

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

**The Hunt for Quasi-periodicities with
Wavelet and
Camera**

Claire Blackman

Thesis Presented for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Astronomy

UNIVERSITY OF CAPE TOWN

October 2008

Abstract

Quasi-periodic oscillations and dwarf nova oscillations occur in dwarf novae and nova-like variables, and have analogues in high-mass X-ray binaries and black-hole candidates. The low coherence of quasi-period oscillations and dwarf nova oscillations can make detection with standard time-series tools such as periodograms problematic.

This thesis develops tools to analyse quasi-periodic brightness oscillations. We review the use of time-frequency representations in the astronomical literature and consider which techniques are best suited to the low signal-to-noise ratio and low coherence of our data. We conduct a detailed analysis of a variety of time-frequency representations using artificial data that mimics our lightcurves. We conclude that the wavelet scalogram is the most effective time-frequency representation for analysing quasi-periodicities in low signal-to-noise data, as it has high time-frequency resolution, and is a minimum variance estimator.

By using both Morlet and Mexican Hat wavelets, we obtain simultaneous high time and frequency resolution. We introduce 'wavelet ridges' to measure instantaneous frequency, and calculate confidence contours based on an OLS fit of the red noise continuum. Empirical Mode Decomposition effectively prewhitens data and gives a detailed time-domain analysis of amplitude evolution. Monte Carlo simulations show that Empirical Mode Decomposition underestimates the amplitude of oscillations in the presence of noise. However, the amplitudes found are larger than those found previously using periodogram analysis, giving a new lower bound on the true amplitude.

We use the wavelet/Empirical Mode Decomposition method to re-analyse archival data from RU Peg, U Gem, TY PsA and VW Hyi, and find 62 new QPOs and 7 new long-period DNOs in VW Hyi. Relative to previous analyses, our method substantially improves the detection rate for QPOs. We show that long-period DNOs

in VW Hyi follow a linear increase in period during outburst and both standard and long-period DNOs decrease in amplitude during the initial non-linear period-increasing phase of outburst.

We show that both randomly driven damped harmonic oscillators and time-varying autoregressive decompositions can fail to detect QPOs and can misclassify DNOs.

Contents

1	Introduction	1
1.1	Cataclysmic Variables	2
1.2	Non-Magnetic Systems	3
1.2.1	Dwarf Novae in Quiescence	5
1.2.2	Dwarf Novae in Outburst	6
1.2.3	Superoutbursts	8
1.3	Magnetic Systems	10
1.3.1	Polars	10
1.3.2	Intermediate Polars	11
1.4	Brightness Fluctuations in Dwarf Novae	12
1.4.1	Flickering	13
1.4.2	Dwarf Nova Oscillations	13
1.4.3	Longer Period DNOs	16
1.4.4	QPOs	17
1.4.5	ZZ Ceti-type Pulsations	19

1.5	QPOs in CVs, HMXBs and BHCs	20
2	Historical Overview of Detection and Characterization of Optical QPOs	23
2.1	Visual Detection of QPOs	23
2.2	Detection through Periodogram Analysis	28
2.3	Statistical Models	30
2.4	Other Methods	32
2.5	Discussion	33
3	Time-frequency Analysis Literature Review	43
3.1	Wavelet Analysis	44
3.1.1	Wavelets for Image analysis	44
3.1.2	Wavelets for Spectroscopy	45
3.1.3	Early Wavelet use in Time Series Work	46
3.1.4	The WWZ	47
3.1.5	Other Wavelet Methods	50
3.2	Alternative Time-frequency Representations	52

3.3	Other Non-stationary Analysis Techniques	53
3.4	Discussion	55
4	Time-frequency Representations	57
4.1	Definitions	57
4.2	Stationarity	58
4.2.1	Stationarity of the Mean	59
4.2.2	Removal of Periodicities	60
4.2.3	Stationarity of the Variance	60
4.3	The Power Spectrum and the Periodogram	62
4.4	Theoretical Spectra	62
4.4.1	Spectra of Periodic signals	62
4.4.2	Spectra of Aperiodic signals	64
4.4.3	Spectra of Stochastic Signals	64
4.5	Estimating Power Spectra	66
4.5.1	Confidence Levels	67
4.6	Instantaneous Frequency	71

4.7	Time-frequency Representations	72
4.7.1	Important Properties	74
4.8	The Wigner-Ville Distribution	79
4.8.1	The Pseudo Wigner-Ville Distribution	80
4.8.2	Interference Terms	82
4.9	The Ambiguity Function & the Correlation Domain	84
4.10	TFRs with Variable IT Attenuation	88
4.10.1	The Cohen Class	88
4.10.2	Smoothed Pseudo WVD	91
4.10.3	Choi-Williams Distribution	97
4.10.4	Born-Jordan Distribution and Cone-shaped Kernel Representation	101
4.10.5	The Affine Class	106
4.10.6	Affine Smoothed Pseudo WVD	110
4.11	TFRs with No ITs	112
4.11.1	Gabor Spectrogram	114
4.11.2	The Wavelet Scalogram	115


