A METHOD FOR OPTIMAL IMAGE SUBTRACTION

C. ALARD[1,2] AND ROBERT H. LUPTON[3]

ABSTRACT

We present a new method designed for optimal subtraction of two images with different seeing. Using image subtraction appears to be essential for full analysis of microlensing survey images; however, a perfect subtraction of two images is not easy, as it requires the derivation of an extremely accurate convolution kernel. Some empirical attempts to find the kernel have used a Fourier transform of bright stars, but solving the statistical problem of finding the best kernel solution has never really been tackled. We demonstrate that it is possible to derive an optimal kernel solution from a simple least-squares analysis using all the pixels of both images, and we also show that it is possible to fit the differential background variation at the same time. We show that point-spread function (PSF) variations can be easily handled by the method. To demonstrate the practical efficiency of the method, we analyzed some images from a Galactic Bulge field monitored by the OGLE II project. We find that the residuals in the subtracted images are very close to the photon noise expectations. We also present some light curves of variable stars and show that despite high crowding levels, we get an error distribution close to that expected from photon noise alone. We thus demonstrate that nearly optimal differential photometry can be achieved even in very crowded fields. We suggest that this algorithm might be particularly important for microlensing surveys, where the photometric accuracy and completeness levels could be very significantly improved by using this method.

Subject headings: methods: data analysis — methods: statistical — techniques: image processing

1. INTRODUCTION

The search for microlensing events toward the LMC (with MACHO [Alcock et al. 1993] and EROS [Aubourg et al. 1993]), the Galactic Bulge (with OGLE [Udalski et al. 1994], MACHO and DUO [Alard & Guibert 1997]), or the M31 galaxy (with AGAPE), has provided us with an impressive database of images of densely crowded fields. The target fields have been monitored for several seasons, providing us with time series containing hundreds of images. Light curves for millions of stars can then be easily obtained with one of the widely used profile-fitting codes, such as DoPHOT (Schechter & Mateo, 1993). The search for variable objects among these huge light curve databases has proved very fruitful for microlensing (with MACHO, OGLE, DUO, and EROS), and also for variable stars (with MACHO, OGLE, DUO, and EROS). However, we would like to emphasize that photometry and detection of variable (including moving) objects should be based on the difference between frames, whereas photometric codes like DoPHOT are designed to perform profile-fitting photometry of stars detected on a reference frame. If a variable object appears but was not seen on the reference, it will not be detected, leading to a serious loss of efficiency for microlensing. The completeness of the variable star catalogue will also be seriously affected. Another concern is photometric accuracy. With multiprofile fitting techniques, the absolute photometry of a given (crowded) star requires perfect point-spread function (PSF) estimation and careful modeling of all other close components, and also a correct estimate of the background value around each star. For the particular problem of finding light curves of variable objects, it is more efficient to estimate only that part of the star's brightness that varies from image to image; this is exactly the problem that image subtraction is designed to solve.

The first attempt at image subtraction was made by Tomaney & Crotts (1996; hereafter TC) for data taken toward the M31 galaxy (Crotts & Tomaney 1996). To make a perfect subtraction of two images, one must match the frames to exactly the same seeing. TC proposed degrading a good seeing image to match a reference frame with bad seeing. The quality achieved in the subtracted image is very dependent on the quality of the kernel determination, and finding the proper kernel is a very delicate operation. TC proposed deriving the kernel by simply taking the ratio of the Fourier transform of a bright star on each image. However, the high frequencies are dominated by noise, and they were forced to use a Gaussian extrapolation to determine the wings (Phillips & Davies 1995). This method provides no guarantee of producing the highest attainable quality of the subtracted image. Even apart from the non-Gaussian wings of the true kernel, and given the limited number of bright, uncrowded stars with sufficient signal-to-noise ratio, this method is not optimal in the sense that it does not use all the information available; in fact, every star, even if extremely crowded, contains information about the kernel, and to get an optimal solution we must use all of that information. In addition, their method has difficulty with rapid, complicated PSF variations, and does not intrinsically handle background subtraction.

