Publications of the Astronomical Society of the Pacific
105: 1342–1353, 1993 November

DoPHOT, A CCD Photometry Program: Description and Tests

PAUL L. SCHECHTER

Physics Department, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
Electronic mail: schech@achernar.mit.edu

MARIO MATEO[1,2]

Carnegie Observatories, 813 Santa Barbara Street, Pasadena, California 91101
Electronic mail: mateo@astro.lsa.umich.edu

ABHIJIT SAHA

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, Maryland 21218
Electronic mail: saha@stsci.edu

Received 1993 May 6; accepted 1993 September 1

ABSTRACT. The design considerations and operational features of DoPHOT, a point-spread function (PSF) fitting photometry program, are described. Some relevant details of the PSF fitting are discussed. The quality of the photometry returned by DoPHOT is assessed via reductions of an "artificial" globular cluster generated from a list of stars with known magnitudes and colors. Results from comparative tests between DoPHOT and DAOPHOT using this synthetic cluster and real data are also described.

1. INTRODUCTION

DoPHOT is a computer program designed to search for "objects" on a digital image of the sky and to produce positions, magnitudes, and crude classifications for those objects. Digital images, like the research programs for which they are obtained, vary enormously, and no single computer program is likely to handle all cases equally well. The particular project for which DoPHOT was written (Schechter and Caldwell 1989) called for star/galaxy classification and stellar photometry using an oddly shaped and poorly sampled point-spread function (PSF), and involved the reduction of a large number of images. DoPHOT was therefore written to be fast, highly automated, and flexible regarding the choice of PSF, with less emphasis on reaching the photon noise limit in photometric and astrometric accuracy.

Since its first incarnation DoPHOT has been revised and expanded, permitting it to deal with a wider range of problems. Some 50 copies of DoPHOT have been distributed by the authors to astronomers on six continents. It has been used to obtain photometry (much of which has proven surprisingly good) for tens of millions of stars (e.g., Udalski et al. 1992), to analyze *HST* images (Sandage et al. 1992), to determine the positions of comparison arc lines on echellograms (Metzger et al. 1991), and to locate spots on Hartmann test data (Schechter and Mack 1993). It has been used to reduce "drift-scanned" data (Schechter and Caldwell 1989; Caldwell et al. 1991) obtained both by turning off the telescope and by tracking the telescope backward. Its users have been encouraged to customize and optimize it for their own purposes. DoPHOT is briefly described by Mateo and Schechter (1989) in the proceedings of the First ESO/ST-ECF Data Analysis Workshop. In the present paper we elaborate upon that description, note recent improvements, and discuss ways in which the DoPHOT approach is different from those of other programs.

The authors believe that astronomers are only beginning to take full advantage of the potential of digital images, and that there is much room for improvement. The reader is therefore urged to keep in mind that while some aspects of DoPHOT might be worth adopting in a next-generation reduction program, there are others which might better be abandoned in favor of alternate approaches. Some idea of the many interesting alternatives which have already been explored may be had by reading the contributions to the ESO workshop proceedings (Grosbøl et al. 1989).

In Sec. 2 we give a general description of DoPHOT. In Sec. 3 we present a fanciful computer program written in an imaginary computer language which ties together the ideas presented in Sec. 2. In Sec. 4 we discuss details.

2. GENERAL DESCRIPTION

A central feature of DoPHOT is the adoption of a *model* for each kind of object which one seeks to identify within a digital image. The model for a star might, for instance, be an elliptical Gaussian. The model for a galaxy might *also* be an elliptical Gaussian, but one which is significantly bigger than those associated with stars. The model for a double star, which DoPHOT treats as a different kind of object, is composed of two single stars. To classify an object one chooses that model which best fits the object. The model for a single star is also used as a filter in searching for new objects. Using the same model for both detection and photometry guarantees a well-defined completeness limit based on signal-to-noise.

With the exception of the model for a cosmic ray, which is modeled as a single high pixel, DoPHOT's models are specified in terms of analytic functions with free parame-

[1]Hubble Fellow.
[2]Presently at the University of Michigan.

ters rather than as lookup tables (which, it should be remembered, are also models). Six parameters are associated with the simplest models: x and y position within the image, total (or central) intensity, and three shape parameters (height, width, and tilt). While one can adopt better models with more parameters at the cost of speed and simplicity, DoPHOT was written on the hunch that one might do reasonably accurate photometry and classification using relatively simple models.

A major consideration in the design of the program was that it operate with minimal user effort. Ideally the user supplies only rough estimates of the seeing and background sky for each digital image. While parameters for a given instrument (e.g., gain and read noise) must also be specified, this need be done only once for a set of images. DoPHOT uses the initial guess for the seeing to identify stars brighter than some fairly high initial detection threshold. The analytic function used to represent a star is then fit to a number of subrasters centered on different objects in order to determine a better estimate of the shape of a typical star. If an object is fit better by a model for a nonstellar object (e.g., a galaxy or a cosmic ray), it is classified as such. The appropriate fitted model is subtracted from the image, producing a "working" image. This object-subtracted working image is then searched for progressively fainter stars by lowering the detection threshold. This procedure—based on an algorithm incorporated in the WOLF photometry program developed by R. Lupton (Lupton and Gunn 1986)—bears considerable resemblance to the CLEAN algorithm (Högbom 1974) used with aperture synthesis data.

After each pass through the working image all objects found on previous passes are once again fit to derive improved estimates of the model parameters. This is done by adding back to the working image the previous best-fitting model. Since neighboring fainter objects have now been subtracted, the shape parameters and magnitude are expected to be better than those obtained in the previous pass. This avoids, in part, the need for fitting several neighboring objects simultaneously.

Throughout the above process, DoPHOT constructs and updates a *noise image* which provides weights for each pixel used in its nonlinear least-squares-fitting subroutine. It is also used to test whether a potential object is significantly above the background. The failings of the analytic PSF are most obvious at the positions of bright stars. Large positive residuals might be expected to trigger the false identification of spurious objects. To avoid such "phantom" stars DoPHOT adds extra noise to the noise array every time it subtracts a star from the working image. A negative consequence of adding this extra noise is that it lowers the likelihood of identifying faint stars near brighter stars.

Given the systematic pattern of residuals in the star-subtracted images, one would expect that the total fluxes derived from fitting the model PSF to the data would likewise suffer systematic errors. But to first order one would expect to make the same systematic error for all stars. As an aid in correcting for such systematic errors, DoPHOT calculates total fluxes inside a suitably chosen subraster. Such "aperture" magnitudes are very much more uncertain than "fit" magnitudes because there is much more Poisson noise from the sky (including regions around poorly fit or faint stars) inside the aperture than there is under the model profile.