4.12 Discussion	121
4.13 Software	126
5 The Wavelet Scalogram	129
5.0.1 Introduction	129
5.1 Choice of Analysing Wavelets	133
5.1.1 The Morlet Wavelet	135
5.1.2 The Mexican Hat Wavelet	136
5.1.3 Maximal Ridges and Instantaneous Frequency	138
5.1.4 Cone of Influence	138
5.1.5 Confidence Contours	139
5.2 Implementation	141
5.3 Wavelet Examples	142
5.3.1 A Simple Sinusoid	142
5.3.2 Gaps	146
5.3.3 Frequency Modulations	146
5.3.4 Phase Jumps	155

5.3.5	Amplitude Modulations	160
5.3.6	Multi-component Signals	160
5.3.7	Noisy Signals	170
5.4	Discussion	186
6	Empirical Mode Decomposition	191
6.1	Introduction	191
6.2	Method	192
6.2.1	Implementation	193
6.3	Testing	195
6.3.1	Gaps	196
6.3.2	Frequency Modulation	196
6.3.3	Phase Jumps	198
6.3.4	Amplitude Modulation	203
6.3.5	Multiple Components	204
6.3.6	Noisy Signals	204
6.4	Uses for EMD	208

6.4.1	Prewhitening using EMD	208
6.5	Discussion	213
7	Application of Time-Frequency Techniques to Four Cata-	
	clysmic Variables	215
7.1	RU Peg	215
7.1.1	s1674	216
7.1.2	s1677	217
7.1.3	s1680	217
7.1.4	s1683	218
7.2	U Gem	227
7.2.1	s1676	227
7.2.2	s1679	228
7.2.3	s1682	228
7.2.4	s1685 and s1691	229
7.3	TY PsA: run s3412	241
7.4	VW Hydri	245

7.4.1	Runs s0018, s0019, s0026 and s0030 (mid-fall to return to quiescence)	249
7.4.2	Runs s0085 and s0093 (Quiescence)	259
7.4.3	Runs s0102, s0105, s0110, s0111, s0112, s0122, s0124 and s0127 (a few days before and spanning superoutburst)	261
7.4.4	Runs s1277, s1307 and s1322 (beginning, middle and end of medium normal outburst)	267
7.4.5	Runs s1571 and s1616 (beginning and end of decline after superoutburst)	272
7.4.6	s0129	277
7.4.7	s2241	277
7.4.8	s2623	277
7.4.9	s6138	278
7.4.10	s6184	278
7.5	Conclusions	288
8	Further Analysis of VW Hydri Data	307
8.1	Introduction	307

8.2	DNOs	309
8.3	lpDNOs	314
8.4	Outburst QPOs	319
8.4.1	DNO-related QPOs	319
8.4.2	Unclassified Outburst QPOs	326
8.5	Quiescent QPOs	331
8.6	Discussion	332
9	Mathematical Models for QPOs	333
9.1	Stationary Auto-regressive Modeling	333
9.2	AR Modelling of RU Peg and VW Hyi	335
9.3	Time-varying Autoregressive Decomposition	337
9.4	TVAR Decompositon of VW Hyi	343
9.5	Discussion	351
10	Conclusion	353



1 Introduction

This thesis develops and tests analytic tools for finding and modeling quasi-periodic phenomena in cataclysmic variable stars. This chapter gives a brief introduction to cataclysmic variables, focusing on dwarf novae and the brightness fluctuations they exhibit, including dwarf nova oscillations and quasi-periodic oscillations. In chapter 2 we review observations and work on quasi-periodic oscillations in cataclysmic variables to date, and discuss the necessary requirements for any method which will enable their objective detection. These requirements include robust statistical detection in low signal-to-noise data and high time and frequency resolution over a broad frequency range. Chapter 3 discusses applications of time series techniques in other areas of astronomy to date. As our data have lower signal-to-noise ratios, and cover broader frequency ranges than the examples in the literature, we conclude that we need to analyse these techniques with appropriate synthetic data in order to choose the best tool. Chapter 4, therefore, is a detailed analysis of several of the techniques introduced in chapter 3, using synthetic data that mimics the key features of our data to test the precision and reliability of the methods. One of the most effective tools discussed in chapter 4 turns out to be wavelet analysis; we conduct an in depth study of this method in chapter 5, using synthetic data to clarify the behaviour of the wavelet in the presence of gaps, amplitude and frequency modulations and noise. In chapter 6 we introduce another useful time-frequency technique, empirical mode decomposition, which provides complimentary amplitude information to the wavelet spectrum. We again use synthetic data to test the behaviour of the technique, as well carrying out Monte Carlo simulations to check the accuracy of the detected amplitudes in the presence of noise. Us-

ing wavelet analysis and empirical mode decomposition, we re-analyse data from RU Peg, U Gem, TY PsA and VW Hyi in chapter 7, both confirming that these new techniques are reliable and informative, and developing a set of objective criteria for the detection of quasi-periodic oscillations. We use the new quasi-periodicities found in VW Hyi in chapter 8 to extend the current model of the occurrence and evolution of quasi-periodicities in VW Hyi. Chapter 9 discusses mathematical models for quasi-periodicities, including stationary and time-varying autoregressive processes. Chapter 10 draws together the key results of the thesis.