An interesting step toward the optimal solution was performed by Kochanski et al. (1996). Their method finds the kernel solution that will minimize the discrepancy between the two images used to subtract. However, this method is a nonlinear least-squares fitting process and has a prohibitive computing time, even for uncrowded fields and when only a fraction of the pixels are fitted. We wish to develop a

[1] Departement d'Astronomie Stellaire et Galactique, Observatoire de Paris-Meudon, 61 Avenue de l'observatoire, F-75014 Paris, France.

[2] Institut d'Astrophysique de Paris, 98 bis Boulevard Arago, F-75014 Paris, France.

[3] Department of Astrophysical Science, Peyton Hall, Princeton University, Princeton, NJ 08544.

method that will also solve the least-squares problem. This method must perform best for the very crowded fields of the microlensing surveys. The challenge will be to achieve a reasonable computing time for processing the huge amount of data from the microlensing experiments. It is essential that the computing time be reduced by at least a factor of 100.

2. THE METHOD

2.1. *Preliminaries*

Before looking for the optimal subtraction, we need to perform some basic operations to register the frames to a reference frame. Usually the frames have slightly different centers and orientation (and possibly scale), and we need to perform an astrometric transform to match the coordinates of the reference frame. We determine this transform by fitting a two-dimensional polynomial using 500 stars on the reference frame, and the same number on the other frame. Using this transform, we then resample the frame on the grid defined by the reference frame. This resampling is performed by interpolation using bicubic splines, which gives excellent accuracy. All the frames are then on the same coordinate system, and we can proceed to match the seeing.

2.2. *The Reference Frame*

Here we emphasize that we choose to take the *best* seeing frame as the reference. We do not wish to degrade the frame to the worst seeing frame, as this will clearly lower the signal-to-noise ratio. Later we will match the seeing in our frames by convolving the reference to the seeing of each other frame. Matching the image quality to that in good seeing frames is more difficult, but we are looking for an optimal result.

2.3. *Seeing Alignment to the Reference*

We now arrive at the fundamental problem of matching the seeing of two frames with different PSFs. We do not want to make any assumption concerning the PSF on the frame, and we plan to use all the pixels. The important point is that most of the stars on a given frame do not have large amplitude variations, varying by at most 1% or 2%. This allows us to say that most of the pixels on any two frames of the same field should be very similar if the seeing were the same. Consequently, one possibility is to try to find the kernel by finding the least-squares solution of the equation

$$\text{Ref}(x, y) \otimes \text{Kernel}(u, v) = \text{I}(x, y), \tag{1}$$

where Ref is the reference image and I is the image to be aligned. The symbol $\otimes$ denotes convolution.

In principle, solving this equation is a nonlinear problem. An attempt to solve this nonlinear problem was made by Kochanski et al. (1996) in order to look for variabilities of active galactic nuclei. However, the computing time of such methods is prohibitive, and they cannot be applied to large data sets, such as those of the microlensing surveys. We need to find a more tractable solution to this problem.

An important consideration is that if we decompose our kernel using some basis of functions, the problem becomes a standard linear least-squares problem. If we use the kernel decomposition of the form

$$\text{Kernel}(u, v) = \sum_i a_i B_i(u, v), \tag{2}$$

solving the least-squares problem gives the following matrix equation for the a_i coefficients:

$$Ma = V,$$

where

$$M_{ij} = \int C_i(x, y) \frac{C_j(x, y)}{\sigma(x, y)^2} \, dx \, dy,$$

$$V_i = \int \text{Ref}(x, y) \frac{C_i(x, y)}{\sigma(x, y)^2} \, dx \, dy,$$

$$C_i(x, y) = \text{I}(x, y) \otimes B_i(x, y).$$

In choosing to solve the problem by least-squares, we have implicitly approximated the images Poisson statistics by Gaussian distributions with variance $\sigma(x, y)^2$, such that

