Photometry - II. Differential and High-Speed Photometry

Dave Kilkenny

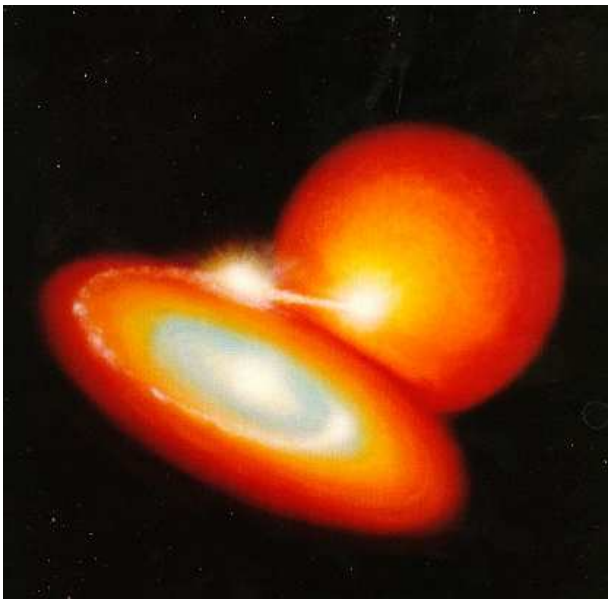

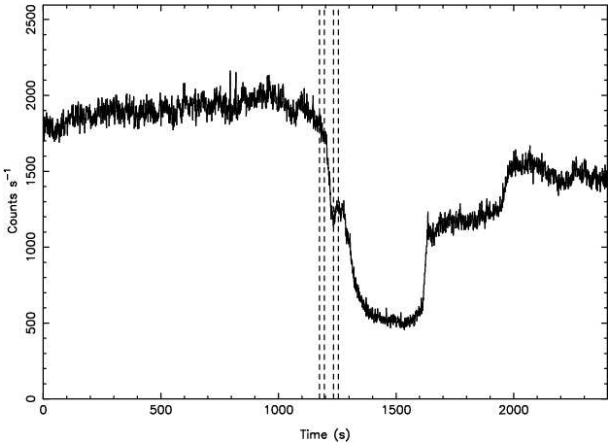

Contents

1 Differential Photometry

If we need more accurate photometry, usually because the amplitude of variability is small – but maybe we just want to do the best possible – then we observe one or more "local comparison" stars. These are stars close to the variable on the sky – and preferably with similar colours to the variable (to minimise any differential effects due to colour). An example of how this might work is shown below and in the figure,

Star	HJD (245 2420+)	V	(B-V)	(U-B)	(V-R)	(V-I)	Residuals ("std" - observed) V	(B-V)	(U-B)	(V-R)	(V-I)
HD123432	5.25803	7.504	1.106	1.069	0.565	1.073	-0.014	0.004	0.011	-0.005	-0.013
V759CEN	5.26303	7.532	0.577	0.092	0.335	0.642					
V759CEN	5.26790	7.542	0.580	0.089	0.336	0.641					
V759CEN	5.27273	7.553	0.576	0.100	0.334	0.645					
V759CEN	5.27759	7.559	0.578	0.093	0.327	0.637					
HD123794	5.28240	7.601	0.366	0.017	0.217	0.436	-0.021	0.004	-0.007	0.003	-0.006
V759CEN	5.28727	7.577	0.579	0.093	0.333	0.639					
V759CEN	5.29215	7.586	0.572	0.099	0.335	0.643					
V759CEN	5.29699	7.593	0.576	0.096	0.335	0.639					
V759CEN	5.30178	7.595	0.586	0.092	0.325	0.637					
HD123432	5.30660	7.505	1.107	1.064	0.569	1.076	-0.015	0.003	0.016	-0.009	-0.016
V759CEN	5.31161	7.604	0.582	0.097	0.331	0.632					
V759CEN	5.31642	7.609	0.584	0.087	0.330	0.640					
V759CEN	5.32118	7.607	0.582	0.092	0.335	0.637					
V759CEN	5.32606	7.606	0.592	0.133	0.298	0.606					
HD123794	5.33090	7.609	0.368	0.018	0.223	0.441	-0.029	0.002	-0.008	-0.003	-0.011
V759CEN	5.33581	7.591	0.582	0.093	0.332	0.638					
V759CEN	5.34057	7.584	0.580	0.092	0.332	0.637					
V759CEN	5.34531	7.573	0.579	0.099	0.336	0.637					
V759CEN	5.35011	7.566	0.580	0.101	0.334	0.639					
HD123432	5.35484	7.502	1.105	1.077	0.569	1.066	-0.012	0.005	0.003	-0.009	-0.006
V759CEN	5.35972	7.549	0.579	0.097	0.337	0.636					
V759CEN	5.36449	7.539	0.578	0.094	0.334	0.635					
V759CEN	5.36937	7.528	0.578	0.099	0.336	0.635					
V759CEN	5.37422	7.519	0.578	0.092	0.336	0.636					
HD123794	5.37899	7.601	0.366	0.020	0.227	0.433	-0.021	0.004	-0.010	-0.007	-0.003
V759CEN	5.38380	7.501	0.581	0.098	0.331	0.630					
V759CEN	5.38859	7.486	0.582	0.095	0.328	0.631					
V759CEN	5.39342	7.480	0.579	0.094	0.330	0.632					
V759CEN	5.39831	7.476	0.576	0.097	0.333	0.636					
HD123432	5.40311	7.502	1.110	1.072	0.571	1.068	-0.012	0.000	0.008	-0.011	-0.008
:	:	:	:	:	:	:	:	:	:	:	:
:	:	:	:	:	:	:	:	:	:	:	:

And the data might look something like:

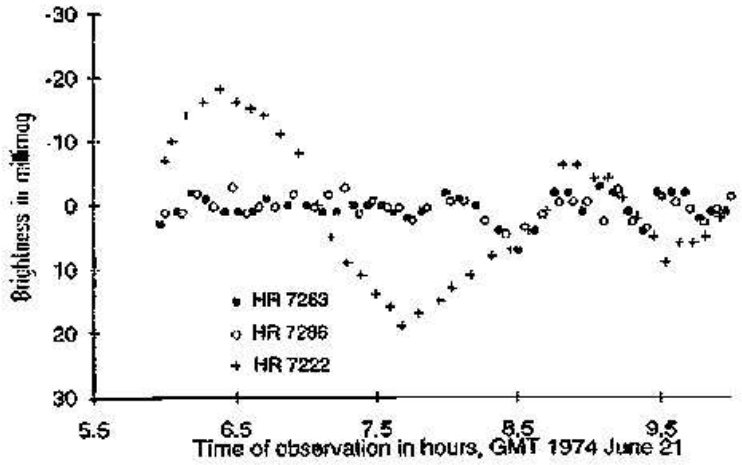

Figure 1: Data for a variable and two comparison stars.

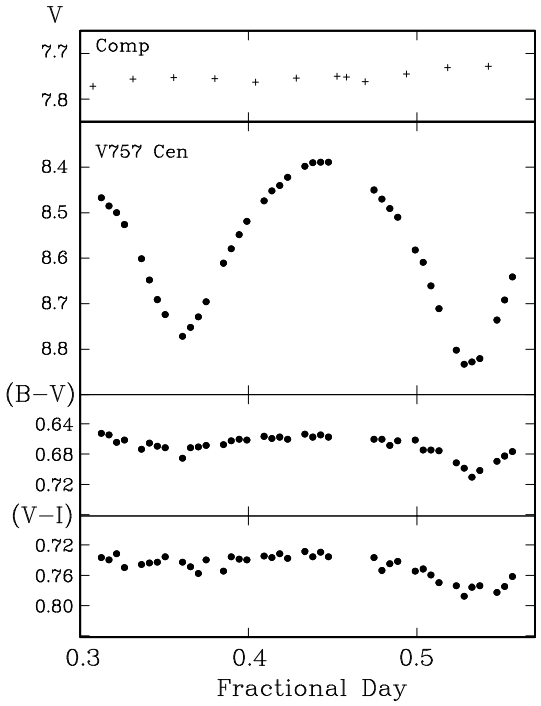

Figure 2: APT data from one night for V757 Cen, a WUMa (W Ursae Majoris) binary system. The plot of magnitude or colour against time is usually called a **light curve**. The upper box shows the magnitude results for one of the comparison stars observed. If we have a well-determined "standard" magnitude (and colours) for this star – and assuming it is itself constant – we can then calculate the differences (*standard minus observed*) and apply interpolated corrections to the variable star photometry. Note: all data have been reduced to the standard system using regularly measured standards (one or two per hour), still, residual effects remain. The "drift" seen in the top panel might be due to small atmospheric transparency changes or to temperature-dependent effects in the instrumentation, for example.

Notes:

- A differential programme might form part of an "all sky" night, with the normal extinction corrections and so on. This frequently happens on the 0.5m telescope.
- The above example shows photomultiplier data. Differential photometry is, in some respects a lot easier with a CCD – you can measure a star, local comparisons and the background sky all at the same time. Of course, you still need separate exposures for the different passbands – and the small field of view can be a problem.
- In most of what follows, I have used real or artificial data with rather short periods, typically a few minutes to a couple of hours. This is partly becasue the effects show up better with many cycles and partly because I have the material immediately to hand. However, the techniques are generally applicable to all variability.

Another two examples of the results of differential photometry are shown. These are both pulsating Hydrogen–deficient stars (sometimes called **Helium stars**, or **Extreme Helium (EHe) stars**) with very short periods (near 0.1 day).

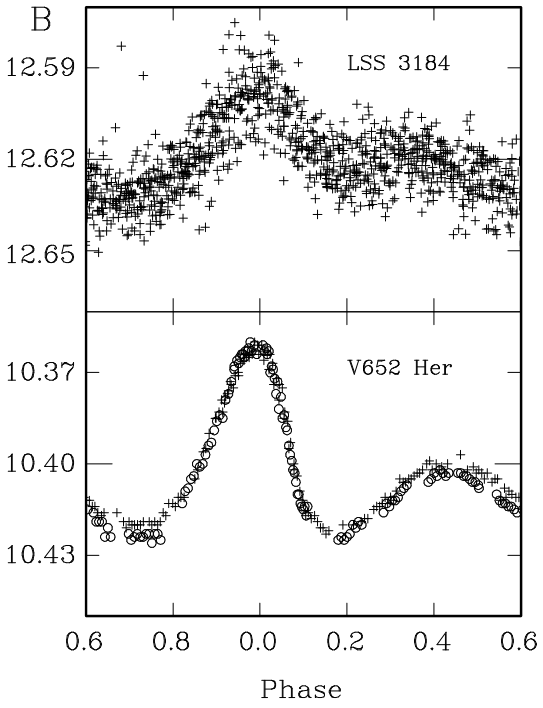

Figure 3: Light curves for (upper) LSS 3184 (BX Cir) and (lower) V652 Her. The upper data are from several runs phased by the pulsation period. The lower data are from two runs about ten years apart, illustrating the potential accuracy of differential photometry. Somewhat curiously, the lower data are using a photomultiplier and the upper data are from a CCD and, although LSS 3184 is substantially fainter than V652 Her, the CCD data should be better. The field of LSS 3184 is, however crowded with many faint stars, so the extraction process might suffer. Alternatively, the star itself might vary slightly (in addition to the very obvious changes) or the stellar variation might change from cycle to cycle.

The next two figures indicate how we might obtain fundamental stellar parameters from the photometry – and in this case using also (Doppler shift) velocities from spectroscopy. LSS 3184 is used as an example; for a more detailed explanation of these processes, see *Monthly Notices of the Royal Astronomical Society,* vol 310, p 1119 (1999).

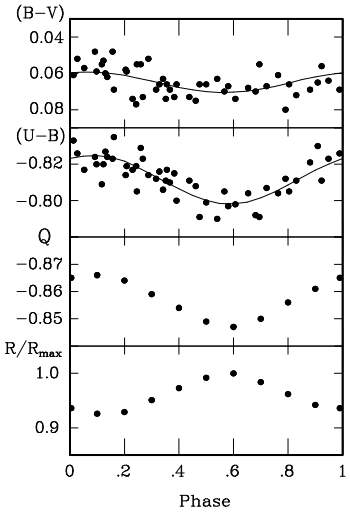

Figure 4: LSS 3184 – the measured quantities $(U-B)$ and $(B-V)$ – which naturally vary as the star pulsates – can be used to derive the "reddening-free" temperature parameter, Q. The relative radius variation is derived from brightness and colour variations.