In fitting a model to an object, the background sky level

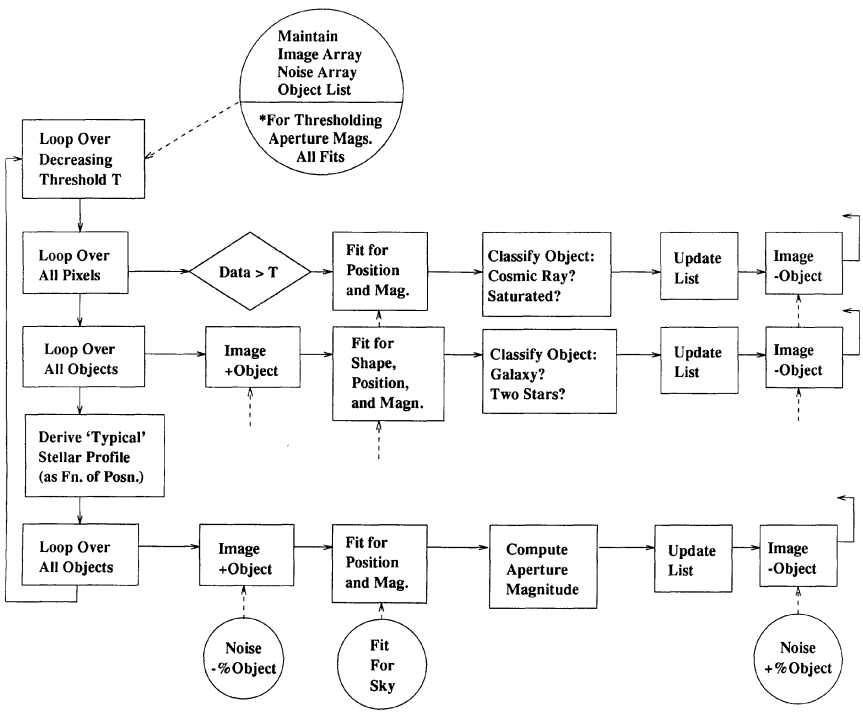

FIG. 1—A flowchart illustrating the operation of DoPHOT.

is treated as an additional free parameter. These sky levels can then be used to determine the background for use in thresholding on subsequent passes.

DoPHOT allows for a "warmstart" in which it reads a list of objects, subtracts them from the working image, and then searches for new objects. The user has the option of either allowing DoPHOT to redetermine the position of an object or insisting that DoPHOT fit a model at the specified position even if there is little evidence for an object. This permits measurement of upper limits on objects too faint to have been detected at a high level of significance. It also results in improved photometry in cases where the positions of the objects are well known.

DoPHOT is composed of a great many relatively short modules. One advantage of this approach is that it permits the rapid substitution of different PSF models—only a few modules need intimate knowledge of the details of the PSF. A second is that it allows for the quick replacement of modules with improved versions. Since no two data sets present exactly the same challenges the modules have evolved considerably since they were first written. It is hoped that they will continue to evolve as new circumstances are encountered. DoPHOT's users are encouraged to improve upon and to customize these routines and to report their experiences back to the authors.

3. A TOUR THROUGH DoPHOT

The following is a schematic version of DoPHOT in an imaginary computer language. The main tasks are identified and their relative locations in the code are correctly illustrated. Figure 1 is a more conventional flowchart for DoPHOT. Reference to this listing of CRYPTO_DoPHOT and the flowchart in Fig. 1 may be useful in following some of the more detailed descriptions of DoPHOT below.

```
CRYPTO_DoPHOT
  INPUT_PARAMETERS
  READ_IMAGE
  MAKE_NOISE_ARRAY
  {  if warmstart then
        READ_OBJECT_LIST
        DETERMINE_TYPICAL_SHAPE_PARAMETERS }
  {  for threshold = highest to lowest
        [ SEARCH_FOR_PIXELS_ABOVE_THRESHOLD
           ( if high_pixel then
                CHECK_FLUX_THROUGH_STAR_MASK )
           ( if significant then
                CLASSIFY_STAR/COSMIC/BADLY_SATURATED
                ADD_OBJECT_TO_LIST ) ]
        [ for all objects in list
             DETERMINE_OBJECT_SHAPE_PARAMETERS
             ( if object_is_big then
                  CLASSIFY_GALAXY/DOUBLE_STAR )]
        DETERMINE_NEW_TYPICAL_SHAPE_PARAMETERS
        [ for all objects in list
             ( if star then
                  DETERMINE_FIT_MAGNITUDE
                  DETERMINE_APERTURE_MAGNITUDE )
           DETERMINE_EXTENDEDNESS ]
        WRITE_OBJECT_LIST_TO_DISK
        WRITE_OBJECT_SUBTRACTED_IMAGE_TO_DISK }
  END
```

After the input image is read and the noise array created (using information on the read noise, electrons per digital number [DN], and Poisson photon count statistics), any "predetermined" objects supplied via the warmstart option are (a) subtracted from the image, and (b) used to compute the values (weighted by signal-to-noise) for the typical shape parameters for a star. In the absence of a warmstart, a good guess for these shape parameters is required, and is provided as part of the start-up parameters.

Before searching for new objects a reasonable idea of the local background sky is also required. On each pass through the data with successively decreasing thresholds, DoPHOT searches for pixels with a value higher than the threshold plus the local background. A first guess of the sky value is needed for the initial pass through the data (unless the median sky model is used). On subsequent passes the program models the variation of the sky across the field using either a uniform gradient model or a modified Hubble profile (Rood et al. 1972; Binney and Tremaine 1987). The latter is useful for stellar photometry in the vicinity of globular clusters. Successfully detecting fainter objects during subsequent passes depends increasingly critically on having a good estimate for the local sky.

When a pixel is encountered that is higher than the threshold plus estimated sky, the flux through a mask representing the current model of a stellar PSF is tested, first against the current sky model, and then against the actual average value of the pixels in the neighborhood. If both tests show adequate signal-to-noise this newly found object is appended to the object list. A test is immediately done to see if a cosmic-ray model is a better match than the model for a star. If so, the pixel is "turned off" and not included in any further calcualtions. If a significant number of pixels in the mask are saturated, DoPHOT concludes that a very bright star is present and a region surrounding the saturated pixels is excised from the image. The best fits of all other objects with the stellar PSF are subtracted from the working image.

After the entire image is searched for new objects, everything in the augmented object list is fitted anew to improve the classifications and the typical shape parameters. This involves fitting the position, brightness, shape, and sky values independently for each object. Since all objects known up to this point have been subtracted from the working image, the previously fitted value (for the object under consideration) must be added back to the image. The shape parameters are examined against those of the current typical stellar PSF, and the object is accordingly classified as a star, galaxy, or double star. Details of this classification are given in Sec. 4.4. An object may not have sufficient signal-to-noise to warrant the "shape" test; such objects remain "unclassified" and are treated like stars, but not used to update the parameters for the current estimate of the typical stellar PSF. There is scope here for the innovative user to devise other tests and object types as needed. After each object is tested its best fit to the appropriate model is again subtracted from the image.

