

WAVELETS FOR PERIOD ANALYSIS OF UNEVENLY SAMPLED TIME SERIES

GRANT FOSTER

American Association of Variable Star Observers, 25 Birch Street, Cambridge, Massachusetts 02138

Electronic mail: gfooster@aavso.org

Received 1995 November 15; revised 1996 June 25

ABSTRACT

The wavelet transform shows great promise as a method for period analysis in time series, particularly for detecting the time evolution of the parameters (period, amplitude, phase) describing periodic and pseudo-periodic signals. However, when applied to unevenly sampled time series, the response of the wavelet transform is often more dependent on irregularities in the number and spacing of available data than on actual changes in the parameters of the signal. Yet by casting the wavelet transform as a projection, we can derive its statistical behavior and devise advantageous rescaled transforms. By treating it as a weighted projection to form the *weighted wavelet Z-transform (WWZ)*, we improve its ability to detect, and especially to quantify, periodic and pseudo-periodic signals. The methods are illustrated by analysis of artificial test data, and of the light curves of the variable stars R Aquilae and FS Comae. © 1996 American Astronomical Society.

1. INTRODUCTION

Fourier analysis is an ideal tool for detecting and quantifying periodic fluctuations in time series, *if* by periodic we mean of truly constant period, amplitude, and phase. Real astrophysical systems rarely exhibit such constancy of fluctuation. Often periodic fluctuations arise intermittently, as transient phenomena. Even for a time series with consistent periodicity we usually see time evolution of the parameters of the fluctuation. Fourier analysis can detect, and to some degree quantify, such behavior, but it is far from *ideal* for such purposes (although recent developments promise to improve its performance, Foster 1995).

The wavelet transform, on the other hand, is well-suited to detect transient periodic fluctuations, as well as changes in their parameters, because it can focus attention on a limited time span of the data. The continuous wavelet transform of a function of time $x(t)$ is defined as (Grossman *et al.* 1989)

$$\begin{aligned} W(\omega, \tau; x(t)) &= \omega^{1/2} \int x(t) f^*(\omega(t-\tau)) dt \\ &= \omega^{-1/2} \int x(\omega^{-1}z + \tau) f^*(z) dz, \end{aligned} \quad (1-1)$$

where f^* is the complex conjugate of f , and the function $f(z)$ is the *wavelet kernel*, also known as the *analyzing wavelet* or *mother wavelet*. The transform depends on two parameters, the *scale factor* ω (ω^{-1} is also called the scale factor, and called the *dilation*; I shall hereafter refer to ω as the *frequency*), and the *time shift* τ (a.k.a. the *shift parameter*, *position*, or *location*; as you can see, there is a wide variety of nomenclature in the literature). By choosing a mother wavelet which is concentrated near $z=0$, we explore the behavior of $x(t)$ near $t=\tau$.

We have a bewildering variety of analyzing wavelets to choose from, but because we are interested in detecting and

quantifying periodic and pseudo-periodic fluctuations, we shall adopt a mother wavelet which fluctuates due to a term of the form e^{iz} . One of the important consequences of using a complex mother wavelet is that it possesses *two* degrees of freedom (its real and imaginary parts) to probe for structure in our time series. The Fourier transform is a special case of (1.1) with mother wavelet $f(z) = e^{iz} = e^{i\omega(t-\tau)}$, for which we do *not* insert the factor $\omega^{1/2}$ in front of the integral. Also, this function is an eigenfunction of the time translation operator, so for Fourier analysis we drop the irrelevant parameter τ .

Recently astronomers have begun to explore new possibilities; the *en vogue* mother wavelet for variable star light curves is the Fourier wavelet with a Gaussian decay profile (or, a Gaussian with harmonic modulation), the so-called *Morlet wavelet* (Grossman & Morlet 1984)¹

$$f(z) = e^{-cz^2} (e^{iz} - e^{-1/4c}) = e^{-c\omega^2(t-\tau)^2} (e^{i\omega(t-\tau)} - e^{-1/4c}), \quad (1-2)$$

where the constant $e^{-1/4c}$ is inserted so that the average value of the analyzing wavelet is zero

$$\int_{-\infty}^{+\infty} f(z) dz = 0. \quad (1-3)$$

The constant c determines how rapidly the analyzing wavelet decays; it is usually chosen so that the exponential term decreases significantly in a single cycle $2\pi/\omega$. For variable star light curves, for example, a popular choice for c is $1/8\pi^2$ (however, c could be treated as a *parameter*, in which case the wavelet transform would be characterized by three parameters instead of two). I shall choose the value

¹The Morlet wavelet is more usually defined as $f(z) = e^{-z^2/2} (e^{i\omega_m z} - e^{-\omega_m^2/2})$, where $\omega_m = 1/\sqrt{2c}$. I have transferred the constant to the exponential factor (which is merely a choice of convention) so that the scale factor ω will exactly correspond to the (radian) frequency of a periodic fluctuation.

$c=0.0125$, which is close enough to $1/8\pi^2$ that the two choices are functionally equivalent. When c is small (as is usually the case when applying wavelets to period analysis of time series), the constant $e^{-1/4c}$ is quite small; for $c<.02$ it is negligible (Grossman *et al.* 1989). Hence it is customary to ignore this term, giving a different wavelet kernel which I will dub the *abbreviated Morlet wavelet*

$$f(z) = e^{iz - cz^2} = e^{i\omega(t-\tau) - c\omega^2(t-\tau)^2}. \quad (1-4)$$

This serves to define the *abbreviated Morlet transform*. When adapted to a discrete time series, it will define the *discrete wavelet transform*, or *DWT*. The abbreviated Morlet transform bears a striking similarity to a windowed Fourier transform, with window $e^{-c\omega^2(t-\tau)^2}$. The crucial difference is the factor ω in the exponential; the size of the “window” is frequency-dependent.

My purpose is twofold: first, to consider the response of the DWT, as defined by the abbreviated Morlet wavelet, to periodic signals which are irregularly sampled; second, to propose modifications to wavelet analysis which improve its performance. Experience shows that adapting even something as thoroughly well-understood as the Fourier transform to a discrete time series presents numerous difficulties. Even with a perfectly regular time spacing the *discrete Fourier transform (DFT)* has nontrivial statistical behavior (Kovacs 1980; Foster 1996a). Uneven time spacing can wreak havoc. False peaks (ghost images of real peaks) can become very strong; in fact two false peaks may overlap and reinforce each other, so that even with a noise-free signal the strongest peak in the DFT is at a *spurious* frequency (Foster 1995). Any signal component we locate in the DFT will be shifted from the true signal frequency (Kovacs 1981), and the amplitude we compute for that frequency component can be very much in error (Ferraz-Mello 1981; Foster 1995).

It should come as no surprise that when we adapt the wavelet transform to uneven time sampling, the same difficulties arise as with the Fourier transform. The good news is, we can apply the same remedies. Many of the problems are overcome by treating Fourier analysis as a *projection* onto the trial functions $\sin(\omega t)$, $\cos(\omega t)$, and an arbitrary constant (Ferraz-Mello 1981; Foster 1995, 1996a, 1996b). Furthermore, we can view the wavelet transform not just as a projection onto a set of trial functions, but as a *weighted* projection. Both strategies enhance our ability to detect periodic fluctuations, and to estimate their parameters with improved precision and consistency. Finally, the statistical behavior of projections is known, and quite simple (Foster 1996a).

To ease the burden to the reader, I have relegated the goriest mathematical details to appendixes, which exist only to establish certain results used herein. There the wavelet transform is formulated in terms of tensor notation, according to the treatment of Foster (1996b). I warn the reader sternly that the notation used in the appendixes is quite different from that of the main body of the paper; those interested should definitely consult the aforementioned paper for details of the notation and conventions.

2. DATA

Pure theory is fine, but no method can be considered satisfactory which has not run the gauntlet of real data with real time sampling. Hence in order to test the efficacy of the various incarnations of the wavelet transform, I will apply these analyses to the light curves of two variable stars, R Aquilae and FS Comae. R Aquilae is a Mira-type long-period variable; according to the GCVS it has a magnitude range from 5.5 to 12.0, period 284.2 days, and spectral type M5E-M9E (Kholopov *et al.* 1985). The data are 10-day averages of visual observations from the AAVSO international database (Mattei 1995a), yielding 1,107 data points covering the time interval JD 2437601 to 2449626. FS Comae is a semiregular (SR B) variable ranging from 5.5 to 6.1, period 58 days, and type M3E spectrum (Kholopov *et al.* 1985). The FS Com data for this study are 511 photoelectric magnitudes taken from the AAVSO international database of photoelectric photometry (Mattei 1995b), spanning the time range JD 2446165 to 2449508. For the most part, these data sets were selected for their time sampling. The sampling of the R Aquilae data is only mildly irregular, especially since we have taken 10-day averages; despite random fluctuations, the data are spaced roughly 10 days apart, with very few data missing. The FS Comae data have a nightmarish time sampling, showing large gaps in the data stream which repeat annually and large fluctuations of data density.

For the purpose of evaluating period analysis techniques, real data do have a drawback: the true signal is not known *a priori*. Therefore in addition to the R Aql and FS Com data, I shall also analyze artificial test data of known signal. Instead of artificially gapping the test data, we have generated test signals with the same time sampling as the real data for R Aql and FS Com. Test data set 1 is a pure sinusoid with perfect time spacing: it has period 300 days, (real semi-) amplitude 1 mag., with 401 data points spaced 10 days apart, covering the (arbitrary) time span JD 2446000 to 2450000. Test data set 2 is a pure sinusoid with period 283.3 days, amplitude 1.9011 mag, and the same time sampling as the R Aql data. Test data set 3 is a pure sinusoid with period 55.98 days, amplitude 0.0598 mag., and the same time sampling as FS Com, while test data set 4 is a pure sinusoid with period 662 days, amplitude 0.0513 mag, and (again) the same time sampling as FS Com. For all data sets we have subtracted the average value from the data prior to analysis. Figure 1 plots the test data sets. While introducing the many different forms of wavelet transform, I shall apply them to the artificial data sets in order to gauge their response to pure signals. After comparative analysis of different wavelet methods, I shall attack the light curves of R Aql and FS Com.

3. DISCRETE WAVELET TRANSFORM (DWT)

An observed time series consists of N data values $x(t_\alpha)$, taken at a discrete set of N times $\{t_\alpha, \alpha=1,2,\dots,N\}$. The most straightforward adaptation of the Fourier transform to a discrete time series is the *discrete Fourier transform (DFT)*

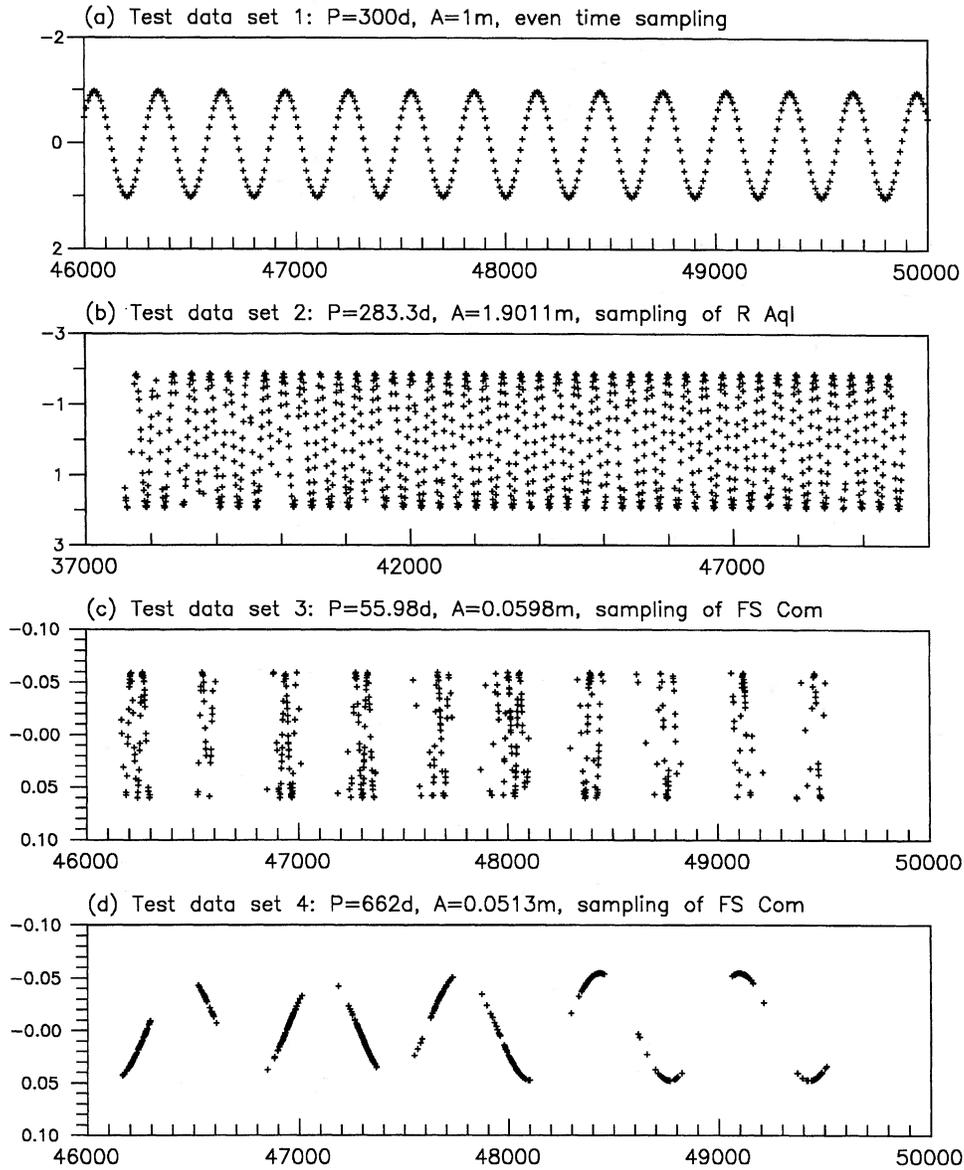


FIG. 1. Test data sets. (a) Period=300, amplitude=1, even time sampling. (b) Period=283.3, amplitude=1.9011, time sampling of R Aql. (c) Period=55.98, amplitude=0.0598, sampling of FS Com. (d) Period=662, amplitude=0.0513, sampling of FS Com.

$$D(\omega; x(t)) = N^{-1} \sum_{\alpha=1}^N x(t_{\alpha}) e^{i\omega t_{\alpha}}. \quad (3-1)$$

We then define the DFT *Power* as

$$P(\omega) = |D(\omega; x(t))|^2, \quad (3-2)$$

which we use to evaluate the frequency ω *statistically*. The (real semi-) amplitude of the periodic fluctuation corresponding to a significant peak in the DFT spectrum is

$$A = 2\sqrt{P}. \quad (3-3)$$

The wavelet analogue of (3.1) is the *discrete wavelet transform (DWT)*

$$W(\omega, \tau; x(t)) = \sqrt{\omega} \sum_{\alpha=1}^N x(t_{\alpha}) f^*(\omega(t_{\alpha} - \tau)). \quad (3-4)$$

For the abbreviated Morlet wavelet (1.4), the real and imaginary parts of the DWT are

$$\text{Re}(W) = \sqrt{\omega} \sum_{\alpha=1}^N x(t_{\alpha}) e^{-c\omega^2(t_{\alpha} - \tau)^2} \cos(\omega(t_{\alpha} - \tau)), \quad (3-5)$$

and

$$\text{Im}(W) = -\sqrt{\omega} \sum_{\alpha=1}^N x(t_{\alpha}) e^{-c\omega^2(t_{\alpha} - \tau)^2} \sin(\omega(t_{\alpha} - \tau)). \quad (3-6)$$

We evaluate the response of the DWT by calculating the *wavelet modulus* $|W|$ (although the wavelet *phase* $\arg(W)$ is also of considerable interest, Goupil *et al.* 1991). The wavelet modulus $|W|$ is also referred to as the *amplitude*, but I shall not do so, reserving that term for another definition (Sec. 3.2).