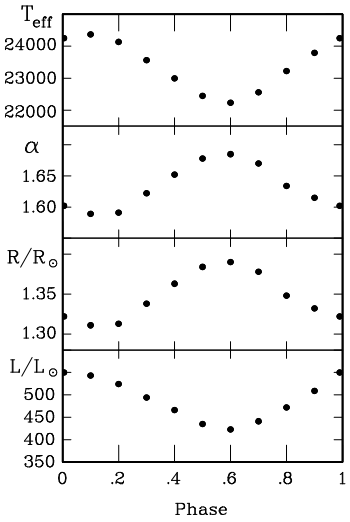

Figure 5: LSS 3184 – $T_{\rm eff}$ is derived from a calibration of Q; the angular size, α, comes from a comparison of observed and expected (from a stellar model) fluxes. Comparing the (Doppler) velocity variations with the angular size variations then gives the distance to the star; the absolute radius and luminosity variations follow from the distance.

2 "High Speed" Photometry

"High speed" or continuous photometry might be done if we are observing a fast variable with a complex light curve where we don't want to lose time by observing standards or even local comparison stars. In this case, we can often remove the extinction effect as a low order polynomial (and any error in this process can be essentially taken out as low frequency "noise" in the Fourier analysis. The example shown below is part of a single light curve for a pulsating sdB star – a class of variable discovered at the SAAO.

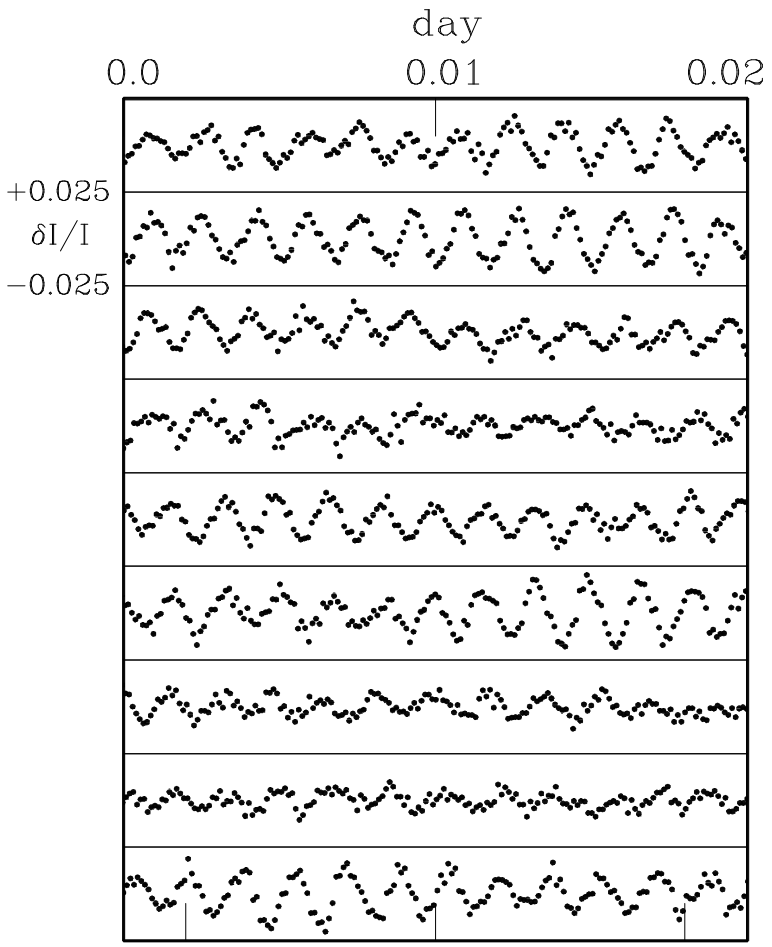

Figure 6: Continuous light curve of PG 1047+003 – read top-to-bottom and left-to-right. Note the very obvious "beating" – a sign that more than one frequency is present.

It turns out that the pulsation characteristics of PG 1047+003 are very complex and in order to have a chance of resolving these, much longer, continuous data trains are required than we can get from one night – or even several nights from one site.

One way around this problem is a "multi–site campaign". One group which organises this is the "Whole Earth Telescope" or WET, a loose grouping of about 50 or more astronomers with access to around 20 observatories. We shall look at the advantage of such an arrangement later, in the context of Fourier analysis

3 Differential high-speed photometry

We can often get good results, even in less–than–perfect conditions, if we do high-speed photometry with a CCD which has a number of stars, as well as the target star, on the frame. The high-speed requirement usually means we have to read out just small areas of the CCD ("windowed mode") or if we cover up half the chip so that only one half is exposed to light, then at the end of the exposure, the exposed half is almost instantaneously shifted to the covered half, which can then be read out less frantically (minimising read-out noise) whilst the next exposure is being made. ("frame transfer mode"). Having other stars on the chip means these can be used as local comparisons to correct for small (and sometimes not so small) changes in the atmosphere.

The next two pages show examples of this applied to the short-period (~ 0.26 day) eclipsing binary, AA Dor.

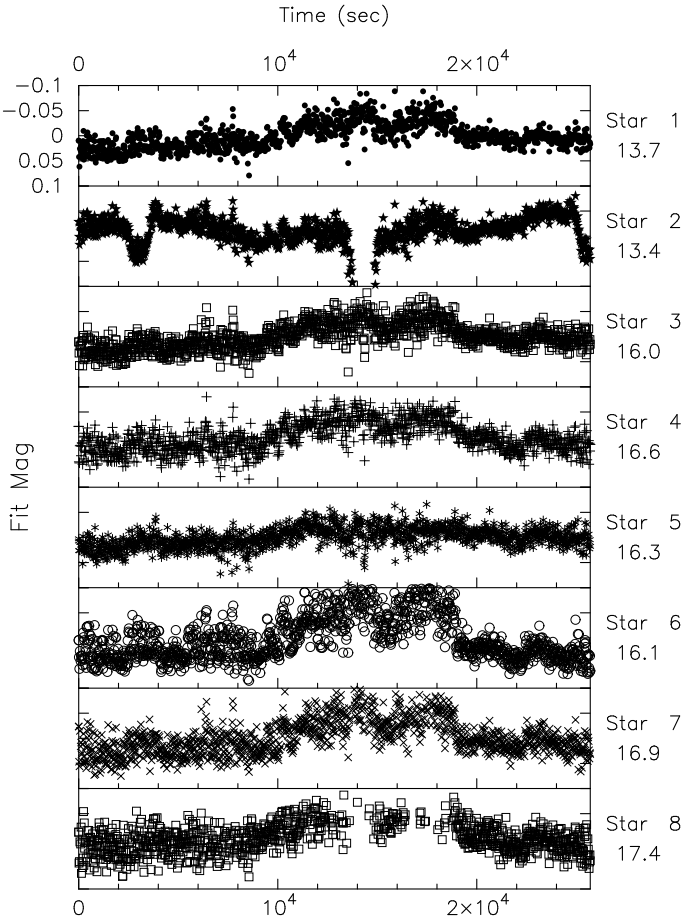

Figure 7: Light curves for AA Doradus and some field stars, uncorrected.

Note that, apart from star 2 – the variable – all the stars follow much the same variation, caused by atmospheric changes (clouds). We can therefore derive an average correction for each data point using several stars – this will reduce somewhat the scatter from each individual star – and then apply this correction to all the stars on the frame. This is shown in the next figure.

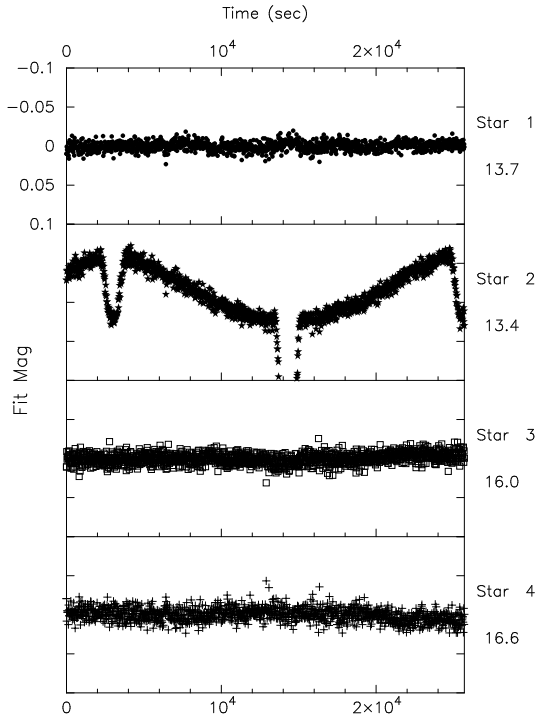

Figure 8: Light curve for AA Doradus and some field stars, corrected using stars 1, 3 and 4.

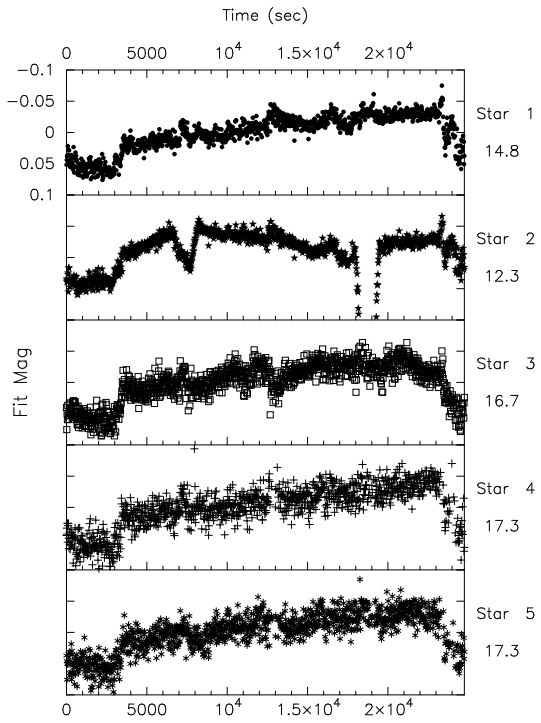

Figure 9: Another example of light curves for AA Doradus and some field stars, uncorrected.

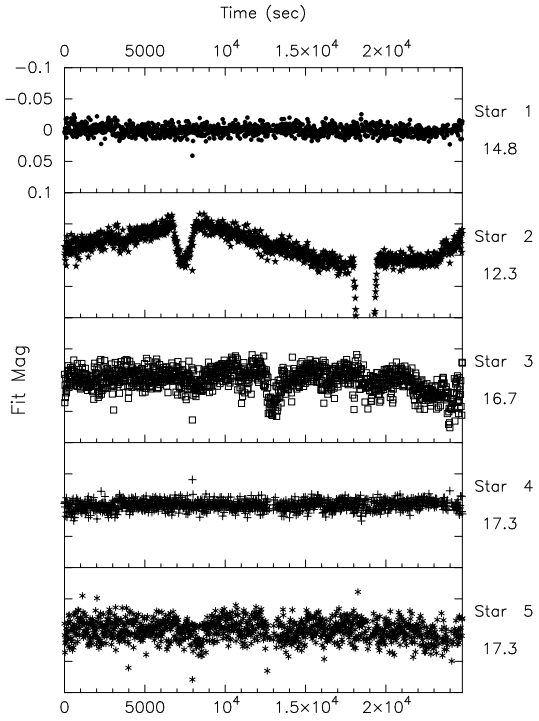

Figure 10: Light curve for AA Doradus and some field stars, corrected using stars 1 and 4.

