly with position in the frame. For the point-spread function to contain a constant brightness volume independent of position in the frame, it is necessary and sufficient that $\Sigma_{\Delta x, \Delta y} C'_X = \Sigma_{\Delta x, \Delta y} C'_Y = 0$. If C'_X and C'_Y as first computed do not average out to zero, then the Gaussian analytic approximation is multiplied by a suitable factor and added to or subtracted from C'_X and C'_Y so that they do.

A PSF-generating routine operating in this fashion is currently in use at DAO to reduce frames obtained with the Kitt Peak No. 1, 0.9-m telescope. With constant model point-spread functions as described in Section III.C, profile-fitting reductions for these frames showed changes in the apparent photometric zero-point amounting to as much as 7%–10% from one corner of the frame to the opposite corner, as determined both by the differences between profile-fitting photometry and aperture photometry for uncrowded stars and by the differences between profile-fitting photometry and published photometry for stars in previously-observed fields. Employment of the linearly-varying PSF described here generally reduced these systematic trends to $< 1\%$.

D. *Matching Stars Between Frames*

If one of the two biggest defects in DAOPHOT is the difficulty of defining the diffuse sky-brightness distribution, the other is the absence of any simple provision for cross-identifying stars from one frame to another. Very few research programs are satisfied with the measurement of a single brightness value or a single position for each object; most require at least a color index or a series of instantaneous magnitudes at a number of epochs. Thus most research programs run quickly into the considerable problem of comparing measurements from different data frames. For users of DAOPHOT, this task is performed after the photometric analyses are complete, using software external to DAOPHOT itself. Generally speaking, after each data frame has been run through the entire reduction procedure individually, and after the positional/photometric catalog for each frame is as complete as possible, transformations relating one frame's coordinate system to another's are derived and stars are cross-identified by positional coincidence. This method is imperfect in at least two ways. First, due to purely random errors, the positions of stellar centroids—particularly those of stars near the faint detection limit—can occasionally be in error by amounts comparable to the half width at half maximum of the stellar profile, which is to say comparable to the minimum separation between "distinguishable" stars. This means that in a certain fraction of cases repeat detections of the same star may seem sufficiently far apart for them to be regarded as distinct objects. Second, in the complex sequence of star-finding/star-rejection inherent in the FIND/NSTAR/SUBSTAR/FIND... loop, DAOPHOT may conclude that a given brightness enhancement in the night sky is a single star in one digital image, a binary in another. When the different sets of results for this enhancement are intercompared, the magnitude differences will be meaningless.

A reasonably good solution to these problems is available within the current state of the art (although outside DAOPHOT as it now exists). After the photometric reductions are complete, the independent positional/photometric catalogs for all available frames of a given field are transformed to a common coordinate system and are intercompared to distinguish probably real from probably spurious detections. The criteria used for this classification may be as simple or as complex as seems appropriate. For instance, when a detection within one frame coincides with one in another within one half width at half maximum, they may be taken as representing a single real object. A detection which occurs in only one frame with no counterpart in any other may simply be rejected as

spurious; or, it may be retained as real if it seems significant at the 5 σ level and rejected otherwise. Positionally-coincident detections whose brightnesses differ by three magnitudes or more might be referred to the astronomer for a professional opinion. In any case, through some procedure a single catalog of "real" astronomical objects in the program field is obtained. This master catalog is then transformed back to the coordinate system of each of the individual frames and is used as the basis for a final round of NSTAR reductions. In some cases it may be desirable to modify NSTAR to remove its capability to adjust the input centroid coordinates of the program stars, allowing it only to scale the model point-spread functions in intensity to match the image data (Crotts 1986). This modification may require some care, however, if the data frames are taken in photometric bandpasses sufficiently different and at air masses sufficiently high that atmospheric dispersion causes the apparent position of a star to depend on its intrinsic color.

The ex post facto imposition of a master star list on a set of data frames still may not represent the best conceivable method for combining multiple data frames. As the technique now stands, the information contained in each stellar image is reduced to a comparatively short list of derived quantities, and it is this condensed information which is then consulted to determine whether the star is real. More complete use of the available information could be achieved with a photometric program which considered an entire series of frames simultaneously. The star-finding algorithm would require that a detection be cataloged only if its existence seemed likely upon consideration of the appropriate patch of sky in all available data frames, and would list only the single (x,y) coordinate pair most appropriate to the totality of the data. Similarly, the profile-fitting routine would consider the data from all available frames together, insisting on mutually-consistent centroids from one frame to the next for every iteration of the reduction. However, each frame would still have its own PSF, the magnitudes derived for a given star would be allowed to differ from frame to frame to permit synoptic reduction of frame sets involving variable stars or a variety of photometric passbands, and frame-to-frame centroid shifts depending upon each star's derived color could also be incorporated.