$$\sigma(x, y) = k\sqrt{\text{I}(x, y)}.$$

We set the constant k by taking into account the gain of the detector (the ratio of photons detected to the analog-to-digital converter unit [ADU]). We also assume the noise in the reference image to be negligible. Note that the matrix M is just the scalar product of the set of vectors C_i, and the vector V is the scalar product of the C_i vectors with I. All we have to do now is to look for a suitable basis of functions to model the kernel. The functions must have finite sums, and must drop rapidly beyond a given distance (the size of an isolated star's image). To solve this problem, we start with a set of Gaussian functions, which we modify by multiplying with a polynomial. These basis functions allow us to model the kernel, even if its shape is extremely complicated. We adopt the following decomposition:

$$\text{Kernel}(u, v) = \sum_n \sum_{d_n^x} \sum_{d_n^y} a_k \, e^{-(u^2 + v^2)/2\sigma_n^2} u^{d_n^x} v^{d_n^y},$$

where $0 < d_n^x \le D_n$, $0 < d_n^y + d_n^x \le D_n$, and D_n is the degree of the polynomial corresponding to the nth Gaussian component. There are a total of $(D_n + 1)(D_n + 2)/2$ terms for each value of n. The value of k is implicit in the values of the other indexes.

In the notation of equation (2),

$$B(u, v) \equiv e^{-(u^2 + v^2)/2\sigma_n^2} u^{d_n^x} v^{d_n^y}.$$

In practice, it seems that three Gaussian components with associated polynomial degrees in the range of 2 to 6 can give subtracted images with residuals comparable to $2^{1/2} \times$ photon noise.

2.4. *Differential Background Subtraction*

Another important issue is that the differential background variation between the frames can be fitted simultaneously with the kernel. In equation (1) we did not consider any background variations between the two frames; to do so, let us modify equation (1) in the following way:

$$\text{Ref}(x, y) \otimes \text{Kernel}(x, y) = \text{I}(x, y) + \text{bg}(x, y). \tag{3}$$

We use the following polynomial expression for bg (x, y):

$$\text{bg}(x, y) = \sum_i \sum_j a_k x^i y^j,$$

where $0 < i \le D_{bg}$, $0 < i + j \le D_{bg}$, and D_{bg} is the degree of the polynomial used to model the differential background variation. The least-squares solution of equation (3) will lead to a matrix equation similar to the previous one, except

that we must increase the number of C_i vectors; our definitions of the matrix M and vector V relative to C_i remain the same as in § 2.3. We then have

$$C_i(x, y) = \begin{cases} x^j y^k, & i = 0 \ldots n_{\text{bg}} - 1 \\ I(x, y) \otimes B_i(x, y), & i = n_{\text{bg}} \ldots n_{\text{bg}} + n, \end{cases}$$

where $n_{\text{bg}} = (D_{\text{bg}} + 1)(D_{\text{bg}} + 2)/2 - 1$ and $n = \sum_j (D_j + 1)(D_j + 2)/2$; note that for $i \geqq n_{\text{bg}}$, the C_i are identical to our previous results.

3. TAKING INTO ACCOUNT THE PSF VARIATIONS

There are two ways to handle the problem of PSF variations. First, most of the time the field is so dense that a transformation kernel can be determined in small areas, small enough that we can ignore the PSF variation. Of course, this would not apply to less dense fields, such as those in a supernova search. This is the great advantage of a method that does not require any bright isolated stars to determine the kernel, but can be used on any portion of an image, provided that the signal-to-noise ratio is large enough to determine the kernel. Indeed, the more crowded the field, the easier it is to model variations of the PSF. A second possibility is to make an analytical model of the kernel variations. We take the kernel model

$$\text{Kernel}(x, y, u, v) = \sum_n \sum_{d_n x} \sum_{d_n y} \sum_{\delta x} \sum_{\delta y} a_k x^{\delta x} y^{\delta y}$$
$$\times e^{-(u^2 + v^2)/2\sigma_n^2} u^{d_n x} v^{d_n y},$$

where $0 < \delta_x < D_k$, $0 < \delta_y + \delta_x \leq D_k$, and D_k is the degree of the polynomial transform that we use to fit the kernel variations. Provided that the kernel variations with x and y are small enough compared to the u, v variations, we can easily calculate new expressions for C_i:

