APPLICATION OF CUBIC SPLINES TO THE SPECTRAL ANALYSIS OF UNEQUALLY SPACED DATA

C. AKERLOF,[1,2] C. ALCOCK,[3] R. ALLSMAN,[3] T. AXELROD,[3] D. P. BENNETT,[1,3] K. H. COOK,[1,3] K. FREEMAN,[4]
K. GRIEST,[1,5] S. MARSHALL,[1,6] H.-S. PARK,[3] S. PERLMUTTER,[1] B. PETERSON,[4] P. QUINN,[4] J. REIMANN,[7]
A. RODGERS,[4] C. W. STUBBS,[1,6] AND W. SUTHERLAND[1]

ABSTRACT

In the absence of a priori information, nonparametric statistical techniques are often useful in exploring the structure of data. A least-squares fitting program, based on cubic B-splines has been developed to analyze the periodicity of variable star light curves. This technique takes advantage of the limited domain within which a particular B-spline is nonzero to substantially reduce the number of calculations needed to generate the regression matrix. By using simple approximations adapted to modern computer workstations, the computational speed is competitive with most other common methods that have been described in the literature. Since the number of arithmetic operations increases as N^2, where N is the number of data points, this method cannot compete with the FFT modification of the Lomb-Scargle algorithm. However, for data sets with $N < 10^4$, it should be quite useful. Examples are shown, taken from the MACHO experiment.

Subject headings: methods: data analysis — methods: numerical — methods: statistical

1. INTRODUCTION

The process of identifying the periodicity of a variable star boils down to finding a particular frequency at which the period-folded data can be adequately described by a curve with many fewer parameters than the number of original data points. If the shape of the curve is known a priori, the problem resolves to finding the frequency which leads to the best least-squares fit. Unfortunately, the folded light curve behavior is almost never known in advance so a large number of methods have been developed which circumvent this problem in various ways. An excellent review of various techniques can be found in a paper by Fullerton (1986).

One approach, the Lomb-Scargle method (Lomb 1976; Scargle 1982; Horne & Baliunas 1986), computes the Fourier power over an ensemble of frequencies. The significant periodicities correspond to the frequencies where the power is maximized. The most attractive feature of this method is that an FFT-like algorithm (Press & Teukolsky 1988; Press & Rybicki 1989) can be used so that the number of computations scales as $N \log N$ rather than N^2, where N is the number of data points. Some of the disadvantages include the need for the light-curve data to be approximately homoscedastic (i.e., uniform variance of all the data points) and the loss of statistical power for nonsinusoidal behavior.

A totally different idea is embodied in the so-called string length method. As the name implies, the statistical measure of the significance of a particular frequency is the sum of the successive absolute differences of adjacent points of the period-folded light curve. Shorter lengths correspond to greater ordering of points and thus indicate the significant oscillation frequencies. A variation on this technique is to compare the interpolation of adjacent points with each actual light-curve value (Friedman 1984; McDonald 1986). The sum of the absolute values of these deviations will be a similar string length statistic. A nice feature of this approach is that the method is unaffected by the light-curve shape, but the requirement that the data points be ordered at each test frequency leads to inherently long computation times.

In the method to be described below, an attempt has been made to retain the computational speed of the Lomb-Scargle procedure while taking advantage of the greater functional degrees of freedom provided by string length algorithms. The basic procedure is to least-squares-fit a sum of cubic spline functions to period-folded light curves. The significant frequencies are those which produce minimum values of the χ^2 statistic. This technique asymptotically requires N^2 arithmetic operations, but some reasonable approximations can keep the computation time modest as long as $N < 10^4$. This technique is particularly suited for modern workstations with several megabytes of memory since efficient evaluation of the χ^2 statistic requires lookup tables which grow linearly with the number of data points. Even if this spline-fitting technique is not used for the initial task of finding the fundamental frequency, the method provides a well-behaved smooth curve which can be used for further analysis. For example, the simple polynomial nature of cubic splines permits the easy evaluation of curve extrema or spectral analysis with Fourier series.