At this point the individual shape parameters of all objects that have been classified unambiguously as single stars are combined with signal-to-noise based weights, to produce an updated model for the typical stellar PSF. (Some users have modified DoPHOT so that this typical PSF is taken to be a function of position on the image). This new

model is then fit to all objects that are *not* definitely nonstellar (i.e., single stars, individual components of double stars, and the unclassified objects). Since the shape parameters are not allowed to vary in this fit, all putative stellar objects have resulting magnitudes that are on the same footing. As before, previous fits are added to the image, fitted, and the new fits subtracted. Ideally, in the final pass through the data, the fits for all significant objects have been subtracted, leaving behind only the residuals. At this time, aperture magnitudes are also obtained for objects with adequate signal-to-noise. Other parameters for each object can be evaluated easily in this phase, in particular an "extendedness" parameter that helps *ex post facto* delineation of stars from galaxies for the "unclassified" objects.

The fitted sky values for definite stars and double stars are used to update the fit to the sky model. If the sky model is not analytic (see Sec. 4.6), the current working image, from which the objects detected thus far have been removed, can be used to create a better low-pass filtered sky lookup table image. The results from this pass are written out, and a new pass with lower detection threshold is begun.

4. SALIENT DETAILS

This section is intended primarily for readers who have given considerable thought to the fine points of PSF fitting and who might be interested in how DoPHOT deals with them.

4.1 Pixels versus Models

A theme which runs through the preceding sections is that of modeling an array of pixels. A small subraster is filtered through a model profile to determine whether or not an object should be added to the list. A somewhat larger subraster is fit with a model to determine the shape of an image. The image is modeled by the entire list of objects. Wherever possible the sky, as sampled by the pixels, is modeled. A goal in the design of DoPHOT was to use model fitting in preference to pixel-based algorithms wherever possible.

An immediate advantage of such a model-based approach is that one can fit the model even when the image is not uniformly sampled. If pixels are missing for some reason (e.g., a cosmic ray, or a bad column, or a saturated pixel at the center of an image), one can still fit a model to the data. Odd-shaped pixels are also easily accommodated. If one had data which were oversampled by a factor of 2, such that the FWHM was of order four pixels, then it should be possible (and sometimes desirable) to compress the data by a factor of 2 in both directions and obtain very nearly the same output object list.

At some point, of course, pixels must be explicitly considered. For example a single high pixel is used to trigger the subsequent test for an object. A cosmic ray is modeled by a single high pixel, and the shape parameters which describe the typical star are dimensioned in pixels. One might choose to model a star with an empirical lookup table, in which case pixels are an obvious possible grid spacing. While this may be expected to lead to better photometry for overlapping stars, we find in practice that analytic PSFs do surprisingly well. In some particular instances, e.g., for undersampled images, analytic PSFs have distinct computational advantages.

4.2 Point-Spread Functions and Accuracy

While an elliptical Gaussian is an obvious possible model for a stellar point-spread function, the surface-brightness profiles of stellar images tend to look more like power laws. The actual model used in DoPHOT consists of similar ellipses of the form

$$I(x,y)=I_0\left(1+z^2+\frac{1}{2}\beta_4(z^2)^2+\frac{1}{6}\beta_6(z^2)^3\right)^{-1}+I_s,$$

where

$$z^2=\left[\frac{1}{2}\left(\frac{x^2}{\sigma_x^2}+2\sigma_{xy}\,xy+\frac{y^2}{\sigma_y^2}\right)\right],$$

and

$$x=(x'-x_0);\quad y=(y'-y_0),$$

with the nominal center of the image at (x_0,y_0). If $\beta_4=\beta_6=1$, Eq. (1) is just a truncated power series for a Gaussian. In practice the user is allowed to specify these two parameters to obtain better looking residuals images. In our experience a wide variety of PSFs can be generated by an appropriate choice of values. There are seven free parameters in this function: the shape parameters σ_x, σ_y, and σ_{xy}; the object center (x_0,y_0); the central intensity I_0; and the background intensity I_s.

The choice of an analytic PSF, as opposed to a tabulated empirical PSF, was driven by several considerations. First was expedience—the authors had experience in fitting analytic functions but none in fitting tabulated functions. Next was speed of calculation. Another was the suspicion that one could do good *relative* photometry within an image even with an imperfect approximation to the point-spread function, since the systematic error arising from the approximate nature of the model should be the same for all stars. This is not strictly correct, since the profiles of bright stars are weighted by the photon statistics of the star itself, while the profiles of faint stars are weighted by the photon statistics of the sky. Results from a limited set of experiments would indicate that the associated scale errors are less than 0.01 mag per magnitude.

4.3 Phantom Stars and the Noise Array

The most obvious shortcoming of an analytic, as opposed to empirical, point-spread function is the larger re-

siduals one gets from the best-fitting model. Since DoPHOT makes multiple passes through the object-subtracted image there is a danger that the program will trigger on a positive residual and identify a spurious "phantom" star. Phantom stars may be checked for in the "synthetic" image, generated by subtracting the object-subtracted image from the original image. Another useful diagnostic is to plot the (x,y) positions of objects using a plotting program, keeping an eye out for unusual grouping of faint objects around bright objects or at the edges of excised regions.

DoPHOT uses its noise array to avoid such phantom stars. Whenever an object is subtracted from an image a user specified fraction of the subtracted signal is added, in quadrature, to the noise array. The larger one makes this fraction the less likely one is to detect a spurious object. In increasing this fraction one pays two penalties: one is less likely to detect faint stars in the vicinity of bright stars and the fits to objects which are detected in the vicinity of another object will be more uncertain than they would otherwise be.

While this scheme seems to work reasonably well at the centers of objects, phantom stars still appear on the peripheries of bright stars when the actual stellar profile has broader wings than the analytic model. Ideally, one ought to try a different analytic approximation. Short of that, one can adjust a factor which expands the shape parameters used in augmenting the noise array by a user specifiable factor.

4.4 Star/Galaxy/Double-Star Classification

If a subraster of pixels produces a significant signal when tested against the stellar model filter one does not know whether this is due to the light from a single star, from a galaxy, or from two or more stars which lie relatively close to each other.

DoPHOT attempts to deal with this classification question economically, taking advantage of the fact that is must determine the shapes of objects to determine the parameters associated with a "typical" star. When those shape parameters are different from the typical image at a level of statistical significance specified by the user and the sense of that difference is that the area of the associated footprint is larger, DoPHOT declares the object to be "big."

The significance of the difference between the shape parameters for an individual object and those for a "typical" star depend first, upon the errors in the fit for the individual star and second, on the star-to-star scatter in those parameters. Given a perfect instrument, and an ensemble of stars of the same brightness, one would expect these to be the same. For many instruments (e.g., wiggly CCDs in fast beams) they are not. DoPHOT computes a star-to-star scatter in the stellar parameters which weights the bright stars, those for which the accidental errors from the fit are likely to be small, more heavily. The scatter for each shape parameter is added in quadrature to the uncertainties for that parameter obtained from the fit to an individual star. The significance of the difference between the measured and expected shape parameters is computed taking this as the expected difference.