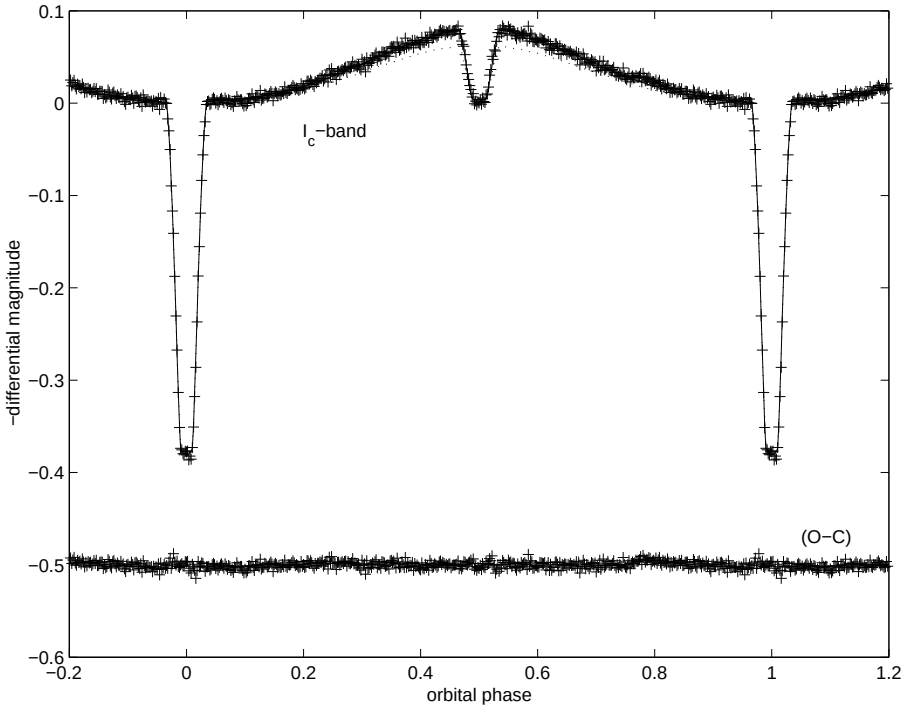

Figure 11: AA Dor: differentially corrected "I" light curve with fitted model.

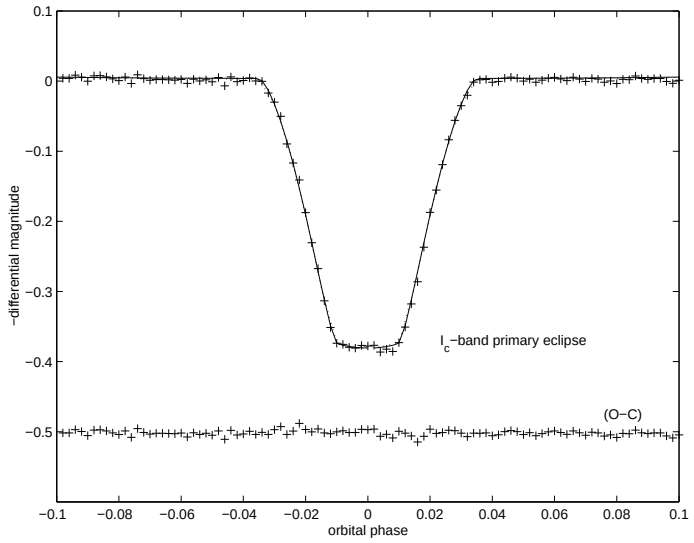

Figure 12: AA Dor: "I" light curve expanded around primary eclipse with fitted model.

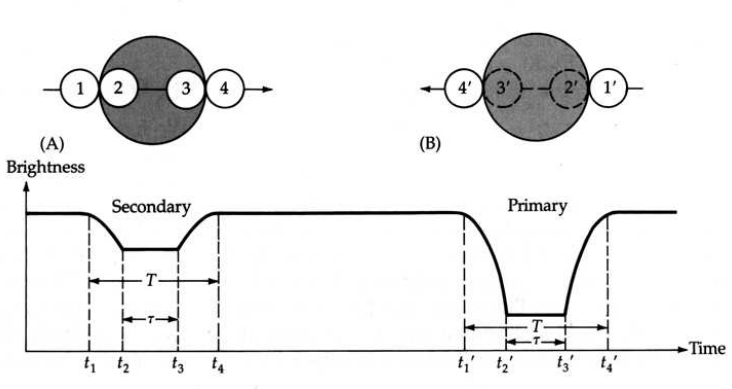

Figure 13: Schematic showing the derivation of stellar radius from an eclipsing binary light curve. The "contact points" t_1, t_2, t_3 and t_4 are clearly related to the relative radii of the stars (as a fraction of the binary orbit) – $R_p = t_3 - t_1$ or $t_4 - t_2$ and $R_s = t_2 - t_1$ or $t_4 - t_3$, assuming a circular orbit in the line-of-sight. The **primary eclipse** of the binary is when the hotter star (which is not necessarily the bigger) is obscured by the cooler, i.e. the "deepest" eclipse.

4 Multi-periodicity: Applications of Fourier Transforms

In the examples we have looked at in differential/"high speed" photometry – eclipsing binaries and singly periodic pulsating variables – phasing the data (or, equivalently, determining the period) is usually relatively easy. When a star (or other object) experiences multi-periodic behaviour, this is much harder.
Fourier analysis is a very powerful tool in variable star research (and many other areas) particularly where observational data contain several (or many) periodicities. A very brief summary of some basic concepts are given below, before we look at some practicalities.

4.1 Integral Transforms

The idea of a transform is to take a difficult problem, transform it into a simpler problem, solve the simpler problem and then apply the inverse transform to recover the solution of the original difficult problem. Logarithms are a simple example of this process – it's easier to add and subtract in "log-space" than it is to multiply and divide in "number space".

An **integral transform** takes a function $f(x)$ **defined on x space** for some interval $a \leq x \leq b$ and multiplies it by $K(p, x)$, a function of x and some parameter p. The integral transform of $f(x)$ is then $F(p)$ where:

$$F(p) \;=\; \int_a^b f(x)\; K(p,x)\; dx$$

$K(p, x)$ is called the **kernel** of the transform. The integral transform $F(p)$ can be regarded as the image of $f(x)$ in **p-space**.

4.2 Fourier Transforms

The Fourier transform of a function $f(x)$ is usually written something like:

$$F(p) \;=\; \int_{-\infty}^{\infty} f(x)\; e^{-i2\pi px}\; dx$$

Since we usually have data as a time series and we transform to frequency space, we can write:

$$F(\nu) \;=\; \int_{-\infty}^{\infty} f(t)\; e^{-i2\pi \nu t}\; dt$$

If we now consider a function

$$f(t) \;=\; A\ \sin\ (\omega t) \qquad \text{over the range } -\infty \;\leq\; t \;\leq\; \infty$$

then the Fourier transform $F(\nu)$ will be zero except where $\nu = \omega$. The Fourier transform of an infinitely long sine wave, will thus be a delta function at frequency ω (Actually two delta functions, one at ω and one at $-\omega$.

In practice, observational data are rarely infinitely long (!) and a finite data set is equivalent to a sine wave (in the case above) convolved with a "top hat" function (more properly called a

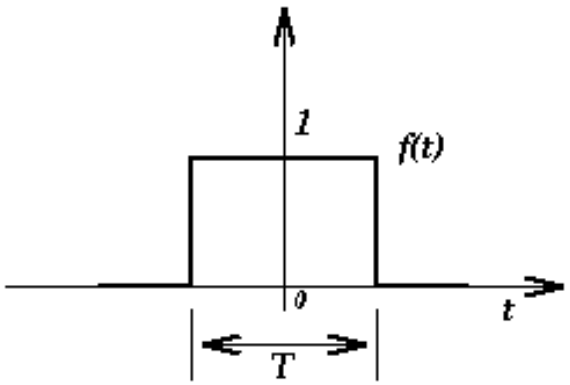

Figure 14: The rectangle function.

rectangle or boxcar function) which has the property that the function is equal to unity in a given range (the range of the data set) and zero everywhere else.
A very powerful property of Fourier transforms is the **convolution theorem** which states that if h(x) is the convolution of two functions f(x) and g(x), i.e.

$$h(x) \;=\; f(x) \otimes g(x)$$

then the Fourier transform of h(x) is the product of the Fourier transforms of f(x) and g(x):

$$H(p) \;=\; F(p)\; G(p)$$

(For a nice graphical demonstration of function convolution (moving gifs, nogal), see http://mathworld.wolfram.com/Convolution.html)

For finite data then, we need to know the Fourier transform of the top hat. Since the function is zero outside $-T/2 \leq t \leq T/2$ and unity within that range, we can write the Fourier transform as:

$$F(\nu) \;=\; \int_{-\infty}^{\infty} f(t)\; e^{-i2\pi\nu t}\; dt \;=\; \int_{-T/2}^{T/2} e^{-i2\pi\nu t}\; dt \;=\; \frac{-\; e^{i2\pi\nu T/2} \;+\; e^{-i2\pi\nu T/2}}{-i2\pi\nu}$$

This reduces to:

$$F(\nu) \;=\; \frac{-2\; i\; \sin\;(\pi\nu T)}{-2\pi i\; \nu} \;=\; T\; \frac{\sin\;(\pi\nu T)}{\pi\nu T} \;=\; T\; sinc\;(\pi\nu T)$$

and the sinc function will appear in transforms of finite data sets; in fact, as noted above, a finite continuous data set will have a Fourier transform which is a product of a delta function and a sinc function:

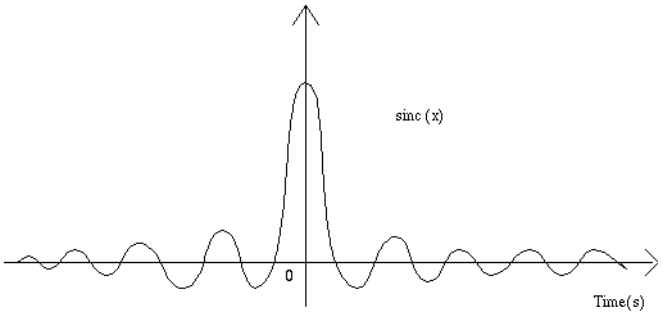

Figure 15: The sinc function.

4.3 Basic concepts

First, we use some very simple artificial data to illustrate some basic ideas.

4.3.1 Periodogram, resolution and spectral window

The figure shows a single, perfect sine wave with period 200 seconds (frequency 0.005 Hz or 5000 μHz) sampled at 0.0001day – or 8.64 seconds – for 1 day (86400 seconds).

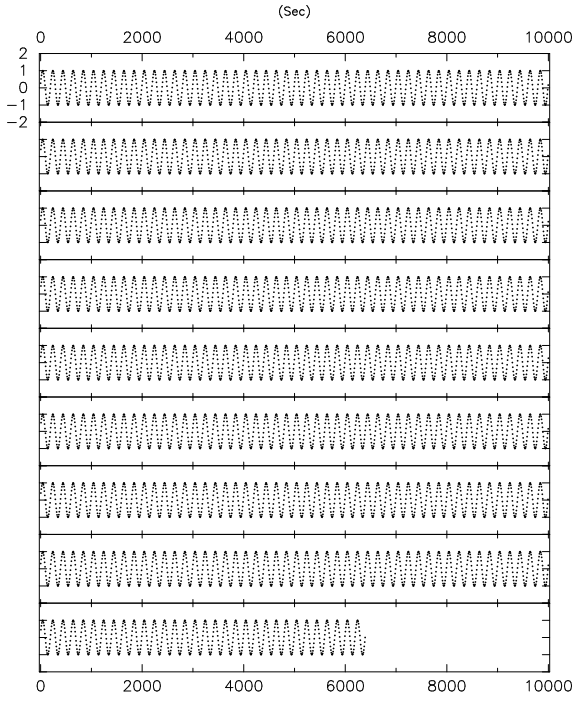

Figure 16: One day's worth (86400 seconds) of perfect sine wave with period of 200 seconds (frequency 5000 μHz) sampled at 0.0001 day (8.64 seconds).