In summary, the periodic cubic spline–fitting procedure described in this paper is particularly attractive for large survey programs such as MACHO or Sloan Digital Sky Survey (SDSS) where a substantial number of objects with

[1] Center for Particle Astrophysics, University of California, Berkeley, CA 94720.
[2] Postal address: Department of Physics, University of Michigan, Ann Arbor, MI 48109.
[3] Postal address: Lawrence Livermore National Laboratory, Livermore, CA 94550.
[4] Mount Stromlo and Siding Spring Observatories, ACT 2611 Weston Creek, Australia.
[5] Postal address: Department of Physics, University of California, San Diego, La Jolla, CA 92093.
[6] Postal address: Department of Physics, University of California, Santa Barbara, CA 93106.
[7] Department of Statistics, University of California, Berkeley, CA 94720.

initially unknown characteristics must be investigated. The beauty of the spline method is that the restricted domain of each individual B-spline plus various other tricks lead to a very fast algorithm for an exhaustive search. At the same time, a spline curve nicely expresses the expected smoothness of the period-folded light curve while retaining the degrees of freedom to conform closely to all of the gross features. To the best of our knowledge, the advantage of using splines in this context has not been previously exploited. The work described below followed a number of conversations with the members of the statistics department at Berkeley and elsewhere. We embarked on this project only after realizing that the traditional methods for searching for periodic behavior could be drastically improved. In the following pages, the mathematical content is fully described, but to use this method, the reader is strongly urged to obtain copies of our own computer code. It took a considerable programming effort to realize the inherent computational efficiency implicit in the fitting procedure which would be senseless to reinvent.

2. CUBIC B-SPLINES

The starting point for the least-squares cubic spline method is the set of functions called B-splines (de Boor 1978; Powell 1981). Assume that a function is to be approximated on the interval $a \leq x \leq b$. A sequence of values, called knots, are chosen in this interval such that $a = x_0^* < x_1^* < \cdots < x_{n-1}^* = b$. A cubic spline curve defined with respect to this knot vector is a piecewise cubic polynomial with coefficients arranged so that the curve and the first two derivatives are continuous in value at each interior knot point. If there are n knot points, $4(n - 1)$ coefficients are required to specify the cubic polynomial in each of the $n - 1$ knot intervals. The continuity conditions provide $3(n - 2)$ constraints, leaving $n + 2$ degrees of freedom to be resolved by other considerations. The definition of B-splines incorporates the continuity constraints directly in the construction of sets of basis spline functions for

$$0 \leq p < n + 2 , \quad -\infty < x < \infty ,$$

$$B_p(x) = (x_{p+1}^* - x_{p-3}^*) \sum_{j=p-3}^{p+1} \left[\prod_{\substack{i=p-3 \\ i \neq j}}^{p+1} \frac{1}{(x_i^* - x_j^*)} \right] (x - x_j^*)_+^3 ,$$

$$u_+ \equiv \max (0, u) . \tag{1}$$

This definition guarantees that

$$B_p(x) > 0 , \quad x_{p-3}^* < x < x_{p+1}^* ,$$
$$= 0 , \quad \text{elsewhere} , \tag{2}$$

and the continuity conditions are satisfied for all x. It follows that any linear sum of these basis functions will be likewise continuous. A spline approximation of a function can thus be written in the form

$$f(x) = \sum_{i=0}^{n+1} c_i B_i(x) \tag{3}$$

by specifying only $n + 2$ coefficients. If the knot points are equally spaced so that

$$x_{i+1}^* = x_i^* + \Delta , \tag{4}$$

the B-splines obey the following trivial recursion relation

$$B_{i+1}(x + \Delta) = B_i(x) . \tag{5}$$

Under these conditions, equation (1) simplifies to

$$B_p(x) = \begin{cases} \frac{1}{6}(4 - 6u^2 + 3u^3) , & 0 \leq u \leq 1 , \\ \frac{1}{6}(2 - u)^3 , & 1 \leq u \leq 2 , \\ 0 , & 2 \leq u , \end{cases} \tag{6}$$

where $u = |(x/\Delta) - p + 1|$. This leads to a significant compression of the computations required to manipulate spline curve representations. To provide a more graphic idea of how B-splines behave, Figure 1 shows the set of spline functions required to approximate a function on the interval $0 \leq x \leq 5$, with knots at the integer values along the x-axis. Eight B-spline functions constitute the basis set.

