We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Analysis of Pulsating White Dwarf Star Light Curves

Denis J. Sullivan

Abstract

Analysis techniques are presented for extracting the frequencies contained in the light curves of pulsating white dwarf stars. In several surface temperature regimes, these astronomical objects are unstable to gravity mode pulsations which result in brightness variations corresponding to the periods of the excited modes. There is a rich array of possible periods with values ranging from about 100 to 1000 seconds. Mode periods present in the light curve are detected by undertaking a Fourier analysis of the time series light curve; theoretical models of the star can be refined with this information. The Fourier analysis needs to take into account such things as finite length, data gaps and the presence of noise.

Keywords: white dwarf stars, pulsators, light curves, Fourier analysis, amplitude power spectra, prewhitening, false alarm probabilities

1. Introduction

The time series and their analyses covered here derive from those white dwarf stars that exhibit light variability as a consequence of their undergoing pulsation. In order to provide a context for this work it is appropriate to outline some key astrophysical details.

Many stars at particular evolutionary stages pulsate due to instabilities created by escaping radiation generated in the hot interiors interacting with partially ionized outer layers in the star. The pulsations cause surface brightness variations; recording and analysing the light curve of a pulsator enables investigators to determine internal properties of the star.

The pulsating white dwarfs are particularly productive targets for these investigations.

1.1 White dwarf stars

White dwarf stars [1] are the cooling remnants of about 98% of all stars. They are extremely dense as they have about the mass of the Sun compacted into a sphere about the size of the Earth. Typical densities are of the order of a million times that of water. Nuclear reactions have ceased and the inward pull of gravity in the object is offset by a pressure created by the quantum behavior of the (separate) dense electron gas.

This (degeneracy) pressure is independent of temperature, and hence provides permanent support down to the final “black dwarf” stage with temperatures

approaching absolute zero. As the object cools the radiation output reduces in intensity and is centered around increasingly longer wavelengths. Eventually, the paucity of radiation emitted from the surface of this extremely cold object will render it invisible.

1.2 Pulsating white dwarf stars

During their long cooling sequence white dwarfs pass through several temperature regimes in which they are observed to pulsate [2]. Something like 80% of white dwarfs have a core consisting of a mixture of carbon-12 and oxygen-16 nuclei (plus the separate electron gas) overlaid by thin shells of helium and then hydrogen on top. Another significant group has only a thin helium surface layer overlaying the core.

At surface temperatures around 12,000 K for hydrogen atmosphere white dwarfs (25,000 K for the helium atmosphere types), the escaping radiation from the core interacts with a partially ionized subsurface layer, leading to the pulsations. The pulsations in most stars are called p-modes as they involve pressure variation propagation that results in radial material motion. This mechanism is the same as that which enables the propagation of sound in fluids.

Due to the extremely dense white dwarf material, the predicted periods for this type of oscillation are in the seconds domain. Although white dwarf p-modes have been searched for they have never been observed, probably due to their very low amplitudes (if they exist) combined with the difficulty of making suitable sub-second contiguous observations (even using 6 m class telescopes).

1.3 Gravity mode pulsations

The pulsations that are observed are nonradial gravity modes (g-modes). These involve largely angular movement of material in the star and have periods in the 100–1000 seconds domain [2]. They involve buoyancy variations in the stellar material; a simplified visual model of the mechanism involves an object floating on the surface of a fluid that will undergo vertical oscillations if it is pushed into the fluid. These oscillations only occur in a gravitational field, and in a spherical star they can only be nonradial.

The significant period difference between the two types of pulsation is directly connected to the larger forces involved in the compressibility (squeezing) of the dense white dwarf material versus the smaller buoyancy forces created by density differences.

The g-mode oscillations in white dwarfs were a serendipitous discovery in the late 1960s; theory predicts a large number of these g-modes with different periods. In the pulsators some of these modes are sufficiently excited such that their presence is detectable. The oscillations in effect result in temperature waves on the surface that produce observed variations at the mode frequencies in the light output of the star. An observed light curve of the white dwarf will therefore contain frequencies that yield information about the mode periods.