3.1 Response of the DWT to Pure Noise

We begin by adopting the null hypothesis that the data are independent random variables with mean 0 and variance σ^2 .² If the time sampling is ideal (perfectly even time spacing and no data missing), and if the density of data ρ is high enough and the frequency ω low enough that the product $\omega^{-1}\rho \gg 1$, and if the position τ is far enough from the beginning and end of the data that edge effects may be ignored (that's a lot of *ifs!*), then the expected value of the squared modulus of the DWT is independent of the frequency and position (Appendix C)

$$\langle |W|^2 \rangle = \sigma^2 \rho \sqrt{\frac{\pi}{c}}. \quad (3-7)$$

This explains the factor $\sqrt{\omega}$ in the definition of the wavelet transform (1.1); it gives a uniform response to random noise. The quantity (3.7) is (almost exactly) chi-square with 2 degrees of freedom.

For a less-than-ideal time spacing, the definition of the DWT indicates that it will depend strongly on the sheer number of data in the vicinity of the position τ . For convenience define the *local data number* as a function of the frequency ω and position τ

$$n(\omega, \tau) = \sum_{\alpha=1}^N e^{-c\omega^2(t_\alpha - \tau)^2}. \quad (3-8)$$

In the continuum limit, in which $\sum_{\alpha} \rightarrow \int \rho dt$, the local data number n turns out to be independent of the position and inversely proportional to the frequency

$$n(\omega, \tau) = \omega^{-1} \rho \sqrt{\frac{\pi}{c}}. \quad (3-9)$$

For irregular sampling, we will *define* the (position- and frequency-dependent) *data density* as

$$\rho(\omega, \tau) = \omega \sqrt{\frac{c}{\pi}} n(\omega, \tau) = \omega \sqrt{\frac{c}{\pi}} \sum_{\alpha=1}^N e^{-c\omega^2(t_\alpha - \tau)^2}. \quad (3-10)$$

The expected value of the squared modulus of the DWT is (Appendix C)

$$\langle |W|^2 \rangle = \sigma^2 \omega n(\sqrt{2}\omega, \tau) = \sigma^2 \rho(\sqrt{2}\omega, \tau) \sqrt{\frac{\pi}{2c}}. \quad (3-11)$$

Note the factor $\sqrt{2}$ in the frequency arguments of the functions n and ρ of (3.11). However, for irregular time spacing

²This is not quite the correct null hypothesis. We should assume the data have unknown mean μ , and that we have subtracted from the data the *average* value; see Appendix C.

this quantity is *not* chi-square, as is often thought (but then, neither is the DFT, Foster 1996a).

We see that the DWT squared modulus varies linearly with the data density ρ . If ρ fluctuates greatly, then the response to pure noise will be strongly τ -dependent. This means that with irregular time sampling, statistical evaluation of the DWT is not so straightforward.

3.2 Discrete Wavelet Power (DWP)

We can rescale the DWT so that it has uniform noise level, even with uneven time sampling. We define the *discrete wavelet power (DWP)* as

$$DWP = \frac{|W|^2}{s^2 \omega n(\sqrt{2}\omega, \tau)}, \quad (3-12)$$

where s^2 is the estimated σ^2 , and is inserted so that if the data are random noise, the expected value of the DWT is 1. The DWP is also not chi-square, but if the time spacing is not too bad it may be treated as approximately so. In addition to compensating for fluctuations in the data density, the DWP has the virtue that it is uniformly scaled for easy statistical evaluation.

3.3 Response of the DWT to a Pure Sinusoid

For an ideal time sampling, and ignoring edge effects, the DWT squared modulus for a pure sinusoidal signal at frequency ω_0 , of (real semi-) amplitude A is (Appendix D)

$$|W|^2 = \frac{\pi \rho^2 A^2}{4c\omega} e^{-(\omega - \omega_0)^2 / (2c\omega^2)}, \quad (3-13)$$

The first thing to note is that the response is inversely proportional to the frequency. As a result, it peaks not at the signal frequency ω_0 but

$$\omega_{max} = \frac{2\omega_0}{1 + \sqrt{1 + 4c}} \approx \omega_0(1 - c). \quad (3-14)$$

For $c=0.0125$ this causes a 1 1/4% shift in the peak frequency. Also note that the response is asymmetric. More important, the response is exaggerated at low frequencies and suppressed at high ones. This is no surprise; for higher frequencies the wavelet function is narrower, so we effectively sample fewer data points. Figure 2(a) shows the response of the DWT to a pure sinusoid with even time spacing. It was computed with a position grid size of 100 days and a frequency grid size of 0.0001 cycle/day. As expected, it is asymmetric in frequency, peaked slightly below the signal frequency, and shows significant edge effects almost three whole periods from either end of the time span.

For any constant location τ , we can use the peak frequency of the DWT to determine the period. For this computation we used a position grid size of 100 days and a frequency grid size of 0.00001 cycle/day, for greater precision. We not only note that this period is 1 1/4% higher than the signal period, we also detect spurious time evolution of the period due to edge effects [Fig. 6(a)]. Nonetheless, the response is constant far from the edges, and the 1 1/4% frequency shift can be compensated easily enough. Likewise, if

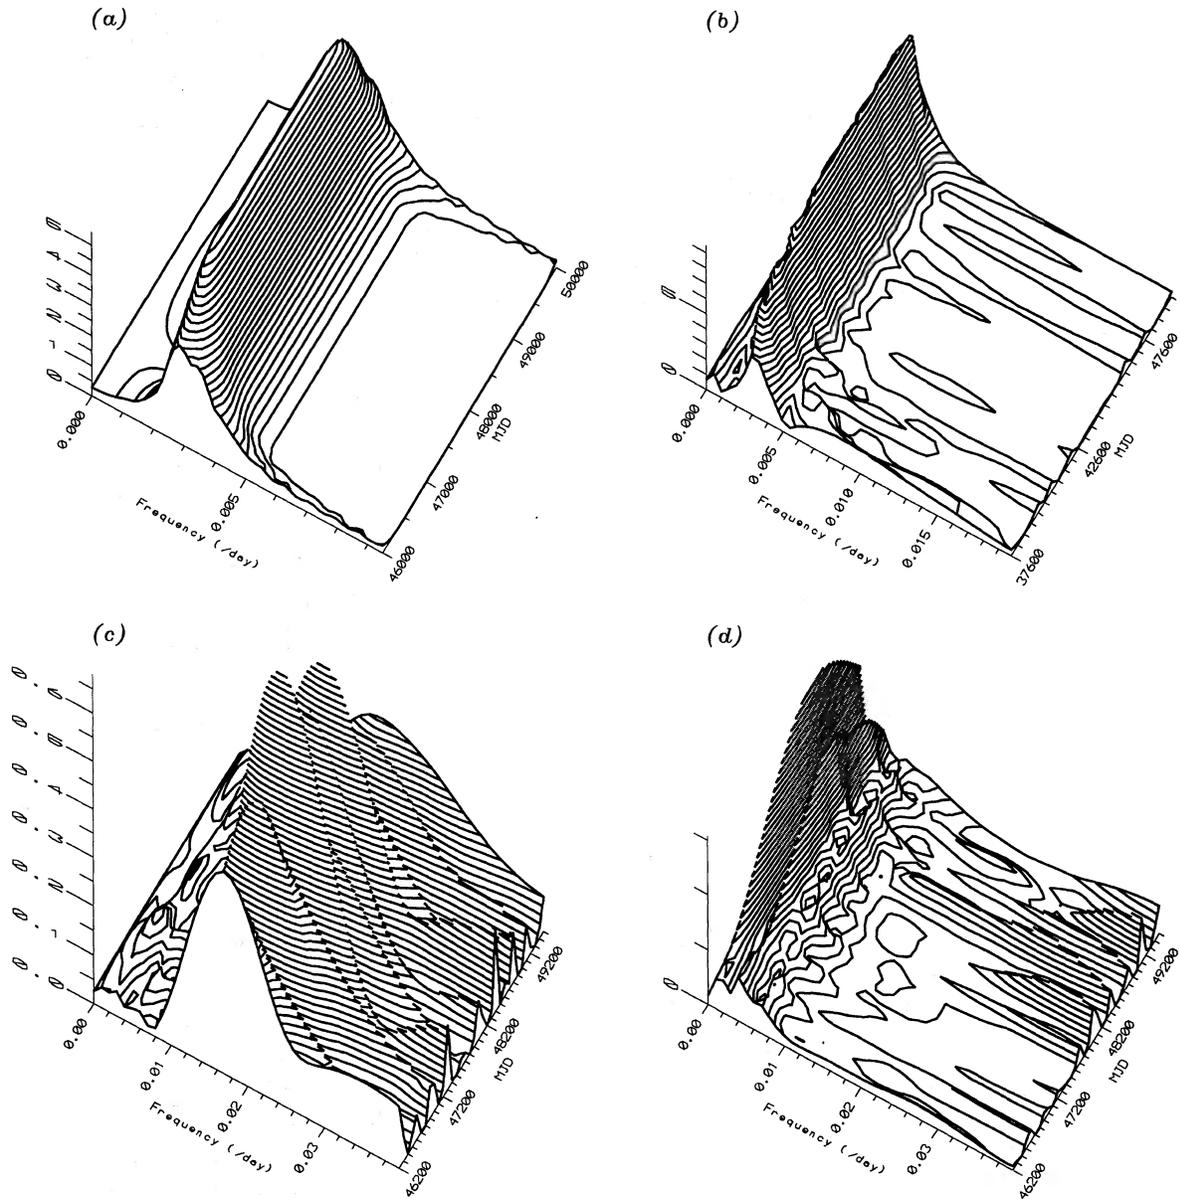


FIG. 2. DWT for test data. (a) Set 1. (b) Set 2. (c) Set 3. (d) Set 4.

we use the wavelet modulus (the DWT) at the peak frequency as an indication of the signal amplitude, we detect spurious amplitude modulation, again entirely from edge effects.

For an irregular time spacing, the response of the DWT to a pure sinusoid (Appendix D) depends so strongly on the irregularities of the data density that we will have a very hard time indeed extracting the physical parameters of a fluctuation (frequency, amplitude, phase) from the DWT. Figures 2(b)–2(d) show the DWT for pure sinusoidal signals with irregular sampling. The peak period varies with time, indicating that it will be difficult to deduce period evolution from the DWT. Still, for mildly irregular time spacing [Fig. 6(c)] the fluctuations are less than 1% above and below the mean

value (more near the edges). For very bad sampling the DWT is a very bad period indicator; with a high-frequency signal [Fig. 6(e)] the fluctuations in peak period are nearly $\pm 20\%$, and for a low-frequency signal [Fig. 6(g)] the DWT indicates a sizeable spurious trend.

The fluctuations in modulus due to varying data density are blatant, making it next to impossible to extract from the DWT alone any information about the time-dependence of the amplitude of these signals. For data set 2 we again see fluctuations, this time about $\pm 10\%$ (larger near the edges). For data set 3, the DWT drops to near zero in the gaps between observing seasons, causing us severely to underestimate the amplitude. And for data set 4, we again find a spurious trend, the DWT modulus quadrupling from mini-