3. FITTING PERIOD-FOLDED LIGHT CURVES

The least-squares technique can be easily applied to the B-spline representation to estimate the coefficients c_i through the normal equations

$$\frac{\partial}{\partial c_j} \sum_{i=1}^{N} w_i [f(x_i) - y_i]^2 = 0 . \tag{7}$$

Since for any $x, f(x)$ is the sum of only four adjacent B-spline functions, each data point, x_i, only contributes to a small number of elements of the regression matrix, independent of the number of knot points. The resulting regression matrix is thus band diagonal with three subdiagonals above and below the main diagonal. Solving such a sparse linear system is considerably faster than for the case when the matrix elements are all nonzero.

For modeling periodic behavior, it is convenient to map the observation time variable t into the unit interval via the prescription

$$x = \omega t \bmod 1 , \tag{8}$$

where ω is the assumed oscillation frequency. This requires a slightly more complex computation of the spline coefficients. Assume that the unit interval, [0, 1], is to be divided equally into n subintervals. The corresponding knot vector contains $n + 1$ points so that the cubic B-spline representation requires $n + 3$ coefficients. This leaves the two ends of the spline com-


FIG. 1.—Set of eight cubic B-spline functions with knots at the integers that approximate a function on the interval $0 \leq x \leq 5$. A periodic function can be represented on this interval by requiring that the spline coefficients obey $c_i = c_{i \bmod 5}$.

pletely independent of each other. To make the representation periodic, the three continuity conditions must be applied at the boundaries. This means that the coefficients obey

$$c_i = c_{i \bmod n} \, . \qquad (9)$$

When this condition is asserted, the regression matrix gains nonzero elements in the upper-right and lower-left corners. Cholesky decomposition can still be used to solve for the coefficients, but the inversion algorithm must deal with a more complicated matrix structure.

To be effective for frequency searches, these computations must be adapted to the capabilities of computers. Two techniques are used to simplify the calculations and reduce the execution time: discretizing the modeling function (in this case, the B-splines) to a finite number of values over the unit phase interval and using a uniformly spaced grid of frequency values.

Truncating the precision of the light-curve phase data reduces the computation of the B-spline functions to a table lookup procedure. If we use only 8 bits to store the phase, a spline function can be calculated for all possible phase values and compactly stored in memory. Since the B-splines are translationally invariant for uniform knot spacing, only one of the spline functions need be stored. We use 16 knot intervals, and construct function tables for the 16 individual subinterval argument values. There are 10 nonzero matrix elements in the regression matrix and four B-spline components which must be calculated for each of the 16 knot subintervals. If these 16×14 values are premultiplied by their appropriate significance weighting for each of the data points of the light curve, the total required storage is $16 \times 14 \times N$ floating point values, where N is the number of photometric observations. This precomputation requires, at most, about 1 Mbyte of memory and results in the innermost program loop, where the regression matrix sums are calculated, consisting of just 14 floating point additions per data point with no multiplications or divisions.

The remainder of the calculations required to find the χ^2 statistic at each sample frequency have a fixed cost which is independent of the number of data points, since the χ^2 statistic can be obtained from a sum of products of the spline coefficients with no further iteration over the number of data points.

The use of an equally spaced grid of frequencies simplifies computation of the light-curve phase. Since the χ^2 statistic must be evaluated at each test frequency, if the search is uniformly gridded in frequency space the light-curve phase for a particular data point can be successively calculated by adding a constant value to some initial value. This process can be mapped to integer arithmetic in such a way that the most significant 8 bits can be interpreted as the light-curve phase. Table lookup can then be efficiently accomplished by simple bit-mask and bit-shift operations.

4. COMPUTATIONAL RESULTS

The basic ideas presented here were encoded in a C program and executed on a Sun IPC workstation. The algorithm was segmented into a large number of subroutines, and each routine was written with both single- and double-precision versions for a total of over 8600 lines of code. The typical performance for analyzing variable star light curves obtained by the MACHO experiment (Alcock et al. 1993) was 85 s to compute χ^2 for 27,000 frequencies with a data set of 77 points. About half of this time was spent by the linear equation solving routines which will not increase for larger data sets. Approx-

imately 500,000 bytes of free storage were required, mostly to store information about the fit at each test frequency. Several hundred light curves have been analyzed to date. Since the MACHO imaging system consists of two separate CCD focal-plane arrays illuminated through both parts of a dichroic beam splitter, two color magnitudes are determined for each observation, corresponding roughly to the photometric V and R bands. The parallel analysis of two independent light curves has provided an important verification that the computer program can identify the correct periodicity. Some examples of fitted light-curve data are shown in Figures 2, 3, and 4, in order of decreasing apparent luminosity.