Detecting the different driven modes in a particular pulsating white dwarf allows comparison with predicted values from a detailed numerical model of the star [2]. Subsequent refinement of the model is then possible. In this way, detecting the various pulsation modes allows the observer to peer inside the otherwise hidden regions. Thus, one can measure such structure quantities as core chemical content and the masses and thicknesses of the overlying material shells. In some cases, the stellar rotation rate can be inferred.

1.4 Light curve time series

The raw materials for the time series are contiguous samples of the light output of a target star of interest. This is achieved using an instrument capable of fast photometry attached to a telescope [3, 4]. Ground-based telescopes dominate these observations, although space-based observations have become increasingly available.

Given the expected mode periods in white dwarfs, sampling rates of 10–20 seconds are typically employed. For the ground-based observations, gaps in the measurements due to the intrusion of clouds and the appearance of the Sun, as well as a signal random noise contribution, need to be accounted for in the analysis, along with the finite data length. Satellite observations can simplify these issues by providing uninterrupted data over extended intervals, although telescope aperture size is often a disadvantage.

The most useful information obtained from a light curve for modeling purposes is the periods of the detected modes. Hence, processing the time series in the optimum way in order to obtain a power spectrum is the primary goal. Since the amplitude of a driven pulsation results from a complicated nonlinear limiting process in the star, and is therefore more difficult to model, the detected amplitudes are primarily of use in differentiating between real and noise peaks. The phases of the different oscillations are in general not very useful, so a power spectrum is usually the ultimate analysis goal.

However, there is one case when the phase of a particular pulsation in an observation set is useful and that is when one is attempting to observe the frequency stability of the mode over an extended (e.g., multiyear) period. As long as one avoids cycle counting errors, tracking the measured phases over successive observation sessions can provide a sensitive test of the frequency stability of the mode.

The phase values for each observation set are best determined by least squares fitting sinusoids to the data. Then a predicted sinusoid maxima in a set is compared with the observed values from set to set to check for stability or any consistent changes. This is referred to as the Observed minus Calculated (O – C) method; for an example, see [5], but this will not be developed further here.

1.5 Time series examples

The observation of two types of pulsating white dwarf star obtained by the author will be used as time series examples in this article.

- a. MY Aps (WD1) is a relatively bright example of the hydrogen atmosphere white dwarf pulsators.
- b. QU Tel (WD2), a fainter target, is a member of the helium atmosphere pulsating white dwarf class.

Henceforth we will use the WD1 and WD2 labels to distinguish between the two objects rather than the astronomical variable names provided.

2. The raw light curve time series

The desirable time series form prior to undertaking a Fourier analysis consists of the light intensity excursions (i.e., modulation) about an average running mean of zero. The absolute detected light values are affected by such things as stellar

distance, telescope aperture and detector efficiency; these quantities are not important in the analysis and certainly not in Fourier transform space.

A “typical” quality observation set is plotted in **Figure 1**. These data were obtained on the night of 12 July 2004 using a three-channel photometer [3] attached to a one meter telescope at the University of Canterbury Mt. John Observatory (UCMJO) in the South Island of NZ.

This photometer enables the measurement of the amount of light passing through three small identical circular apertures in the vicinity of the particular target object (WD2 here). Three identical photomultipliers operated in photon counting mode were employed in each “channel” to digitize the light intensities, and each plotted point in the graph corresponds to a 10 seconds integration of the counts.

The blue points correspond to the variable white dwarf, the green points to a nearby (brighter) constant comparison star, and the red points to the sky background light from a clear patch of sky. The comparison star data show the effect of observing through a changing amount of airmass as the object transits across the sky, along with the occasional intrusion of cloud, primarily around 15.5 hours.