Periodograms of the whole data set, half of the data and one quarter of the data (i.e. 24, 12 and 6 hours of continuous data) are shown in the next figure:

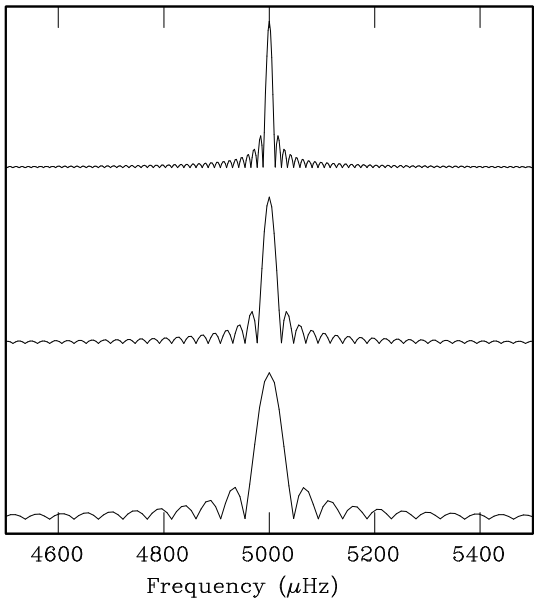

Figure 17: Periodograms of the perfect sine wave for 24, 12 and 6 hours of continuous data.

It is clear that:

- the function displayed in each of the three plots is not *sinc* but looks like $(sinc)^2$. This occurs because in the process of calculation of the Fourier transform it is usual to examine the (purely real) **power spectrum** which is proportional to the **amplitude spectrum** squared. The power spectrum is also called the **periodogram**. In fact, the amplitude spectrum (= Fourier transform) is often then calculated and is the more normal way of displaying astronomical results – but taking the square root, obviously does not recover the negative values in the *sinc* function. To confuse things, the amplitude spectrum is also often called a periodogram. Hereafter, the amplitude spectrum will be displayed uniformly (and also referred to as the periodogram).
- it is evident that as the data set decreases in length, the FWHM of the main peak increases in inverse proportion. This is a result of the *sinc* function, but also makes sense because as the data set gets smaller, the frequency is less accurately determined.

From the last point, it can be seen that the length of the data set affects the **resolution** in the frequency domain, and although the theoretical value for resolution is 1/T – the half-width of the $(sinc)^2$ peak – a more conservative value of 1.5/T for "complete resolution" is often adopted (Loumos & Deeming Astrophys. & Sp. Sci. 56, 285, 1978).

In the examples given above, data sets of length 86400, 43200 and 21600 seconds, the resolutions, 1.5/T, will be about 17, 35 and 69 μHz.

The above figure is also instructive because often you will find that papers show the **spectral window** for a data set. This is the Fourier transform of a perfect sine wave, sampled in an

identical way to the actual data set and therefore representing the best that the data can produce in the absence of noise. Since these Fourier transforms are of a perfect sine wave, they are also the spectral window of the data.

4.3.2 Noise

Next we take the 24 hour perfect sine wave and introduce some noise into the data – each point has a small, random artificial "error" added or subtracted, with the resulting data set (or part of it) shown below. Clearly, it would be a lot harder to be confident there was a signal in these data just by looking at them.

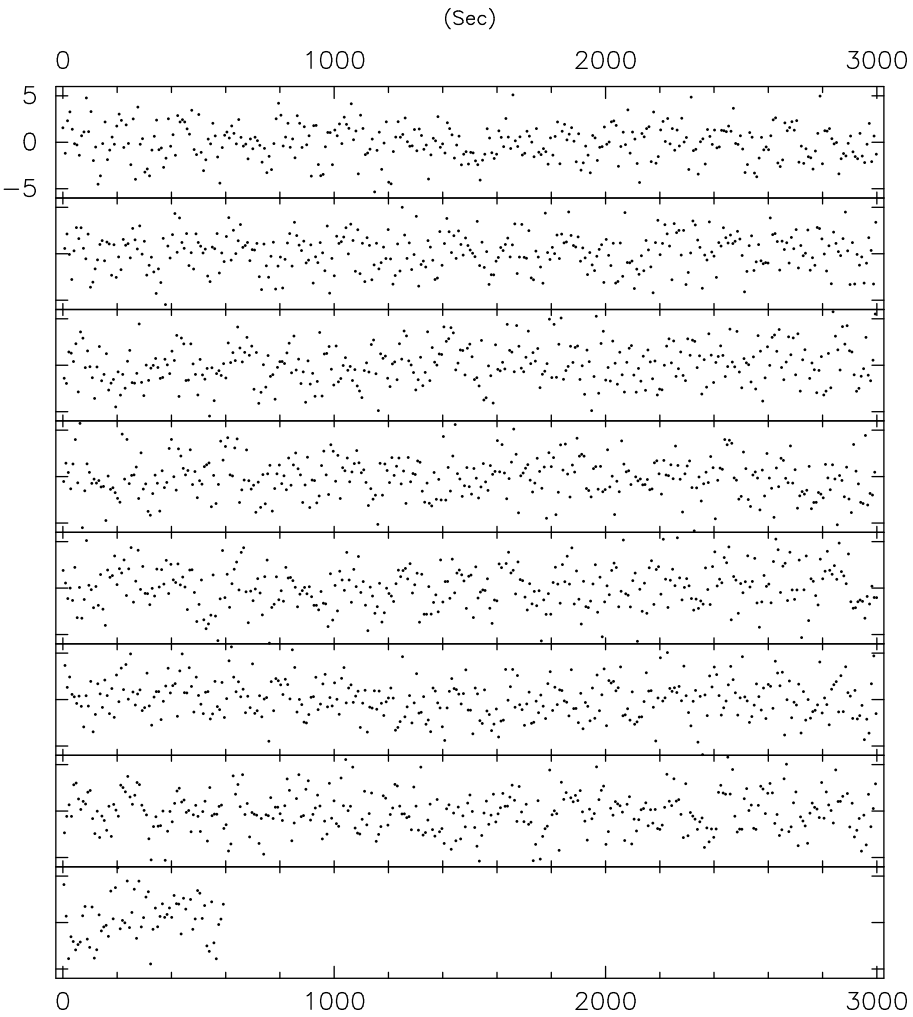

Figure 18: Same 5000 μHz data sampled at 0.0001 day but with added noise.

However, the Fourier transforms (shown below) are very convincing – we recover the original periodicity to a high degree of accuracy.

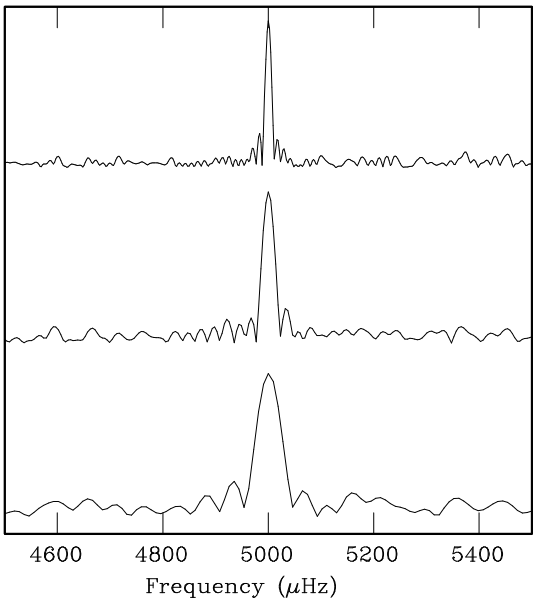

Figure 19: Periodograms for noisy data (24, 12 and 6 hour data sets.

4.3.3 Multiperiodicity

Next, we create an artificial data set which includes 5 frequencies at 5000, 6000, 7000, 5600 and 3300 μHz (corresponding to approximate periods 200, 167, 143, 179 and 303 seconds) with respective amplitudes of 1.0, 0.6, 0.3, 0.2, and 0.1. With these frequencies all at phase zero at time zero, we get a light curve which looks like this:

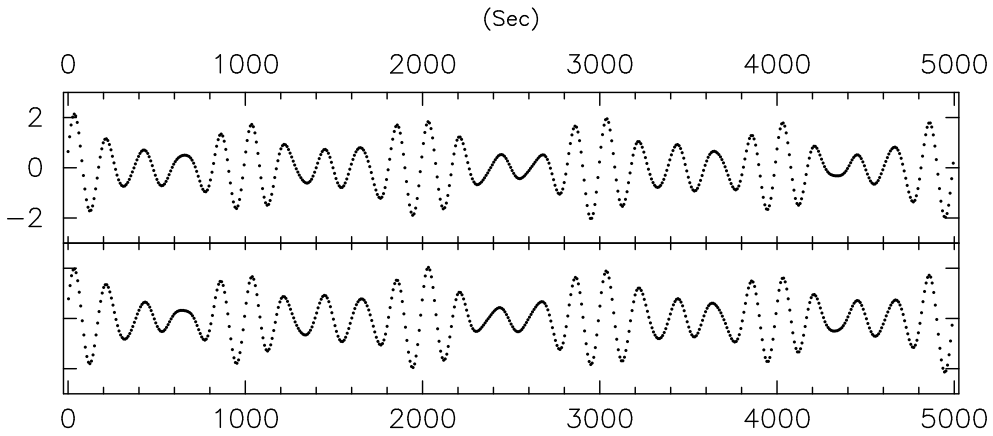

Figure 20: Synthetic light curve generated by the five frequency/amplitude combinations noted in the text.

It looks like a simple two-frequency beating at first glance, but remember that some of the input amplitudes are rather small – the light curve will be dominated by the larger amplitude variations. Also, you can see that the pattern does not repeat exactly.

The next figure shows a series of Fourier transforms of the data. The top plot is the spectral window for the data set. The second plot is the Fourier transform or periodogram (actually the **amplitude spectrum**) of the data.

The third plot is the periodogram of the data with a sine wave of the period, amplitude and phase of the dominant (largest amplitude) frequency determined in the Fourier analysis removed from the data. This is known as **"pre-whitening"** the data by the sine wave. Subsequent plots are pre-whitened by two, three, four and five frequencies

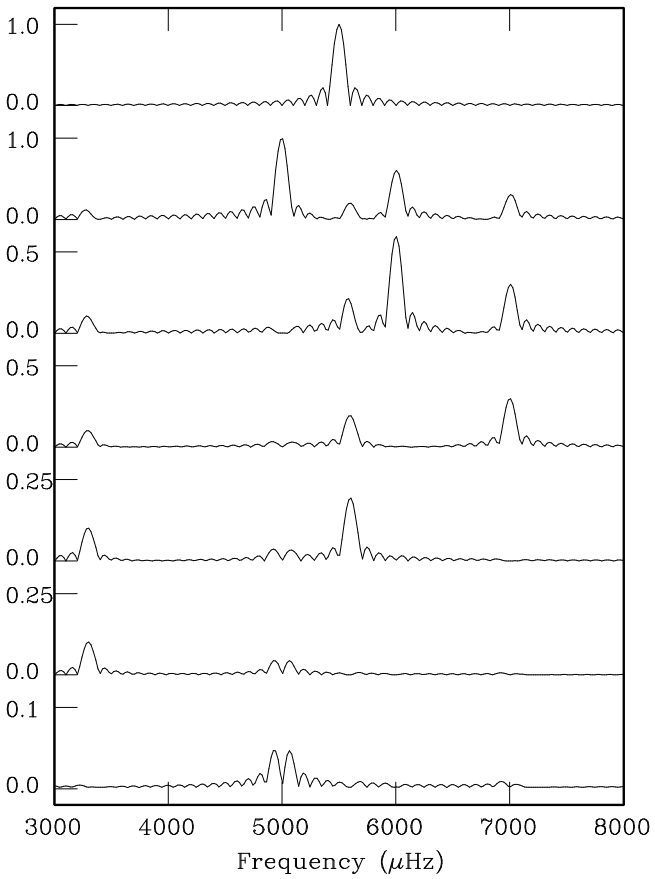

Figure 21: The top plot is the spectral window for the "five frequency" light curve. The second plot is the periodogram of the synthetic data; the third plot is the periodogram of the data with the first (largest amplitude) sinusoid removed; the fourth plot has two sinusoids removed, and so on. Note that the amplitude scale increases down the figure.

Note that the lower plot shows some power near the frequency of the first (largest amplitude) frequency. This is due to the fact that the derived frequency was not a perfect match to the "real" frequency, so some residual periodic variation remains.

With real data, having located five frequencies with some confidence, we would then solve for all five simultaneously.

The next two plots show the same thing but with noise added. A summary table compares the results.