As an additional check, 180 pairs of light curves were analyzed for periodicity using both the least-squares spline technique described above and the Stellingwerf method (Stellingwerf 1978). Most of the data (75%) in this particular sample corresponded to eclipsing binaries with the remainder exhibiting totally aperiodic behavior. The frequencies obtained by the two methods were compared and in most cases were the same or else just differed by exactly a factor of 2. For 3% of the light curves, the Stellingwerf code failed to find the appropriate periodicity. In some cases, the selected frequency was a higher harmonic or an alias, and in the rest, the value had no obvious relation whatsoever to the correct number. No example was found in which the Stellingwerf algorithm found a valid frequency that was substantially different from the least-squares spline estimate. From this sample we conclude that the least-squares spline technique is at least as robust as a more conventional method that has been widely adopted for variable star data analysis.

5. FREQUENCY ERROR ESTIMATION

An important question about any determination of periodic behavior is the confidence limits that should be associated with the measurement. This aspect of spectral analysis seems to be only rarely mentioned in the literature. One approach is to explore the shape of the χ^2 distribution in the neighborhood of the appropriate minimum. If the measurement errors are all Gaussian and independent, the 1 σ confidence limits for the frequency are set by the points on either side at which the χ^2 curve increases to the value, $\chi^2_{\min} + 1$.

To obtain an estimate of frequency error, assume that the light-curve data is approximated by the function $f(x)$. Averaging over all possible experiments with data points obtained from the same variable star, the expectation value for the χ^2 statistic is

$$\langle \chi^2 \rangle = \langle \Sigma w_i [y_i - f(x_i)]^2 \rangle \, , \qquad (10)$$

which can be rewritten as

$$\begin{aligned}
\langle \chi^2 \rangle &= \langle \Sigma w_i \{ y_i^2 - \langle y_i \rangle^2 - 2(y_i - \langle y_i \rangle) f(x_i) \\
&\quad + [\langle y_i \rangle - f(x_i)]^2 \} \rangle \\
&= N + \Sigma w_i [\langle y_i \rangle - f(x_i)]^2 \, , \qquad (11)
\end{aligned}$$

where N is the total number of data points and $\langle y_i \rangle$ is the true mean of the observations, which might not necessarily be identical to $f(x_i)$.

As a simple example, assume homoscedasticity and sinusoidal light-curve behavior. We explore the situation where the approximating function, $f(x_i)$, differs slightly in frequency from

FIG. 2.—Light curves for a Cepheid variable in the Large Magellanic Cloud (LMC). The upper two plots show the blue and red data as a function of observation date. The lower two plots show the same data folded at the estimated oscillation frequency and extended over two full cycles.

the true value so that

$$y(x) = A \sin (2\pi\omega_0 t) ,$$
$$f(x) = A \sin (2\pi\omega_1 t) . \qquad (12)$$

As long as the data is relatively uniformly sampled in time, then summation can be replaced by an integral so that

$$\langle \chi^2 \rangle = N + \frac{NA^2}{T\sigma^2} \int_{-T/2}^{+T/2} [(\sin 2\pi\omega_0 t) - (\sin 2\pi\omega_1 t)]^2 \, dt \qquad (13)$$

For $\omega_1 \simeq \omega_0$ this expression is approximately given by

$$\langle \chi^2 \rangle \cong N \left[1 + \frac{(2\pi)^2(\omega_0 - \omega_1)^2 T^2 A^2}{24 \, \sigma^2} \right] . \qquad (14)$$