When an object is found to be "big," DoPHOT attempts to fit two typical stellar profiles to the subraster. The goodness-of-fit parameters returned from the fits to the single object and the two typical objects are compared, and based on a user specifiable parameter a decision is made on its classification. If the object is double, it is "split" and two entries are made in the object list. These entries are then subject to further testing and splitting on subsequent passes through the image. DoPHOT treats split stars as a distinct class of object, and reports them as such. Split stars can be later reclassified as galaxies or split again, but they cannot be reclassified as single stars.

The criterion for deciding whether a "big" object is a galaxy or two stars clearly depends to some extent upon the nature of the field observed. At high galactic latitudes, in regions of low star densities, the user would want to adjust the parameter which controls the galaxy/double-star decision to favor galaxies. At low galactic latitudes and in star clusters the decision should favor double stars. The synthetic image is useful for diagnosing how this parameter should be set.

One difficulty with modeling objects is that at progressively lower signal-to-noise ratios one can support fewer and fewer free parameters. Since the fitting in DoPHOT is carried out iteratively, the symptom of too low a signal-to-ratio is a failure to converge. There is a user specifiable signal-to-noise limit such that only objects with higher signal-to-noise are fitted for shape parameters. Objects which are fainter than this cannot be reliably classified as either stars or galaxies and are treated as a unique class of object. These are fit only for sky, (x,y) position, and central intensity while adopting the current best estimates of the stellar shape parameters.

A large fraction of such faint objects may be galaxies, especially in fields at high galactic latitudes. To provide at least some handle on whether or not such objects are galaxies/double stars or single stars, an "extendedness" parameter is computed for every object. Using the typical stellar parameters for an initial guess, the fitting program is allowed to determine the first step toward improving the shape parameters. The change in the goodness-of-fit parameter predicted by its second derivative matrix is then estimated. If the sense of the change is to decrease the area of the footprint, this change is reported as a negative number; otherwise it is reported as positive.

Limited tests indicate that while this quantity (PGAL) does help to separate stars from galaxies (and from globular clusters in nearby galaxies), its value is not independent of the brightness of the object. In plots of magnitude versus PGAL, the PGAL values for bona-fide stars lie over a narrow range of PGAL which is relatively independent of magnitude, whereas the PGAL values for galaxies fall along a broad, magnitude-dependent swath, with values *higher* than those for stars. The separation of the ridge line for stars and galaxies is good at brighter magnitudes, but gets smaller as one goes fainter. At the very faint end, within 1.5 mag of the detection limit, the range of PGAL

values for stars increases rapidly to the point that star-galaxy separation is much less reliable.

A number of other PSF-related parameters are also available from DoPHOT processing (even though they may not be currently reported in the standard output). Some of these may be used in conjunction with PGAL in a cluster/discriminant analysis in multidimensional space to provide more robust star-galaxy separation for fainter objects. This is an area that requires some experimentation.

4.5 Completeness and Faint-End Bias

The model which is used as the filter for measuring the signal-to-noise ratio for a potential object is identical to that of the typical stellar profile used for fitting the object. It should not come as a surprise, therefore, that when the number of objects identified is plotted as a function of apparent magnitude obtained from the fit, one finds that it drops to zero quite abruptly. The detection and the measurement are carried out consistently.

The naive user might be deceived into thinking that the sample of detections is remarkably complete. It is not. Consider several stars for which the expected signal-to-noise ratio is exactly the limiting value. Some of those will be superposed on positive noise fluctuations, and will be detected, with measured fluxes brighter than their true fluxes. Others will be superposed on negative noise fluctuations, and missed. Not only will the sample be incomplete, the estimated magnitudes will be biased by an amount that depends upon the logarithmic slope of their number-magnitude distribution and upon the limiting signal-to-noise ratio.

4.6 Model Sky and Average Sky

For every star in the object list DoPHOT produces an estimate of the sky brightness. It is therefore possible to construct a "model" of the variation of the sky brightness over the face of the chip. This model can then be used to *estimate* the sky value when a potential object is filtered through the typical stellar profile.

If the object passes through the filter using the model sky, it is tested once again, this time using a weighted average of the sky computed from the fit subraster. This second test avoids spurious detections where the background level is changing in a way which is not modeled.

The model for the sky need not be a constant. Typically it is a plane, which allows for a small gradient across the field. An option which has proved quite useful is modeling the sky as a "Hubble" profile plus a constant. In fitting this model (and the much simpler plane) the data points are the values of the sky derived from fits to individual stars. The "Hubble" profile requires a relatively large number of points (of order 50) to converge.

In another option the working image is median filtered to obtain an image which closely matches the sky background. This image serves as a lookup table model of the sky, and is useful when the background varies erratically over the image but on scales larger than the FWHM of a stellar PSF.

4.7 Warmstarts and Fixed Positions

There are a variety of circumstances under which it is helpful to give DoPHOT a starting list of objects. For example, having done a first pass on the data, one might want to try lowering the threshold without having to start over from the beginning. Or one might want to insert an object by hand into the list, or to change the parameter that controls whether objects are stars or galaxies. This feature allows the designation of entirely new classes of objects—stars (and galaxies) whose positions (and shapes) are taken as given. For these objects, DoPHOT computes fluxes and uncertainties using these specified positions. This can be helpful when one has several images taken under different seeing conditions: the positions obtained with good seeing can be used to constrain the photometric fits to the poorer data, where the blending of images might otherwise produce larger uncertainties.

4.8 Saturated Stars

Bright stars present several problems for automated photometry. Since detectors often go nonlinear and then saturate at high surface brightness levels, it is difficult to subtract badly overexposed stars from the image. Rather than try to fit and subtract such an object, DoPHOT automatically excises a rectangular subraster around the brightest objects. However, because different detectors saturate more and less gracefully, no general algorithm will do as well as masking pertinent portions of the image.

4.9 Trouble

A number of circumstances arise where, try as it might, DoPHOT cannot come up with magnitudes or shapes for objects. The most common of these is when there are too few samples to do an adequate job of fitting the model. There are two user adjustable parameters which specify the fraction of the fit subraster that must be present for an object to be fit. Special object types are assigned for particular problems: if too few pixels are present for the shape of an object to be determined, and/or for objects where the shape fitting does not converge.

Similarly special object types are also given to cases where there are not enough pixels to do even a fit with the typical stellar PSF, or if such a fit does not converge. In these two final cases it is not clear what to subtract, and such objects are therefore *not* subtracted—they are said to have been "deactivated." They are left in the object list so that a record is left that something out of the ordinary has happened. If the signal-to-noise ratio of an object which was previously identified is found to be less than the user specified value for the identification of objects, the object is given a special type and "deactivated."

4.10 Internal Parameter Representation

DoPHOT fits for the parameters σ_x^2 and σ_y^2 rather than σ_x and σ_y so that the program can recover more easily from iterations which might otherwise make σ_x and σ_y negative. The quantity σ_{xy} is measured in *inverse* square pixels so that round images produces parameters which behave well. When fitting models DoPHOT divides the central intensity and sky values by 100 to avoid overflows.