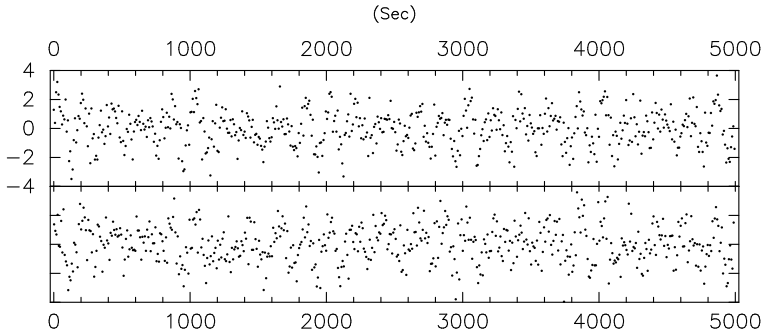

Figure 22: Same as the "five frequency" light curve but with random noise added to each point.

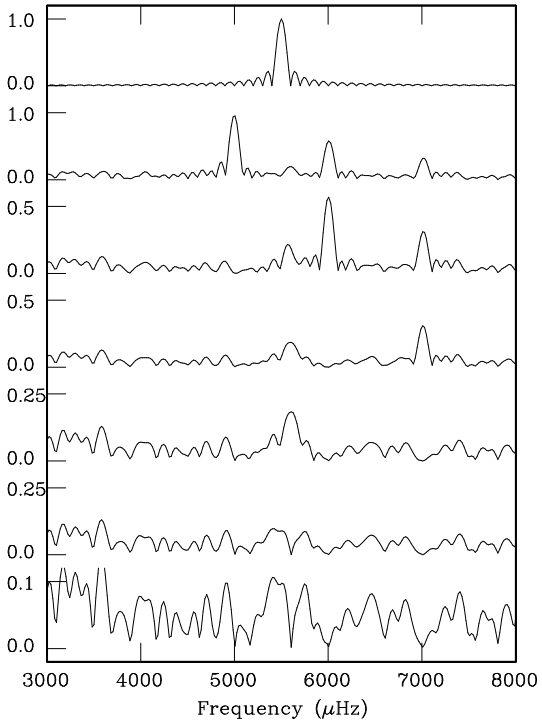

Figure 23: Periodograms of the noisy "five frequency" light curve. Note that we do NOT recover the weakest of the five frequencies – it is lost in the noise.

Table 1: Summary of FT results for the "five frequency" data.

Input		Noiseless results		Noisy results	
Freq. (μHz)	Amplitude	Freq. (μHz)	Amplitude	Freq. (μHz)	Amplitude
5000.0	1.0	4996.9	1.001	4996.9	0.961
6000.0	0.6	6000.1	0.597	6004.1	0.571
7000.0	0.3	7001.3	0.298	7011.3	0.311
5600.0	0.2	5599.6	0.194	5605.6	0.184
3300.0	0.1	3298.8	0.101		

4.3.4 Aliasing

Now we look at what happens when the data are not continuous. The first figure shows the sort of thing you might get in three observing nights – runs of about 6 hours with gaps of about 18 hours (actually, it should be possible to get somewhat longer runs than 6 hours, but go with me on this ...). The data are now singly periodic but with some noise added.

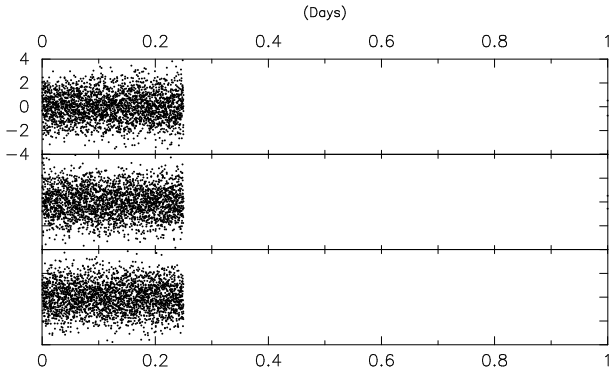

Figure 24: Simulation of three 6-hour data sets on the same object, singly periodic at a frequency of 5000 μHz.

Now the periodogram shows a distinctive pattern, due to the gaps in the data.

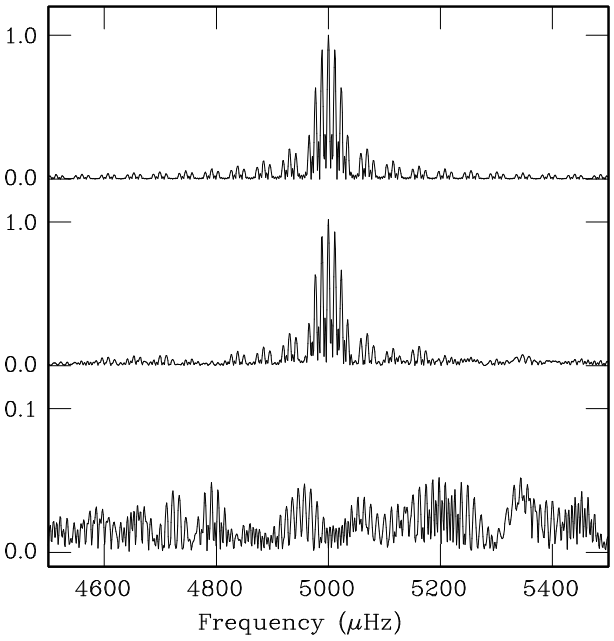

Figure 25: The top plot is the spectral window; the middle plot is the periodogram (almost identical to the spectral window; the bottom plot is the periodogram of the data pre-whitened by the frequency derived from the middle periodogram.

This occurs because when there is a gap in the data, the best frequency which fits the data on both sides of the gap will have (say) n cycles missing (in the gap). If the gap is sufficiently large,

then it will be possible to fit a frequency which has $n-1$ or $n+1$ cycles into that gap. With a less good fit, it might be possible to fit $n-2$ or $n+2$ cycles into the gap – and so on.

With astronomical data, there is often (as in the synthetic example given above) a one day frequency built into the observing because the star will be observed at about the same time each night. The peaks in the above figure will then be separated by a day – therefore will have a frequency of 1/86400 or 11.57 μHz (and twice that, three times, etc). This effect is called **aliasing**.

If the data sets are short and the gaps are large, it might not be easy – or even possible – to decide which peak is the real frequency. One way around this is to use **multi-site data**. Although quite difficult to organise, a programme of observing one star from several observatories can yield excellent results. One such group – the "WET" (Whole Earth Telescope) runs two campaigns a year in this way in which SAAO/UCT people usually participate.

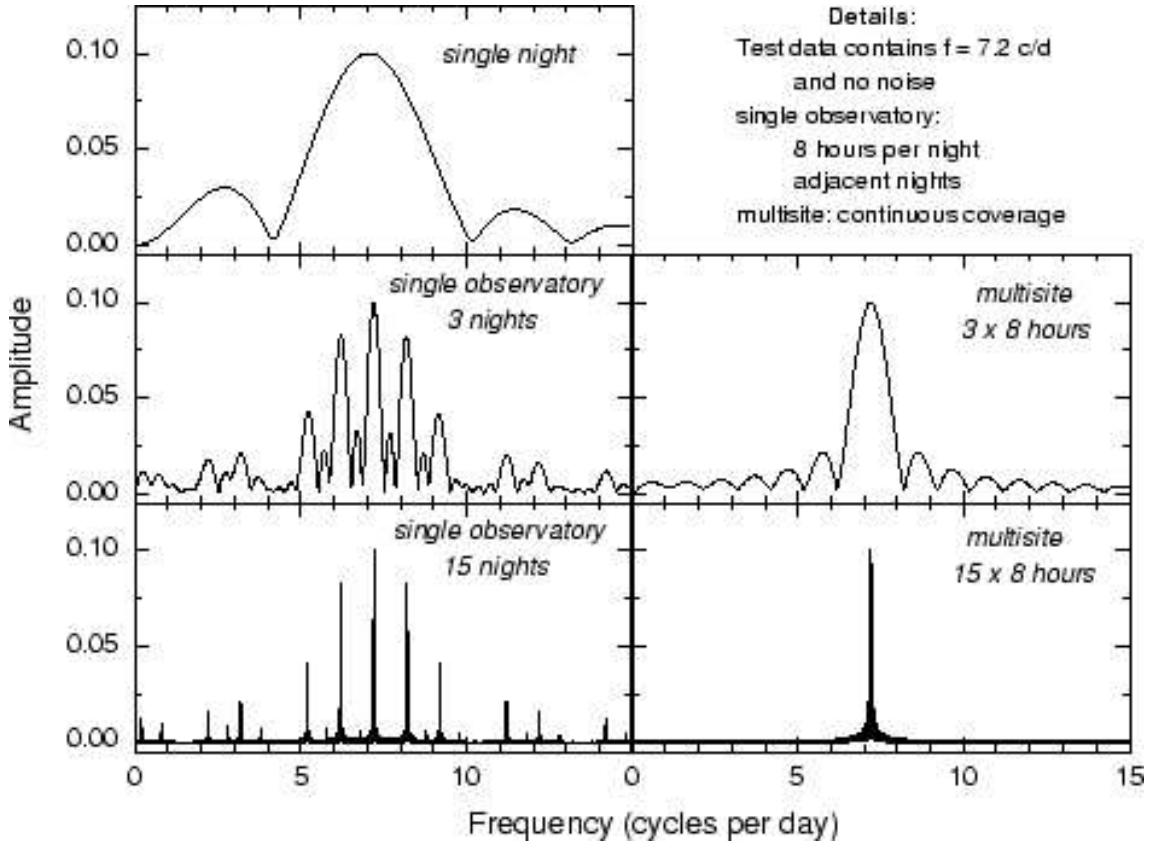

Figure 26: Aliasing effects with different data sets.

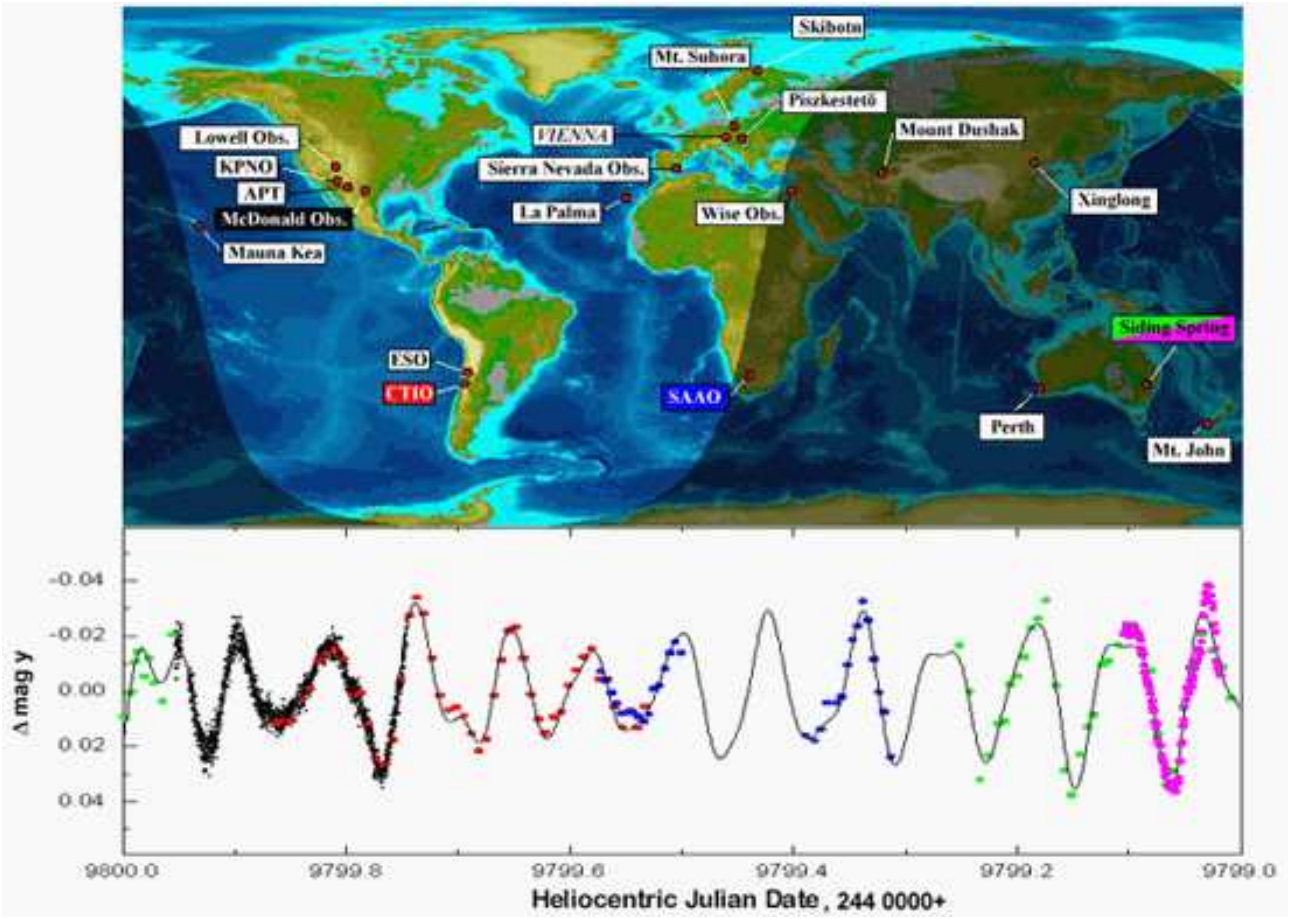

Figure 27: Some of the sites involved in the WET organisation.