Thus, the variance for the frequency estimation is

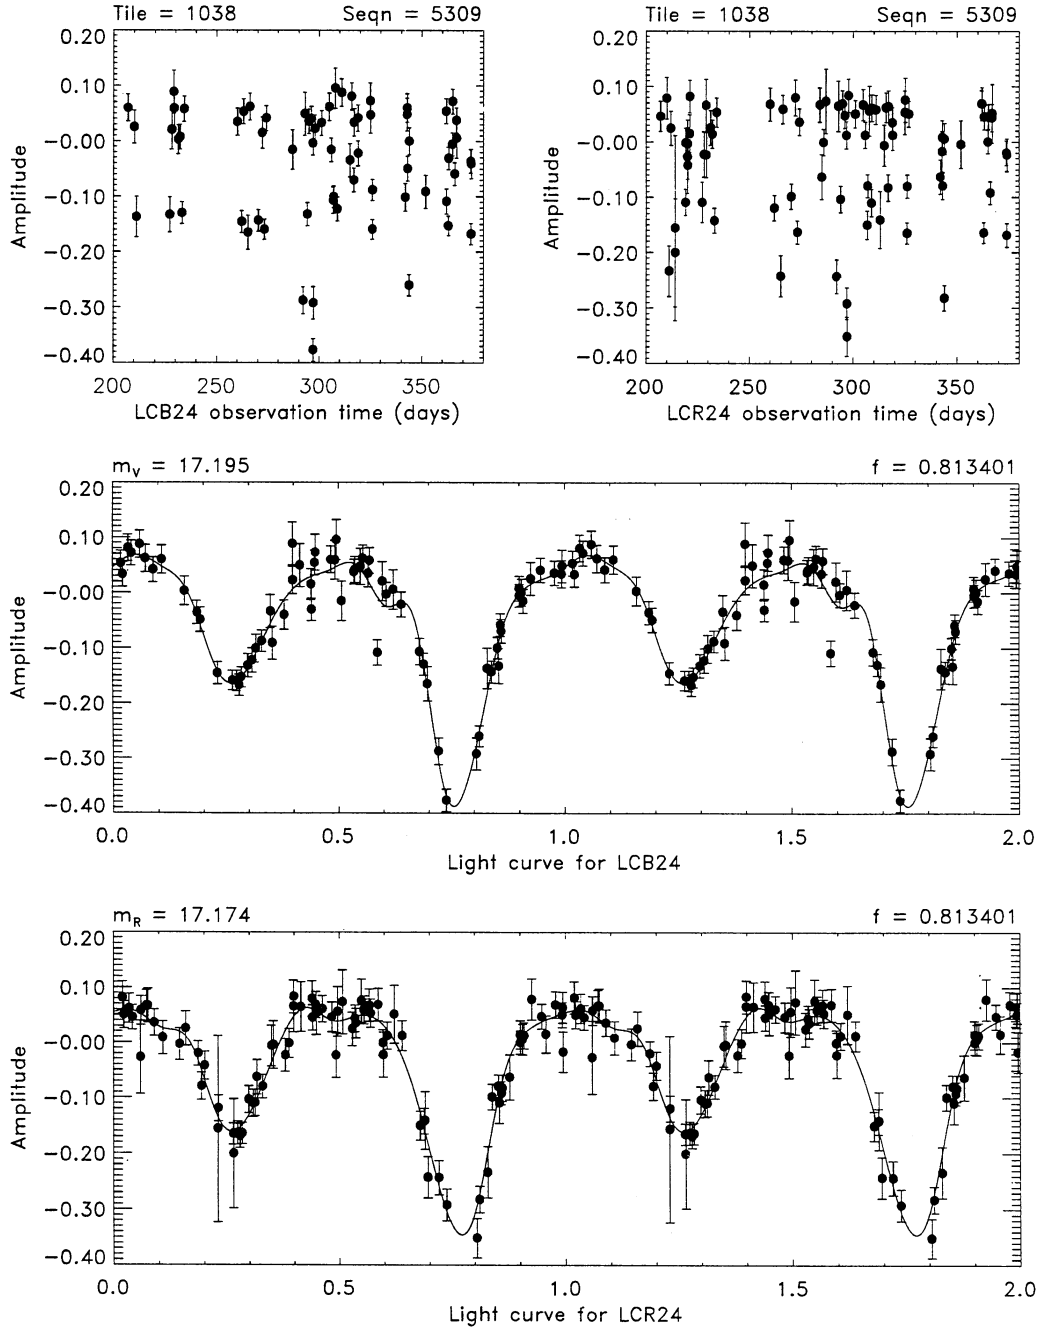$$\sigma_\omega^2 = \frac{24}{N} \left(\frac{\sigma}{2\pi T A} \right)^2 . \qquad (15)$$

FIG. 3.—Light curves for an eclipsing binary star in the LMC. The upper two plots show the blue and red data as a function of observation date. The lower two plots show the same data folded by the estimated eclipse period and extended over two full cycles.

A similar expression has been previously derived by Lucy & Sweeney (1971, eq. [20]).