The sky data show the effect of changing background light (on a moonless light in this case). Note that the white dwarf and sky points are plotted using the same scale. The target star here is rather faint and about half the light intensity is contributed by the sky background. For optimum measurements it is important to use an aperture of an adequate size so that part of the image formed by the telescope optics (the point spread function) is not obscured.

Prior to processing the variable star data using Fourier techniques, one requires a time series of light curve modulation values, so conversion of the variable star data to fractional excursions about a running mean is required.

The first step, as detailed in the figure, is to divide the sky-corrected variable data by the sky-corrected comparison data and multiply by an appropriate constant

Figure 1.

Raw time series data obtained on the pulsating white dwarf WD2 on 12 July 2004 from the University of Canterbury Mt. John one meter telescope. The data were obtained using a photometer attached to the telescope that enabled the monitoring of three regions of the sky defined by small circular apertures [3]; identical photomultiplier tubes operated in photon counting mode measured the light intensities. The blue, green, and red plots correspond to the target white dwarf, a nearby (brighter) comparison star and a blank section of the sky, respectively. The mauve plot is the result of the indicated transform aimed to eliminate observing artifacts.

to yield a time-series curve with the correct mean value. This yields the mauve points plotted in **Figure 1**.

The second step is to convert light curve intensities to fractional deviation values about the running mean. A 3 hours section of the resulting data is presented in the middle panel (2) of **Figure 2**. A visual inspection of this plot reveals that along with a noise contribution there are oscillatory signals present.

Two intensity deviation measures are commonly utilized in this work. They are fractional or percent deviations, and millimodulation intensities (mmi) such that 10 mmi is equivalent to a 1% deviation. The latter measure is useful given the fact that real signal variations below 1% are detectable and often feature.

The light curve data displayed in **Figure 1** is of sufficiently high quality that almost no transformation steps beyond those listed above were required. However, useable time series light curves can be recovered from observations obtained in less than ideal photometric conditions by using a number of other techniques.

Such things as deleting obviously bad integrations and dividing the observations into separate groups to be “stitched together” later can be used. These manipulations are best accomplished interactively using a graphical display. As part of this process, the author developed a program [6] that enabled the interactive fitting of cubic splines to the light curves in order to remove “unwanted” low frequency variations. By implementing before and after Fourier transforms while manipulating the data, the user can directly see the impact of the changes and either accept or reject them.

The analysis will then of course be blind to the presence of possible low frequency signals, but ensuring that there is no impact on mode periods of ~ 100 seconds or smaller will not impact the analysis. Some judgment is required in the whole process of extracting an optimum time series from the raw observations that can then be Fourier analysed.

The bottom panel (3) in **Figure 2** is a segment of observations made of WD2 in July 2003 using a charge coupled device (CCD) instrument camera attached to a 6.5 m aperture telescope in Chile (the Magellan Clay telescope). The CCD takes electronic images of the small field of view and in this case both the white dwarf and a nearby comparison star were captured, allowing reductions to proceed similarly to those described above.

The CCD exposure times were 20 seconds followed by a 10 seconds gap in order to allow the (serial) reading out of the individual CCD pixel charges and digitizing

Figure 2.
 Three hours sections of reduced light curve time series suitable for undertaking a Fourier analysis. The top panel (1) corresponds to WD1 and the middle (2) and bottom (3) panels to WD2. Note that the vertical scale is the same for the three panels and the major difference between the two WD2 plots is the telescope aperture size.

each amount. The shuttering of the CCD is required to prevent contamination from incoming photons while reading out. This means the system only received light from the stars for 2/3 of the observing time (see below). However, the much larger “light grasp” of this telescope provided significantly enhanced signal to noise measurements as can easily be seen by examining the plot.

It is clear that there are at least two closely spaced frequencies in the light curve by observing the *beating* (sinusoidal variation of the amplitudes). One can use both the period of the oscillations and the beating period to estimate the two frequencies and their separation in frequency space. We will leave this as an exercise and simply refer to the appropriate (amplitude) power spectrum in **Figure 3**.

