

Reduction of WFC Images with DAOPHOT III

Peter B. Stetson¹

Abstract

New additions to the DAOPHOT family of stellar-photometry software are described, and results of their application to WFC imagery are presented.

I. Introduction

For almost exactly ten years, I have devoted a major fraction of my professional efforts to the development of software for extracting stellar photometry from digital images. The agglomeration of code that has resulted may be referred to by the generic name DAOPHOT, but that term includes a number of generations and a myriad of modifications. DAOPHOT Classic (Stetson 1987) was the first photometry package—as far as I know—to incorporate the concept of the hybrid point-spread function (PSF): the model PSF of an image is first approximated by a continuous analytic Gaussian function, and the brightness residuals from that fit are stored as a look-up table of corrections from the analytic first approximation to the true model PSF. When the brightness value for a given pixel at a particular point in the stellar profile is to be predicted, the analytic first approximation is numerically integrated over the area of that pixel, and then a correction to the true PSF is obtained by interpolation within the look-up table. The hybrid PSF succeeds because the look-up table provides the flexibility to cope with asymmetric or irregular PSFs, while the analytic first approximation, representing most of the flux in the profile, provides the high-order spatial derivatives needed for accurate interpolation in critically sampled or slightly undersampled grids. PSFs which varied with position in the digital image were soon encountered; this was dealt with by replacing the one look-up table of corrections from analytic to true PSF with three tables, which allowed the empirical correction at each point in the profile to be represented by a first-order Taylor expansion as a function of position in the frame.

As we moved into the *HST* era, it became necessary to deal with even more severely undersampled and spatially complex PSFs than we had seen before. DAOPHOT II: The Next Generation (Stetson, Davis, & Crabtree 1990) was written before we learned of the spherical aberration in *HST*, but it has fortuitously turned out to be comparatively effective in dealing with the aberrated PSF as well (Stetson 1991, 1992). DAOPHOT II: TNG allows the user a choice of analytic first approximations — a Gaussian function (as before), two different Moffat functions, a Lorentz function (this is the best for *HST*), and the sum of a Gaussian function with a Lorentz function

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

(this is seldom used). In addition, six look-up tables of corrections allow for a PSF which varies as a quadratic function of position in the frame. Various other differences in detail, most notably in how the average PSF is estimated from a large number of stars at various positions and magnitudes in the science frame, were incorporated at about the same time. For instance, I made it possible to include centrally saturated stars in the PSF; the saturated pixels near the peak of the profile are ignored, but the unsaturated outskirts can improve the signal-to-noise ratio in the extended wings of the model PSF.

II. DAOPHOT III: This Time It's Personal

More recently, I have undertaken some further refinements of the DAOPHOT approach to stellar photometry; these have been found helpful in reducing images obtained with the aberrated *HST*, and I would like to describe them here.

First and most trivially, the maximum number of look-up tables containing the Taylor expansion of the corrections from the analytic first approximation of the PSF to the true PSF has been increased from 6 to 10, to allow encoding the stellar profile as a cubically varying function of position.

Second, in developing the empirical PSFs to be used in the reduction of the images of M81 obtained for the *HST* Extragalactic Distance Scale Key Project (Freedman, et al. 1994; Hughes, et al. 1994), I found that there were not enough bright, unsaturated, isolated stars in any one field to define a good empirical PSF. In a first crude attempt to deal with the problem, I summed median-averaged images of the M81 major-axis fields observed with chips WF1 and WF2, to produce a single image with twice the surface density of bright-ish stars. From this image I derived a single quadratically varying PSF which was then used in the reduction of the images obtained with all four detectors. In my second attempt, for each of the four WFC chips I summed the median-averaged image of the M81 major-axis field with the averaged-image of the so-called V30 field on a chip-by-chip basis. These summed images permitted me to estimate a separate quadratically varying model PSF for each of the four chips.

This approach has some advantages and some drawbacks. The first advantage is obvious: it doubles the number of stars that can go into the definition of each PSF. Since the image of each field is itself an average of individual exposures obtained at different epochs, the derived PSF is appropriate for some sort of average of the various focus settings and jitter histories of the individual exposures. It is not clear whether this is an advantage or disadvantage. (I can say, however, that in my experience, the changes in the PSF due to the breathing of the telescope and the tracking wander during the integration are not dominant sources of photometric error, when compared to other obvious problems of *HST* photometry.) The averaging of the various images of each field reduces the effect of readout and Poisson noise, but the unavoidable subpixel offsets of the various exposures does broaden the core of the resulting PSF. And at the same time as the summing of two different fields doubles the number of available PSF stars, it also doubles the degree of crowding they are subject to.