A more general and elegant derivation can be developed from the methods of statistical regression analysis (Draper & Smith 1981). Consider a statistical model of the form

$$y_i = f_\beta(x_i) + e_i , \quad i = 1, \ldots, N , \tag{16}$$

where f is a function of the x variable and a parameter vector β of dimension p; the errors e_i have mean zero and constant variance and are independent. Then, if N is large, the variance-covariance matrix of the estimated parameters is approx-

imated by[8]

$$V(\beta) \approx \sigma^2 [Z'Z]^{-1} , \tag{17}$$

where Z is an $N \times p$ matrix defined by

$$Z_{ij} = \frac{\partial f(x_i; \beta)}{\partial \beta_j} , \quad i = 1, \ldots, N, j = 1, \ldots, p . \tag{18}$$

As a simple example, assume sinusoidal light-curve behavior, with $\beta = (A, B, \omega)$ and $f(x_i; \beta) = A \cos (2\pi\omega x_i) + B \sin (2\pi\omega x_i)$. We now need to calculate $Z'Z$; as an example, we present the

_____

[8] This approximation requires that the model can be locally approximated as a linear function of the β parameters.
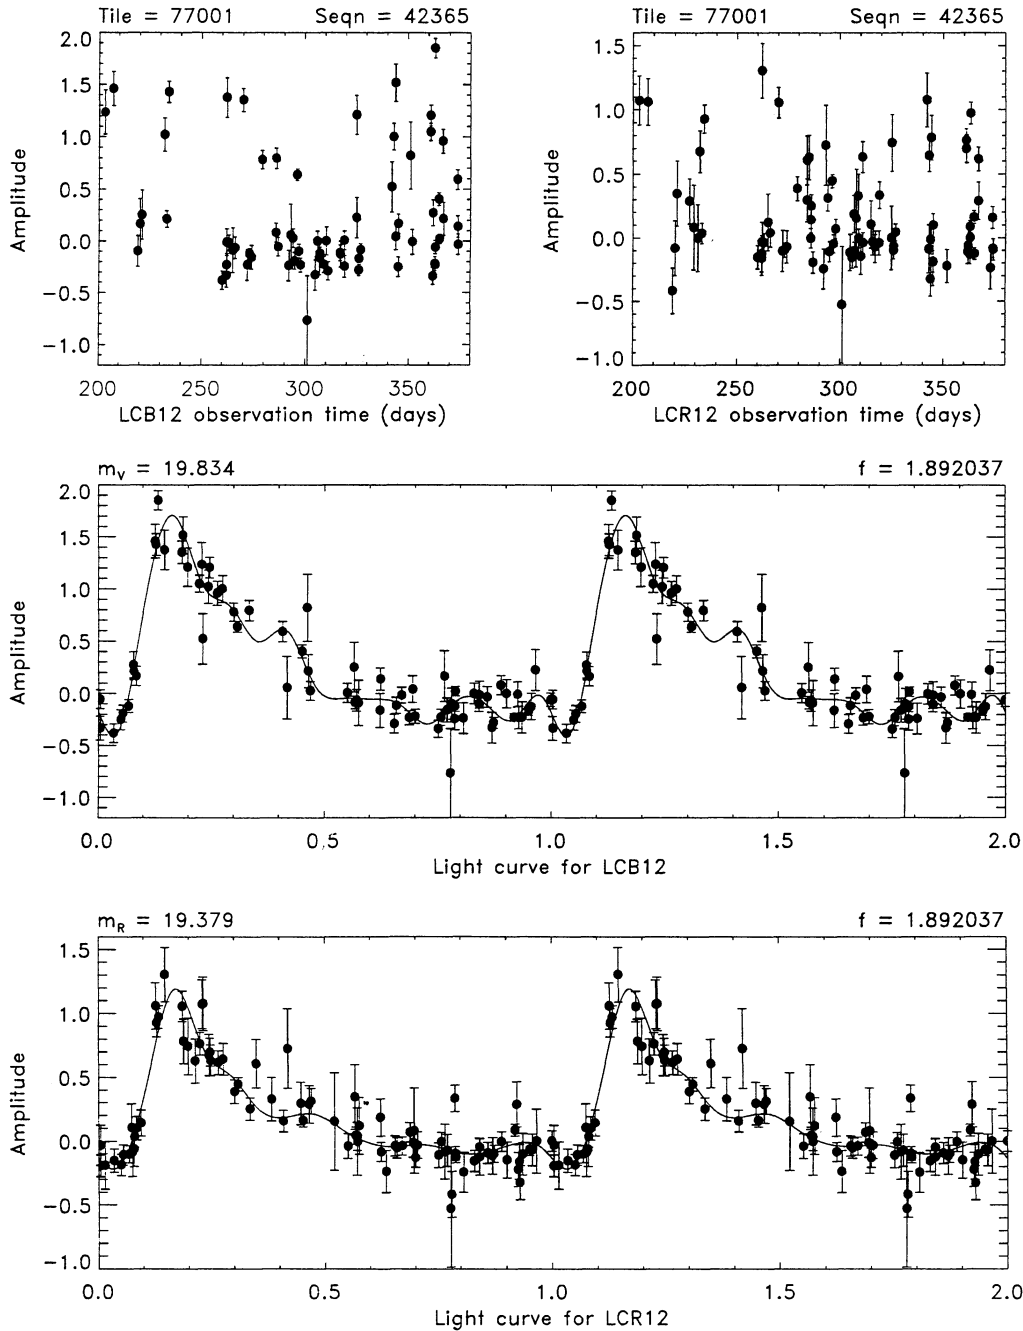
Fig. 4.—Light curves for an RR Lyra variable star in the LMC. The upper two plots show the blue and red data as a function of observation date. The lower two plots show the same data folded at the estimated oscillation frequency and extended over two full cycles.