The time series in the top panel (1) of **Figure 2** is a segment of the light curve of WD1 obtained in May 2013 using a CCD photometer attached to the Mt. John one meter telescope. This CCD photometer [4] is a frame transfer CCD that obviates the necessity of shuttering the device while reading out the previous exposure. This is accomplished by using a double sized CCD chip with one section masked off and used as a storage area. The analogue charges in the exposed CCD area can be rapidly moved to the storage area, and the next exposure proceed while the pixels in the storage area are being read out and digitized.

Each point in the plot corresponds to a 20 seconds exposure. The analogue charge shifting time is a small fraction of the exposure time so no shuttering is required. Minimal photon contamination occurs due to the rapid movement of the analogue charge.

It is clear from the time series plot that WD1 exhibits at least one pulsation mode with both a smaller period and a smaller amplitude compared with the WD2 oscillations.

Figure 3.

Four amplitude power spectra (DFTs) resulting from a Fourier analysis of the full time series light curves corresponding to the segments depicted in **Figure 2**. The observing window functions (see text) are presented alongside the DFTs. These functions show the impact of time series length and gaps on the DFT. Panel 1 results from observations of WD1 over a six night interval. Panels 2a and 2b correspond to observations of WD2, with 2a resulting from the light curve in **Figures 1** and 2b the full six nights of observation. Panel 3 corresponds to two successive nights of the WD2 light curve plotted in the bottom panel of **Figure 2**.

The introduction of CCD-based photometers (see for example, [4]) in the last few decades has made the acquisition of quality time series data a lot simpler; among other things all the star intensities and background measurements are made with effectively identical instruments. For sparsely populated observing fields the relevant intensities can be extracted from the images using synthetic aperture techniques [4], and if all images are recorded, the parameters can be varied offline to obtain optimum results.

3. Fourier analysis of the reduced light curves

Having suitably prepared the light curve time series, the next step is to perform a Fourier transform. As mentioned previously, the primary targets of the analysis are both the frequencies and amplitudes of the detected modes, with the phases mostly being of no interest. Therefore we require a power spectrum of the data calculated at specified frequency values.

It is much more enlightening to work in amplitude space, so the square root of a power spectrum is what is used in this work. Sometimes this result is called a periodogram. These amplitudes then match more closely the values one gets by least squares fitting sinusoids to the data; this will be discussed in the prewhitening section below.

Henceforth, we will simply describe the amplitude power spectrum used in this work as a discrete Fourier transform (DFT). Clearly we need to choose a set of frequencies at which to determine the DFT values.

It is informative to view each amplitude Fourier transform frequency point as a cross correlation of the data with a sine function defined by the chosen frequency and a variable phase factor. One obtains the correct amplitude value by varying the phase factor so as to obtain the maximum value. The standard algorithm does not actually do this, but instead performs the cross correlations separately with sine and cosine functions with fixed (zero) phases, and computes the square root of the sum of the squares of the two amplitudes thus determined. This is then the value of the (amplitude) DFT at the chosen frequency.

The complete DFT of interest is the set of amplitude values calculated for all the chosen frequencies.

This raises the interesting question of the choice of frequencies, the sampling rate in the time domain and the Nyquist frequency. For a given (uniform) sampling interval (δt) in the time domain, in principle there is frequency information in the DFT up to the Nyquist frequency defined by $f_N = 1/2\delta t$. This result follows from the fact that only two contiguous samples are required to define each sinusoid.

So for example, the choice of a light curve sampling value of 10 seconds yields a Nyquist frequency of $f_N = 1/20 \text{ Hz} = 5 \times 10^4 \text{ } \mu\text{Hz}$. Note that with periods of 10^2 – 10^3 seconds of interest in the light curves, it is convenient to use the $\mu\text{Hz} = 10^{-6} \text{ Hz}$ frequency scale. Thus a typical mode period of 200 seconds corresponds to frequency of 5000 μHz .

Given current computational speeds it is more informative in the Fourier analysis procedures to use light curve sampling intervals that yield a Nyquist frequency significantly higher than the highest expected frequency of interest in the light curve.