I have attempted to deal with the drawbacks while retaining the advantages of determining the model PSF from a multiplicity of images, by creating a new stand-alone module which I have named MULTIPSF. MULTIPSF is merely the old PSF routine from within DAOPHOT, with the addition of a third dimension. While the PSF routine derived a model point-spread function from stellar images recorded in a two-dimensional digital image, MULTIPSF derives a single model PSF from stellar images recorded in a stack of two-dimensional images. As was always the case with DAOPHOT, a provisional model PSF is used to fit the stars in each image, and then the neighbor stars are subtracted from each image leaving selected bright stars more or less isolated in their frames, for the derivation of an improved model PSF. If the spacecraft has been moved between exposures, or if entirely different fields have been imaged, the various frames will contain stars sampling the PSF in different portions of the focal plane. The spatial variation of the PSF will therefore be better constrained than it could have been by any one of the input images. Nevertheless, the stars used to define the PSF are not more crowded than before, because the model is derived from the individual images in the stack, rather than from their sum. Readout and Poisson noise are still beaten down by the inclusion of numerous stars in a multitude of frames in the model PSF, but the core radius of the derived PSF is not spuriously broadened, again because the individual exposures are employed, not their average.

The third and largest component of DAOPHOT III is a program which I call ALLFRAME. As described in Stetson (1994), ALLFRAME is the culmination of a sequence of increasingly sophisticated model-profile fitting packages consisting of the DAOPHOT routines PEAK and NSTAR, and the stand-alone programs ALLSTAR and ALLFRAME. PEAK performs fits of the model PSF to stars contained in a digital image one star at a time. NSTAR performs simultaneous profile fits to small groups (≤ 60 stars) of mutually blended star images. ALLSTAR extends the scope of NSTAR to the simultaneous derivation of photometric parameters for all stars contained in a given digital image. ALLFRAME carries this process to the logical limit: it performs simultaneous profile fits for all stars contained in *all* the images of a given patch of sky. In doing so it maintains a single, self-consistent list of program objects and of their positions, and solves for an independent magnitude for each star at each epoch. At present, ALLFRAME also determines an independent value for the sky brightness underlying each target, but it is conceivable that in the future the sky-brightness models for the different frames could be coupled in some way.

ALLFRAME offers several distinct advantages for stellar photometry:

- It maintains a consistent star list for all frames. When different exposures of a given field are reduced independently, and their photometric results are then combined *ex post facto*, it often happens that a blob of light is reduced as a single star in some frames and as a blended double in others. When the results are combined, the single star is identified with one component of the double in the other, and the result is a spurious variable star (if the frames are in the same filter) or a ludicrous color (if the frames are in different filters). With ALLFRAME, a blended double is always a blended double, always with the same position angle and separation.
- It uses all available data for all stars. PEAK, NSTAR, and ALLSTAR are

instructed to discard any detection which is significant at less than a three-sigma level. They must do this to prevent individual bright stars from being reduced as tight clusters of stars of dubious reality, and to prevent the diffuse sky brightness from being represented by millions of faint stellar profiles, one for each noise peak in the background. This has the obvious consequence that if a star appears as a 3.1σ detection in one frame and a 2.9σ detection in the next, it is retained in the former and discarded in the latter. The final average magnitude for the star is based on only half the data—predominantly the overestimated half. ALLFRAME retains every object that is a three-sigma detection in the combined data for all frames, and uses all data for that object from every frame in which it lies.

- ALLFRAME has fewer degrees of freedom. Rather than deriving an independent position for each star in each frame, it solves for a single position per star, and transforms that to the coordinate system of each frame. This results in higher photometric precision, especially in crowded fields.
- ALLFRAME does a better job of recognizing and then disregarding blemishes and cosmic rays. A defect in the wings of a stellar image would normally influence the determination of the star's centroid. By requiring the centroid of the blemished profile to be consistent with that star's position in all of the other frames in which it appears, the contamination is made more apparent and can be ignored.

III. Test 1: IC 4182

Besides the aforementioned work on the M81 data for the Key Project, I have analyzed two public-domain datasets from WFC. The first of these, the IC 4182 data from program 2547, *Calibration of Supernovae of Type I as Standard Candles* (A. Sandage, PI), is an example of a comparatively simple problem in relative *HST* photometry: the telescope pointing and roll angle were virtually identical (to within about four pixels, peak-to-peak) at all epochs. This means that systematic errors in the model PSF and in the flat-field corrections will cancel out (I used data that had been subjected only to the pipeline calibration procedures): apparent epoch-to-epoch variations in stellar magnitude should be real, or should reflect the response of the reduction software to the inherent precision of the data.