$$C_i(x, y) = \begin{cases} x^j y^k, & i = 0 \ldots n_{\text{bg}} - 1 \\ I(x, y) \otimes B_k(u, v) x^{\delta x} y^{\delta y}, & i = n_{\text{bg}} \ldots n_{\text{bg}} + n, \end{cases}$$

where the values of δ_x, δ_y, and k are implicit in the index i, and now $n = \sum_j (D_j + 1)(D_j + 2)/2 \times (D_k + 1)(D_k + 2)/2$. Unfortunately, these equations do not guarantee the conservation of flux. Consequently, we need to add the condition that the sum of the kernel must be constant. To simplify the equation, we also normalize the B_i functions so that each sums to one. We can then rewrite the kernel decomposition as

$$\text{Kernel}(x, y, u, v) = \sum_{k=0}^{n-1} a_k x^{\delta x} y^{\delta y}$$
$$\times [B_i(u, v) - B_n(u, v)] + \text{norm} \times B_n$$

We can calculate the norm (the sum of the kernel) by making a constant PSF fit in several small areas. The different values will then be averaged to get the constant norm.

The solution of the system for the coefficients a_n is very similar to the previous case of a constant PSF. We do not bother to give all the details here.

4. APPLICATION OF THE METHOD TO OGLE DATA

The OGLE team has kindly provided us with a stack of images of a field situated $2°$ from the Galactic Center, in order to experiment with our method. In these particular images the optimal kernel has a complicated shape, which would probably be very difficult to compute reliably with a simple Fourier division; we consider this field an excellent test of our method. The data were taken in drift-scan mode (TDI), so the form of the PSF can vary rapidly with row number on the CCD. We extracted a small (500×1000) subframe from the 2048×8192 original images. One of the images has quite outstanding seeing, and we took it as a reference. All frames were resampled to the reference grid using the method described above. To model the kernel, we took three Gaussian components with associated polynomials. We took $\sigma = 1$ pixels for the first Gaussian, and $\sigma = 3$ and $\sigma = 9$ for the two others. The degree of the associated polynomials were 6, 4, and 2, respectively. We divided the subframe into 128×256 pixel regions. We applied our method to each of these regions, providing us with one subtracted image per region. We reconstructed the subtracted image of our whole subframe by mosaicing the subtracted images obtained for each region. In this set of 86 images, the seeing varies from $0\overset{''}{.}7$ to $2\overset{''}{.}5$, and some of the frames have elongated stellar images. We started by making an initial residual image using all unsaturated pixels. We then made a 3 σ rejection of the pixel list, to get rid of the variables. We usually used four iterations of the method, to be completely unbiased by large-amplitude variables. We found that for all images, the final residual calculated from the subtracted image was very close to that expected from Poisson statistics. To illustrate this result, we plot in Figure 1 the initial images and the subtracted image for a small field containing a variable star at its center.

The stellar images are sharply peaked on the reference, while they look quite fuzzy and asymmetric on the other image. This is confirmed by the shape of the best convolution kernel, which looks elongated and has a complicated shape. This example clearly illustrate the ability of our method to deal with any kernel shape. We can imagine that in this case, any Gaussian approximation of the kernel itself or of its Fourier transform would not be satisfactory. For illustrative purposes, we also normalized the subtracted image by the sum of the photon noise expected from the two images (see Fig. 2). Once this normalization is applied, we see that the larger deviations visible at the location of the bright stars disappear, suggesting that the subtraction errors correspond to Poisson noise. This is confirmed by calculating the reduced χ^2, for which we find $\chi^2/\nu = 1.05$ (before doing this calculation, we removed a small area around the variable star at the center of the image). We also plot the histogram of the normalized deviations in Figure 2. This histogram is very close to a Gaussian with zero mean and unit variance [i.e., $N(0, 1)$]. We observe deviations significantly larger than the Poisson expectations only for very bright stars (about 5 to 10 times brighter than the brightest stars in the small field we present). We believe that these residuals are due to seeing variations (see § 5 for more details); the number of such bright stars in an image is very small.