So for the white dwarf pulsators with frequencies up to $\sim 10^4 \text{ } \mu\text{Hz}$ of interest, 10 seconds ($f_N = 5 \times 10^4 \text{ } \mu\text{Hz}$) and 20 seconds ($f_N = 2.5 \times 10^4 \text{ } \mu\text{Hz}$) are suitable sampling times.

3.1 DFT versus FFT

A version of the very efficient Fast Fourier transform (FFT) algorithms was published by Cooley and Tukey in 1965 (see for example, [7]), and these computational methods became widely used in the following decades.

In effect the FFT algorithm trades increased algorithmic complexity and a certain lack of flexibility for (often vastly) improved computational times when calculating a Fourier transform. The algorithm is most efficient for a uniform sequence of samples with a number that is a power of 2, e.g., 2^M with M an integer; it removes many redundant DFT calculations by successively dividing the sequence into halves such that at the core one has a 2×2 transform; efficient computation steps are invoked to regain the full transform.

Compared with a direct calculation using the DFT method for the same number of frequency points, the computational speed difference between the two methods increases rapidly with sample number. For a sequence of length $N = 2^M$, the computing time for the FFT increases as $\sim N \log_2 N$ compared with $\sim N^2$ for the DFT.

However, in spite of this speed advantage, where practical it is advantageous to use the direct DFT for the Fourier analysis discussed here. And, the large advances in computing speeds have made this choice nearly always practical. The FFT algorithm requires uniform samples, and the frequency values and separations are defined by the number of time domain samples. This can be quite limiting and there is only one case in the analysis discussed here in which the author introduced use of the FFT (see conclusion section).

3.2 DFTs and the window function

DFTs of the time series light curves obtained for WD1 and WD2 are presented in the four panels in **Figure 3**, along with the associated window functions in the panels on the right. The vertical axes use the millimodulation amplitude (mma) scale, such that 10 mma corresponds to a signal amplitude of 1%, while the horizontal axes use the measure $\mu\text{Hz} = 10^{-6} \text{ Hz}$.

The dominant pulsation mode in WD1 (panel 1) has a peak amplitude of less than 1% at 6 mma. The two primary pulsation modes in WD2 (panel 3) have larger amplitudes ~ 13 mma. (note that the smaller amplitudes for WD2 displayed in panels 2a and 2b result from light contamination from a nearby faint star that could not be separated from the target image during the aperture photometry measurements made with the smaller aperture Mt. John telescope.)

The window function used here is an extended version of the window function that appears in the standard Fourier literature, which depicts the effect of a finite segment length on either the Fourier series or Fourier transform.

Using the simplest example, one can view a finite length time series of length T as being an infinite time series multiplied by a “top hat” function with a value 1 between $t = -T/2$ and $t = +T/2$ and otherwise zero. This will abruptly turn the signal on and off and therefore represent the actual time series of finite length T .

Fourier theory shows that multiplying two signals in the time domain corresponds to a convolution in the Fourier domain (and vice versa). The Fourier transform of a top hat is a sinc function of the form $\sin x/x$, so any frequency delta function peak in the spectrum will be expanded to the shape of a sinc function. In turn squaring these shapes to produce an amplitude power spectrum will result in peak shapes like the window function in panel 5a in **Figure 3**—as the time series here is that reduced from **Figure 1**, which is essentially a continuous section of 10 seconds samples over an ~ 9 hour interval.

What this result is telling us is the fact that in Fourier frequency space we are endeavoring to decompose the signal into sine and cosine functions that extend to $\pm\infty$, so we require a range of frequencies to interfere and turn the time series signal on and off.

For the real white dwarf time series that can have multiple and varied gaps, the shape of the window function is most readily determined by calculating values at all the data points of a (noiseless) sinusoid with some frequency (e.g., 5000 μHz here) and then computing the DFT of this function in the same way as for the data.