When fitting large objects for two "typical" stars, DoPHOT treats the logarithm of the central intensity as the free parameter; this allows recovery from overzealous attempts to decrease the brightness of one of the two elements of the double star.

4.11 Error Estimation

The nonlinear least-squares model fitting is done by minimizing χ^2. The covariance or curvature (of the χ^2 surface) matrix used in solving this problem contains the necessary information for estimating errors in the fit parameters. The fit magnitudes are derived directly from the height parameter of the fitted stellar model, so the errors in the magnitudes are the same as the errors in determining this particular fit parameter.

Under the assumption that the fitted parameters are not strongly correlated, a diagonal term of the curvature matrix is the reciprocal of the variance (σ_a^2) of the corresponding fitted parameter (a), provided that the noise and object models are both perfect so that the minimum χ^2 per degree of freedom (henceforth the "reduced" χ^2) is unity. Thus σ_a corresponds to the amount by which the parameter a must be changed in order to change χ^2 by unity, while allowing the other parameters to change as they will in order to minimize χ^2. Unfortunately, if the model being fitted is less than perfect, or if the real noise is not modeled accurately, the minimum reduced χ^2 will not be unity. The approach taken in DoPHOT is to scale σ_a^2 given above by the value of the minimum reduced χ^2 for each fitted object. In effect, the errors are being increased (or decreased) on the assumption that the error estimates for each pixel have been overestimated (or underestimated) by the same factor.

From the analysis of repeat observations of a given field the formal errors derived by DoPHOT have been found to be in excellent empirical agreement with direct estimates made by comparing photometry on repeat observations (Caldwell et al. (1991).

5. DoPHOT PERFORMANCE

5.1 DoPHOT Reduction of Tonry 1

As a check on the quality of the stellar photometry returned by DoPHOT in a crowded field, we have systematically reduced simulated CCD frames of a dense globular cluster. These frames were generated for us by John Tonry, and we shall refer to this artificial cluster as "Tonry 1;" the I-band image of the cluster is shown in Fig. 2. Figure 3 shows the adopted color–magnitude diagram (CMD; in I, $V-I$) for the cluster. Only the brightest 87173 stars are plotted; however, a total of 231439 stellar images were generated to make the final images of Tonry 1. The faintest stars are about 2 mag fainter than the limit shown in Fig. 3 and were included to simulate the unresolved background light present in real clusters. The CMD for Tonry 1 is meant to be similar to that of 47 Tuc with the exception that the "cosmic scatter" about the principal sequences was assumed to be negligible in order to provide

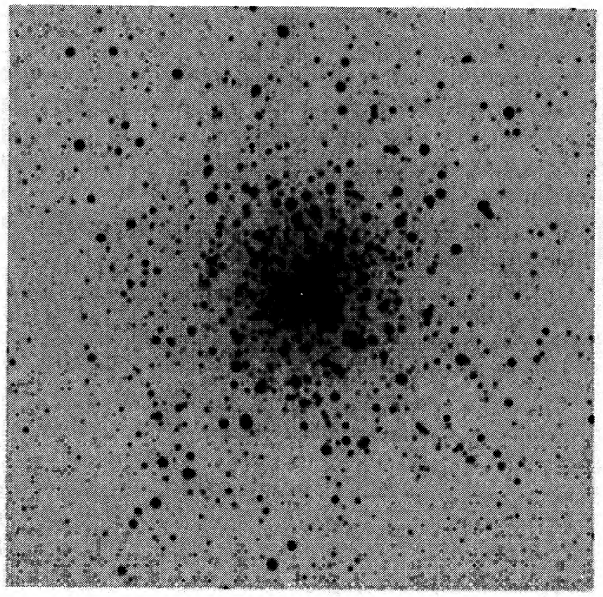

FIG. 2—An I-band "image" of Tonry 1.

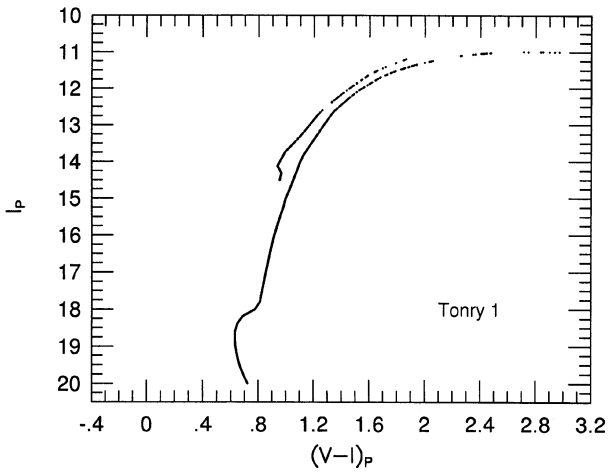

FIG. 3—The "true" I vs. $(V-I)$ color–magnitude diagram of Tonry 1. Only the brightest 87173 stars are plotted. A total of 231439 stars were generated to produce the synthetic CCD images of Tonry 1. This diagram is based on a fiducial sequence meant to mimic the color–magnitude diagram of 47 Tuc.

an easy visual check on the quality of the measured photometry. For both the artificial V- and I-band frames, the stellar profiles are scaled copies of the observed PSF on a set of CCD images obtained by John Tonry using the prime-focus camera on the KPNO 4-m telescope. The distribution of stars follows a relatively concentrated modified Hubble profile (Rood et al. 1972) profile with a core radius ($R_{\rm core}$) of about 80 pixels. The cluster center was placed in the middle of both of the 1024×1024 frames. Photon and readout noise were added to the simulated images, although there was no attempt to mimic detector flaws or cosmic ray events. The counts were scaled to minimize the number of saturated stars.

DoPHOT possesses a large number of user-adjustable parameters. In general, these parameters can be fine-tuned to tailor the reductions to the specific task faced by the program, though in general the final results are often remarkably robust to even rather large changes in the values of most of these parameters. For our reductions of Tonry 1 we specifically did not try to fine-tune DoPHOT's parameters; rather, we adopted standard values suitable for any general reduction of a crowded field. This is also true of the DAOPHOT reductions we describe below. Our aim was to perform a *representative* reduction using DoPHOT rather than try to achieve the very best possible results. In this spirit, many of the numeric parameters used by DoPHOT were determined in an automatic manner based only on the FWHM and sky estimate. The background model assumed for the Tonry 1 reductions was the median-smoothed version of the star-subtracted image (see Sec. 4.6). The total time required to generate this parameter file was a few minutes. Details of the parameters and how they are input to the program are described in the DoPHOT user's manual which is distributed with the program.

The left-hand panels in Fig. 4 show the measured CMDs for Tonry 1 using DoPHOT. Results are shown for three regions: within 100 pixels of the cluster center, from 100 to 300 pixels of the center, and beyond 300 pixels from the cluster center. Not surprisingly, the photometry improves significantly as one works further from the cluster center. The tightness of the brighter portions of the principal sequences in the CMD of the outer region demonstrates that DoPHOT has produced very good relative photometry.