calculation of the (1, 3) element of $Z'Z$:

$$[Z'Z]_{1,3} = \sum_{i=1}^{N} \cos{(2\pi\omega x_i)}[2\pi B x_i \cos{(2\pi\omega x_i)}$$
$$- 2\pi A x_i \sin{(2\pi\omega x_i)}]$$
$$\approx \frac{N}{T} \int_0^T [2\pi B x \cos^2{(2\pi\omega x)}$$
$$- 2\pi A x \cos{(2\pi\omega x)} \sin{(2\pi\omega x)}]dx$$

for large N, assuming observations are approximately uniformly sampled in time. Thus

$$[Z'Z]_{1,3} \approx \frac{N}{T} \int_0^T \frac{2\pi B x}{2} \, dx$$

$$= \frac{2\pi B N T}{4} \, . \tag{19}$$

TABLE 1

SYNOPSIS OF FREQUENCY ERROR ESTIMATES FOR THE B-SPLINE SEARCH PROCEDURE

Stellar Type (1)	f (2)	χ^2/dof (3)	$f_{\text{red}} - f_{\text{blue}}$ (4)	$\sqrt{\sigma_{\text{red}}^2 + \sigma_{\text{blue}}^2}$ (5)	$\sqrt{\delta f_{\text{red}}^2 + \delta f_{\text{blue}}^2}$ (6)
Cepheid variable	0.198192	0.26	+0.000023	0.000011	0.000335
Eclipsing binary	0.813401	1.05	−0.000070	0.000056	0.000111
RR Lyra variable	1.892037	1.84	+0.000048	0.000024	0.000028

NOTE.—The values in each column are described in the text.

Similar calculations give

$$(Z'Z) \approx \begin{bmatrix} \dfrac{N}{2} & 0 & \dfrac{2\pi}{4} BNT \\[2mm] 0 & \dfrac{N}{2} & -\dfrac{2\pi}{4} ANT \\[2mm] \dfrac{2\pi}{4} BNT & \dfrac{-2\pi}{4} ANT & \dfrac{(2\pi)^2}{6}(A^2 + B^2)NT^2 \end{bmatrix}.$$

(20)

Thus, the appropriate variance-covariance matrix for all three parameters is

$$V(A, B, \omega) \approx \frac{2\sigma^2}{(A^2 + B^2)} \begin{bmatrix} \dfrac{A^2 + 4B^2}{N} & \dfrac{-3AB}{N} & \dfrac{-6B}{2\pi NT} \\[2mm] * & \dfrac{4A^2 + B^2}{N} & \dfrac{6A}{2\pi NT} \\[2mm] * & * & \dfrac{12}{(2\pi)^2 NT^2} \end{bmatrix},$$

(21)

and the variance for the frequency estimate is

$$\sigma_\omega^2 = \frac{24}{N}\left[\frac{\sigma}{2\pi T(A^2 + B^2)^{1/2}}\right]^2.$$