Only the shape and frequency scale is important, so the window functions for the four DFTs in **Figure 3** are plotted about a central value of zero. The expanded detail in the vicinity of peaks in a DFT can be inspected with the window function shape in mind.

For ground-based observing under generally good conditions, the basic frequency scale and shape of the symmetrical window function is dictated by two factors. The overall length of the time series determines the width of the central peak and the daily gaps dictate the accompanying satellite peaks. This is nicely depicted in the three window functions 5a, 6 and 5b in **Figure 3**.

The 5a window corresponds to a basically uninterrupted time series of length ~ 10 hours = 3.6×10^4 seconds, so the width of the central peak is the inverse of this which is about 28 μHz . The window in panel 6 results from two approximately 9 hours observing sessions on successive nights, so the central peak is reduced accordingly with added alias peaks offset by the frequency $(\text{day})^{-1} = 11.6 \mu\text{Hz}$. Panel 5b continues this trend with observations over six successive nights along with daily gaps.

The panel 4 window function derives from observations over six successive nights so it has a narrow central peak related to the total interval of ~ 5.5 days, the daily 11.6 μHz aliase peaks and additional peaks created by data gaps due to observing an alternative target.

4. Prewhitening the light curves

A very useful technique for identifying low amplitude frequencies in a light curve, especially when close to a larger amplitude one, is the procedure of prewhitening. To prewhiten a light curve of a given frequency one uses least squares fitting to determine the amplitude and phase of a given frequency, or more generally these quantities for a set of frequencies. One then computes a light curve based on these derived parameters and subtracts it point by point from the original light curve.

In this way, the frequency (or frequencies) are “removed” from the light curve. A DFT of this prewhitened light curve will reveal more details about the frequencies remaining.

Although least squares fitting can be used to determine the frequencies of interest, it is much simpler to use the frequency values determined from the DFT as then linear least squares fitting can be used to provide direct results. If the frequency values are also found using least squares then this unavoidably leads to a nonlinear least squares fitting problem. This is more complex and requires iterative techniques.

On the other hand, a sine function with a known frequency but an unknown phase can be expanded into the sum of sine and cosine functions, both with zero phases and unknown amplitudes. Such a combination of functions are suitable for linear least squares fitting.

Figure 4.

An expanded section of the WD1 DFT presented in **Figure 3** showing the structure around the dominant pulsation mode: $f_1 \sim 5190 \mu\text{Hz}$, period ~ 192 seconds.

Figure 4 demonstrates the utility of prewhitening using a section of the WD1 DFT displayed in the top panel of **Figure 3**. The top panel in **Figure 4** is an expanded version of the DFT around the dominant frequency f_1 . The impact of the observing window function is clearly evident with the additional peaks surrounding the central one. This particular structure results from the daily breaks, together with additional peaks created by observing interruptions during the night due to the monitoring of another target star.

The red curve in the bottom panel of **Figure 4** shows the DFT of the light curve prewhitened by f_1 . This reveals two frequencies, f_2 and f_3 , of smaller amplitudes either side of the main peak. These peaks are also affected by the window function. The DFT with no prewhitening has been added to the plot as a dotted blue curve.

The green curve displaying essentially noise is a DFT of the light curve prewhitened by all three frequencies and clearly demonstrates the reality of the two smaller peaks. The horizontal dashed lines marked FAP 0.36 represent a significance level for the DFT peaks that is explained in the next section.

This analysis along with other WD1 observations will be published in Ref. [8].

5. Significance levels

As one examines peaks in a DFT that approach the noise level, it is important to adopt a quantitative criterion that facilitates differentiating between real and noise peaks. The False Alarm Probability (FAP) concept [9] provides such an intuitively straight forward approach.

Since we are dealing with statistical quantities, a particular FAP value (e.g., 0.001) will allow one to assert that there is a 1/1000 chance of a noise peak being as large as this value. So a peak exceeding this amount can be asserted to be real with a false alarm probability of 1/1000.