1.1 Cataclysmic Variables

For a general introduction to cataclysmic variable stars, see Hellier (2001), or Warner (1995a), which is also the main reference for this chapter.

Cataclysmic variables (CVs) are close binary systems with orbital periods typically of a few hours, consisting of a late main sequence secondary which is transferring mass onto a white dwarf primary. Due to their close proximity, the gravitational influence of the primary greatly distorts the shape of the secondary, and the secondary fills its Roche lobe. If the secondary expands to overflow its Roche lobe, or the orbit shrinks through loss of angular momentum, the secondary transfers mass into the Roche lobe of the primary through the inner Lagrangian point (L_1).

1.2 Non-Magnetic Systems

The presence of a magnetic field in the primary can influence the way in which gas is accreted onto the primary; we begin by discussing non-magnetic CVs - those with a primary field strength less than about 10^5 G.

As the stream of gas flows from L_1 towards the primary it is deflected by the Coriolis force and begins to orbit within the primary's Roche lobe. The gas forms a ring which, due to viscous shear, gradually spreads into a differentially rotating disc. The stream impacts onto the outer rim of the disc at supersonic velocities, creating a shock-heated bright spot which at optical wavelengths may radiate as much or more energy than all the other components (primary, secondary and disc) combined (see figure 1). Viscous shear between adjacent annuli causes most of the particles in the disc to move inward towards the primary, while a few particles move outward carrying the excess angular momentum.

Classical Novae (CN) have only one observed eruption, and are modeled as thermonuclear runaways of gas accreting onto the surface of the primary. **Recurrent Novae** are previously recognised CN that are seen to repeat their eruptions. **Dwarf Novae (DN)** have outbursts every four days to tens of years during which they brighten by between 2 and 5 mag above their quiescent level. These outbursts are modeled as a release of gravitational energy caused by a temporary increase in the rate of mass transfer through the disc. DN are further divided into three subtypes, according to the nature of their outbursts:

Z Cam stars show prolonged standstills about 0.7 mag below maximum




Figure 1: Schematic of a non-magnetic CV. The size of the arrows on the disc illustrate the fact that material orbits faster near the disc centre. Adapted from Hellier (2001).

brightness, during which outbursts cease for intervals of tens of days to years. **SU UMa stars** have occasional superoutbursts, in which the star achieves a brighter state by about 0.7 mag at maximum. Superoutbursts last about 5 times longer than ordinary outbursts. Superhumps - periodic brightenings of up to 0.4 mag with a period a few per cent longer than the orbital period of the system - are seen in all SU UMa stars during superoutburst.

U Gem stars include all remaining dwarf novae.

Nova-like Variables (NL) include all the non-eruptive CVs, for example pre-novae, post-novae and perhaps Z Cam stars effectively in permanent standstill.

Mass transfer is only possible if the secondary star fills its Roche lobe. When the secondary transfers mass to the primary, the transferred material loses angular momentum, so the binary separation increases to compensate, causing the secondary to detach from its Roche lobe and hence stop further

mass loss. However, gravitational radiation and magnetic braking extract angular momentum from the binary, ensuring steady, long-lived mass transfer. Most systems are thought to evolve from longer period binaries losing angular momentum through magnetic braking to shorter period systems governed by gravitational radiation.

Systems with orbital period $P_{\text{orb}} < 2$ h have mass-transfer rates characteristic of gravitational radiation, while systems with $P_{\text{orb}} > 3$ h have higher mass transfer rates characteristic of magnetic braking (Hellier, 2001). Between 2 and 3 h, however, there is a minimum in the orbital period frequency distribution. It is thought that while a system evolves across this so-called period gap the binary is detached, and is consequently too faint to be seen, as there is no mass transfer. While most U Gem stars have periods above the period gap, all SU UMa stars (with the exception of TU Men and several others recently found) have $P_{\text{orb}} < 2.1$ h.

1.2.1 Dwarf Novae in Quiescence

Apart from the periodic brightness fluctuations due to the rotational geometry of the system, quiescent CVs also exhibit stochastic flickering on a variety of shorter time scales, from less than a second to tens of minutes, with amplitudes of a few tenths of a magnitude (Fritz & Bruch, 1998). In a few extreme cases the amplitude may be over a magnitude, requiring a source of considerable energy; Bruch (1992) argues that it is only in regions close to the primary that there is sufficient energy to power the flickering.

In a few rare cases, quasi-periodic oscillations are seen in quiescent CVs. These oscillations typically appear for a few cycles, dying away to be re-

placed with another oscillation at a slightly different phase or period. They generally have a time scale of several hundred seconds, and an amplitude ranging from 0.025 mag in VW Hyi (Woudt & Warner, 2002a) to 0.5 mag in MV Lyr (Pavlenko, 1998).