To spot the variables stars, we created a "deviation image" by coadding the square of the subtracted images. We normalize the deviation image by normalizing by the pixels' standard deviations. We found many variables at very significant levels. Most of these seem to be bright giants with small amplitudes. Some of the variables appeared to be periodic; we found a few RR Lyraes and some eclipsing variables. We computed the flux variations for these stars by making simple aperture photometry. In

Fig. 1.—Example of subtracted image. The two bottom figures of the panel are the original images. On the right is the reference image, and on the left is the image to be fitted by kernel convolution. The two upper figures show the best kernel solution on the right, and the subtracted image on the left. Note the complicated shape of the kernel.

Figures 3 and 4 we give an illustration of the result we have obtained.

5. COMPUTING TIME

One might think that a method that fits all the pixels in an image (even if the fit is linear) would be much more time consuming than conventional methods. But the actual cost of the calculations is much lighter than might appear at first glance. Most of the computing time is taken by the calculation of the matrix we define in § 2.3, an N^2 process (where N is the number of basis function we used). The rest of the calculation is an N process. However, the matrix could be calculated once and then used to fit the kernel solution for all images. One problem with this approach is that we reject different pixels on each frame (due to new saturated pixels or variable stars), so consequently the matrix elements change. But in practice, we find that we reject no more than 1% percent of the total number of pixels, so that all that we

have to do is to calculate the matrix elements for the rejected pixels and subtract them from the original values. This process costs very little CPU, and once the original matrix has been built, the kernel solution can be fitted very quickly even if we use several clipping passes. The rest of the operation requires about the same computing time. By applying this method, we can process a 1024 × 1024 frame in about 1 minute with a 200 MHz PC; this could certainly be improved further by using better numerical algorithms for the solution of the linear system.

6. SOURCES OF NOISE IN THE RESIDUAL IMAGE

As discussed above, the variance of the residual image is approximately equal to the sum of the variances of the input images. If we created a reference image by coadding a large number of images with good seeing, we could remove the contribution of the noise in the reference; we would, of course, have to be careful about variability between the

FIG. 2.—Noise in the subtracted image. The image on the left is the subtracted image normalized by the Poisson deviations of both the reference and the image for the small field presented in the previous figure. On the right we show the histogram of the pixels in this image. We superimposed on this histogram a Gaussian of variance 1 (*dashed line*). Note that the deviations due to the bright stars are no longer visible. The variable star at center clearly stands out at very significant level. The dark pixel in the image is a cosmic ray.

different reference frames. Upon inspection of the residual images, however, some showed significantly larger residuals than expected from Poisson statistics near the position of bright, but nonsaturated, stars (representing less than 1% of the stars visible on the frame). These showed the characteristic signature of centering errors, with equal positive and negative residuals even for stars that show no evidence of variability (i.e., the sum of all residuals within a few arcseconds is zero). We believe that this effect is produced by the turbulent atmosphere modulating our kernel on the scale of our subregions.

Shao & Colavita (1992) quote the variance in the angle

between two stars separated by θ as

$$\sigma_\delta^2 \approx 5.25(\theta/\mathrm{rad})^{2/3}(t/\mathrm{s})^{-1}\int C_n^2(h)h^{2/3}V^{-1}(h)dh$$

for the regime in which we are interested (their eq. [2]). In this equation, C_n is the structure function of the refractive index, h is the height above the telescope, V is the wind speed, and t is the integration time. They evaluate the integral using data from Roddier et al. (1990) for a night on Mauna Kea with $\approx 0''.5$ seeing to give

$$\sigma_\delta \approx 1.1(\theta/\mathrm{rad})^{1/3}(t/\mathrm{s})^{-1/2} \; \mathrm{arcsec} \; .$$