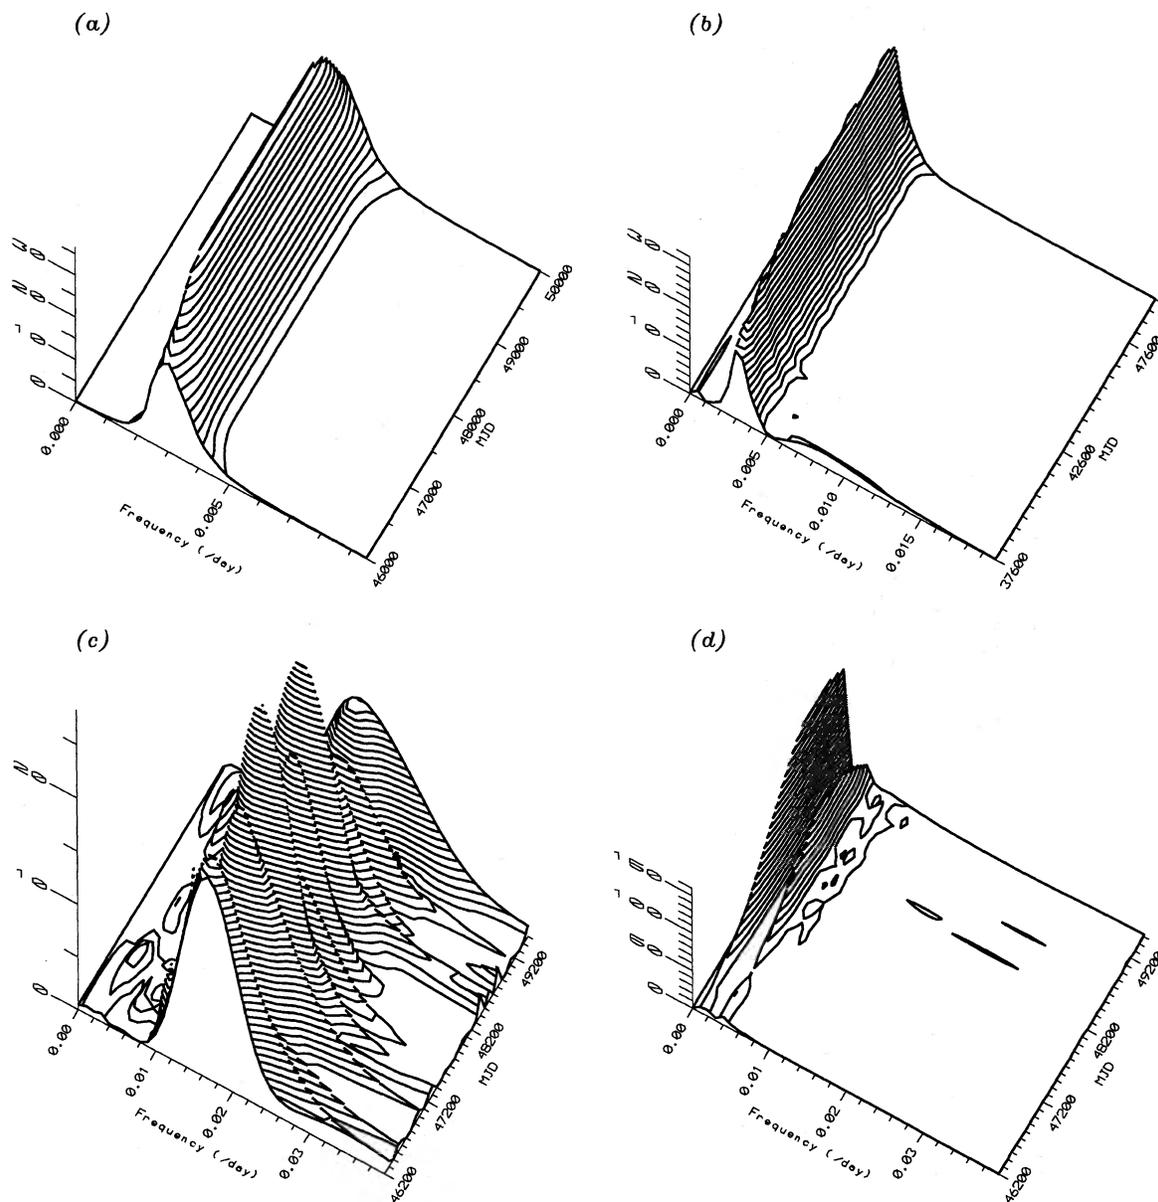


FIG. 3. DWP for test data. (a) Set 1. (b) Set 2. (c) Set 3. (d) Set 4.

num to maximum. These problems are confirmed by experience in applying the DWT to irregularly sampled variable star light curves (Szatmary & Vinko 1992; Szatmary 1994) and to artificial signals with artificial irregular sampling (Szatmary *et al.* 1994).

We could, of course, estimate the signal frequency by locating the peak of the DWP rather than the DWT. Figure 3 shows the DWP for our test data sets. Note that it is indeed an improved indicator of *statistical significance*, suppressing the high-frequency noise. But it is no better at locating the precise signal frequency than the DWT. Again, even for perfect time spacing [Fig. 6(a)] the peak frequency is too low (peak period too high), and edge effects are visible nearly

three periods from the edge, both in peak period and magnitude of the DWP. For irregular time sampling [Figs. 6(c), 6(e), 6(g)], the performance of the DWP is little different from that of the DWT.

3.4 Discrete Wavelet Amplitude (DWA)

Another problem with the DWT and DWP which is not illustrated by these test data sets is that their response to pure sinusoids is explicitly frequency-dependent. If we have a truly multiperiodic signal, then unless we compensate for the frequency-dependence we can't even tell the *relative* amplitudes of the various components. We can, however, define a

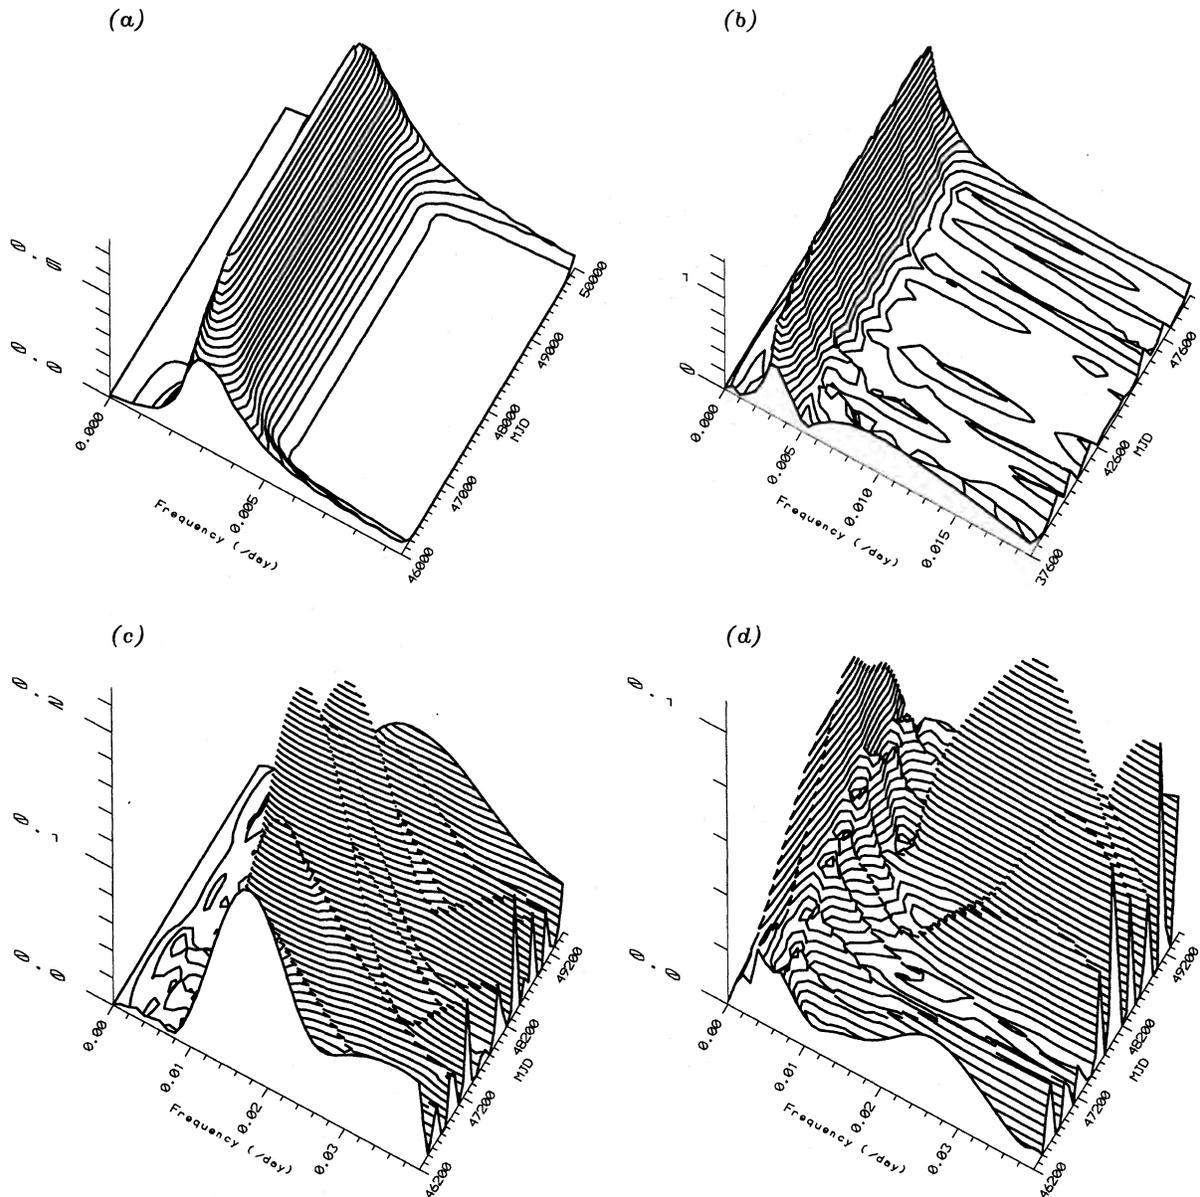


FIG. 4. DWA for test data. (a) Set 1. (b) Set 2. (c) Set 3. (d) Set 4.

function whose response to a pure sinusoid is better behaved. Define the *discrete wavelet amplitude (DWA)* as

$$DWA = \omega \left| \sum_t e^{-c\omega^2(t-\tau)^2} e^{-i\omega(t-\tau)} \right| = \sqrt{\omega} |W(\omega, \tau)|. \quad (3-15)$$

For a sinusoidal signal and ideal time spacing [Fig. 4(a)], the DWA peaks at the signal frequency, with peak value proportional to the amplitude and independent of frequency (except for edge effects). If we use the peak of the DWA (for constant τ) to determine the signal frequency, we get the correct value [Fig. 6(a)]. Unfortunately, for random noise the DWA is not so well behaved; the noise level grows with increasing

frequency. This bad statistical behavior is apparent in the DWA of unevenly sampled test data [Figs. 4(b)–4(d)]; especially in Figs. 4(c) and 4(d) showing many large spurious peaks at high frequencies. If we ignore the spurious peaks, and use the real ones to determine the period and amplitude (Fig. 6), we find that the DWA shows about the same false time evolution as the DWT and DWP.

3.5 Wavelet Amplitude Function (WAF)

Even if we know the exact true frequency of a periodic fluctuation, we cannot get the amplitude of that signal component unless we properly rescale the DWT. For this purpose

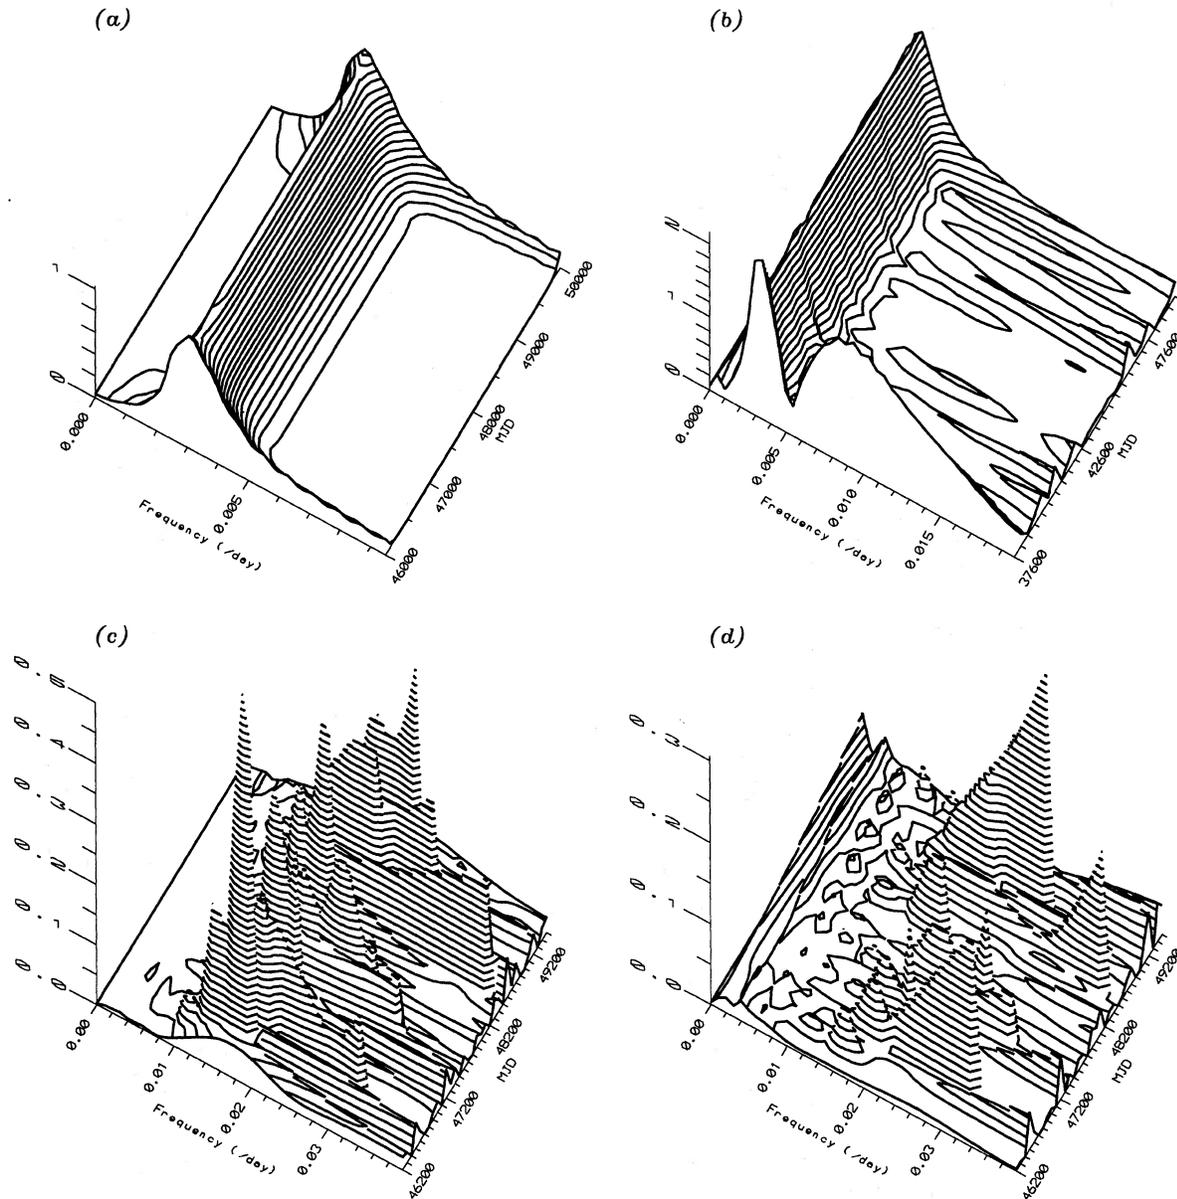


FIG. 5. WAF for test data. (a) Set 1. (b) Set 2. (c) Set 3. (d) Set 4.

we define the *Wavelet Amplitude Function (WAF)*

$$WAF = \frac{|W|}{\rho(\omega, \tau)} \sqrt{\frac{4c\omega}{\pi}} \quad (3-16)$$

If the signal is a pure sinusoid with amplitude A , then for even time spacing, and ignoring edge effects, the WAF is equal to that amplitude at the true signal frequency. Figure 5 shows the WAF for test data sets 1–4. For even time spacing [Fig. 5(a)], the WAF is well-behaved. But for uneven sampling, the WAF turns out to be a good estimate of the amplitude *only* when evaluated at the signal frequency; it is a dreadful statistical measure of the likelihood of periodic fluctuation in the data, or of the location of the signal frequency.