1.2.2 Dwarf Novae in Outburst

DN outbursts are characterised by a rapid rise to maximum, followed by a more or less gradual decline back to minimum. While no two outbursts in a given DN are identical, one of the most striking features of DN outbursts is a bimodality in the *duration* of outbursts. Most pronounced for SU UMa stars, the presence of 2 possible outburst durations is nonetheless seen in both U Gem and Z Cam type stars (van Paradijs, 1983).

Each DN has a characteristic recurrence time scale for outbursts which may vary by a factor of 2 or 3. The shortest recurrence times are about 4 d (Kato & Kunjaya, 1995), while the longest recurrence times are probably restricted to $\lesssim 80$ y by the observational baseline (Warner, 1995a).

DN outbursts are thought to be due to the release of gravitational energy caused by the passage of matter through the disc. The model, originally proposed by Osaki (1974), depends on there being two equilibrium mass transfer rates at the surface temperature of the disc, both of which are unstable, and differ from the mass transfer rate from the secondary, \dot{M}_2 . The lower disc mass transfer rate, \dot{M}_Q , is associated with quiescence and the higher, \dot{M}_O , with outburst. During quiescence, when the hydrogen in the disc is un-ionised, if the rate of mass transfer arriving at the outer edge of an annulus is greater than \dot{M}_Q matter will build up in the annulus, in-

creasing its surface density Σ and viscosity. This in turn heats the annulus until hydrogen ionisation sets in, at which point any further rise in Σ results in a runaway rise in temperature, since opacity during ionisation increases $\propto T^{10}$. Σ stays roughly constant throughout the increase in viscosity and temperature until the hydrogen is fully ionised, when a new hot, highly viscous equilibrium is established at \dot{M}_O . The increased luminosity during this outburst phase is sustained by the higher inward flow of material caused by the higher viscosity of the ionised state. However, since \dot{M}_O is higher than \dot{M}_2 , the surface density drops as matter flows out of the annulus, dropping the temperature until the hydrogen ions begin to recombine and the lower opacity returns. The heat necessary to maintain the high surface temperature escapes, so the temperature falls, speeding the ion recombination and concomitant drop in viscosity. The annulus returns to the cold, quiescent state, and the cycle begins again.

Outbursts start at one annulus of the disc and propagate in both directions from the triggered annulus. Once the entire disc is in outburst, matter flows through the disc onto the primary until Σ drops sufficiently to stop the outburst in the outer annulus, after which a 'cooling wave' propagates inward returning the disc to quiescence. As the cooling wave sweeps inward, matter is transferred back into the larger radii; only about 10% of the matter in the disc is accreted onto the white dwarf during the outburst. The rest simply moves inwards during the heating phase, and outwards again during the cooling wave. The distribution of matter in the disc at the end of an outburst determines the evolution of the following outburst.

During quiescence the radius of the disc shrinks as more and more low specific angular momentum matter from the secondary is stored in the disc.

During outburst, as matter and angular momentum are accreted onto the primary, the disc expands, reaching maximum radius at outburst maximum.

The thermal instability causing the quiescence-outburst cycle in dwarf novae can only occur if the mass-transfer rate into the disc corresponds to a temperature in the partially ionized zone. If the flow is so high that it corresponds to the temperature of the outburst state then the disc will be in permanent outburst, as is the case for the nova-like variables. Z Cam stars straddle the border between DN and NLs: they show periods of outburst activity and prolonged standstills at slightly below outburst luminosity. For Z Cam stars the mass-transfer rate varies from just below \dot{M}_O (during the outburst phase) to just on or above it (during the nova-like standstill phase).

1.2.3 Superoutbursts

The SU UMa subclass of dwarf novae, apart from exhibiting the normal outbursts discussed in the previous section, also show superoutbursts, which are about 0.7 mag brighter than normal outbursts and have an extended plateau of brightness which lengthens the superoutburst duration by factors of 5-10 in comparison to normal outbursts.

Superoutbursts profiles are almost identical for a given star, although the beginning of the superoutburst may be disguised by a triggering normal outburst. Normal outbursts never, however, appear at or near the end of a superoutburst, confirming that normal and superoutbursts are linked phenomena. The recurrence time scale for superoutbursts (T_S) is generally longer than the recurrence time scale for normal outbursts (T_N), with $T_N < T_S \lesssim 14 T_N$. The number of normal outbursts occurring between successive

superoutbursts in a given DN can vary widely.

Superhumps, with a period a few percent longer than the orbital cycle, are seen near every superoutburst maximum. They have a period about 3% longer than the orbital period, and generally appear at maximum light, fading until they disappear near the end of the plateau. At their maximum amplitude of 0.3 to 0.4 mag they have an almost triangular profile, and are equally prominent in all SU UMa stars independent of inclination.

If the mass transfer rate of the secondary remains constant, and each normal outburst drains less material from the disc than has been fed into it since the last outburst, then the disc will grow in mass and radius over successive outbursts. While material in the inner disc orbits the white dwarf in circular orbits, the gravitational pull of the secondary distorts the orbit of material near the outer edge. The result is a tidal bulge lying slightly ahead of the line joining the primary and secondary. Usually, the tidal torque produced by the secondary on the bulge slows the material in the bulge, providing a drain of angular momentum and stopping the growth of the disc.