The data set consisted of 19.5 cosmic-ray splits in the F555W filter, and two C-R splits in F785LP, totalling 43 separate exposures with each of the four WFC chips. All exposures were in the range of 1900-2100 sec, and in the four subfields I determined magnitudes for a total of some 48,000 stellar objects. The total number of individual stellar models actually fit was 1,790,780. The root-mean-square repeatability of the individual stellar magnitudes for stars measured in at least 19 of the 39 F555W-band images (after the removal of a single zero-point constant for each frame intended to allow for the epoch-to-epoch sensitivity variations in the camera) is approximately 0.03–0.05 mag from the magnitude level at which saturation sets in to a level roughly three magnitudes fainter. Below that, the standard deviation increases as an exponential function of apparent magnitude, passing through 0.3 mag maybe five magnitudes below saturation.

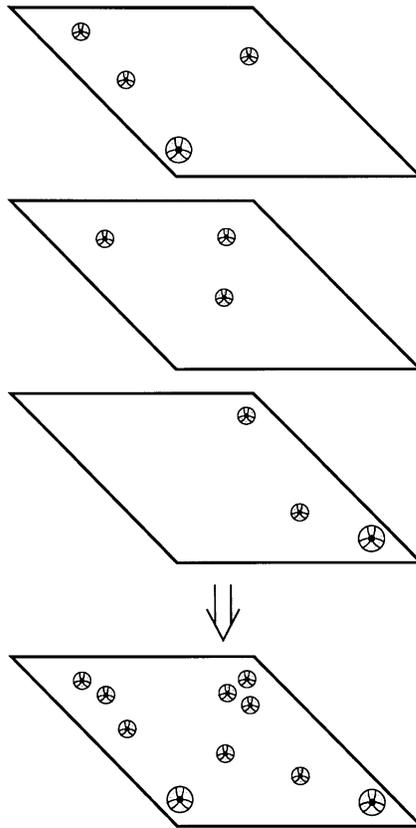


Figure 1: MEDPIC

However, it must be remembered that in 39 exposures of order 2000 seconds apiece, most of the stars in the field will have been involved in a cosmic-ray event at least once. The root-mean-square estimator of the photometric repeatability will be greatly inflated by the inclusion of the spurious magnitude estimates obtained from these contaminated images. A careful astronomer with a lot of time to spare would sort through the results and identify and eliminate those measurements that are obviously and hopelessly incorrect, and would derive the mean magnitudes and the estimates of precision from only those observations that seem to be valid. For my present purposes, I will estimate the photometric precision of my results using a statistic that is more robust against the presence of extreme outliers, and is more sensitive to the spread in the residuals of more typical observations: 1.2533 times the mean absolute residual. If the error distribution were truly Gaussian, this would have the same expectation value as the root-mean-square statistic, but since each the first power of each residual is used rather than the square, the rare extreme outliers have much less influence on the result.

When this more robust estimator of precision is used, the magnitude level at which a particular photometric repeatability is achieved is at least a magnitude fainter than when the root-mean-square residual is employed. The frame-to-frame repeatability is in the range 0.03–0.05 mag to about five magnitudes below the saturation level,

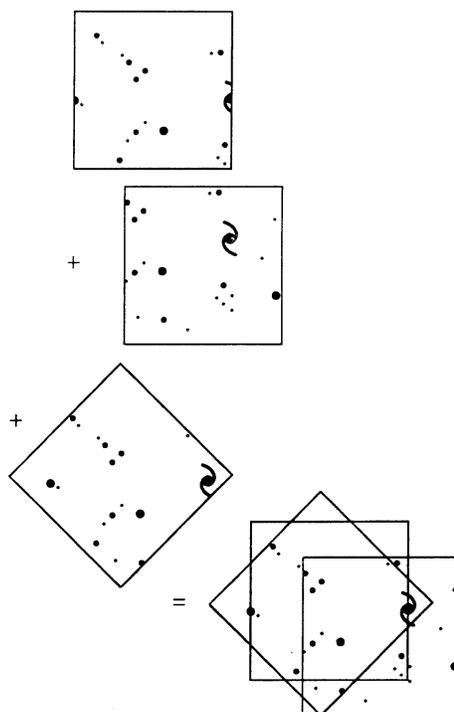


Figure 2: ALLFRAME

and the data are repeatable to ± 0.3 mag more than six magnitudes below saturation.

The use of ALLFRAME also permits the measurement of stars nearly two magnitudes fainter than the older one-frame-at-a-time packages. ALLSTAR retains only those stars that are three-sigma or greater detections in each input frame; as just mentioned, using the robust estimator this precision is achieved roughly 6.5 magnitudes below saturation in the F555W images of IC 4182. ALLFRAME is able to retain stars that have lower precisions in individual frames, provided that they are significant to of order 3σ in the aggregate of all frames. Adding the additional criterion that stars must have been recoverable in at least 19 of the 39 F555W exposures, I find that the star list extends to somewhat more than eight magnitudes below the saturation level. Stars at the detection limit are precise only to ± 0.8 mag per frame, but by the time 19–39 independent magnitude determinations have been averaged, the standard error of the mean magnitude at the detection limit is ~ 0.2 mag.