In fact, especially in the gaps between observations, the WAF can take on ridiculously huge values. The tall thin peaks in Figs. 5(c), 5(d) occur not in the observing seasons, but in the large gaps between them. This is essentially due to the shape of the abbreviated Morlet wavelet. Figure 7 shows the real part of the wavelet (1.4), together with data which can fool the WAF. Note that in order for the wavelet function to have any appreciable value at its extremities (where the data are to be found), it must be immensely tall in the gap between data clusters—hence the enormous value of the WAF.

This is indicative of a more general problem with the DWT, namely that the shape of the abbreviated Morlet

wavelet does not closely resemble the signal shape we're looking for. Because the wavelet transform in many ways behaves as a *projection*, we are projecting onto essentially the wrong trial functions. To search for subtle changes in period and amplitude of an oscillation, we should use trial functions which do not decay with time. Even if we are looking for transient fluctuations, the Morlet wavelet does not resemble the sought-after signal unless we fine-tune the constant c to mimic the correct decay rate (in which case we should treat c not as a constant but as a variable parameter).

Hence we must bear in mind that the WAF is to be used only to find the amplitude of a periodic fluctuation, after we have determined the period by some other method, and even at that, it is only to be trusted in the observing seasons. We might, for instance, locate the peak of the DWT, DWP, or DWA to determine the period, then use the WAF to find its amplitude. Figure 6 shows the time evolution of the period and amplitude of our test data (which of course have constant period and amplitude), when determining the period by DWT, DWP, or DWA, and the amplitude from the WAF. Clearly, for bad time sampling all these wavelet transforms are unsatisfactory for tracking either period or amplitude.

4. DATE COMPENSATION BY PROJECTION

The poor performance of these transforms reflects the fact that for uneven time sampling, the expression (3.4) is a naive approximation to the continuous integral (1.1). Furthermore, we should bear in mind that the continuous integral (1.1) itself is only one (of many) means to an end; what we really seek is to discover whether our analyzing wavelet, when optimally scaled in time and magnitude, bears any statistically significant resemblance to the data. If so, we can conclude that it also bears a resemblance to the *signal*, in addition to concluding that a signal actually exists, i.e., the data are not simply noise.

We must, therefore, undertake *date compensation*, i.e., find the optimal match of analyzing wavelet to data, accounting for the times of observation. Adopting the usual definition of goodness-of-fit in terms of a sum of squared residuals, the direct solution is by least-squared regression. An equivalent approach is to treat the analysis as a *projection* (Foster 1996a, 1996b). In my opinion, this method has many benefits; it gives us powerful mathematical tools to apply, a very concise notation, and a good intuitive insight into the operations we apply.

For Fourier analysis, for example, we perform a projection onto the two *trial functions* $\sin(\omega t)$ and $\cos(\omega t)$. This leads to the *modified periodogram* (Scargle 1982; Foster 1996a), and essentially solves the problem that with irregular time spacing, the Fourier transform is no longer properly normalized. We reap further improvement by including a third trial function, the *constant function* $\mathbf{1}(t)=1$ for all t , which produces the *date-compensated discrete Fourier transform (DCDFT)* (Ferraz-Mello 1981; Foster 1996a). It solves the problem that with irregular time spacing, the functions sine and cosine no longer have average value zero, causing an error in determining the zero point of the signal.

It is difficult to overemphasize the importance of includ-

ing the constant function in the analysis. Not doing so amounts to assuming that either the zero point of the signal is the data average (when for irregularly sampled data they are almost always significantly different), or that the trial functions all have average value zero (likewise a rare occurrence). Even with lots of data and a very long time span, this can still be a *major* problem (Foster 1995). Wavelets are *especially* vulnerable to this; for high frequencies the mother wavelet narrows, we effectively sample fewer data points, so the likelihood rises dramatically that the Morlet wavelet no longer has average value zero, and our data average deviates from the signal's zero point.

Projection computes the coefficients y_a of a set of r *trial functions* $\phi_a(t), a=1,2,\dots,r$ for which the *model function*

$$y(t) = \sum_a y_a \phi_a(t), \quad (4-1)$$

best fits the data (in the sense that it minimizes the sum of the squared residuals). We can compute the projection by defining the *inner product* of two functions $f(t)$ and $g(t)$ as

$$\langle f|g \rangle = \frac{\sum_{\alpha=1}^N w_{\alpha} f(t_{\alpha}) g(t_{\alpha})}{\sum_{\beta=1}^N w_{\beta}}, \quad (4-2)$$

where w_{α} is the statistical weight assigned to data point α (in most cases the data points have equal statistical weights). We then compute the *S-matrix*, which is the matrix of inner products of the trial functions

$$S_{ab} = \langle \phi_a | \phi_b \rangle. \quad (4-3)$$

We determine the best-fit coefficients of our trial functions by multiplying the inverse of the *S-matrix* by the vector of inner products of the trial functions with the data

$$y_a = \sum_b S_{ab}^{-1} \langle \phi_b | x \rangle. \quad (4-4)$$

For the Fourier transform, these coefficients will be different from those determined by the simple DFT; they are much better estimates of the actual amplitude of a physical fluctuation.

We define the *power* (which we use to evaluate the projection statistically) not as the squared modulus of a projection as in (3.2), but by (Foster 1996a)

$$P = \frac{N}{(r-1)s^2} \left(\sum_{a,b} S_{ab}^{-1} \langle \phi_a | x \rangle \langle \phi_b | x \rangle - \langle \mathbf{1} | x \rangle^2 \right), \quad (4-5)$$

where N is the number, and s^2 the estimated variance of the data. P is a chi-square statistic with $r-1$ degrees of freedom and expected value 1 (a more usual scaling for a chi-square statistic is $(r-1)P$, having expected value $=r-1$).

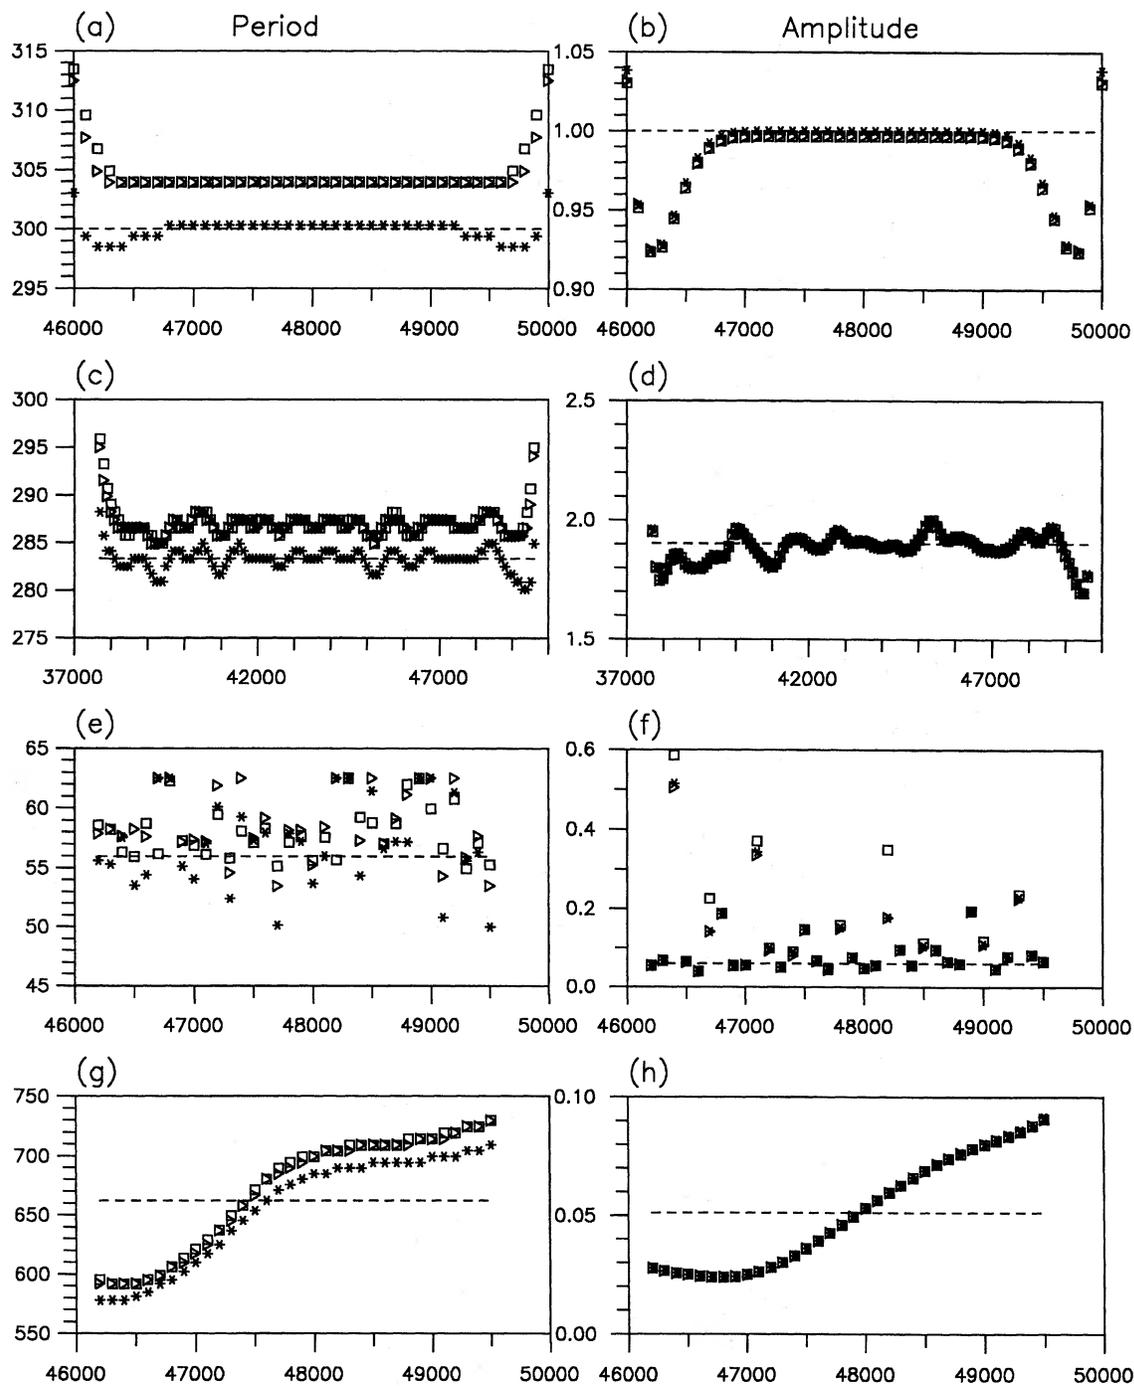


FIG. 6. Period from maximum of the DWT (triangles), DWP (squares), and DWA (asterisks), with amplitude determined by the WAF at those frequencies (plotted using the same symbols). Dashed lines: actual periods, amplitudes. (a) Set 1, period. (b) Set 1, amplitude. (c) Set 2, period. (d) Set 2, amplitude. (e) Set 3, period. (f) Set 3, amplitude. (g) Set 4, period. (h) Set 4, amplitude.

The natural choice would be to use as trial functions the constant function

$$\phi_1(t) = \mathbf{1}(t) = 1, \quad (4-6)$$

and the wavelet functions

$$\phi_2(t) = e^{-c\omega^2(t-\tau)^2} \cos(\omega t), \quad (4-7)$$

$$\phi_3(t) = e^{-c\omega^2(t-\tau)^2} \sin(\omega t). \quad (4-8)$$

However, we will still have the problem illustrated by Fig. 7;

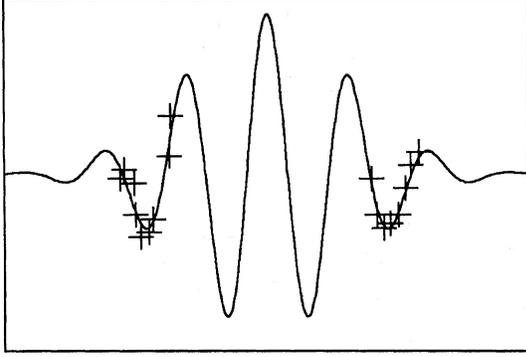


FIG. 7. Real part of the abbreviated Morlet wavelet (solid line), together with data which can cause a high value of the WAF (plus signs).

essentially, we are projecting onto trial functions which do not resemble the signal shape we hope to detect. We would prefer to use trial functions which do not decay with time, and to keep the decay term $e^{-c\tau^2}$ so we can focus on a limited time span (which is why we adopted wavelets in the first place). Fortunately, there is a way to do both.

5. WEIGHTED PROJECTION

When treating the wavelet transform (3.4) as a projection, we have *two* possible interpretations. One is that it is a projection onto the complex trial function

$$\phi(t) = e^{i\omega(t-\tau) - c\omega^2(t-\tau)^2}. \quad (5-1)$$

But we can also adopt the viewpoint that it is a *weighted* projection onto the trial function

$$\phi(t) = e^{i\omega(t-\tau)}, \quad (5-2)$$

with statistical *weights* chosen as

$$w_\alpha = e^{-c\omega^2(t_\alpha - \tau)^2}. \quad (5-3)$$

Although (3.4) is the same in both cases, the projection coefficients (4.4), statistical significance (4.5), and physical parameters determined by the two approaches will be very different.