If, however, the radius of the disc is such that the material in the outer annulus is moving with one third of the orbital period, then the tidal bulge is amplified, rather than damped, every time the secondary passes. When the secondary sweeps past a tidal bulge, material in the outer elliptical orbits collides with material closer to the disc center in circular orbit, causing the disc to brighten. This resonance effect is only possible for systems with mass ratio $M_2/M_1 = q \lesssim 0.3$, i.e. short-period systems. SU UMa stars, which lie below the period gap, have a mass ratio small enough to allow superhumps; U Gem stars, almost all of which lie above the period gap, have $q \gtrsim 0.3$,

and consequently never exhibit superhumps.

In SU UMa stars, then, the disc grows during successive normal outbursts, until the radius is large enough to support the 3:1 resonance with the orbital period. The disc then becomes elliptical, and precesses, initiating superhumps. The enhanced tidal stresses increase the angular momentum drain on the disc, enhancing the inward flow of matter and sustaining the disc in its hot state. This enhanced accretion rate exceeds the rate of mass transfer, so the disc begins to shrink. The disc can however maintain its elliptical shape over a significant drop in radius, allowing most of the material to accrete out of the disc. Since superoutbursts involve so much of the disc's mass, they last much longer than normal outbursts.

1.3 Magnetic Systems

If the primary has a strong magnetic field (greater than 10^5 G) then accretion flow near the primary will be controlled by the magnetic field, and formation of the accretion disc may be partially or completely inhibited. The magnetic field will ensure that the primary rotates as a solid body. In some CVs the magnetic axes of the white dwarf primary may be inclined to their rotation axes. Magnetic CVs fall into two subclasses based on the strength of the magnetic field:

Polars, also known as AM Her stars, have magnetic fields so strong that disc formation is totally inhibited; the accretion stream is threaded onto the field lines close to L_1 , and accretes directly onto the primary near the magnetic poles (see figure 2). Depending on the location of the threading




Figure 2: Schematic of a Polar, from Hellier (2001), figure by Jens Kube.

region in the magnetic field, matter may be fed towards either one or both polar regions (Warner, 1995a). As in non-magnetic systems, accretion will tend to speed up the rotation of the primary, but counter-torques induced by the magnetic field usually hold the primary phase-locked with the secondary (Warner, 1995a). Brightness modulations at P_{orb} are observed in inclined systems as radiation from the accretion shock zones moves in and out of view.

Intermediate Polars (IP) have weaker fields, resulting in a truncated disc from which the inner disc region is absent. Matter is swept from the inner edge of the disc onto the magnetic field lines, accreting onto the primary in extended accretion arcs (Woudt & Warner, 2002a) (see figure 3).

Within the region of the disc on which the magnetic field acts, the gas flow drags the field lines with it, producing a torque on the primary. For a slowly rotating primary all the magnetic torque acts to spin the primary up, whereas for a faster rotator the change in rotation of the primary is determined by the net torque and may even be zero. While the magnetic field is generally not strong enough to synchronise the primary, the primary rotates as a solid body as in Polars. The complicated geometry of the system,




Figure 3: The pattern of field lines feeding from the inner edge of a disc onto the white dwarf. Straight lines from the white dwarf mark the spin axis and the offset axis of the magnetic dipole. From Hellier (2001), figure by Andrew Beardmore.

with the oblique rotator primary rotating asynchronously to the rest of the system, results in IPs display a range of periodic phenomena. IPs include DQ Her stars, which are rapidly rotating IPs lacking strong hard X-ray emission.

1.4 Brightness Fluctuations in Dwarf Novae

With these models in mind, we turn now to the short ($\lesssim 5000$ s) time scale oscillations seen in dwarf novae and nova-likes during outburst and super-outburst, and occasionally, quiescence. These brightness fluctuations come in several different types, ranging from almost coherent dwarf nova oscillations (DNOs), through quasi-periodic oscillations (QPOs), to stochastic flickering.

1.4.1 Flickering

Magnetic instabilities in the turbulent inner disc, and, in a few systems, the bright spot, give rise to stochastic flickering during quiescence and outburst (Bruch, 1992). Time scales for flickering range from a few seconds to several hours; the only pattern that is discernable in flickering is that longer fluctuations (up to several thousand seconds in length) have larger amplitudes. At much longer timescales, the amplitude of the flickering decreases. Flickering systems thus have red noise periodograms (Hellier, 2001), often characterized at low frequencies by several significant discrete peaks (Elsworth & James, 1982, 1986) which are, however, not reproducible (Bruch, 1992), and are due to the random structure of the flickering (Papadaki et al., 2006; Kozhevnikov, 2003).