There are theoretical methods to estimate this number [9], but with the speed of modern computers a Monte Carlo approach is both enlightening and preferable (as with other uncertainty estimate requirements).

For the DFT FAP the method is as follows. First, the time series is prewhitened by the identified frequencies (in practice, the key requirement is to prewhiten by

Figure 5. Histograms of the maximum noise peaks obtained from the Monte Carlo noise analysis described in the text. The red histogram corresponds to the analysis of the WD1 DFT presented in panel 1 of **Figure 3**. The blue histogram corresponds to the analysis of the WD2 DFT in panel 2b of **Figure 3**, and the mauve histogram to the DFT analysis of a similar multi-night WD2 observation set.

the frequencies with larger amplitudes). Next, the data time stamps are randomly shuffled, a DFT calculated for this randomized time series and the largest noise peak selected. If this randomizing and DFT calculation process is repeated N times, with the largest DFT peak each time recorded, an ensemble of largest peaks is obtained. It can then be asserted that there is only a $1/N$ chance of random noise conspiring to produce a DFT peak as high as the largest peak in the ensemble.

It is instructive to plot histograms of these ensembles of largest peaks. Using a value for N of 1000, three examples are plotted in **Figure 5**. The red histogram corresponds to the May 2013 observing programme of WD1 yielding a value of 0.36 mma as the maximum value. Thus, there is a 0.001 probability of a random noise peak exceeding this peak value for the WD1 DFT. One can then assume with reasonable certainty that any peaks above this threshold are real with a false alarm probability of 0.001.

The blue and mauve histograms correspond to DFTs for WD2: the blue histogram relates to the DFT for WD2 plotted in panel 2b of **Figure 3** (see Ref. [6]), while the mauve histogram derives from multi-night photometry of WD2 obtained in June 2004 under similar conditions to that of July 2004. Although their shapes are different, both these histograms yield the same 0.001 FAP peak height of 0.93 mma.

Clearly the WD1 and WD2 histograms show that the WD1 data set is affected by significantly less noise; this is primarily the result of the white dwarf being a brighter target, but also in part due to the observations being obtained by a superior instrument. A comparison of the panels 1 and 2 in **Figure 2** qualitatively demonstrates a similar conclusion.

It should be emphasized that the FAP procedures outlined here relate only to random uncorrelated noise estimates. Nonrandom “noise” problems are mostly clearly identifiable on a case by case basis.

6. Conclusions

The analysis techniques described here relate specifically to those white dwarf pulsators that undergo coherent oscillations which are stable in frequency and phase

over long periods. This is the class of isolated white dwarfs that exhibit gravity mode pulsations, which provide a potentially rich spectrum of modes that may be excited.

The amplitudes of some of the frequencies in a few pulsators are variable over intervals of weeks or months, with modes appearing and disappearing. However, the underlying frequency values indicating the mode structure remains stable. And, as mentioned above, it is this mode structure that is used to improve the theoretical models of the target white dwarf.

A class of pulsator that does not present stable coherent pulsations involves a pulsating white dwarf in a close binary configuration such that material is flowing from the companion on to the white dwarf. Although these are interesting objects, the Fourier techniques presented here are not as powerful [10].

The two pulsator examples used here are relatively low amplitude pulsators, but other white dwarf pulsators display large and nonsinusoidal light curve variations. This means the DFT can contain both harmonic and combination frequencies. The WD2 DFT shows some low amplitude combination frequencies [6], but none of these effects are present in the WD1 DFT.

The mode frequencies deduced from the WD2 time series observations enabled improved modeling of the white dwarf as well as an estimate of its rotation period of close to 2 hours [6]. Using the theory of “rotational frequency splitting,” analysis of the WD1 data revealed a rotation period close to 12 hours [8]. In addition, when combined with single-site observations over several decades, at least one mode frequency was stable enough to enable the detection of an evolutionary cooling timescale [11].