IV. Test 2: NGC 1850

A large number of frames of the Large Magellanic Cloud cluster NGC 1850 were obtained during the course of programs 3008, *WFPC SAT Observation: Young Cluster Photometry*, (J. Westphal, PI); 3367, *WFPC Astrometric Calibration, Plate Distortion, Pointing Assistance Calibration*, (R. Gilmozzi, PI); and 4161, *Mapping the Position Dependence of the WFPC PSF — Verification*, (R. Gilmozzi, PI). These observations provide a far more stringent test of photometric precision with *HST* and whatever

software package you may be using, because they were taken with a wide variety of center positions and roll angles — any given star can occur at different places on different chips at the various epochs. This will test the quality of both the flat-fielding (again, I reduced pipeline-calibrated data) and the model PSF. Unfortunately, at the present time I have no way to evaluate the relative importance of these two sources of systematic error, but future analyses based on new flat fields and utilizing comparisons with ground-based results may eventually resolve the ambiguities.

The data set for NGC 1850 includes 17 individual exposures in the F555W filter, with exposure times ranging from 10 sec to 1100 sec; 10 exposures in F785LP, again from 10 sec to 1100 sec; and two 1100 second exposures in F439W. Of course, each exposure consists of actual data frames from each of the four WFC detectors. Final photometric data were extracted for a total of 16,441 stars in the NGC 1850 field appearing in at least two of the 116 individual frames; 353,598 individual model profiles were actually fit.

The robust error estimates for the NGC 1850 results were considerably poorer than for IC 4182: at no magnitude level was the frame-to-frame repeatability consistently better than 0.10 mag. Since the exposure times ranged over more than two orders of magnitude, it is not possible to reference the precision to any particular saturation magnitude. Nevertheless, in the global combination of results the limiting precision of 0.10 mag was obtained over about a four-magnitude range. Since for each star this figure includes results from 10 second frames where everything is faint to 1100 second frames where many stars are saturated, a better figure of photometric merit might be derived from frames with a single exposure time. The data set for NGC 1850 includes seven 300 second exposures, each one taken with a different pointing of the telescope. If I then take the photometric results for each star that appeared in at least 4 (so that the sigmas will be reasonably well-defined) and not more than 7 (to eliminate stars split between frames by the edges of the pyramid mirror) of the 300 second exposures, and compute their photometric scatter with respect to the mean of all available data, the precision is still no better than 0.10 mag at any magnitude level. However, this level of precision is achieved over about a five-magnitude range, from about one magnitude above saturation to four magnitudes fainter than saturation. Fainter than this, the photometric error grows exponentially with magnitude, passing through ± 0.3 mag some six magnitudes below the saturation level. Again, at this point I have no way of knowing whether the errors are dominated by the flat fields or by the model PSF, but I suspect the former.

Acknowledgments

I am grateful to the goofy bunch of guys at the Canadian Astronomy Data Centre, located at the Dominion Astrophysical Observatory in Victoria, for their assistance in acquiring some of these public-domain data; and to Jeremy R. Mould, lately of Caltech, for his moral and financial support in carrying out this work.

References

- Freedman, W. L., Hughes, S. M., Madore, B. F., Mould, J. R., Lee, M. G., Stetson, P. B., Kennicutt, R. C., Turner, A., Ferrarese, L., Ford, H., Graham, J. A., Hill, R., Hoessel, J. G., Huchra, J., & Illingworth, G. D. 1994, to appear in ApJ
- Hughes, S. M. G., Stetson, P. B., Turner, A., Kennicutt, R. C., Jr., Hill, R., Lee, M. G., Freedman, W. L., Mould, J. R., Madore, B. F., Ferrarese, L., Ford, H. C., Graham, J. A., Hoessel, J. G., & Illingworth, G. D. 1994, submitted to ApJ
- Stetson, P. B. 1987, PASP, 99, 191
- Stetson, P. B. 1991, in Third ESO/ST-ECF Data Analysis Workshop, eds. P. J. Grosbøl & R. H. Warmels (Garching: ESO), 187
- Stetson, P. B. 1992, in Astronomical Data Analysis Software and Systems I, eds. D. M. Worrall, C. Biemesderfer, & J. Barnes, PASP Conf. Ser. 25, 297
- Stetson, P. B. 1994, submitted to PASP
- Stetson, P. B., Davis, L. E., & Crabtree, D. R. 1990, in CCDs in Astronomy, ed. G. H. Jacoby, PASP Conf. Ser. 8, 289