By using the weighted definition of the inner product (4.2), we can compute the projection according to the prescription of Sec. 4. A modification imposed by a weighted projection is that in the definition of the power (4.5) we must replace the number of data points N by the *effective number* (Foster 1996a, 1996b)

$$N_{\text{eff}} = \frac{(\sum w_\alpha)^2}{(\sum w_\alpha^2)} = \frac{[\sum e^{-c\omega^2(t-\tau)^2}]^2}{\sum e^{-2c\omega^2(t-\tau)^2}} = \frac{n^2(\omega, \tau)}{n(\sqrt{2}\omega, \tau)}. \quad (5-4)$$

5.1 Weighted Wavelet Transform (WWT)

Therefore we perform a weighted projection onto the three trial functions

$$\phi_1(t) = \mathbf{1}(t), \quad (5-5)$$

$$\phi_2(t) = \cos(\omega(t-\tau)), \quad (5-6)$$

$$\phi_3(t) = \sin(\omega(t-\tau)), \quad (5-7)$$

with statistical weights assigned according to (5.3). We evaluate the response by computing the power according to equation (4.5), using the effective number of data points (5.4). However, in (4.5) for the variance in the dominator we use the *weighted* estimated variance

$$s_w^2 = \frac{N_{\text{eff}} V_x}{N_{\text{eff}} - 1}, \quad (5-8)$$

where V_x is the *weighted variation* of the data

$$V_x = \frac{\sum_\alpha w_\alpha x^2(t_\alpha)}{\sum_\lambda w_\lambda} - \left[\frac{\sum_\alpha w_\alpha x(t_\alpha)}{\sum_\lambda w_\lambda} \right]^2 = \langle x|x \rangle - \langle \mathbf{1}|x \rangle^2. \quad (5-9)$$

We can define the weighted variation of the model function similarly

$$V_y = \frac{\sum_\alpha w_\alpha y^2(t_\alpha)}{\sum_\lambda w_\lambda} - \left[\frac{\sum_\alpha w_\alpha y(t_\alpha)}{\sum_\lambda w_\lambda} \right]^2 = \langle y|y \rangle - \langle \mathbf{1}|y \rangle^2. \quad (5-10)$$

We can then define the *weighted wavelet transform* (WWT) as

$$WWT = \frac{(N_{\text{eff}} - 1)V_y}{2V_x}. \quad (5-11)$$

For fixed parameters ω and τ , the WWT may be treated as a chi-square statistic with two degrees of freedom and expected value 1.

5.2 Weighted Wavelet Z-Transform (WWZ)

Of course, the above expected values for the WWT only hold when the data are random noise. For a sinusoidal signal, we might expect the WWT to peak at the signal frequency, but instead it tends to peak at a lower frequency, due to an effect which strongly impacts the WWT. At lower frequencies, our “window” is wider, so we effectively sample more data points; the effective number N_{eff} is larger. This can make the WWT increase with decreasing ω , even though the fit to the data is poorer, simply because the factor N_{eff} in (5.11) increases. We would get a better estimate of the frequency of a significant peak if we had a test statistic which was less sensitive to the effective number of data. Fortunately, there is such a statistic for projections: we apply the Z-statistic of Foster (1996a) which we dub the *weighted wavelet Z-transform* (WWZ)

$$Z = \frac{(N_{\text{eff}} - 3)V_y}{2(V_x - V_y)}. \quad (5-12)$$

It follows the F-distribution with $N_{\text{eff}} - 3$ and 2 degrees of freedom, and expected value 1.

The last sentence is nonsense; because the effective number N_{eff} is not an integer, there is no F-distribution with $N_{\text{eff}} - 3$ and 2 degrees of freedom (although we can construct it by analytic continuation). The statistical behavior (5.12) is derived for a projection where the statistical weights are inversely proportional to the variances of the data, not chosen for convenience of the analysis. The implication is that this purely theoretical approach to estimating the statistical be-

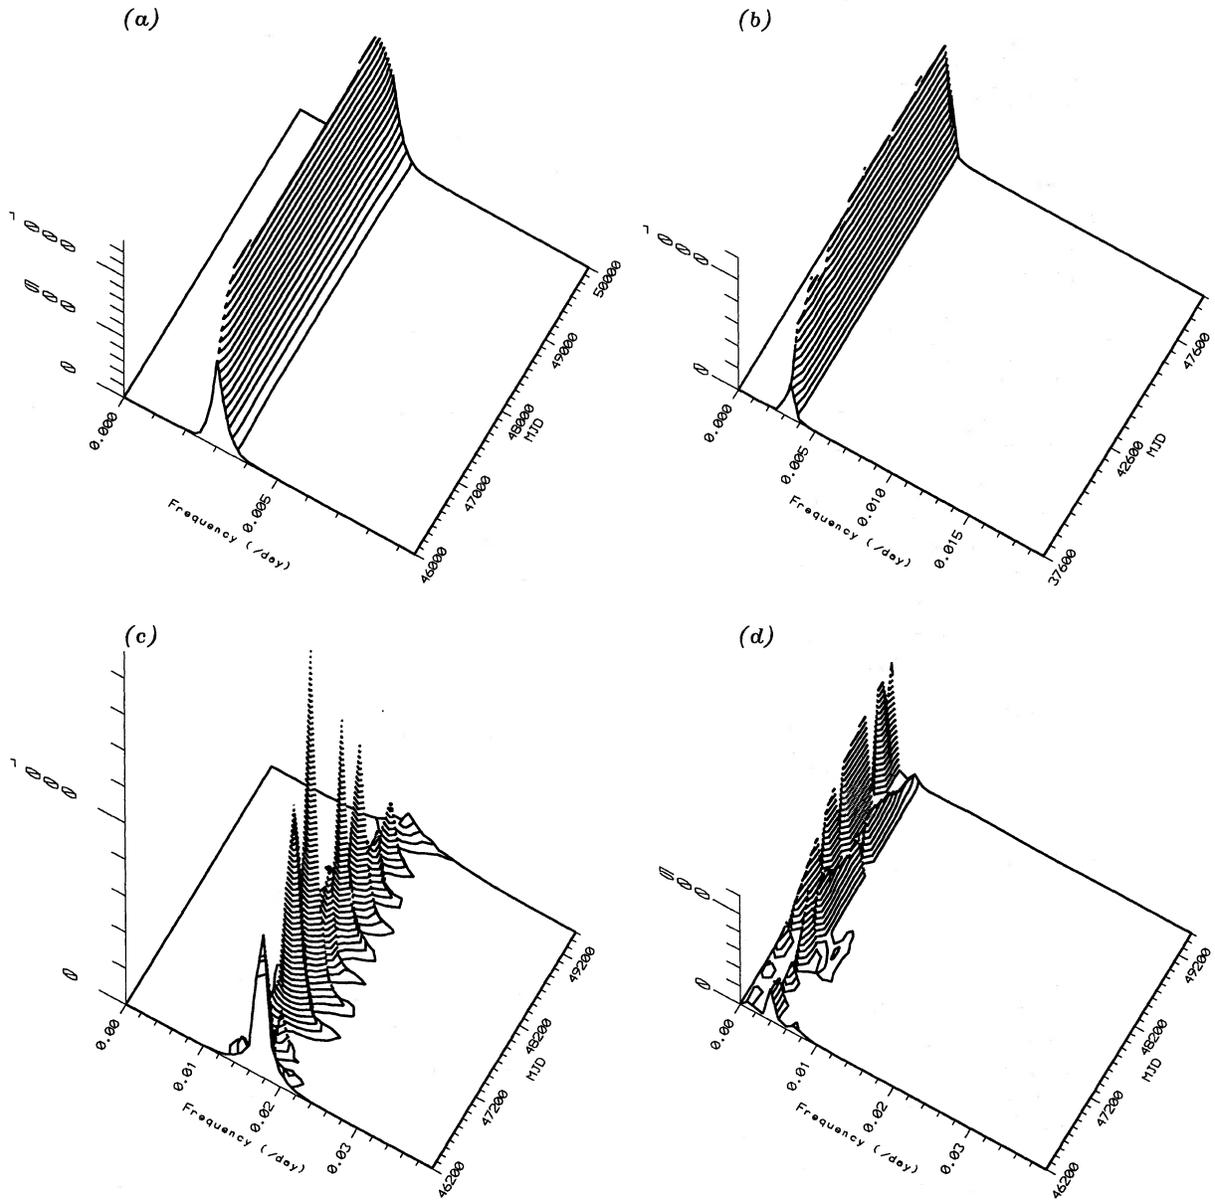


FIG. 8. WWZ for test data. (a) Set 1. (b) Set 2. (c) Set 3. (d) Set 4.

havior of the WWZ calls for a heavy dose of caution (see also Sec. 6).

Figure 8 shows the WWZ for our test data. Except for data set 4 [Fig. 8(d)], which shows a false overtone frequency, spurious peaks have been entirely eliminated. Considering the abysmal time sampling (especially for test data sets 3 and 4), this is a powerful testimony to the effectiveness of the WWZ. If we use the peak of the WWZ (for constant τ) to determine the period [Figs. 9(a), 9(c), 9(e), 9(g)], we find stunning agreement with the true period (Fig. 9 is plotted on the exact same scale as Fig. 6, for direct comparison of the WWZ with the DWT, DWP, and DWA).

If we take the ratio of the WWZ to the WWT, we get

$$\frac{Z}{(WWZ)} = \frac{(N_{\text{eff}} - 3)V_x}{(N_{\text{eff}} - 1)(V_x - V_y)} = \frac{(N_{\text{eff}} - 3)}{(N_{\text{eff}} - 1)}(1 + (S/N)), \quad (5-13)$$

where S/N is the estimated signal-to-noise ratio. When S/N is much less than 1, the WWZ therefore is little different from the WWT, even in the limit of large data size N . Therefore at low S/N , the WWZ shares the drawbacks of the WWT. But for S/N even moderate (≥ 1), the WWZ is a significant improvement on the WWT.

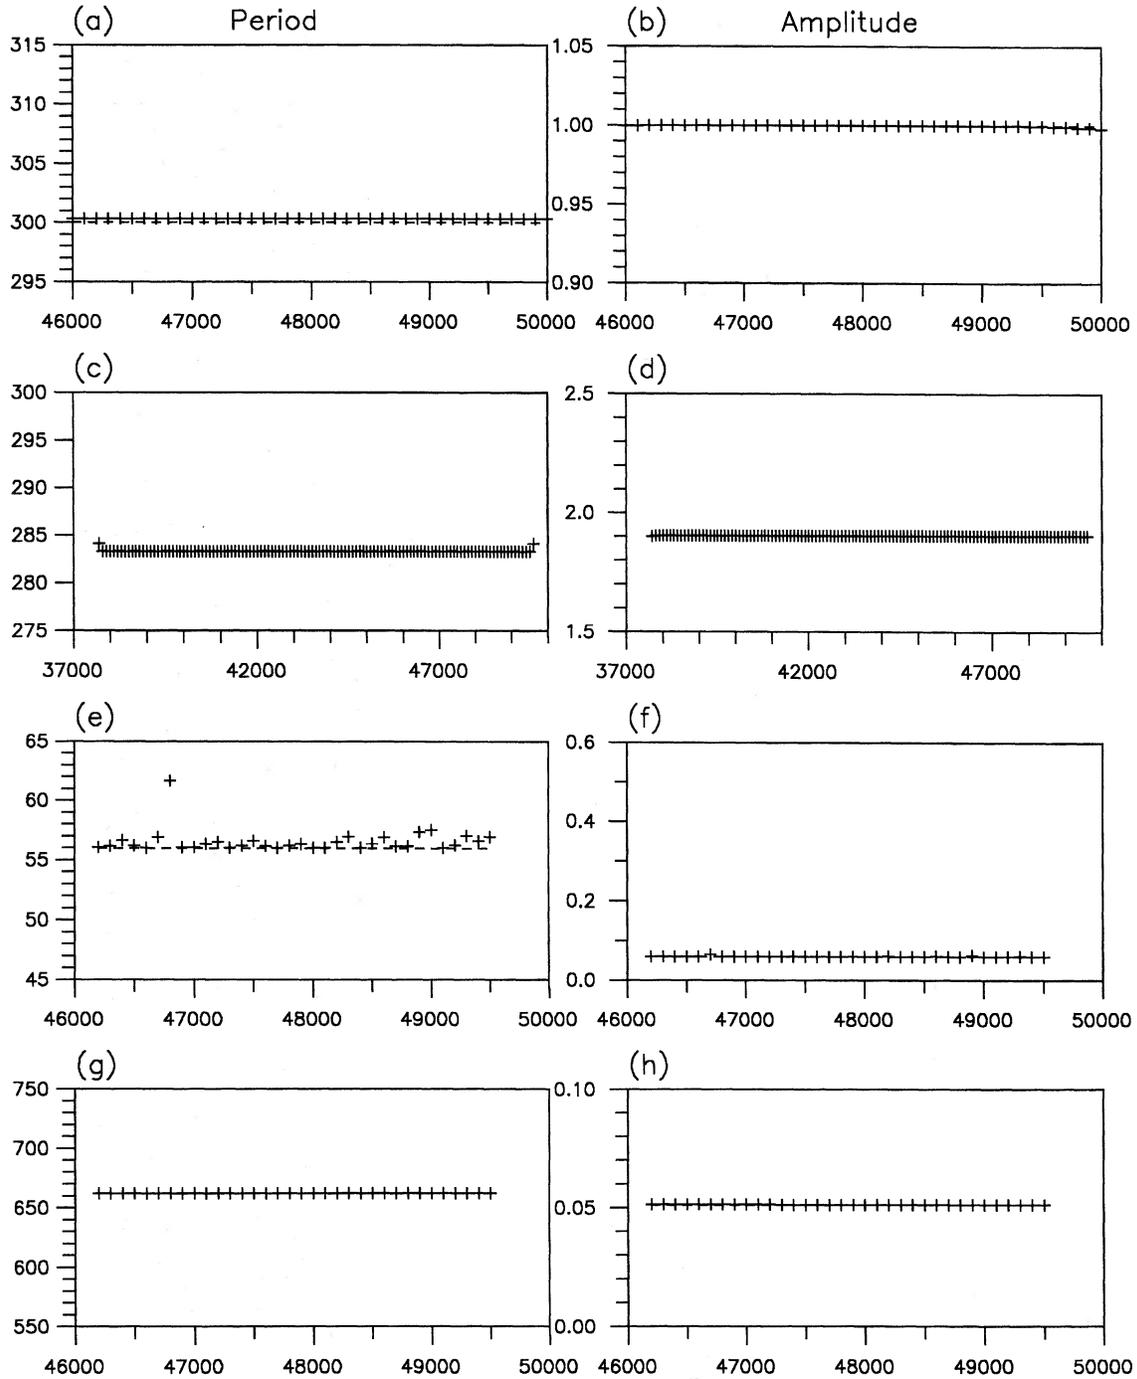


FIG. 9. Period from peak of the WWZ, amplitude from the WWA. The dashed lines, indicating the actual periods and amplitudes, are invisible in many of the graphs because the WWZ/WWA estimates are nearly identical to the true values. (a) Set 1, period. (b) Set 1, amplitude. (c) Set 2, period. (d) Set 2, amplitude. (e) Set 3, period. (f) Set 3, amplitude. (g) Set 4, period. (h) Set 4, amplitude.