1.4.2 Dwarf Nova Oscillations

Seen usually only during outburst, dwarf nova oscillations (DNOs) have periods of ~ 3 -100 seconds, with varying degrees of coherence ($Q = |dP/dt|^{-1} \approx 10^4 - 10^6$) (Woudt & Warner, 2002a). Their amplitude is generally small enough that they are only detectable through Fourier analysis, but they are occasionally visible directly in the light curve (Warner, 2004). DNOs typically appear during the rise to outburst, decreasing in period until outburst maximum, after which their period lengthens again, until they disappear during outburst decline (Woudt & Warner, 2002a). This correlation between DNO period and optical brightness suggests that higher accretion rates produce shorter periods (Warner, 2004).

DNOs are modeled as an extension of the intermediate polar model to white dwarfs with weaker magnetic fields of less than about 10^5 G. In the Low-Inertia Magnetic Accretor (LIMA) model the magnetic field is not strong enough to lock the envelope to the interior of the white dwarf, but is powerful enough to cause channeled accretion flow (Warner, 1995b; Woudt & Warner, 2002a). During accretion, only the surface equatorial belt of the white dwarf is spun up as higher angular momentum material accretes onto it (Paczynski, 1978), resulting in accretion curtains and shocks in the same manner as IPs, but with variable frequency as the belt is spun up during the high-accretion phase, and decelerated afterwards. As the primary spins the beams of high-energy radiation caused by the accretion shocks, reprocessed from the accretion disc, are alternately seen and hidden, producing periodic DNOs which decrease in period during the spin up phase, and increase in period during the spin down phase. Estimated rates for DNO period decrease and increase during a normal outburst are $\dot{P} \sim 10^{-4}$ and $\dot{P} \sim 3 \times 10^{-5}$ respectively, the different rates being due to the smaller inertia of the equatorial belt during the rise to maximum compared to the almost fully accreted mass of the equatorial belt during the fall from maximum (Warner & Pretorius, 2008).

Some CVs, notably OY Car, SS Cyg, TY PsA and VW Hyi, show considerably larger values of \dot{P} late in decline than that predicted above, requiring a mechanism for extracting additional angular momentum from the equatorial belt (Warner & Woudt, 2002). The inner radius of the disc r_0 is determined by a competition between the mass transfer rate through the disc, \dot{M} , and field strength of the primary. \dot{M} first increases and then decreases during outburst, decreasing and then increasing the inner radius of the disc. If the

rapid reduction of \dot{M} at the end of outburst causes r_0 to increase faster than the co-rotation radius r_{co} (i.e. the radius at which the primary and disc annulus angular velocities match), then as r_0 exceeds r_{co} gas will be centrifuged outwards on the field lines in a propeller action, instead of being accreted onto the equatorial belt (Warner & Pretorius, 2008). As gas is centrifuged out, angular momentum is extracted from the equatorial belt, increasing its rotation period and hence the period of the DNOs (Warner & Pretorius, 2008).

The classical DNOs, with period determined by the rotation period of the equatorial belt of the primary, are referred to as having a sidereal period. Occasionally, double DNOs are seen with their difference frequency equal to the observed QPO frequency (Warner & Pretorius, 2008). The additional DNOs are explained as reprocessing of the beam of high-energy radiation from the accretion zone by a prograde travelling wave in the inner disc (which will be discussed in more detail in section 1.4.4), and are referred to as having a synodic period (Warner & Woudt, 2006a).

Frequency doubling and tripling of DNOs, which is seen in several systems, is discussed in detail in Warner & Woudt (2006a). Briefly, transition from single-pole to two-pole accretion, and a change in the visibility of the accretion curtain due to the change in the radius of the inner disc, are two possible mechanisms that could produce the observed frequency doubling. Frequency tripling is explained as a rotational sideband, rather than as a physical frequency of the system.

On a shorter time scale than the rise and fall of outburst, DNOs exhibit stretches of stability followed by sudden small jumps in period. In general,

the coherence time scale of DNOs has a maximum at outburst maximum, decreasing as the system brightness decreases (Warner & O'Donoghue, 1989). Two models have been suggested as possible explanations of the period jumps seen in DNOs.

Magnetic field lines connecting the equatorial belt to an annulus of the accretion disc with different velocity will be wound up (Warner & Woudt, 2002). A 2π rotation of the primary relative to the disc connection annulus will break the field lines and it may be the breaking and subsequent reconnection of these field lines which causes the jumps in period observed in DNOs. Near maximum, \dot{M} is most stable so P_{DNO} will tend to be stable, and the time between reconnections becomes very long. At the end of outburst, however, \dot{M} decreases rapidly, increasing the rate of reconnection events, and decreasing the coherence of the DNOs.

Recently, it has been noted that both VW Hyi (Warner & Woudt, 2006a) and TY PsA show DNOs that alternate between the sidereal and synodic periods on a time-scale close to the beat period (or a multiple thereof) between the alternating values (Warner & Pretorius, 2008); it is possible that the short time scale period changes seen in all DNOs are simply the due to this switching.