FIG. 3.—Examples of bright giant light curves. The x-axes are days, the y-axes are percentage variation (with respect to reference frame). Errors bars are derived from the Poisson deviations associated with each image; we do not include the deviations associated with the reference image.
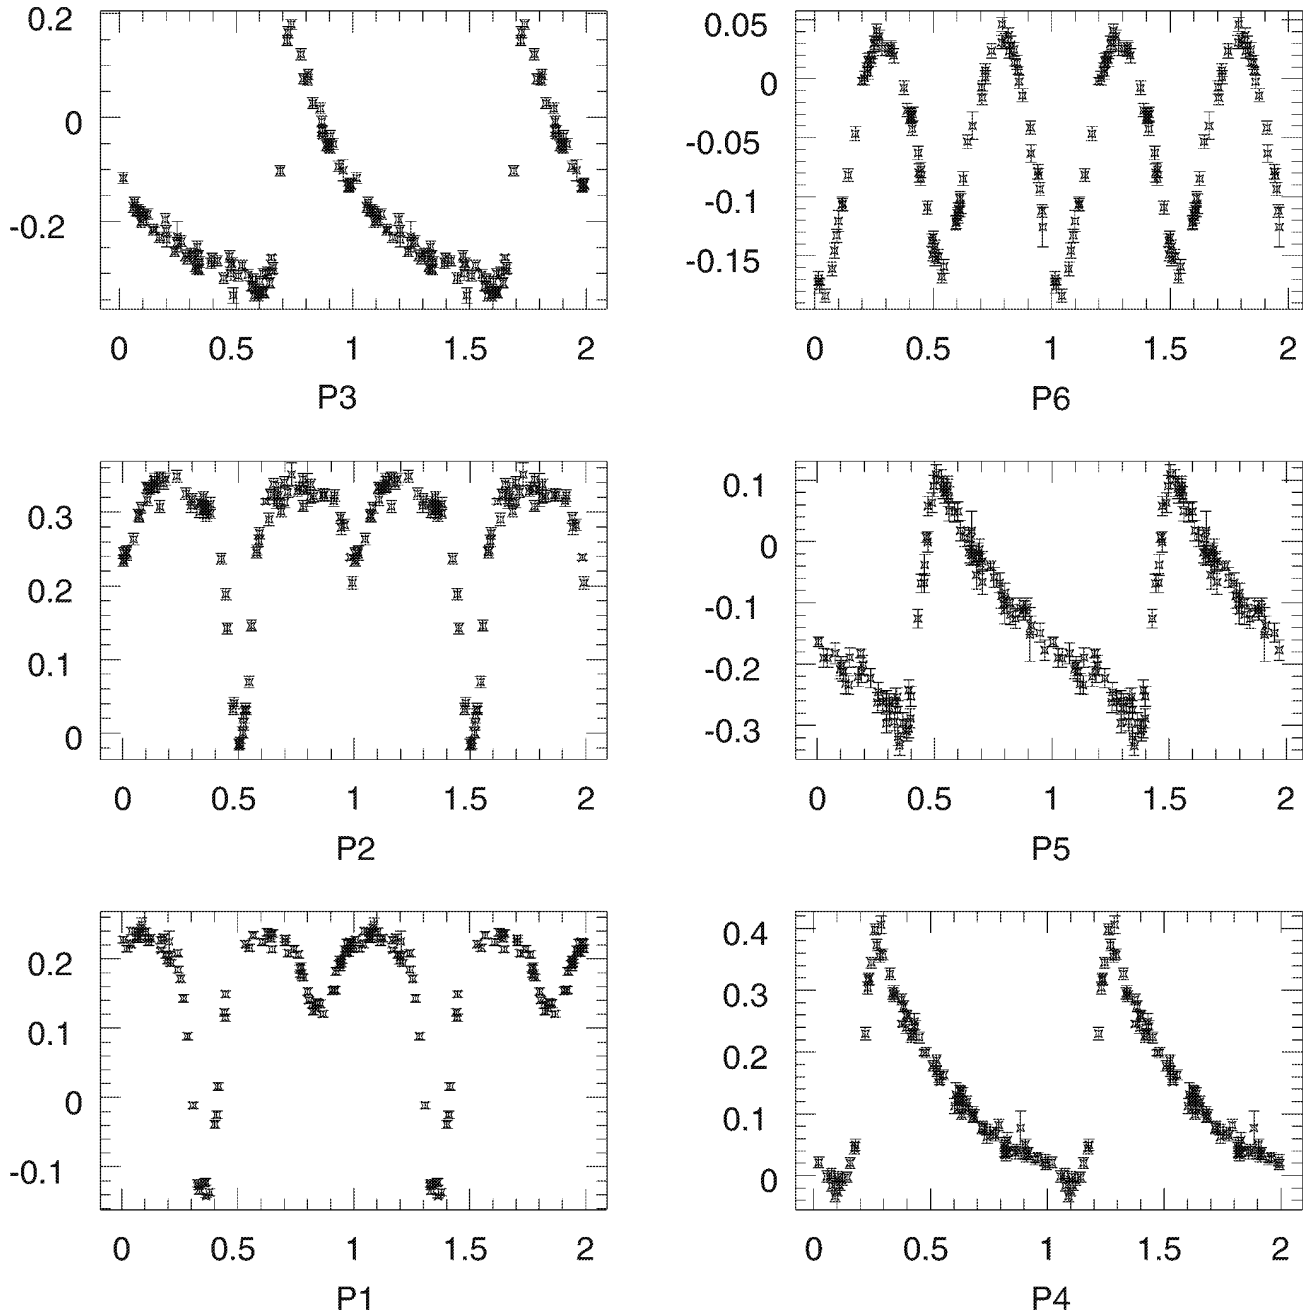
FIG. 4.—Periodic variables. The x-axes represent the phase, the y-axes are percentage variation (with respect to reference frame).