5.3 Weighted Wavelet Amplitude (WWA)

Although the WWZ may be an excellent locator of the signal frequency, it is a poor measure of amplitude. With a projection it is easy to define the amplitude of the corresponding periodic fluctuation: it is the square root of the sum

of the squares of the expansion coefficients for the sine and cosine functions, which I call the *weighted wavelet amplitude (WWA)*

$$WWA = \sqrt{(y_2)^2 + (y_3)^2}. \tag{5-14}$$

Therefore the weighted wavelet transform enables us to

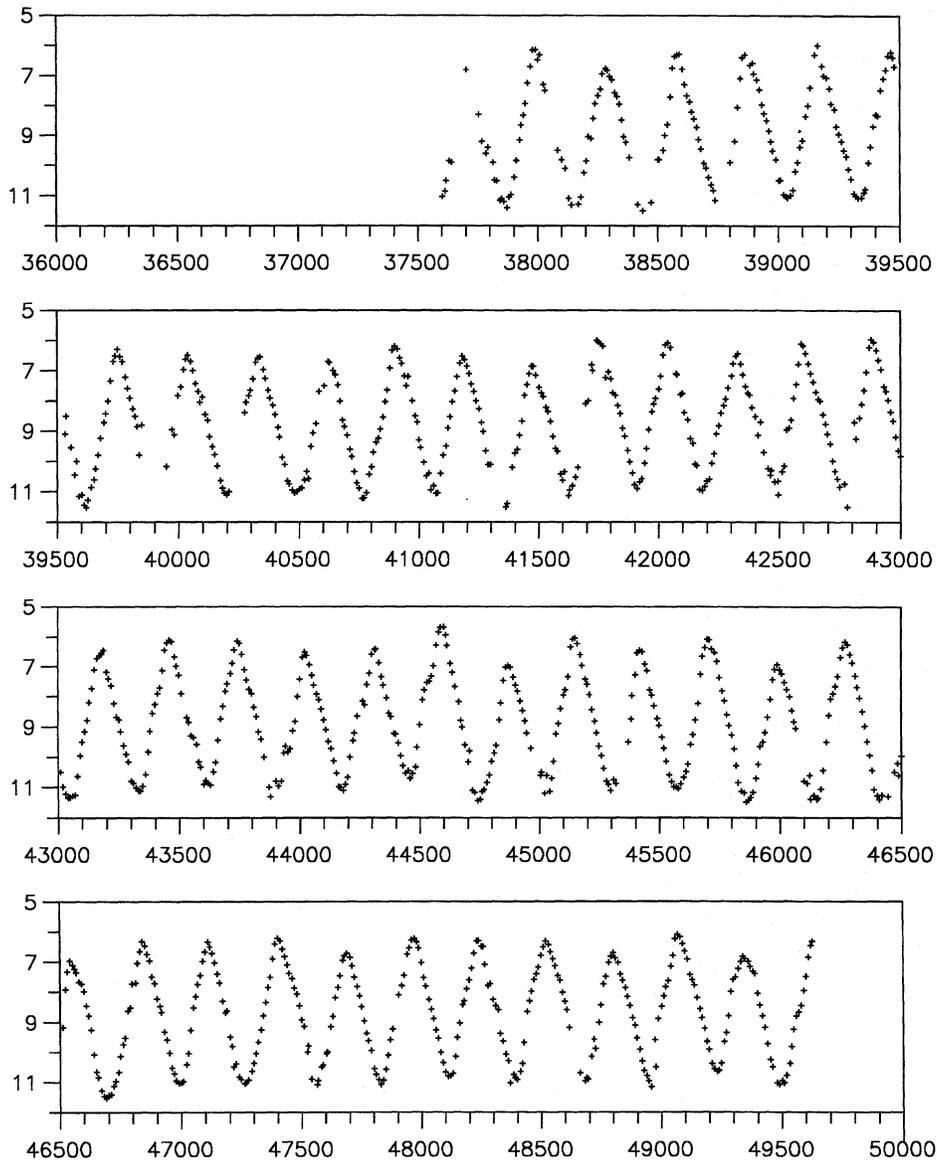


FIG. 10. Visual light curve of R Aquilae.

compute the amplitude by the WWA, after we have determined the period from the WWZ. Figures 9(b), 9(d), 9(f), 9(h) show the time evolution of amplitude, as determined by the WWZ/WWA, for the test data sets. It is readily apparent that this method is vastly superior to its wavelet rivals (Fig. 6) for identifying the true nature of these test signals, despite their highly undesirable time samplings. Hence with the WWZ/WWA combination, we have a wavelet analysis technique which performs well for pure signals, even with very bad time sampling. We may therefore have a new and more powerful weapon in the battle against unevenly spaced data.

6. FORMAL ERRORS

The problem of computing the formal errors of period and amplitude as computed by the WWZ/WWA, is an extraordi-

narily complex one. However, an estimate is available (Foster 1996a, 1996b). Having located the peak frequency ω , and determined the constant, cosine, and sine coefficients y_1, y_2, y_3 , we define *four* trial functions ϕ_A , $A = 1, 2, 3, 4$ as

$$\phi_1(t) = 1(t), \quad (6-1)$$

$$\phi_2(t) = \cos(\omega(t - \tau)), \quad (6-2)$$

$$\phi_3(t) = \sin(\omega(t - \tau)), \quad (6-3)$$

$$\phi_4(t) = (t - \tau)[y_3 \cos(\omega(t - \tau)) - y_2 \sin(\omega(t - \tau))]. \quad (6-4)$$

We then compute the *super-S-matrix*, which is the matrix of inner products of these trial functions

$$\Omega_{AB} = \langle \phi_A | \phi_B \rangle, \quad (6-5)$$

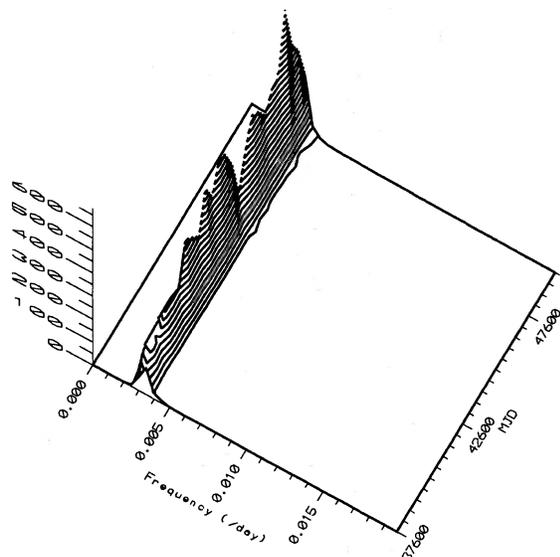


FIG. 11. WWZ of R Aql.

using the weighted inner product. We also compute the estimated residual variance

$$s_{\text{res}}^2 = \frac{N_{\text{eff}}(V_x - V_y)}{N_{\text{eff}} - 3}. \quad (6-6)$$

Then the variances of the projection coefficients are proportional to the diagonal elements of the inverse of the super- S -matrix

$$\text{Var}(y_A) = N_{\text{eff}}^{-1} \sigma_{\text{res}}^2 \Omega_{AA}^{-1}, \quad (6-7)$$

and the variance of the estimated frequency is given by the same equation, letting the index $A=4$. If, for example, this procedure is applied to Fourier analysis with even time sampling, it reproduces the usual formulas for uncertainty in frequency and amplitude (Foster 1996a).

However, the purely formal errors so computed are not to be trusted for the WWZ, for several important reasons. First, the WWZ (like Fourier analysis) is a *parametric* projection, i.e., the trial functions depend on a variable parameter, the frequency. In this case the exact statistical behavior cannot be deduced (the same is true for Fourier analysis of unevenly sampled data, Foster 1996a). Second, (6.7) was derived for a standard projection, in which the statistical weights are inversely proportional to the variances of the data. For a weighted projection, (6.7) is only an approximation, a good one if the effective number N_{eff} is large, not so good if it is small. Third, the entire derivation is predicated on the usual null hypothesis, that the data are a pure sinusoidal signal with constant frequency and amplitude, plus random noise. We know that this null hypothesis is false; if the signal were a pure sinusoid, we would not need wavelet analysis. In light of these difficulties, perhaps the most trustworthy method to

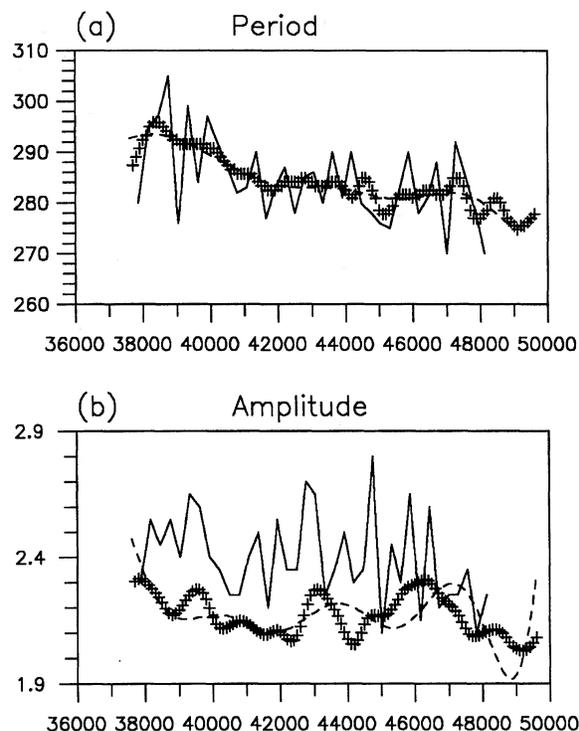


FIG. 12. Period and amplitude of R Aql, as determined by the WWZ (plus signs), CLEANEST Fourier spectrum (dashed line), and times and magnitudes of maxima and minima (solid line). (a) Period. (b) Amplitude.

estimate uncertainties for the coefficients and parameters would be Monte Carlo simulations.

7. R AQUILAE

Figure 10 shows the visual light curve of R Aquilae. The most obvious feature is a consistent oscillation with period, according to the GCVS, of 284.2 days. The GCVS also lists the range as 5.5 to 12.0. Fourier analysis using the date-compensated discrete Fourier transform (DCDFT, Ferraz-Mello 1981) identifies the oscillation as having period 283.3 days, amplitude 1.9011 mag. However, the period of R Aql is clearly not constant. Use of the CLEANEST Fourier spectrum, together with the method of complex amplitude reconstruction (Foster 1995), establishes that the period exhibits a trend, and provides estimates of the period and amplitude as functions of time. In addition, we have direct estimates of the times and magnitudes of maxima and minima for R Aql (Campbell 1926, Campbell 1955, Mattei *et al.* 1990), enabling us to compute individual periods as the time from one maximum to the next, and amplitudes as the magnitude change from maximum to minimum, or from minimum to maximum.

Figure 12 compares the period and amplitude of R Aql as determined by the times and magnitudes of maxima and minima, to those determined by the CLEANEST spectrum. The agreement between period from these two methods is as

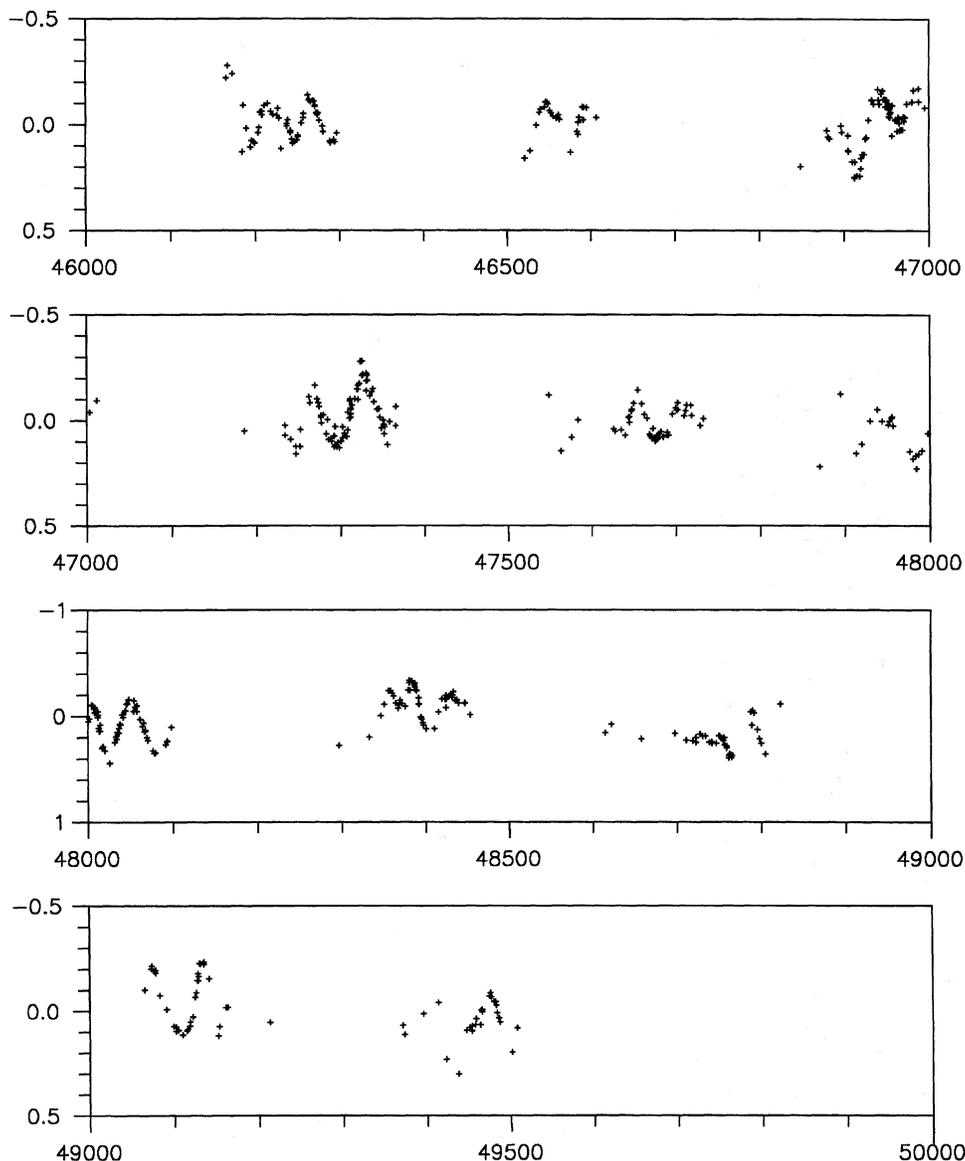


FIG. 13. Photoelectric photometry of FS Comae.

good as could be expected [Fig. 12(a)]. We see that the periods determined by the MAX/min data exhibit considerable scatter. Much of the scatter is due to errors in determining the times of maxima; however, there are genuine cycle-to-cycle fluctuations. Regardless of the scatter, we have a good picture of the time evolution of the period.