1.4.3 Longer Period DNOs

A second class of DNOs, with periods about four times longer than the DNOs of the previous section, have been dubbed 'longer period DNOs' (lpDNOs) (Warner, Woudt & Pretorius, 2003). While showing similar coherence levels to the DNOs of the previous section, these DNOs appear to be indepen-

dent of the system luminosity, and can appear during quiescence (Warner & Woudt, 2003). These lpDNOs may be caused by two-pole magnetic accretion onto the bulk of primary rather than the equatorial belt, as the periods of the lpDNOs appears to be roughly twice the estimated spin rate for the white dwarfs in which this phenomenon is seen (Warner & Woudt, 2003). Alternatively, Warner & Pretorius (2008) suggest that lpDNOs may be due to shocks located at regions just outside the corotation radius, where gas threaded during propellering will be accelerated azimuthally and radially outwards, eventually reaching supersonic velocity and forming shock waves.

1.4.4 QPOs

These oscillations typically have periods of 50-5000 s, and are of much lower coherence than DNOs ($Q \approx 1 - 10$). There appear to be at least two different types of QPOs. The first type, with periods in the range 50-700 s, are linked to DNOs in that their periods tend to be about 15 times the period of DNOs observed at the same time, and their periods exhibit a similar period-luminosity relationship. The second group have longer periods, typically of the order of several thousand seconds, and have so far shown no connection to system luminosity. QPOs are seen mainly during outburst, but occasionally in quiescence.

The longest QPOs have periods close to the expected rotation of the outer rim of the disc (Woudt & Warner, 2002a). However, it is possible that all QPOs are caused by inner disc oscillations, either by intrinsic luminosity variations, or by variation of intercepted radiation from the central highly luminous regions of the disc by vertical oscillations in disc thickness. It is

the latter interpretation that we will focus on here.

A possible model for the DNO-related QPOs is that the QPOs are the result of reprocessing and/or obscuration of radiation from the white dwarf by a prograde travelling wave in the inner disc (Woudt & Warner, 2002a). In the inner disc, the disc flow speed is greater than the speed of the magnetic field lines, resulting in the deceleration of material to match the speed of the field lines. If, as is assumed in the DNO model, the magnetic field of the primary is not aligned with its rotation axis, there will be an azimuthal variation in the field strength around the inner disc. Material attaching to the strongest field lines will be decelerated more than the material attached to weaker field lines in the same annulus, resulting in a pile up of material in strong field regions. As material from the weaker field lines collides with the material ahead of it, it will heat and thicken the disc in this region, also possibly lifting gas out of the disc as field lines are bent over. The field lines will be dragged around the disc until they snap and reconnect to a lower energy position, releasing the blob of gas, which may perturb the disc sufficiently to start a travelling wave. Although the travelling wave may damp out quickly, it will be regenerated by the next reconnection event; the period of reconnection events will determine the recurrence period of the wave generation. Resonance could generate waves of considerable amplitude.

QPOs not related to DNOs have several possible origins. It is possible that some of the systems which show only long period QPOs are IPs from which large amplitude coherent oscillations and X-Ray modulation are suppressed (Patterson et al., 2002a). Since some systems showing these QPOs do however also show DNOs and lpDNOs, and consequently cannot be IPs, other mechanisms must be explored. It has been noted by Warner & Woudt (2003)

that the time scales of these QPOs is comparable to the rotation period at the outer edges of high mass transfer rate systems, although no models have been developed around this point. Variations in the mass transfer rate of the secondary due to non-radial oscillations similar to the ~ 5 min oscillations of the sun are another possible QPO source (Warner & Woudt, 2003).

1.4.5 ZZ Ceti-type Pulsations

Townsley & Bildsten (2003) show that the effective temperature, T_{eff} , of the white dwarf primary in a cataclysmic variable is determined by the accretion rate \dot{M} . There are thus a few systems that have sufficiently low \dot{M} for T_{eff} to fall in the instability strip, which for isolated DA white dwarfs is $\sim 11\,000$ to $12\,000$ K (Bergeron et al., 1995). These primaries exhibit ZZ Ceti-type non-radial pulsations, which are visible in optical observations against the accretion luminosity flux because \dot{M} is low (Warner & Woudt, 2005); for higher \dot{M} the accretion flux dominates the optical spectrum.

XX Ceti-type pulsation periods range from about 500 to 1600 s, with amplitudes up to 40 mmag (Warner & Woudt, 2005). The power spectra of these pulsations are very complicated, often consisting of a linear combination of the principle pulsation frequency as well as eigenfrequencies of other principle driving modes (Nilsson et al., 2006), which require long observations to deconvolve. Individual modes can vary greatly in amplitude over days, months and years.

While non-radial pulsations and QPOs occur on similar time scales, they are unlikely to occur in the same system, since the former are low \dot{M} phenomena, while QPOs are typically high \dot{M} phenomena. Non-radial pulsations

and QPOs also have very different appearances. Non-radial pulsations are on-going phenomena, while QPOs are intermittent and generally only observed during outburst. The periodograms of non-radial oscillations are often described by sequences of frequencies (e.g. Woudt & Warner (2004)); apart from period doubling and tripling, periodograms showing QPOs obey no such formulae. Low \dot{M} also suggests a long outburst interval; since DNOs and QPOs are usually outburst phenomena, there are relatively few opportunities to hunt for them. About 10 CV/ZZ Ceti hybrids have been discovered to date (Nilsson et al., 2006) and while one, HS2331, appears to exhibit permanent quiescent superhumps (Araujo-Betancor et al., 2005), none have so far been seen to exhibit DNOs or QPOs.