Figure 5 is a plot showing the differences in the I magnitudes and $V-I$ colors for the stars measured by DoPHOT in the outer region and the true magnitudes and colors of the stars (these are known from the list used to generate Tonry 1). The matching program used to generate Fig. 5 considers only stellar positions to identify a star in the photometry lists. The matching radius used here was 1 pixel. There was no attempt to "optimize" the matching radius by determining the number of matches as a function of this radius (e.g., see Bailyn et al. 1992); however, the results described here are not significantly changed from the case where a matching radius of 0.75 pixels was adopted. It is clear from Fig. 5 that the residuals vary such that the measured magnitudes are brighter than the true magnitudes. Given the high density of stellar images throughout Tonry 1, this trend is not surprising; it undoubtedly results from blending of faint stellar images. Some of the most extreme residuals in Fig. 5 clearly correspond to mismatched stars where the measured stellar position was sufficiently in error that the star was matched with a fainter object in the list containing the true photometric results. At $I\lesssim16$ (where blending begins to occur frequently), the differences in the measured minus true magnitudes and colors have a standard deviation of $\lesssim0.03$ mag for each annulus; for the brightest stars, $\sigma\sim0.002$ mag.

5.2 DAOPHOT Reductions of Tonry 1

For the sake of comparison, we have also reduced the Tonry 1 data using a version of the commonly used CCD reduction program DAOPHOT (Stetson 1987) supplied by its author to astronomers at the Observatories of the Carnegie Institution of Washington. As with DoPHOT, we did not in any way attempt to perform an "exceptional" reduction of the cluster using DAOPHOT, although a conscientious effort was made to construct a reasonable point-spread function from stellar images in the outer region of the cluster. Also, as discussed in more detail in Sec. 5.3 below, the DAOPHOT reduction was performed in two passes; after running ALLSTAR on the original frame, the subtracted image was searched for additional stars and a second pass with ALLSTAR was performed using the combined object list. The CMD of Tonry 1 generated by DAOPHOT are shown in the right-hand panels of Fig. 4, while the difference of the DAOPHOT photometry and the true photometry is plotted in Fig. 5. We shall defer the discussion comparing the DoPHOT and DAOPHOT results to Sec. 5.3 below. For now, we stress that the DAOPHOT photometry clearly also suffers from the same blending errors apparent for DoPHOT (also in Fig. 5); this effect is clearly worse in the middle and inner annuli. That these systematic deviations are not intrinsic to either program is emphasized by the overall good agreement (in the mean) of the DoPHOT and DAOPHOT results (Fig. 6). The implication is that all crowded-field photometry probably suffers from significant faint-end errors due to blending. The magnitude of this effect cannot be easily be determined solely from comparing actual data obtained under similar seeing conditions; to ascertain the magnitude of this effect one must perform extensive false-star experiments.

5.3 A Comparison of the DoPHOT and DAOPHOT Results

5.3.1 Tonry 1

Because DAOPHOT is well known in the community as a CCD reduction program for crowded fields, it is particularly instructive to compare the DAOPHOT and DoPHOT results for Tonry 1. We re-emphasize that no attempt was made to maximally optimize either code for these reductions; rather, our goal was to simulate results corresponding to a typical reduction using the two programs. There is no doubt that expert DoPHOT and DAOPHOT users could

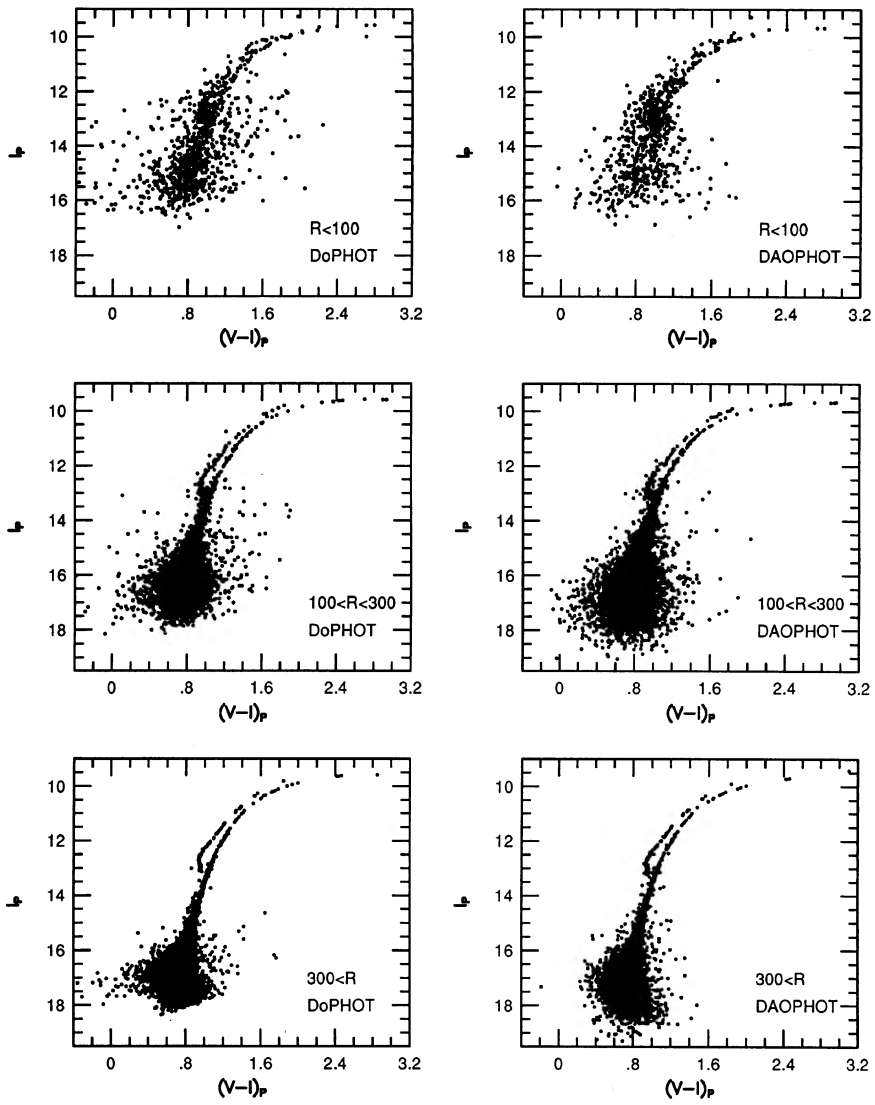

FIG. 4—(Left panel; top) The I vs. $(V-I)$ color–magnitude diagram of the inner region of Tonry 1 (i.e., for stars within 100 pixels from the cluster center) measured using DoPHOT. (Left, center) The same for stars located in a middle annulus (100–300 pixels). (Left, bottom) The same for stars located more than 300 pixels from the cluster center. (Right panels) The same as the left-hand panels except for DAOPHOT reductions of Tonry 1. the zero point of the ordinate scales on all of these plots is the same, but arbitrary.

significantly improve on the results described here by adjusting many of the tuneable parameters in both programs.