If we assume that the integral over the atmosphere scales with seeing in a way similar to the integral

$$\int C_n^2(h)dh \, ,$$

which enters into the definition of the Fried parameter, r_0, we may expect that this result will scale as $(r_0\lambda^{-6/5})^{-5/6}$ (a result that is independent of λ as a result of the wavelength dependence of r_0). In 1″ seeing, therefore, we may expect

$$\sigma_\delta \approx 2(\theta/\text{rad})^{1/3}(t/s)^{-1/2} \text{ arcsec} \, .$$

On the typical scale of our 128 × 256 regions, and for 128 s exposures, this corresponds to an rms image motion of about 0″.011. If we model the PSF as a Gaussian with width parameter α ($\alpha \approx 0.424$ for 1″ FWHM images), this would produce a maximum residual of σ_δ/α exp $(-\frac{1}{2})$, or 1.6%. This is on the same order as the residuals that we see in our frames.

7. HARMONIC FITTING TO THE PERIODIC VARIABLES

We expect the periodic variables' light curves to be well approximated by truncated Fourier series. We calculate the period using the Renson method (Renson 1978), and we fit Fourier series with different numbers of harmonics. The errors are calculated from the photon noise in each image. We do not include the noise associated with the reference image because it only produces an error in the total magnitude, and to first order does not affect the variable part of the object's flux. We estimate at each time the χ^2 per degree of freedom (χ_d^2), and we look for the best χ^2 with the

TABLE 1

Harmonics Fitting to the Periodic Variables

Variable	χ_d^2	Mean Residual (%)	Δ_{mag}
P1	2.01	1.0	−1.117
P2	1.43	1.1	−0.4402
P3	1.55	1.6	0
P4	1.46	1.2	0
P5	1.17	1.3	0
P6	1.1	0.6	−0.4348

Note—In the third column we give the value of the mean residual to the fit. It is useful to compare this residual to possible flat-fielding errors. We also give an estimate of the star magnitude difference to the RR Lyrae Δ_{mag}. We assume that the RR Lyrae all have the same mean magnitude. A crude estimate for the RR Lyrae mean magnitude in this field is $I \simeq 17$.