Unfortunately, the two methods do *not* agree on amplitude; the amplitudes from CLEANEST are consistently about 0.2 magnitudes less than those determined from MAX/min magnitudes. Close inspection of the visual light curve reveals the reason; the shape of the oscillation is distinctly non-sinusoidal, with very sharp maxima and minima. Thus the amplitude, as defined by half the difference between maximum and minimum brightness, will necessarily be consistently greater than the amplitude as defined by the size of

a best-fit sinusoid. Even accounting for this offset due to different definitions of amplitude, we see no real correlation between the amplitudes from MAX/min and those from CLEANEST. It is difficult to trust CLEANEST when so much experience demonstrates that Fourier methods in general are more precise indicators of period than of amplitude, and my experience is that CLEANEST is a superior tracker of period evolution but a poorer indicator of amplitude evolution. It is also difficult to trust the MAX/min amplitudes, due to significant scatter in the determination of magnitudes at maximum and minimum, although in my opinion they are a better indicator of amplitude evolution than CLEANEST. We expect the WWZ to yield amplitudes on the same scale as CLEANEST, because wavelet methods are, after all, cousin to Fourier analysis.

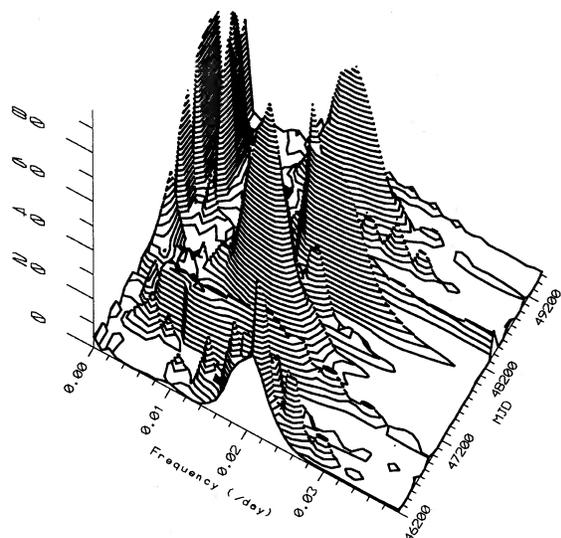


FIG. 14. WWZ of FS Comae.

Figure 11 shows the WWZ computed from the R Aql data. Using the peak, for a constant location τ , to determine the period and amplitude (Fig. 12), we see that the WWZ detects essentially the same trend in period as indicated by CLEANEST and MAX/min. The small deviations of the WWZ period from that of CLEANEST correlate with the MAX/min periods; it seems that the WWZ is better at tracking period evolution than CLEANEST. Figure 12(b) compares the amplitude from the WWA, to that from CLEANEST and the amplitudes from MAX/min. We see only weak correlation with the amplitude as determined by CLEANEST, but a hint of correlation with the amplitude from MAX/min. The peaks in WWA amplitude, for instance, at MJD 39500 and MJD 43100, correspond (roughly) to the occurrence of a consecutive pair of high-amplitude cycles. Still, the agreement is rough at best, and considering the errors inherent in all the methods, it is difficult to choose which one to believe.

8. FS COMAE

In Fig. 13 we see the photoelectric light curve of FS Comae. The (approximately) 58-day period is plainly visible, but it is also evident that the period and amplitude of this fluctuation are highly variable. Furthermore, the annual gaps in the data are much longer than this period; hence we *cannot* expect to draw reliable conclusions about the main fluctuation during these unobserved gaps. CLEANEST analysis does not paint a realistic picture of the time evolution of period and amplitude for FS Com; neither are MAX/min data available to establish the true behavior of this star. Therefore I have analyzed each year's data separately by the DCDFT, to establish the *average* period and amplitude for each observing season. Nonetheless, these estimates are far from ideal, since in several cases significant changes in period and (especially) amplitude are visible within a single observing season.

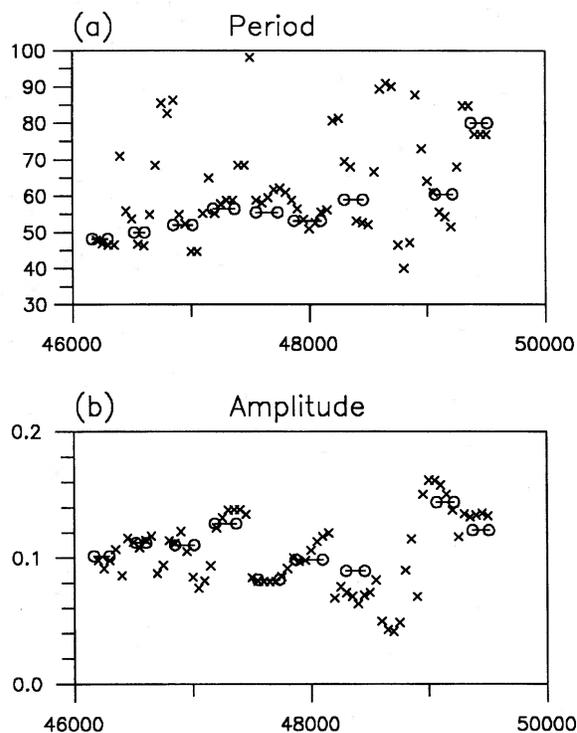


FIG. 15. Period and amplitude of FS Com, as determined by the WWZ (x's), and Fourier analysis of each observing season (circles joined by solid line). (a) Period. (b) Amplitude.

Figure 14 shows the WWZ of FS Com. The main periodicity is evident, but it is also clear that the period is not constant. Using the WWZ to estimate the period and amplitude of the main fluctuation, we get highly variable results (Fig. 15). However, we note that there is excellent agreement between the WWZ and DCDFT, both for period and amplitude, *during the observing seasons*. In fact the WWZ is to be preferred, since it allows for changing period and amplitude within a single observing season. In the unobserved gaps, the amplitude and (especially) period show large fluctuations, confirming what we knew all along: that we cannot expect to draw conclusions about the period or amplitude of a cycle when the gap in the data is longer than the period of fluctuation.

9. CONCLUSION

The wavelet transform is known to be an excellent period analysis method for evenly sampled time series. For uneven sampling, although the DWT is very susceptible to false fluctuations, the WWZ is highly resistant to such spurious response. Our preliminary investigation indicates that it performs better than other methods for determining period and amplitude of unevenly sampled time series. Indeed, considerable practical experience between the first writing of this paper and its revision, bears out the remarkable utility of the WWZ/WWA with only modestly good data density and signal-to-noise ratio.

Furthermore, it has a significant advantage over such methods as CLEANEST and estimating maxima and minima: it is very simple and fast to apply. In particular, it is much faster than CLEANEST, which can take a very long time due to a large number of iterations. Despite its superb performance most of the time, these investigations confirm that any attempt to quantify period and amplitude during a gap in the data which is longer than the period in question, is folly.

The WWZ/WWA turns out to be especially good at discerning the time evolution of period and amplitude from visual observations of long-period variables (LPVs). Most LPVs for which AAVSO has visual data have moderate to high signal-to-noise ratio, good but uneven data density (often with sizeable periodic gaps), and exhibit considerable fluctuations in their period and amplitude from one cycle to the next. In fact, recent results indicate that for LPVs with good data coverage, there is advantage in experimenting with lower values for the decay constant c , so as to focus attention on not a small number of cycles, but a single cycle.

One other thing is clear: the wavelet transform in general is a remarkably versatile tool for extracting information from time series, even with uneven time sampling. The Morlet wavelet is only one of many choices, and it is likely that several different techniques for compensating an irregular time sampling remain to be discovered. This paper has explored only one tiny aspect of the application of wavelets to time series; we can look forward to much ingenuity yet to come in this field.

I am very grateful to Dr. Janet A. Mattei, who first pointed me in the direction of wavelets. I am also most grateful to the army of amateur variable star observers, whose unevenly sampled time series have frustrated many analysis methods while acting as a spur to much invention on my part.

APPENDIX A: TENSOR FORMULATION OF THE DWT

We will treat the wavelet transform as a projection, according to the tensor treatment of Foster 1996b. I caution the reader **strongly** that the notation used in these appendixes is markedly different from that of the main body of the paper. A particular source of possible confusion is the fact that some symbols have one meaning in the body of the paper, and another quite different meaning in the appendixes. For example, the symbols y_a , in the body of the paper, refer to the expansion coefficients for the trial functions, as in equation (4-1). But, in the appendixes they refer to the *covariant* components of the projected data vector; the expansion coefficients are given by the *contravariant* components, which

in the notation of the appendixes is y^a . There are many such pitfalls in the change of notation, but the notation of the appendix is a *natural* adaptation to tensor analysis of the notation in the main body of the paper. I strongly urge those interested in these appendixes to consult Foster 1996b for details of the notation and conventions used from here on out.

Hence the data $x(t_\alpha)$, $\alpha=1,2,\dots,N$ define an N -dimensional *contravariant vector* in *sampling space* as

$$x^\alpha = [x(t_1), x(t_2), \dots, x(t_N)]. \quad (\text{A1})$$

We give geometrical structure to this vector space by defining the *metric tensor* as

$$g_{\alpha\beta} = \begin{cases} w_\alpha / (\sum w) & \text{if } \alpha = \beta \\ 0 & \text{otherwise} \end{cases}, \quad (\text{A2})$$

where w_α is the statistical weight applied to data point α . In most cases the data have equal statistical weights, in which case the metric tensor is

$$g_{\alpha\beta} = N^{-1} \delta_{\alpha\beta}, \quad (\text{A3})$$

where N is the number of data points. We will adopt the *Einstein summation convention*, that any repeated index is to be summed over all possible values.

The DWT corresponds (roughly) to a projection onto the two trial functions (real and imaginary parts of the Morlet wavelet) $\phi_1(t) = e^{-c\omega^2(t-\tau)^2} \cos(\omega(t-\tau))$ and $\phi_2(t) = e^{-c\omega^2(t-\tau)^2} \sin(\omega(t-\tau))$. Therefore we have the two trial *vectors*

$$\phi_1^\alpha = e^{-c\omega^2(t_\alpha-\tau)^2} \cos(\omega(t_\alpha-\tau)), \quad (\text{A4})$$

$$\phi_2^\alpha = e^{-c\omega^2(t_\alpha-\tau)^2} \sin(\omega(t_\alpha-\tau)). \quad (\text{A5})$$

The inner products of the trial vectors with the data vector give the *covariant* projection coefficients

$$y_a = \langle \phi_a | x \rangle = \phi_a \alpha x^\alpha = N^{-1} \sum_\alpha \phi_a(t_\alpha) x(t_\alpha). \quad (\text{A6})$$

Note that because we are now using tensor notation, these are *not* the same as the expansion coefficients for the model function defined in (4.1). The *S-matrix* is the matrix of inner products of the trial functions

$$S_{ab} = \langle \phi_a | \phi_b \rangle = N^{-1} \sum_\alpha e^{-2c\omega^2(t_\alpha-\tau)^2} \begin{bmatrix} \cos^2(\omega(t_\alpha-\tau)) & \sin(\omega(t_\alpha-\tau))\cos(\omega(t_\alpha-\tau)) \\ \sin(\omega(t_\alpha-\tau))\cos(\omega(t_\alpha-\tau)) & \sin^2(\omega(t_\alpha-\tau)) \end{bmatrix}, \quad (\text{A7})$$

and is the *covariant* form of the metric tensor for the subspace spanned by the trial functions. It can be put in the form

$$S_{ab} = \frac{1}{2} N^{-1} \sum_{\alpha} e^{-2c\omega^2(t_{\alpha}-\tau)^2} \begin{bmatrix} 1 + \cos(2\omega(t_{\alpha}-\tau)) & \sin(2\omega(t_{\alpha}-\tau)) \\ \sin(2\omega(t_{\alpha}-\tau)) & 1 - \cos(2\omega(t_{\alpha}-\tau)) \end{bmatrix}. \quad (\text{A8})$$

For convenience I will define the *local data number*, which depends on the parameters ω and τ , as

$$n(\omega, \tau) = \sum_{\alpha} e^{-c\omega^2(t_{\alpha}-\tau)^2}, \quad (\text{A9})$$

and the *local data density* as

$$\rho(\omega, \tau) = \omega \sqrt{\frac{c}{\pi}} n(\omega, \tau) = \omega \sqrt{\frac{c}{\pi}} \sum_{\alpha} e^{-c\omega^2(t_{\alpha}-\tau)^2}. \quad (\text{A10})$$

I will also define the *modified cosine* and *sine factors*

$$\tilde{C}(\omega, \tau) = \frac{1}{n(\sqrt{2}\omega, \tau)} \sum_{\alpha} e^{-2c\omega^2(t_{\alpha}-\tau)^2} \cos(2\omega(t_{\alpha}-\tau)), \quad (\text{A11})$$

$$\tilde{B}(\omega, \tau) = \frac{1}{n(\sqrt{2}\omega, \tau)} \sum_{\alpha} e^{-2c\omega^2(t_{\alpha}-\tau)^2} \sin(2\omega(t_{\alpha}-\tau)). \quad (\text{A12})$$

Now the S -matrix may be written

$$S_{ab} = \frac{1}{2N} n(\sqrt{2}\omega, \tau) \begin{bmatrix} 1 + \tilde{C}(\omega, \tau) & \tilde{B}(\omega, \tau) \\ \tilde{B}(\omega, \tau) & 1 - \tilde{C}(\omega, \tau) \end{bmatrix}. \quad (\text{A13})$$