A comparison of the results shown in Fig. 4 do not reveal any striking differences in the overall quality of the photometric results generated by the two programs. In the inner field, DoPHOT does appear to have produced "tighter" sequences than DAOPHOT, though at the same time there are more extremely deviant stars in the DoPHOT CMD (in fact, there are more stars overall in all of the DoPHOT diagrams than in the corresponding CMDs produced by DAOPHOT). The results for the middle and outer fields are quite similar, though again DoPHOT does seem to produce a relatively larger number of extreme outliers; this is especially evident near the cluster turnoff in the DoPHOT CMD of the outer annulus of Tonry 1.

The mean difference in the DAOPHOT and DoPHOT photometry is reasonably small in the mean (Fig. 6). The slight systematic trends that are apparent clearly depend on the average crowding. For example, in the outer annulus the DoPHOT and DAOPHOT results differ by up to about 4% within 1–2 mag of the limits of the photometry in the sense that DAOPHOT returns brighter results. In the middle annulus, this effect is more pronounced, reaching about 16% at the limits of the photometry with DAOPHOT again returning brighter values. Because of poorer statistics, the overall trend in the inner annulus is more difficult to de-

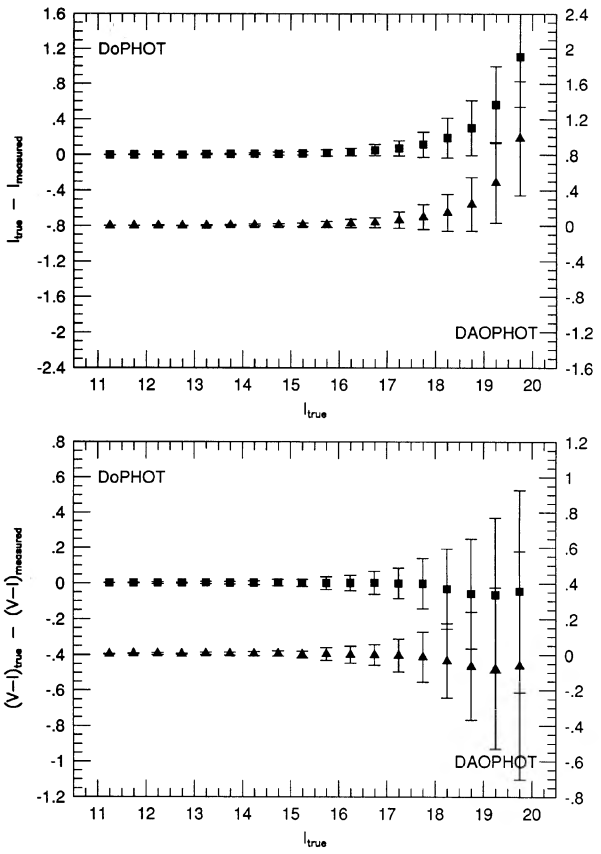

FIG. 5—(Top panel) The difference of the I magnitudes (in the sense "true" *minus* measured) as a function of the "true" I magnitude for DoPHOT (squares) and DAOPHOT (triangles). The left-hand axis is for the DoPHOT results. Both sets of differences have been shifted to zero for the brightest magnitude bins. Only stars measured on both the V and I images have been include in this figure. (Bottom panel) A similar plot of the differences in the $(V-I)$ colors as a function of the "true" I magnitude. The symbols and axes are defined in the same way as in the upper panel.

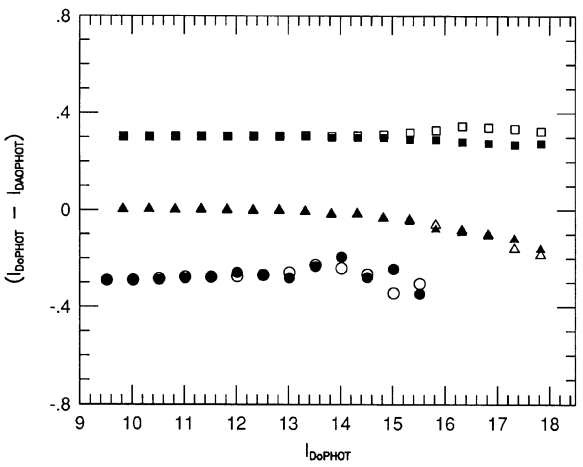

FIG. 6—A plot of the mean differences in the DoPHOT and DAOPHOT I magnitudes as a function of DoPHOT magnitude for the outer (squares), middle (triangles), and inner (circles) annuli. The solid symbols represent the actual differences in the mean magnitudes of stars measured in common by both programs. The open symbols are the calculated magnitude differences based on the mean difference in the sky values computed by each program; see Sec. 5.3.1 for details. The plots for each annulus have been offset by 0.3 mag for clarity.

termine; however, there is a general tendency for the DAOPHOT results to become relatively brighter than the DoPHOT results as one goes to fainter stars.

These results reveal true differences in the output of the two programs. One possible source of this effect is the systematic differences in sky values calculated by DoPHOT and DAOPHOT. These differences can be summarized as follows. For bright stars, DoPHOT returns higher sky values than DAOPHOT. This probably results from contamination by the wings of bright stars in the outer parts of the fit boxes used by DoPHOT. For fainter stars, DAOPHOT systematically returns higher sky values than DoPHOT, with the effect worse for stars in more crowded fields. The source of this discrepancy is not understood. In terms of the final photometry, the second effect is clearly more important since even a small change in the sky value is relatively much more severe for the faintest stars.

In an attempt to quantify this we have estimated the expected photometric variations caused solely by these differences in the sky values. For each magnitude bin shown in Fig. 6, a mean sky difference Δ_s was determined for stars measured by both programs. This was translated into a brightness excess (or deficit) using the relation $\Delta_b \equiv 2\pi\sigma^2\Delta_s$, where σ is the Gaussian scale length of the stellar profile; in Fig. 6 Δ_b is plotted as a function of magnitude for each annulus measured in Tonry 1 (open symbols) along with the actual magnitude differences (closed symbols). It is apparent that the differences in the sky values returned by DoPHOT and DAOPHOT do account for some of the systematic differences in the final photometry returned by the two programs. This is especially true in the inner and middle annuli. In contrast, in the outer annulus the observed magnitude differences and Δ_b diverge for the faintest magnitude bins. Apparently, factors other than just the differences in the sky values—for example the detailed manner in which DoPHOT and DAOPHOT fit their PSFs to the stellar profiles—contribute significantly to the observed differences in the photometric results of the two programs.

5.3.2 Independent Empirical Comparisons

Further comparisons between DoPHOT and DAOPHOT using real data have been carried out and described in the literature, most recently by Janes and Heasley (1993), who compared results obtained with DoPHOT and DAOPHOT with results from a program of their own. One such study was that of Friel and Geisler (1991) who compared DoPHOT and DAOPHOT reductions of the low-latitude globular clusters NGC 5927 and NGC 6496. The main observations from their comparison are:

(1) The principal sequences in the color–magnitude diagrams that they obtain for their clusters using DoPHOT

"appear to be more tightly defined" than those obtained with DAOPHOT.