The multiple-frame reduction would be especially valuable in the case of seriously undersampled data, such as will be obtained with the wide-field camera of the *Hubble Space Telescope*, for example. In undersampled data, weak cosmic rays will be difficult to distinguish from faint stars. A standard procedure will be to obtain two data frames in each photometric bandpass, and retain only those objects which are independently discovered in both frames of a pair. But some gain in limiting magnitude may be achieved by basing an object's acceptance on simultaneous consideration of more than two frames, and some gain in observing efficiency may be realized by allowing data from one photometric passband to assist in evaluating the reality of an object in a frame taken in a different passband. A further problem with undersampled data is that a star's position cannot be established to better than half a pixel when its light is nearly all contained in a single picture element. When multiple frames with slight translational shifts are considered simultaneously, however, the positional information that becomes available when a star's light is divided among two or more pixels can be exploited in the analysis of those frames where it is mostly contained in one.

I am extremely grateful to those who contributed to the development of DAOPHOT by attempting to use the innumerable versions of the code, conveying their likes, dislikes, and bug discoveries to me. Principal among them are Jim Hesser, Bob McClure, and Ed Olszewski; but others too numerous to mention individually here have been a tremendous help as well. We at DAO are all grateful to Jeremy Mould and Keith Shortridge for donating their software to us, and to Linda Stryker and Ed Olszewski for bringing it up. In addition, Phil Hodge, Jim Hesser, Bob McClure, and Graeme Smith read preliminary drafts of this paper and made many valuable suggestions. Thanks are also due to Dave Duncan for his usual impeccable work in producing the plates and figures. Finally, I would like to thank an anonymous referee for a careful reading of the manuscript and for many very constructive suggestions.

REFERENCES

Boyle, W. S., and Smith, G. E. 1970, *The Bell System Technical Journal*, **49**, 587.
Brown, D. C. 1955, Ballistics Research Laboratories, *Report No. 937*, Aberdeen Proving Grounds, Maryland.
Buonanno, R., Buscema, G., Corsi, C. E., Ferraro, I., and Iannicola, G. 1983, *Astr. Ap.*, **126**, 278.
Buonanno, R., Coluzzi, R., Corsi, C. E., De Biase, G. A., and Ferraro, I. 1979, *Mem. Soc. Astr. Italy*, **50**, 451.
Crotts, A. P. S. 1986, *A.J.*, **92**, 292.
Eichhorn, H., and Clary, W. G. 1974, *M.N.R.A.S.*, **166**, 425.
Hamilton, W. C. 1964, *Statistics in Physical Science* (New York: The Ronald Press Company).
Jacoby, G. 1985, *N.O.A.O. Newsletter*, No. 4, 15.
King, I. R. 1971, *Pub. A.S.P.*, **83**, 199.
———. 1983, *Pub. A.S.P.*, **95**, 163.
Kristian, J., and Blouke, M. 1982, *Scientific American*, **247**, No. 4 (October), 66.
Lauer, T. R. 1983, Ph.D. thesis, University of California, Santa Cruz.
Lupton, R. H., and Gunn, J. E. 1986, *A.J.*, **91**, 317.
McClure, R. D., Hesser, J. E., Stetson, P. B., and Stryker, L. L. 1985, *Pub. A.S.P.*, **97**, 665.
Massey, P. 1985, *N.O.A.O. Newsletter*, No. 2, 13.
Monet, D. G., and Dahn, C. C. 1983, *A.J.*, **88**, 1489.
Penny, A. J. 1976, In *IAU Colloquium 40, Astronomical Applications of Image Detectors with Linear Response*, ed. M. Duchesne and G. Lelievre (Paris: Paris-Meudon Observatory), p. 49-1.
———. 1979, *M.N.R.A.S.*, **187**, 829.
Penny, A. J., and Dickens, R. J. 1986, *M.N.R.A.S.*, **220**, 845.

Peterson, B. A., Murdin, P., Wallace, P. T., Manchester, R. N., Penny, A. J., Jordan, A., Hartley, K. F., and King, D. 1978, *Nature*, **276**, 475.
Ratnatunga, K. U., and Newell, E. B. 1984, *A.J.*, **89**, 176.
Smith, G. H., McClure, R. D., Stetson, P. B., Hesser, J. E., and Bell, R. A. 1986, *A.J.*, **91**, 842.
Stryker, L. L. 1981, Ph.D. thesis, Yale University.
Tody, D. 1981, *RICHFLD Photometry Program Users Guide* (Tucson: Kitt Peak National Observatory).
Walker, A. R., 1984, *M.N.R.A.S.*, **209**, 83.
Walker, G. A. H., Johnson, R., Christian, C. A., Waddell, P., and Kormendy, J. 1984, *SPIE*, **501**, 353.