6.1 Software techniques

The software tools described here have been developed in a Fortran 77/90 environment using the GFortran compiler provided by the GNU project. This personalized software has gradually evolved over a number of years, and the author has always found it more flexible to operate in a general programming environment rather than be limited by packaged products.

This certainly does not mean all software has been developed ab initio. Of considerable assistance has been access to the software routines both described and provided in the Numerical Recipes book series by Press et al. [12]. These authors in their last offering have also produced a version using the C++ programming language.

The Numerical Recipes routines, along with the comprehensive descriptions, provide an excellent base with which to develop appropriate software. The direct DFT described here is quite straight forward, but for other software developments the author has utilized Numerical Recipes routines (usually modified to suit) or ideas. Examples are: least squares fitting (both linear and nonlinear), randomizing routines, cubic spline fitting and FFT routines.

As mentioned previously, there was only one situation where the author decided the additional effort in using an FFT algorithm was worth it.

A multi-site multi-day observing campaign [6] on WD2 produced a total of 46,289 10 seconds samples over a 9 day interval. Calculating a false alarm probability at the 0.001 confidence level required 1000 DFT computations of a time series of this size. In order to exploit the power of the FFT, missing sample values were padded with zeros and the time series was extended to the size of $65,536 = 2^{16}$ by adding zeros. The overall computation time was vastly faster than using the direct DFT (which was also undertaken as a comparison and check).

Acknowledgements

The observational material discussed here was obtained at either the University of Canterbury Mt. John Observatory or the Magellan telescope facility in Chile, and the author thanks the relevant committees for the awards of observing time. The NZ Marsden Fund provided financial assistance for some of this work.

Author details

Denis J. Sullivan
School of Chemical and Physical Sciences, Victoria University of Wellington,
New Zealand

*Address all correspondence to: denis.sullivan@vuw.ac.nz

IntechOpen

© 2019 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/3.0>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

References

- [1] Tayler RJ. *The Stars: Their Structure and Evolution*. Cambridge University Press; 1994
- [2] Winget DE, Kepler SO. Pulsating white dwarf stars and precision asteroseismology. *Annual Review of Astronomy and Astrophysics*. 2008;**46**
- [3] Sullivan DJ. The New Zealand WET three channel photometer. *Baltic Astronomy*. 2000;**7**:100
- [4] Chote P, Sullivan DJ, Brown R, Harrold ST, Winget DE, Chandler DW. Puoko-Nui: A flexible high-speed photometric system. *Monthly Notices of the Royal Astronomical Society*. 2014; **440**:1490
- [5] Kepler SO et al. Measuring the evolution of the most stable optical clock G 117–B15A. *The Astrophysical Journal*. 2005;**634**:1311
- [6] Sullivan DJ et al. Whole earth telescope observations of the hot helium atmosphere pulsating white dwarf EC 20058–5234. *Monthly Notices of the Royal Astronomical Society*. 2008;**387**: 137
- [7] Fast Fourier Transform [Internet]. 2018. Available from: https://en.wikipedia.org/wiki/Fast_Fourier_Transform [Accessed: November 22, 2018]
- [8] Sullivan DJ et al. The hydrogen atmosphere pulsating white dwarf L 19–2 I: The 1995 WET campaign. *Monthly Notices of the Royal Astronomical Society*
- [9] Scargle JD. Studies in astronomical time series analysis II: Statistical aspects of spectral analysis of unevenly spaced data. *The Astrophysical Journal*. 1982; **263**:835
- [10] Chote P, Sullivan DJ. The post-outburst pulsations of the accreting white dwarf in the cataclysmic variable GW Librae. *Monthly Notices of the Royal Astronomical Society*. 2016;**458**: 1393
- [11] Chote P, Sullivan DJ. The hydrogen atmosphere pulsating white dwarf L 19–2 II: Detection of an evolutionary timescale. *Monthly Notices of the Royal Astronomical Society*
- [12] Press WH, Teukolsky SA, Vetterling WT, Flannery BR. *Numerical Recipes in Fortran*. 2nd ed. Cambridge University Press; 1992