(2) The derived magnitudes show no systematic difference at the bright end, but for faint objects DoPHOT magnitudes are brighter than DAOPHOT magnitudes. This difference is greater for more crowded fields.

(3) The sky (background value) determinations from DoPHOT are systematically lower than those obtained with DAOPHOT. The amount of the discrepancy increases with crowding, and in the sparse limit the difference is negligible. This is consistent with the findings for Tonry 1 (cf. Fig. 6).

Friel and Geisler attribute the systematic differences in magnitude primarily to the difference in sky determinations. To verify and understand these effects, we performed the following experiment. DAOPHOT was run on a crowded field in the usual way. The aperture preprocessing was then done again, but on the *subtracted* image. DAOPHOT predetermines the sky values around each object from this preprocessing, and does not change the sky values during fitting. The sky values obtained from the subtracted image are negligibly different from those obtained with DoPHOT. This indicates that in crowded fields, the far wings of stars contribute significantly to the ambient "background." The subtracted image has this effect reduced. Thus iterative processing with DAOPHOT (i.e., rerunning DAOPHOT on the original image but with sky estimates from the subtracted image) results in the same sky background as that given by DoPHOT, which already does the equivalent iteration in its CLEAN-like algorithm.

The systematic discrepancy in the faint stars from DAOPHOT and DoPHOT are greatly reduced by the DAOPHOT iteration described above, confirming Friel and Geisler's conclusion that the discrepancy in the background contributes greatly to this effect. However the discrepancy does not disappear entirely. The source of the remaining discrepancy lies elsewhere: the selection criteria for objects is very different in the two programs, and the measurement of limiting faint objects is biased in different ways. DAOPHOT finds objects in a separate preprocessing pass. If the "find" parameters are shallow, the object selection is signal-to-noise limited. If the "find" parameters are too deep, specious detections result, and objects are rejected during the fitting process. The selection effects are no longer well defined. On the other hand, DoPHOT identifies objects in on a strict signal-to-noise basis and fits *all* objects that are identified. Thus the output object list has a well-defined signal to noise limit with its associated faint end bias (see Sec. 4.5). Thus in both cases the systematic magnitude shifts at the faint end must be characterized from synthetic star experiments that precisely reproduce the quality of the data at hand as well as the parameters used in object identification and photometry.

5.4 Comparisons of Running Time

Both the DoPHOT and DAOPHOT reductions of Tonry 1 were carried out on Sparc 2 workstations at Carnegie Observatories. Whereas DoPHOT required approximately 1.8 CPU hr to measure the stars on the two frames, DAOPHOT took about 15 hr for a comparable number of stars. In terms of the number of stars reduced per hour, DoPHOT is clearly superior. Much of this gain in speed is probably attributable to the use of an analytic PSF, and the fact that no more than two stars are ever fit simultaneously in DoPHOT. In contrast, DAOPHOT allows simultaneous reductions of large numbers of stars; each of these solutions can be very time consuming because of the need to invert $3n+1 \times 3n+1$ matrices, where n is the number of stars in the group being fit. Because of its relative speed, DoPHOT is currently being used by two groups attempting to identify microlensing events by performing near-real-time photometry in crowded fields in the LMC and Galactic Bulge (Alcock et al. 1992; Udalski et al. 1992). It is also very useful as a tool to reduce large sets of CCD images of individual fields (e.g., Balona and Jerzykiewicz 1992; Freedman et al. 1992; Kubiak et al. 1992; Carney et al. 1993; Kaluzny and Krzeminski 1993).

We would like to thank the several colleagues who have tested and supplied useful feedback about DoPHOT, in particular Scott Anderson for sharing with us his experiments with the extendedness parameter. Ian Thompson provided MM with DAOPHOT assistance; we thank him for this help. We also gratefully acknowledge John Tonry for creating and letting us use Tonry 1 to test DoPHOT, and John Caldwell and Christopher Naylor for producing Fig. 1. This work was supported in part by NSF grant AST83-18504, and by grant HF-1007.01-90A awarded by STScI which is operated by AURA for NASA under contract NAS5-26555.

REFERENCES

Alcock, C., Axelrod, T. S., Bennett, D. P., Cook, K. H., Park, H.-S., Griest, K., Perlmutter, S., Stubbs, C. W., Freeman, K. C., Peterson, B. A., Quinn, P. J., and Rodgers, A. W. 1992, in Gravitational Lenses, ed. R. Kayser, T. Schramm, and L. Nieser (Berlin, Springer), p. 156

Bailyn, C. D. 1992, ApJ, 392, 519

Balona, L. A., and Jerzykiewicz, M. 1993, MNRAS, 260, 782

Binney, J., and Tremaine, S. 1987, Galactic Dynamics (Princeton, Princeton University Press), p. 39

Caldwell, J. A. R., Keane, M. J., and Schechter, P. L. 1991, AJ, 101, 1763

Carney, B. W., Storm, J., and Williams, C. 1993, PASP, 105, 294

Friel, E. D., and Geisler, D. 1991, AJ, 101, 1338

Freedman, W. L., Madore, B. F., Hawley, S. L., Horowitz, I. K., Mould, J. R., Navarete, M., and Sallmen, S. 1992, ApJ, 396, 80

Grosbøl, Murtaugh, and Warmels 1989, Proceedings of the First ESO/ST-ECF Workshop on Data Analysis

Högbom 1974, A&AS, 15, 417

Janes, K. A., and Heasley, J. N. 1993, PASP, 105, 527

Kaluzny, J., and Krzeminski, W. 1993, MNRAS, in press

Kubiak, M., Kaluzny, J., Krzeminski, W., and Mateo, M. 1992, Acta Astron., 42, 155

Lupton, R. H., and Gunn, J. E. 1986, AJ, 91, 317

Mateo and Schechter 1989, Proceedings of the First ESO/ST-

ECF Workshop on Data Analysis, ed. Grosbøl, Murtaugh, and Warmels), p. 69
Metzger, M. R., Caldwell, J. A. R., McCarthy, J. K., and Schechter, P. L. 1991, ApJS, 76, 803
Rood, H. J., Page, T. L., Kintner, E. C., and King, I. R. 1972, ApJ, 175, 627
Sandage, A., Saha, A., Tammann, G. A., Panagia, N., and Macchetto, D. 1992, ApJ, 401, L7
Schechter, P., and Caldwell, J. 1989, in The Gravitational Force Perpendicular to the Galactic Plane, ed. A. G. Davis Phillip (Schenectady, L. Davis), p. 143
Schechter, P., and Mack, P. 1993, in preparation
Stetson, P. B. 1987, PASP, 99, 191
Udalski, A., Szymanski, M., Kaluzny, J., Kubiak, M., and Mateo, M. 1992, Acta Astron., 42, 253