4.3.5 Harmonics

So far, we have looked at sinusoidal variations. The figure below shows a non-sinusoidal – but stll strictly periodic – variation. RR Lyrae variations look something like this - as do the Helium stars in section 1.

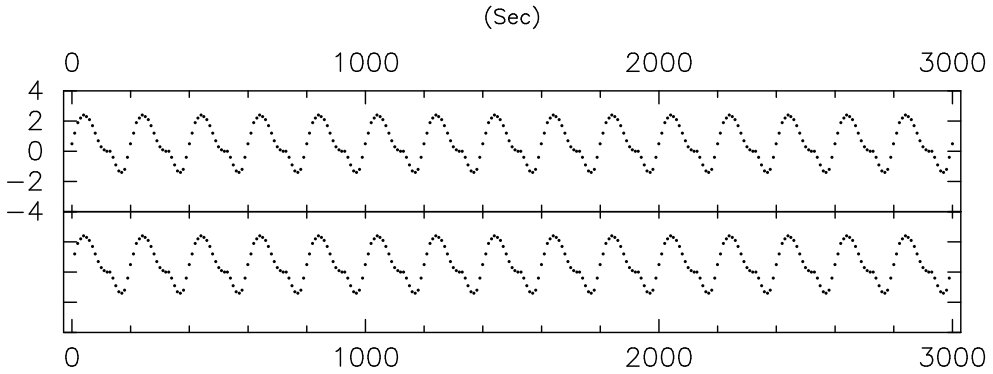

Figure 28: A non-sinusoidal, strictly periodic variation. Note the faster "rise" and slower "fall" – somewhat similar to an RR Lyrae star.

The Fourier transform is a continuous time version of the Fourier series so, of course it will try to represent a non-sinusoidal variation by a sum of **harmonics**, so even though the variation is singly-periodic, we get what looks like a multi-periodic solution:

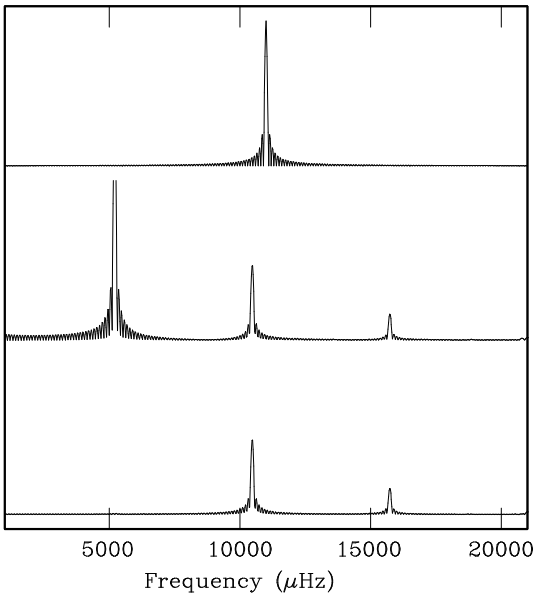

Figure 29: Periodograms of the above data. The top plot is the spectral window; the middle plot is the Fourier transform (actually the **amplitude** spectrum; and the bottom plot is the same but pre-whitened by the dominant frequency.

In this simple case, it is easy to see that the weaker frequencies are simply $2f$, $3f$, ... of the dominant (in fact only) frequency, f.

4.4 The Nyquist frequency

An important consideration of period-finding techniques is the sampling rate. This is illustrated in the figure below.

The top panel is sampled at twice the frequency of the variation and it is clear that one could recover that frequency from the "observational data" using that sampling rate.

The second panel shows a curve with a higher frequency which is sampled at the same rate as the top curve. In this case, it is clear that the recovered frequency - illustrated in the third panel - is quite wrong. The obtained data in fact **undersample** the variation.

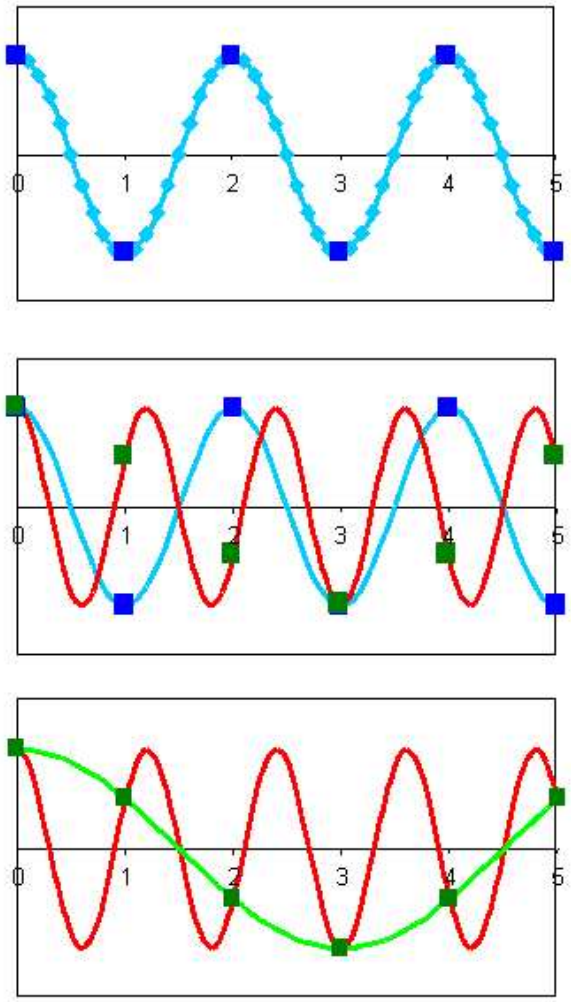

Figure 30: Illustration of undersampling. The top panel shows a frequency, f, which is barely adequately sampled at a rate $2f$. Sampling a higher frequency (red line in the middle panel) at the same rate, results in a frequency determination which is wrong (green line in the bottom panel).

Undersampling can also lead to errors in the determination of the amplitude of variation, even at $2f$, as shown in the next figure

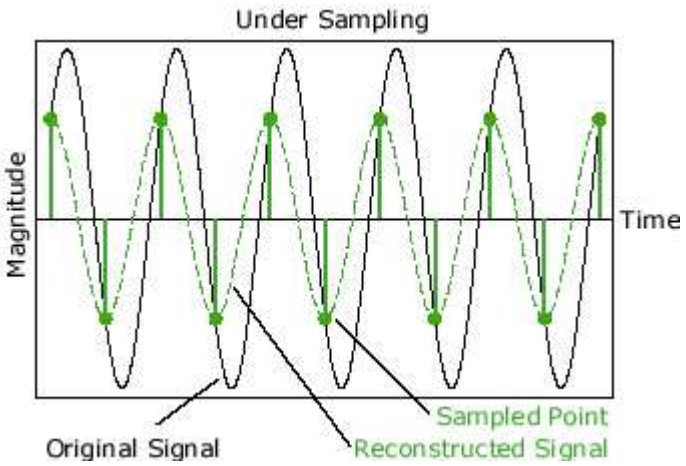

Figure 31: Undersampling resulting in a wrong amplitude determination.

So, if a signal is sampled at a frequency $2f$, you cannot reliably recover information which is periodic at greater frequency than f. This is known as the **Nyquist frequency** from the Nyquist–Shannon sampling theorem. Actually, it is necessary to sample at somewhat higher frequency than the Nyquist frequency, as indicated above.

The higher frequencies are not lost but, as the above figure suggests, are merged with lower frequencies from which they cannot be separated. It is important, therefore to sample at as high a frequency as possible. Often S/N considerations conflict with this – with faint objects, longer integration times (exposures) will be required, so it is important to realise that this will set a high frequency limit to what you can detect.

Amongst the fastest pulsating stellar objects are the DA white dwarf pulsators (ZZ Ceti stars) and the hottest subdwarf pulsators (EC14026 stars or sdB variables). These can have periods as short as $\sim$ 100 seconds, so typically we try to observe them at 10 – 20 seconds resolution, and faster if possible.

5 Some case studies

We now look at some particular stars to see application of the basic ideas just covered:

5.1 LSS 3184 = BX Cir

LSS 3184 is an Extreme Helium star which pulsates – probably in radial fundamental mode – with a period of 0.1065784 days, about 2.55 hours. As indicated earlier, although now less massive than the Sun, and not a whole lot bigger ($R_* \sim 1.35\ R_\odot$), it has evolved into a hot star with $L_* \sim 500\ L_\odot$. It has essentially no Hydrogen in its atmosphere ($< 0.001\%$). (See, for example, *Monthly Notices of the Royal Astronomical Society,* vol 310, p 1119 (1999)). A phased light curve was shown in section 1. The figure below shows some individual light curves that went into that composite plot.

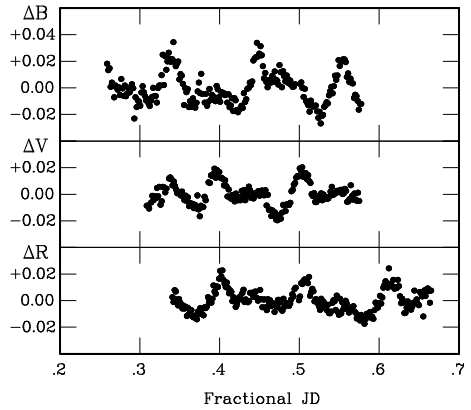

Figure 32: Samples of differentially corrected CCD light curves from 1995.

A fair amount of data of this kind was obtained in April and May 1995 and the periodogram of the R band data is shown below:

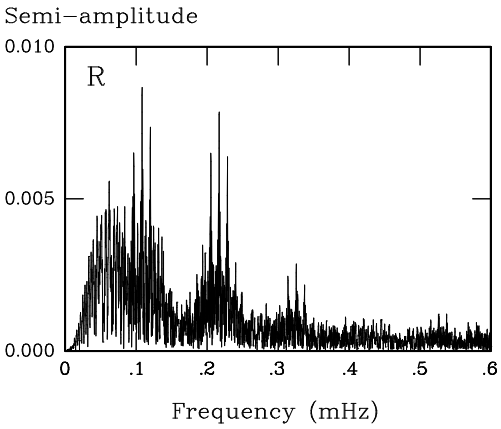

Figure 33: Periodogram of LSS 3184 data from 1995.

One can see:

- the fundamental frequency, f at ∼0.108 mHz (∼9200 sec ∼0.106 day) along with (at least) the $2f$ and $3f$ harmonics and their associated alias patterns;
- the obvious alias spikes are separated by about 0.01 mHz – they are our old friend the 11.57 μHz, daily aliases;
- the much less obvious alias spikes with a separation of about 1 μHz, which are due to a gap of about 9 days between the April and May data;
- the big blob of power below 0.1 mHz which is probably due to residual sky variations - but could be low frequency changes in the star itself – we can't tell from these data.

All well and good. Now see what happens when we try to include some earlier data from 1994. The top plot shows the 0.108 mHz peak, as before.