Note that the S -matrix satisfies

$$\delta_{ab} S_{ab} = N^{-1} \sum_{\alpha} e^{-2c\omega^2(t_{\alpha}-\tau)^2} = N^{-1} n(\sqrt{2}\omega, \tau). \quad (\text{A14})$$

Finally, note that the *contravariant* form of the metric is the matrix inverse of the S -matrix

$$S^{ab} = \frac{2N}{[1 - \tilde{C}^2(\omega, \tau) - \tilde{B}^2(\omega, \tau)] n(\sqrt{2}\omega, \tau)} \times \begin{bmatrix} 1 - \tilde{C}(\omega, \tau) & -\tilde{B}(\omega, \tau) \\ -\tilde{B}(\omega, \tau) & 1 + \tilde{C}(\omega, \tau) \end{bmatrix}. \quad (\text{A15})$$

The contravariant form of the S -matrix defines the contravariant projection coefficients in terms of their covariant counterparts

$$y^a = S^{ab} y_b, \quad (\text{A16})$$

and these *are* the expansion coefficients of Eq. (4.1), defining the *model function*

$$y(t) = y^a \phi_a(t). \quad (\text{A17})$$

Of fundamental importance is the *constant vector* $\mathbf{1}^{\alpha}$, defined to take the value 1 everywhere. $\hat{\mathbf{1}}_a$ is the covariant representation of the *projection* of the constant vector $\mathbf{1}^{\alpha}$, i.e.,

$$\begin{aligned} \hat{\mathbf{1}}_a &= N^{-1} \sum_t e^{-c\omega^2(t-\tau)^2} \begin{bmatrix} \cos(\omega(t-\tau)) \\ \sin(\omega(t-\tau)) \end{bmatrix} \\ &= N^{-1} n(\omega, \tau) \begin{bmatrix} C(\omega, \tau) \\ B(\omega, \tau) \end{bmatrix}, \end{aligned} \quad (\text{A18})$$

which serves to define the (*unmodified*) *cosine* and *sine factors*

$$C(\omega, \tau) = \frac{1}{n(\omega, \tau)} \sum_t e^{-c\omega^2(t-\tau)^2} \cos(\omega(t-\tau)), \quad (\text{A19})$$

$$B(\omega, \tau) = \frac{1}{n(\omega, \tau)} \sum_t e^{-c\omega^2(t-\tau)^2} \sin(\omega(t-\tau)). \quad (\text{A20})$$

The wavelet transform is not actually treated as a projection onto the trial vectors (A.4) and (A.5). Instead we define the complex wavelet transform as

$$W = N\omega^{1/2}(y_1 - iy_2), \quad (\text{A21})$$

and we evaluate the strength of the transform by computing its squared norm

$$|W|^2 = N^2 \omega [(y_1)^2 + (y_2)^2] = N^2 \omega \delta_{ab} y_a y_b. \quad (\text{A22})$$

The squared norm can be referred to as the *wavelet power*, while the (unsquared) norm $|W|$ is called the *wavelet modulus*.

APPENDIX B: CONTINUUM LIMIT

As important as it is to understand the quirks of the wavelet transform due to irregular time spacing, we must also establish its behavior for an ideal time spacing. Hence I consider the *continuum limit*, in which we replace \sum_{α} by $\int_{-\infty}^{\infty} \rho dt$ (where ρ is the data density). This corresponds to a large number of data with even time spacing, which we might call the “ideal time sampling limit.” Noting that factors of the form $e^{-c\omega^2(t-\tau)^2}$ have a “width” proportional to ω^{-1} , we see that in the continuum limit the local number n will also be proportional to ω^{-1} (and will be independent of τ). Also in the continuum limit, the sine factors B and \tilde{B} vanish (they become integrals of odd functions), but the cosine factors C and \tilde{C} do not; rather they approach constants which depend only on the constant c , not on the parameters ω and τ . We can evaluate these factors using the identity

$$\int_{-\infty}^{+\infty} e^{-cz^2} \cos(\lambda z) dz = e^{-\lambda^2/4c} \sqrt{\frac{\pi}{c}}. \quad (\text{B1})$$

In the continuum limit we have

$$n \rightarrow \omega^{-1} \rho \sqrt{\frac{\pi}{c}}. \quad (\text{B2})$$

As already said, the sine factors vanish

$$B \rightarrow 0, \quad \tilde{B} \rightarrow 0, \quad (\text{B3})$$

while the cosine factors approach

$$C \rightarrow e^{-1/4c}, \quad \tilde{C} \rightarrow e^{-1/2c}. \quad (\text{B4})$$

APPENDIX C: RESPONSE OF THE DWT TO PURE NOISE

If the data are random noise with mean 0 and variance σ^2 , then the expected value of the product of two projection coefficients y_a and y_b is³ (Foster 1996b)

$$\langle y_a y_b \rangle = N^{-1} \sigma^2 S_{ab}, \quad (\text{C1})$$

so the expected value of the squared norm of the wavelet transform is

$$\begin{aligned} \langle |W|^2 \rangle &= N^2 \omega \delta_{ab} N^{-1} \sigma^2 S_{ab} \\ &= \sigma^2 \omega \sum_t e^{-2c\omega^2(t-\tau)^2} = \sigma^2 \omega n(\sqrt{2}\omega, \tau) \\ &= \sigma^2 \rho(\sqrt{2}\omega, \tau) \sqrt{\frac{\pi}{2c}}. \end{aligned} \quad (\text{C2})$$

It is more usual that the data do *not* have mean value zero. Then we assume that they have unknown mean μ and variance σ^2 , and we subtract the *average* value from the input data, applying the DWT not to the data vector x^α but to the *zeroed* data vector

$$z^\alpha = x^\alpha - \mathbf{1}^\alpha \mathbf{1}_\beta x^\beta. \quad (\text{C3})$$

Then the expected value of the product of two projection coefficients is (Foster 1996b)

$$\langle y_a y_b \rangle = N^{-1} \sigma^2 (S_{ab} - \hat{\mathbf{1}}_a \hat{\mathbf{1}}_b), \quad (\text{C4})$$

so the expected value of the DWT is

$$\langle |W|^2 \rangle = N^2 \omega \delta_{ab} N^{-1} \sigma^2 (S_{ab} - \hat{\mathbf{1}}_a \hat{\mathbf{1}}_b) \quad (\text{C5})$$

$$= N \sigma^2 \omega [N^{-1} n(\sqrt{2}\omega, \tau) - N^{-2} n^2(\omega, \tau) \times [C^2(\omega, \tau) + B^2(\omega, \tau)]] \quad (\text{C6})$$

$$= \sigma^2 \omega [n(\sqrt{2}\omega, \tau) - N^{-1} n^2(\omega, \tau) \times [C^2(\omega, \tau) + B^2(\omega, \tau)]]. \quad (\text{C7})$$

In the continuum limit, these extra terms are negligible, but with highly irregular time sampling they can be quite significant.

APPENDIX D: RESPONSE OF THE DWT TO A PURE SINUSOID

If the data are a pure sinusoid of frequency ω_0 , amplitude A , and phase θ ,

$$x^\alpha = A \cos(\omega_0 t_\alpha + \theta), \quad (\text{D1})$$

then we can compute the wavelet transform most easily by using the complex form

$$x^\alpha = \frac{1}{2} A (e^{+i(\omega_0 t_\alpha + \theta)} + e^{-i(\omega_0 t_\alpha + \theta)}). \quad (\text{D2})$$

Note that x^α defined by (D.2) is most definitely *real*; we cannot simply model a periodic fluctuation as a complex signal $x = A e^{i(\omega_0 t + \theta)}$. The wavelet transform is thus

$$\begin{aligned} W &= \frac{1}{2} A \sqrt{\omega} \sum_t e^{-c\omega^2(t-\tau)^2} [e^{i((\omega_0 - \omega)(t-\tau) + \theta + \omega_0 \tau)} \\ &\quad + e^{-i((\omega_0 + \omega)(t-\tau) + \theta + \omega_0 \tau)}] \\ &= \frac{1}{2} A e^{i(\omega_0 \tau + \theta)} \sqrt{\omega} \sum_t e^{-c\omega^2(t-\tau)^2} [e^{i(\omega_0 - \omega)(t-\tau)} \\ &\quad + e^{-i((\omega_0 + \omega)(t-\tau) + 2\theta + 2\omega_0 \tau)}]. \end{aligned} \quad (\text{D3})$$

Its squared norm is therefore

$$\begin{aligned} |W|^2 &= \frac{1}{4} A^2 \omega \left\{ \left[\sum_t e^{-c\omega^2(t-\tau)^2} (\cos[(\omega - \omega_0)(t-\tau)] \right. \right. \\ &\quad \left. \left. + \cos(2\omega_0 \tau + 2\theta) \cos[(\omega + \omega_0)(t-\tau)]) \right]^2 \right. \\ &\quad \left. + \left[\sum_t e^{-c\omega^2(t-\tau)^2} \sin(2\omega_0 \tau + 2\theta) \sin[(\omega + \omega_0)(t-\tau)] \right]^2 \right\}. \end{aligned} \quad (\text{D4})$$

This motivates me to define the *major* and *minor response functions*, the *M-function* and the *m-function*, as

$$M(\omega, \Omega, \tau) = \sum_t e^{-c\omega^2(t-\tau)^2} \cos[\Omega(t-\tau)], \quad (\text{D5})$$

$$m(\omega, \Omega, \tau) = \sum_t e^{-c\omega^2(t-\tau)^2} \sin[\Omega(t-\tau)]. \quad (\text{D6})$$

I note that

$$M(\omega, 0, \tau) = n(\omega, \tau), \quad (\text{D7})$$

$$M(\omega, \omega, \tau) = n(\omega, \tau) C(\omega, \tau), \quad (\text{D8})$$

$$m(\omega, \omega, \tau) = n(\omega, \tau) B(\omega, \tau), \quad (\text{D9})$$

$$M(\sqrt{2}\omega, 2\omega, \tau) = n(\sqrt{2}\omega, \tau) \tilde{C}(\omega, \tau), \quad (\text{D10})$$

$$m(\sqrt{2}\omega, 2\omega, \tau) = n(\sqrt{2}\omega, \tau) \tilde{B}(\omega, \tau). \quad (\text{D11})$$

Now we can write the squared modulus as

$$\begin{aligned} |W|^2 &= \frac{1}{4} A^2 \omega \{ [M(\omega, \omega - \omega_0, \tau) + \cos(2\omega_0 \tau \\ &\quad + 2\theta) M(\omega, \omega + \omega_0, \tau)]^2 + \sin^2(2\omega_0 \tau \\ &\quad + 2\theta) m^2(\omega, \omega + \omega_0, \tau) \}, \end{aligned} \quad (\text{D12})$$

which gives the response of the DWT to a pure sinusoid.

In the continuum limit, the major and minor response functions converge to

$$M(\omega, \Omega, \tau) \rightarrow \omega^{-1} \rho \sqrt{\frac{\pi}{c}} e^{-\Omega^2/(4c\omega^2)}, \quad (\text{D13})$$

$$m(\omega, \Omega, \tau) \rightarrow 0. \quad (\text{D14})$$

In most astronomical applications, c is chosen small enough, and ω restricted to low enough frequencies that the product $\omega^{-1} \rho$ is $\gg 1$, in which case the term $M(\omega, \omega + \omega_0, \tau)$ is

³Here is a fine example of possible confusion due to switching notation from body to appendix. The contravariant form is $\langle y^\alpha y^\beta \rangle = N^{-1} \sigma^2 S^{\alpha\beta}$, which in the notation of the main body of the paper becomes $\langle y_a y_b \rangle = N^{-1} \sigma^2 S_{ab}^{-1}$.

negligible. Then the response of the DWT to a pure sinusoid is approximately

$$\begin{aligned} |W|^2 &\approx \frac{1}{4} A^2 \omega M^2(\omega, \omega - \omega_0, \tau) \\ &= \frac{1}{4} A^2 \omega (\omega^{-2} \rho^2 \pi c^{-1} e^{-(\omega - \omega_0)^2 / (2c\omega^2)}) \\ &= \frac{\pi \rho^2 A^2}{4c\omega} e^{-(\omega - \omega_0)^2 / (2c\omega^2)}. \end{aligned} \quad (\text{D15})$$

Therefore the wavelet modulus, for ideal time sampling, plenty of data, and small enough constant c , is

$$|W| = A \rho \sqrt{\frac{\pi}{4c\omega}} e^{-(\omega - \omega_0)^2 / (4c\omega^2)}, \quad (\text{D16})$$

which is the formula usually given for the response of the DWT (Goupil *et al.* 1991).

REFERENCES

- Campbell, L. 1926, *Harvard Ann.* 79, no. 2
 Campbell, L. 1955, *Studies of Variable Stars*, AAVSO
 Ferraz-Mello, S. 1981, *AJ*, 86, 619
 Foster, G. 1995, *AJ*, 109, 1889
 Foster, G. 1996a, *AJ*, 111, 541
 Foster, G. 1996b, *AJ*, 111, 555
 Goupil, M. J., Auvergne, M., & Baglin, A. 1991, *A&A*, 250, 89
 Grossman, A., & Morlet, J. 1984, *SIAM J. Math. Anal.*, 15, 723
 Grossman, A., Kronland-Martinet, R., & Morlet, J. 1989, in *Wavelets: Time-Frequency Methods and Phase Space*, edited by J. M. Combes, A. Grossman, and Ph. Tchamitchian (Springer, Berlin), p. 2
 Kholopov, P. N., *et al.* 1985, *General Catalogue of Variable Stars*, Fourth edition, Moscow
 Kovacs, G. 1980, *ApSS*, 69, 485
 Kovacs, G. 1981, *ApSS*, 78, 175
 Mattei, J. A., Mayall, M. W., & Waagen, E. O. 1990, *Maxima and Minima of Long Period Variables 1949-1975*, AAVSO
 Mattei, J. A. 1995a, *International database of the American Association of Variable Star Observers* (private communication)
 Mattei, J. A. 1995b, *International database of photoelectric photometry of the American Association of Variable Star Observers* (private communication)
 Scargle, J. D. 1982, *ApJ*, 263, 835
 Szatmary, K., & Vinko, J. 1992, *MNRAS*, 256, 321
 Szatmary, K., Vinko, J., & Gal, J. 1994, *A&ASS*, 108, 377
 Szatmary, K. 1994, *Wavelet Analysis of Deficient and Gapped Data Series of Variable Stars*, Poster at IAU 22nd General Assembly