(22)

We expect that the variance for nonsinusoidal light curves will be functionally similar, differing only by dimensionless factors of order unity. This provides a useful guide to selecting the frequency interval required to explore neighborhoods of χ^2 minima. As an example, the periodicity analysis for the data graphed in Figures 2, 3, and 4 is summarized in Table 1. Since the light-curve frequency was independently searched in both red and blue wavelength bands, the difference between these two values gives a primitive measure of the overall error (col. [4]). This can be compared with the frequency error computed from equation (15) (col. [5]). Finally, an independent estimate is obtained by determining the frequency shift required to increase the minimum χ^2 value by one unit (col. [6]). For the Cepheid variable, the photometric errors are quite small due to the relative brightness of the object. In fact, the error bars for bright stars were systematically overestimated in the early MACHO photometric reductions, leading to a ratio of χ^2 to degrees of freedom substantially less than unity. This artificially broadens the frequency range computed in the last column. The difference in frequency between the red and blue light curves, 0.000023 cycles day^{-1}, is at the granularity limit of the digital gridding algorithm, $1/256T$. A more accurate deter-

mination of the χ^2 minimum could easily be explored if warranted by data accuracy and number of observations. Since the requisite number of computations is quite limited, calculational efficiency ceases to be an issue. For the eclipsing variable, the ratio of χ^2 to degrees of freedom is close to one and the frequency error estimates are fairly comparable. The fit for the RR Lyra variable is not so good, probably due to intrinsic stellar fluctuations from the smoothed light curve. It should be expected that the frequency errors determined from equation (15) are close to a lower bound. That calculation was based on the assumption that the original data exactly followed a sinusoidal shape. In the B-spline analysis, the light curve is allowed an arbitrary continuous form and consequently the loss of specificity increases the uncertainty of the location of the χ^2 minimum. Once the light-curve shape has been accurately parameterized and the functional ambiguity reduced, a second-pass analysis can optimally determine the most probable frequency, should that be necessary.

In this general area of confidence interval determination, the B-spline least-squares approach is likely to be superior to the Lomb-Scargle method. The later technique is not amenable to the type of analysis discussed above since no attempt is made to represent the light-curve shape. It would be an interesting exercise in statistics to compare the asymptotic relative efficiencies of these two techniques (i.e., the ratio of frequency variances) as the light-curve shape departs from simple harmonic. One expects that the purely harmonic situation will be handled better by Lomb-Scargle but that this will shift toward the spline method for curves with increasingly higher Fourier components. It is unclear what number of knot intervals for the periodic splines is optimal for reducing the frequency variance. These kinds of questions may only be susceptible to answer by massive Monte Carlo simulations.

6. SUMMARY

In summary, a new method of analyzing the periodicity of unequally spaced data has been described which can be implemented efficiently on modern computer processors. Copies of the code can be obtained via the Internet from akerlof@mail.physics.lsa.umich.edu.

REFERENCES

Alcock, C., et al. 1993, in Robotic Telescopes in the 1990s, ed. A. V. Filippenko (San Francisco: ASP), 193
de Boor, C. 1978, A Practical Guide to Splines (New York: Springer)
Draper, N. R., & Smith, H. 1981, Applied Regression Analysis (2d ed.; New York: Wiley)
Friedman, J. H. 1984, A Variable Span Smoother, SLAC Preprint PUB-3477
Fullerton, A. W. 1986, The Study of Variable Stars Using Small Telescopes, ed. J. R. Percy (New York: Cambridge Univ. Press), 201
Horne, J. H., & Baliunas, S. L. 1986, ApJ, 302, 757

Lomb, N. R. 1976, Ap&SS, 39, 447
Lucy, L. B., & Sweeney, M. A. 1971, AJ, 76, 544
McDonald, J. A. 1986, SIAM J. Sci. Stat. Comput., 7, 665
Powell, M. J. D. 1981, Approximation Theory and Methods (New York: Cambridge Univ. Press)
Press, W. H., & Rybicki, G. R. 1989, ApJ, 338, 277
Press, W. H., & Teukolsky, S. A. 1988, Comput. Phys., 2, 77
Scargle, J. D. 1982, ApJ, 263, 835
Stellingwerf, R. F. 1978, ApJ, 224, 953