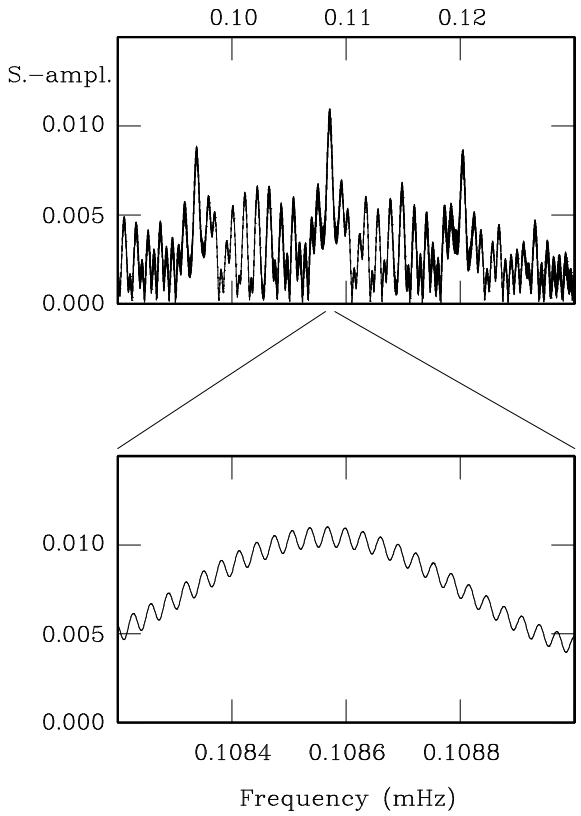

Figure 34: Upper plot: Periodogram of combined 1994 and 1995 data. Lower plot: Main peak at 50 times the resolution.

The bottom plot shows that peak expanded by a factor of 50. Now you can see aliasing at about 0.032 μHz (equivalent to one cycle in about 31 million seconds – or one cycle in about a year.

But what if you want to get the period accurately ? Which of those peaks is the right one ? It might not be possible to tell.

Peter Martinez and Chris Koen of the SAAO have developed a method of harmonic analysis (*Monthly Notices of the Royal Astronomical Society,* vol 267, p 1039 (1994)) which fits harmonic

components to non-sinusoidal light curves and produces a "residualgram" analogous to the periodogram in Fourier analysis.

For non-sinusoidal light curves the harmonic analysis is often better as it uses more information in fitting the harmonic components simultaneously. Equivalently, it is trying to fit a light curve which is nearer to the actual light curve shape. The figure shows a comparison between the Fourier and harmonic analyses for LSS 3184.

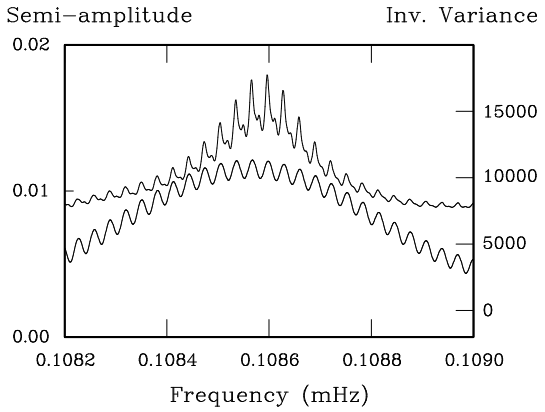

Figure 35: Comparison of the periodogram from Fourier analysis and the residualgram from the harmonic analysis method of Martinez & Koen.

5.2 V652 Her

V652 Her is similar in many ways to LSS3184 but, unlike the other EHe stars, it has about 2% Hydrogen. It has a mean radius $\sim 2\ R_\odot$ and luminosity $\sim 1100\ L_\odot$ and is a radial fundamental mode pulsator (see: *Monthly Notices of the Royal Astronomical Society,* vol 209, p 387 (1984)). The light curve below is reproduced from section 1.

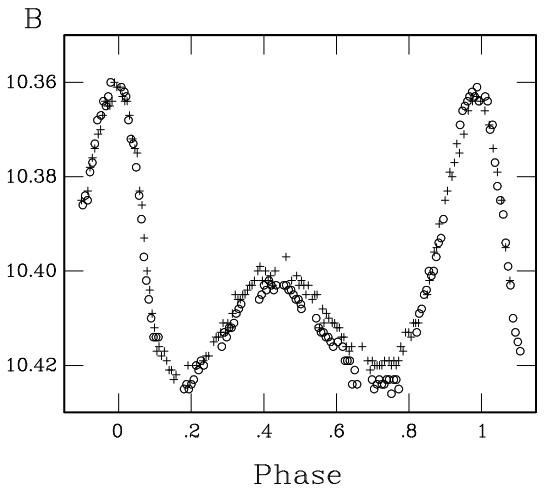

Figure 36: B light curve of V652 Her.

Our interest here – and in LSS 3184 – is that careful measurements over many years have shown the pulsation period of V652 Her to be decreasing (and we wanted to look for the same effect in LSS 3184) :

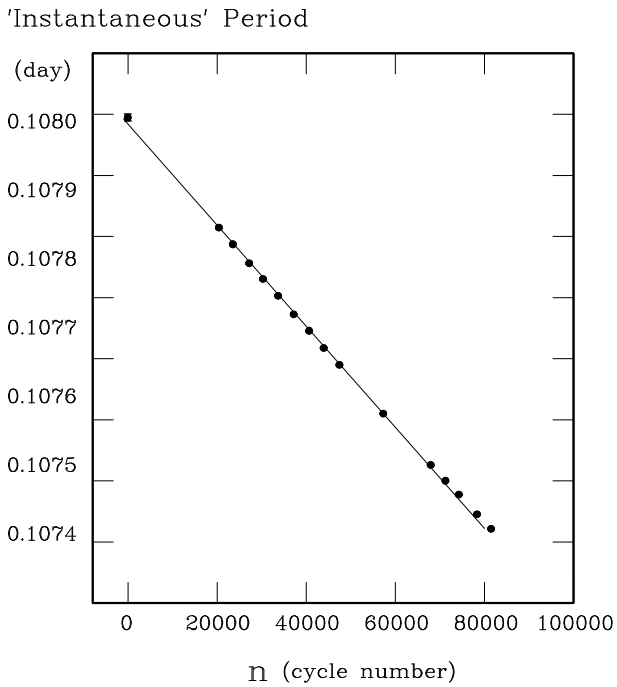

Figure 37: The pulsation period of V652 Her as a function of time. Each point is determined by observing several maxima per year - over 3 months, this would be a baseline of ~ 1000 cycles - and calculating a mean period.

The interest is because, for a pulsating star, there is a simple relation between the pulsation period, P, and the mean density of the star, $\bar{\rho}$, namely:

$$P\,\sqrt{\bar{\rho}} \;=\; Q$$

where Q is a constant dependent on the internal structure of the star. Thus:

$$P\,\sqrt{\frac{M_*}{R_*^3}} \;=\; \text{constant}$$

so, if the mass of the star remains constant (a moderately fair assumption), and the period decreases, then the radius must be decreasing – at a measurable rate – a rate which cannot be measured in any other way. Cool.

It's worth introducing at this stage (but don't tell Chris Koen) the concept of an $(O-C)$ diagram. If you measure, say, a few timings of maxima of a (singly-periodic) pulsating star, or mid-points of eclipses of an eclipsing binary, then you can derive an **ephemeris** for predicting eclipses (or maxima):

$$T_{min} \;=\; T_0 \;+\; n \;\times\; P$$

where T_{min} (or T_{max}) is the predicted time of the recurrent event, T_0 is some start point (eg. the first eclipse observed), n is an integer and P is the period. The ephemeris then gives a "calculated"

value for any future event which can be compared to an "observed" event; any difference is the $(O - C)$ residual.

Clearly a plot of residuals against time should show a scatter around zero. If there's a **linear** trend then you've simply got the period slightly wrong and you can correct this to form a new ephemeris. If, however, there's a curve, then the period is increasing or decreasing with time which might have interesting consequences, such as that noted above.

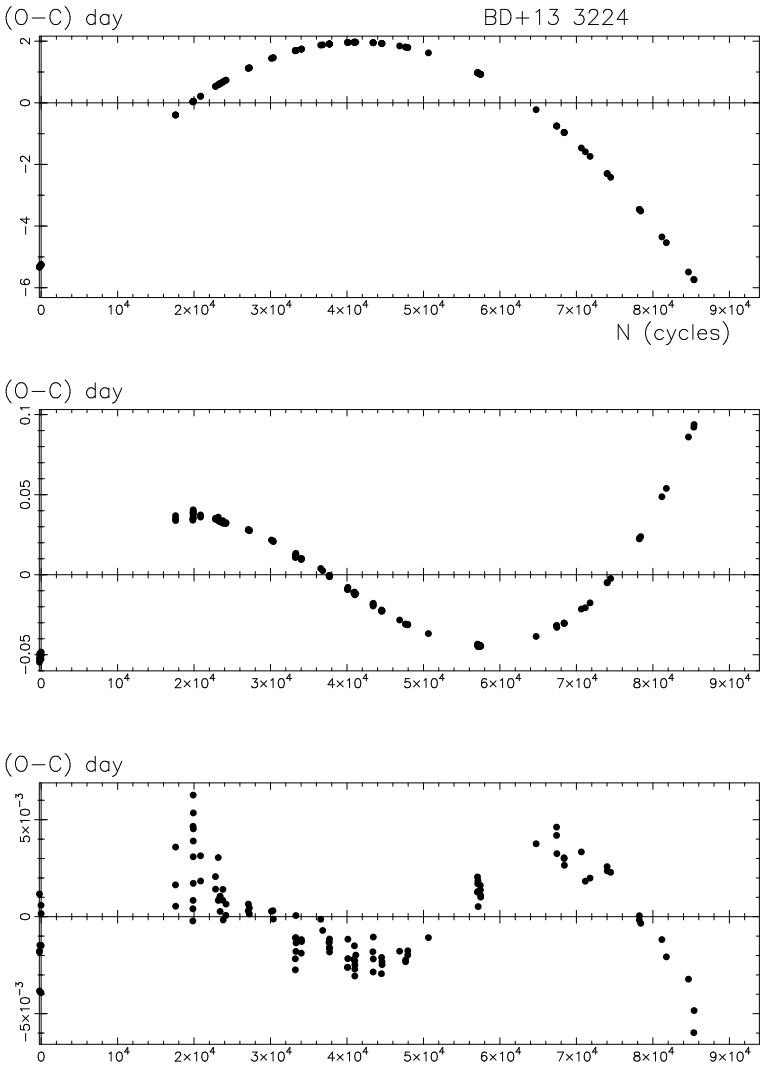

Figure 38: $(O - C)$ diagram for V652 Her. Upper: residuals from a linear least-squares fit to times of maxima. Middle: residuals from a quadratic fit. Lower: residuals from a cubic fit.

It is the case that very few (if any) stars have such a clear and well-defined (O–C) diagram. Chris Koen's dislike of $(O - C)$ diagrams is that many people have over-interpreted the results, attributing physical significance to changes that might well be random. So you have to be careful.

For V652 Her, however, the change is real and clear. It has not been fully explained, but the first order term (the obvious linear term) has been explained in terms of the star contracting as it evolves.

5.3 PG 1336–018 = NY Vir

The figure shows a sample light curve for the exotic PG 1336–018, a pulsating sdB star in an eclipsing binary with a very cool companion.