TABLE 2

Harmonics Fitting to the Light Curves Obtained with the New Reference

Variable	χ_d^2	Mean Residual (%)
P1	2.0	1.0
P2	1.16	1.0
P3	1.27	1.45
P4	1.45	1.2
P5	1.15	1.3
P6	1.03	0.6

Note—In the third column we give the value of the mean residual to the fit. It is useful to compare this residual to possible flat-fielding errors. We also give an estimate of the star magnitude difference to the RR Lyrae Δ_{mag}. We assume that the RR Lyrae all have the same mean magnitude. A crude estimate for the RR Lyrae mean magnitude in this field is $I \simeq 17$.

minimum number of harmonics. The results are given in Table 1, where we see that the resulting value of χ_d is close to unity for most variables. Except for variable P1, our mean error is at most only 25% larger than the Poisson expectation (i.e., $\chi_d^2 < 1.56$); of course, this χ_d^2 excess is significant. In the case of the variable P1, the χ_d^2 is very inconsistent with the Poisson expectation. This variable has about the same brightness as P6. We checked the quality of the subtracted images, but could not identify any defects. The quality of the image subtraction is as good for P1 as for P6, and they have about the same brightness; so what is wrong? Considering that the mean error is fairly small (about 1%), we might suspect some residual error due to flat-fielding. However, we get a mean residual of only 0.6% for P6 and $\chi_d^2 = 1.1$, showing that the flat-fielding errors are much smaller than 0.6%. This is not surprising, because these images were taken in drift-scan mode, and consequently we average the sensitivity of many pixels. We conclude that there must be some intrinsic reason for P1's bad χ_d^2. It is possible that variables do not repeat perfectly from cycle to cycle. This kind of variable star is well known to have spots that are likely to induce variability at the subpercent level. It is also possible that the RR Lyraes do not repeat perfectly; they are known to show the Blashko effect, and we can explain some of the χ_d^2 as an effect of cycle-to-cycle variations. Although estimating the χ_d^2 of periodic variable stars is not an absolute test, we conclude that on average we are only about 20% above the Poisson error, and consequently there is not much to be gained from improving our method. However, we must note that the errors attributable to the reference frame are the same for the integrated flux of a star on each image only at the first order of approximation. By convolving the reference each time to fit the seeing variations, we slightly change the noise distribution around the star. Especially for the case in which a bright star is close to our object, convolving with the kernel might spread some noise into our photometric aper-

ture. This effect will be negligible for good seeing frames, but noticeable when the seeing is bad. An obvious solution is to construct a reference with as good a signal-to-noise ratio as possible by stacking the best seeing images (see § 8).

Another approach with potential for improving the signal-to-noise ratio would be to use a matched filter to measure the variability of our stars. Unfortunately, simply applying the usual PSF filter leads to problems with aperture corrections, and we do not investigate this approach here.

8. IMPROVING THE REFERENCE FRAME

We averaged the 20 best seeing images to build a reference frame with an excellent signal-to-noise ratio. The resulting seeing is of course not as good as it was in our previous reference, which was the best image. But the seeing variations are much reduced, as is the noise amplitude. All the images were reprocessed using this new reference. We found that all the subtracted frames were improved. Even for the good seeing frame, where the seeing quality of the reference is critical, we found some improvements. The light curves of the variables stars were also improved; we give the result of harmonic fitting in Table 2.

REFERENCES

Alard, C., & Guibert, J. 1997, A&A, 326, 1
Alcock, C., et al. 1993, Nature, 365, 621
Aubourg, E., et al. 1993, Nature, 365, 623
Crotts, A., & Tomaney, A. 1996, ApJ, 87, L473
Kochanski, J., Greg, P., Tyson, A., & Fischer, P. F. 1996, AJ, 111, 1444
Phillips, A. C., & Davis, L. E. 1995, in ASP Conf. Ser. 77, Astronomical Data Analysis and Systems IV, ed. R. Shaw Payne & J. J. E. Hayes (San Francisco: ASP), 297

Renson, P. 1978, A&A, 63, 125
Roddier, F., Cowie, L., Graves, J. E., & Songaila, A. 1990, Proc. SPIE, 1236, 485
Schechter, P., & Mateo, M. 1993, PASP, 105, 1342
Shao, M., & Colavita, M. M. 1992, A&A, 262, 353
Tomaney, A., & Crotts, A. 1996, AJ, 112, 2872 (TC)
Udalski, A., et al. 1994, Acta Astron., 44, 165