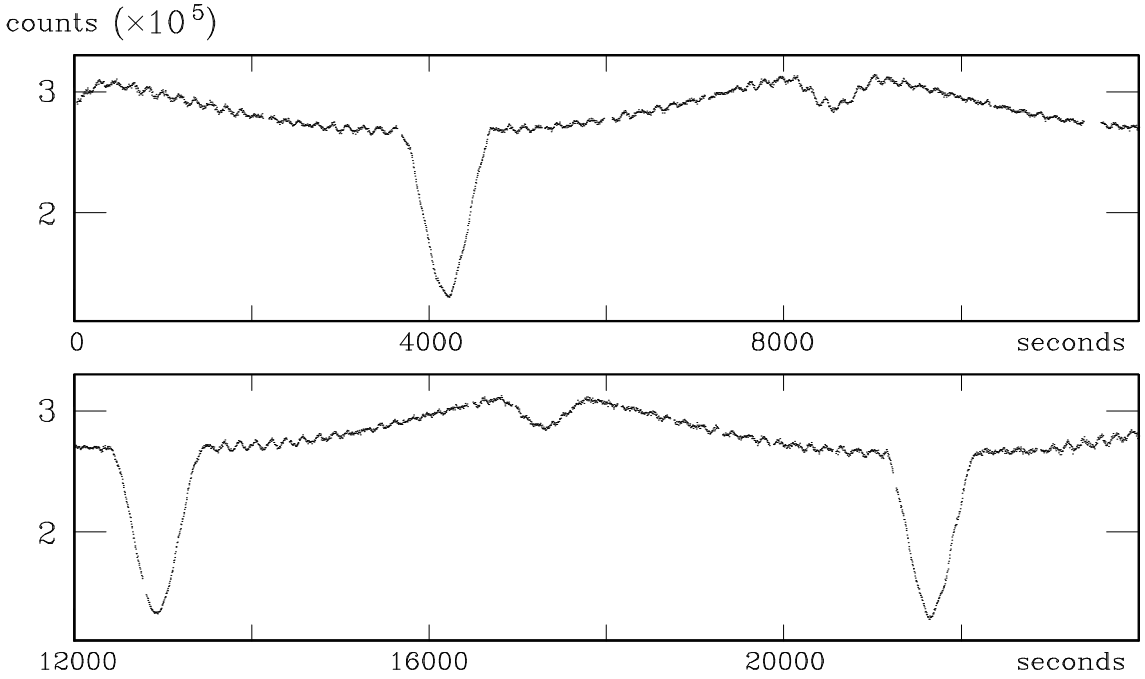

Figure 39: Beautiful light curve for the eclipsing binary PG1336-018, obtained at 3 second time resolution by Francois van Wyk using the SAAO 1.9m telescope. Note the very obvious primary and secondary eclipses and the substantial reflection effect, caused by the hot sdB primary heating one side of the cool M dwarf. Note also, the weak variations in the system light, caused by the sdB star pulsating with periods near 3 minutes. The orbital period of this system is very short at close to 0.1 day.

Because the variations in this star are so complex and because it is potentially vital to studies of pulsation in evolved stars, PG 1336 was the target of a "WET" campaign in 1999. A detailed description of this work can be found in *Monthly Notices of the Royal Astronomical Society,* vol 345, p 834 (2003).

The next figure shows the entire data set – the actual data is all plotted, except that the eclipses have been snipped out and the 'reflection effect" variations removed by fitting them with a double sine wave, because in this case we are interested in the pulsational variations.
With so much data, it is possible to divide the data into two approximately equal parts and analyse them separately. This has the advantage that frequencies which are found in both sets of data will have a strong chance of being real. (Of course, if some artificial frequency is being introduced into the data at one or more sites, this could still persist in both data sets).

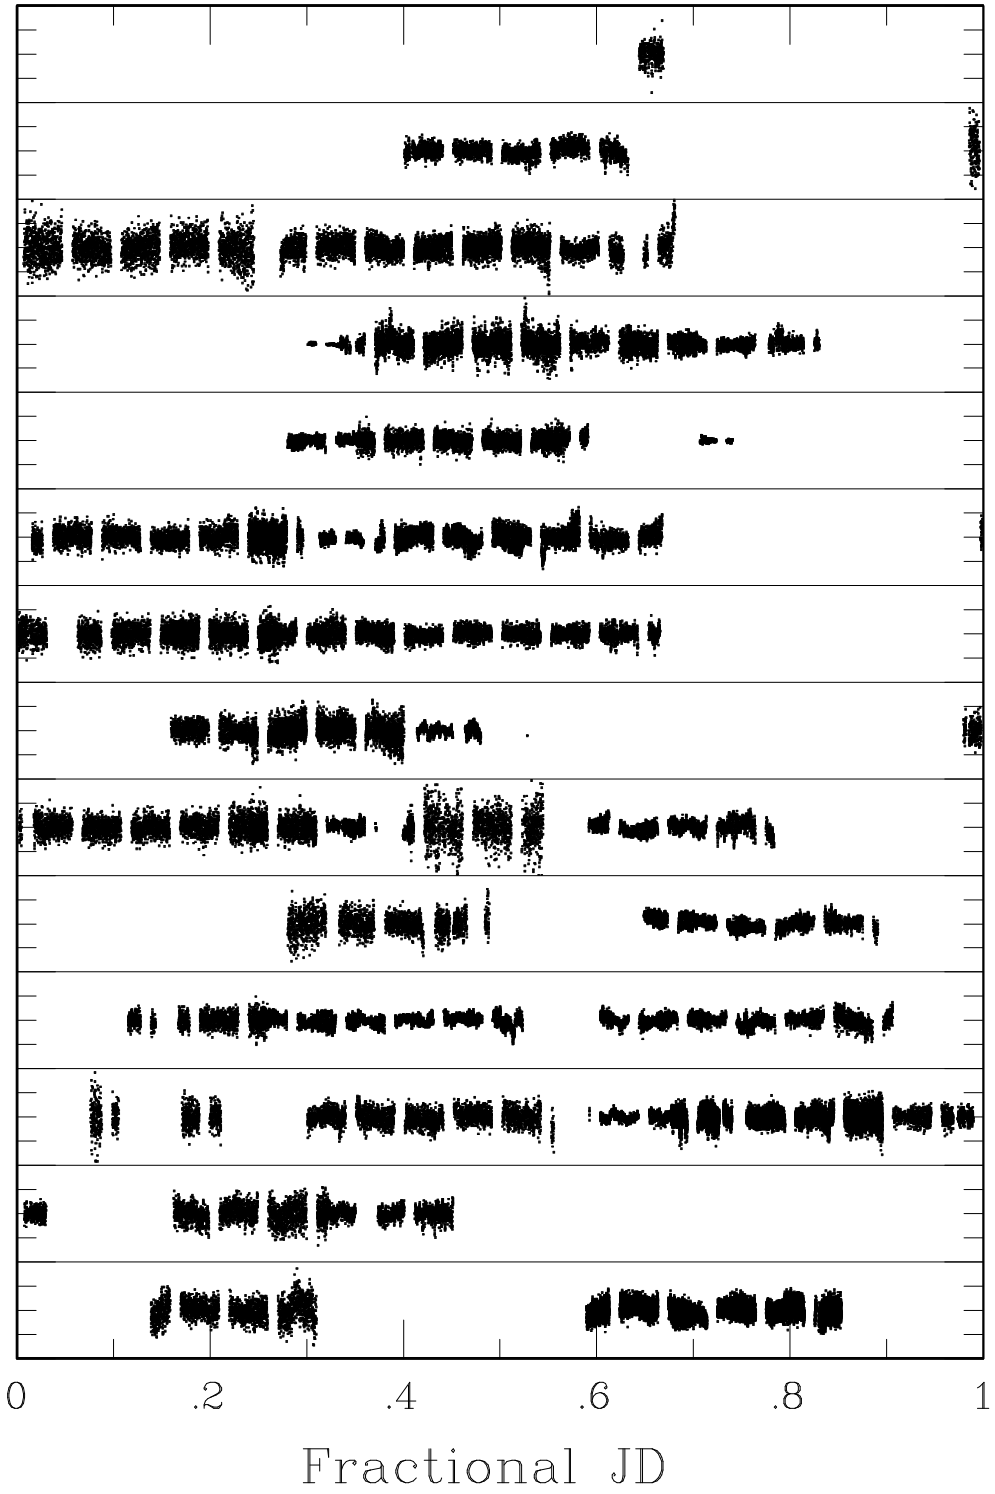

Figure 40: Data from a WET run on PG1336–018. Each panel represents 24 hours of data. Extinction corrections have been made, the eclipses have been removed and the reflection effect taken out by fitting a double sine wave, so that we can investigate the pulsation structure.

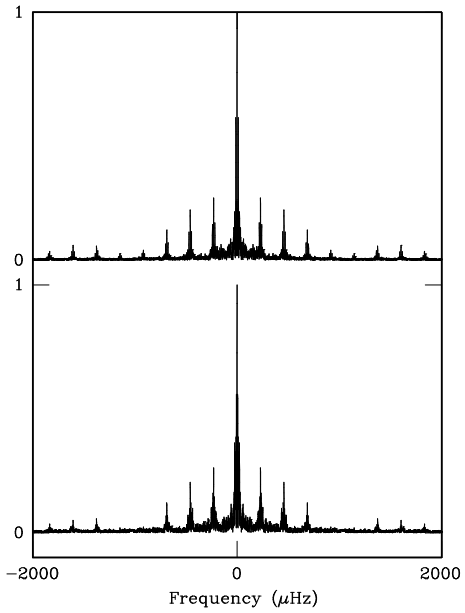

Figure 41: Spectral window of two halves of the "WET" data for PG1336-018.

The spectral windows of the two data sets are shown above. One can see:

- the very obvious aliasing at $\Delta f = 230$ μHz which corresponds to $\sim$4350 seconds, or about 1.2 hours, and is clearly half of the orbital period – or the frequency of the gaps in the data due to the eclipses being removed;
- and the effect of the one-day aliases (11.57 μHz) is just about discernible.

The first figure on the next page shows the periodograms for the two halves of the data – from which we were able to extract over 20 frequencies.

The second figures shows a comparison between the WET data ($\sim$ 2 weeks) and two earlier runs of a few hours only. Apparent are:

- the enormous improvement in resolution with the WET data;
- the appearance of frequencies in the amplitude spectrum is not constant; in some way we do not understand, the pulsation of the star is varying !

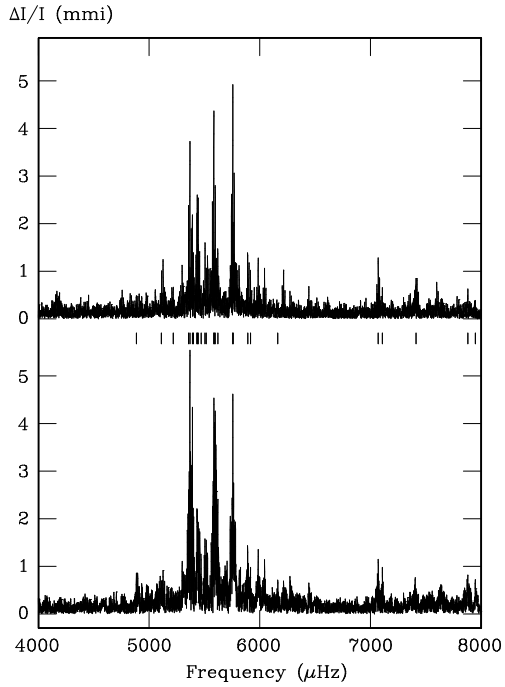

Figure 42: Periodogram of two halves of the "WET" data for PG1336-018. The units ΔI/I are fractional intensity in milli-magnitudes – where a millimag is about 0.1% variation (only approximately because the magnitude scale is logarithmic, not linear).

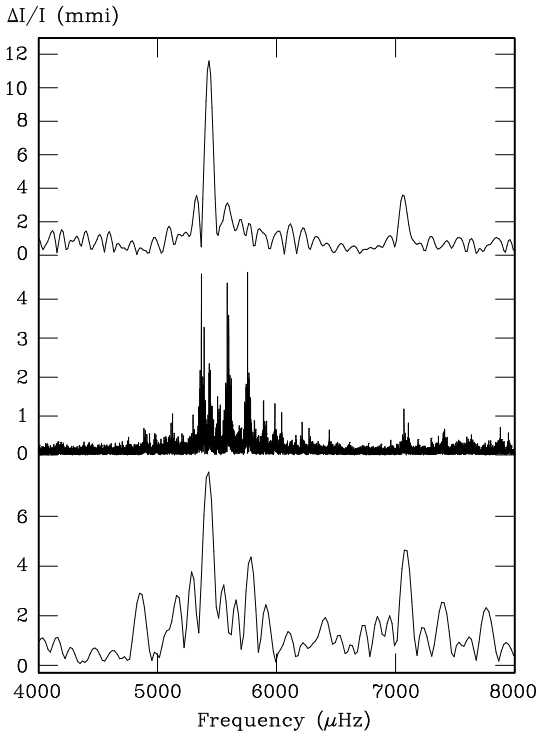

Figure 43: Periodogram from WET run (centre) compared with periodograms from single nights.