Warner & Woudt (2005) notes that shorter outburst intervals (associated with higher \dot{M}) may inhibit the growth of non-radial pulsations, and DB variables are unlikely to be detected since the high \dot{M} required to maintain the $\sim 25\,000$ K temperature of these objects gives a high accretion luminosity which swamps the primary flux component. There are however a few AM CVn stars which occasionally show low \dot{M} periods, so a few DBV primaries may be found, although none are currently known.

1.5 QPOs in CVs, HMXBs and BHCs

The ratio $P_{QPO}/P_{DNO} \approx 15$ seen in dwarf novae and nova-likes is also seen in neutron stars and black hole high-mass X-ray binaries (HMXBs) in the form of the ratio of low to high frequency QPOs at kHz frequencies (Woudt & Warner, 2002a; Warner & Pretorius, 2008). This correlation is seen to extend over nearly six orders of magnitude in frequency. During X-

ray bursts oscillations lying between the high and low frequency QPOs are seen, which resemble lpDNOs in that they are linked to the rotation period of the primary (Warner & Woudt, 2003). A high time resolution study of the neutron star system 4U1608-52 by Barret, Olive & Kluzniak (2005) has shown that the QPOs previously thought to be low coherence ($Q = dP/dt \approx 10$) oscillations are in fact short trains of high coherence ($Q \approx 10^3$) with jumps in frequency between bursts and amplitudes ~ 8 times that previously thought. This increases the similarity between high frequency X-Ray QPOs and DNOs in CVs even further. The similarities between oscillations in CVs and X-Ray binaries may have profound implications if, as suggested by Robertson & Leiter (2003), magnetic stars collapsing towards the BH end point have more in common with magnetic neutron stars and white dwarfs than has previously been realised (Warner & Woudt, 2003). A detailed study of QPOs in dwarf novae and nova-like is therefore not only of relevance to the study of cataclysmic variables, but to close binary systems with neutron star and black hole primaries as well. The advantage to studying CVs is that individual cycles in trains of DNOs and QPOs can be studied with relatively good signal to noise ratio, while the low fluxes from X-ray binaries require statistical analysis (Warner & Woudt, 2008).

2 Historical Overview of Detection and Characterization of Optical QPOs

This chapter is an historical overview of the existing work on optical QPOs in non-magnetic CVs: how, when and where they have been detected, and what attempts have been made to model them. Table 1 is a summary of these results: DNO/lpDNO related QPOs are shown in bold in the table, while other QPOs are shown in normal type. Historically, many oscillations have been misidentified. For example, DNOs detected in VW Hyi by Haefner et al. (1977) later became QPOs (Robinson & Nather, 1979), and subsequently were re-interpreted as lpDNOs (Warner & Woudt, 2003). We present those oscillations which are currently thought to be QPOs in non-magnetic systems, while being aware that at least some of these will be reclassified at some point. We also include in table 1 the method(s) used to detect and analyse the QPOs. We conclude the chapter with a discussion of the features that any method capable of reliably detecting QPOs must have.

2.1 Visual Detection of QPOs

The first QPOs, detected in the dwarf nova RU Peg during outburst by Patterson et al. (1977), were clearly visible in the light curve as an approximately 50 s modulation. The 50 s oscillations were dubbed *quasi-periodic* as it was clear that the while the interval between maxima was not constant from cycle to cycle, it varied only between about 40 and 65 s, exhibiting a preferred time scale (see figure 4). The magnitude of the oscillations varied




Figure 4: Portion of the light curve of RU Peg from run 1677. The individual maxima of the 50 s QPOs are marked to illustrate the lack of short-term coherence. The inset scale corresponds to a period of 51.7 s, the mean period during the illustrated portion of the light curve. From Robinson & Nather (1979)

from undetectable to 0.02 mag, although they never disappeared for more than 10 minutes at a time.

Since then QPOs have been detected visually in a number of CVs during outburst. Robinson & Nather (1979) reported QPOs in U Gem with a mean amplitude of 0.008 mag and a period of 79 s changing abruptly to 71 s. Robinson & Warner (1984) found visible DNOs and QPOs at 23 s and 253 s respectively during a normal outburst of VW Hyi, an SU UMa type CV, the amplitude of the 23 s DNO being largest near peaks in the 253 s QPO. SW UMa was observed by Robinson et al. (1987) to have 5 min QPOs during a superoutburst, while Kuulkers et al. (1991) found 19 and 26 min oscillations in WX Hyi during separate normal outbursts. Tovmassian (1988) observed QPOs in SS Aur in quiescence. VW Hyi has also been observed to have QPOs in quiescence (Woudt & Warner, 2002a), as have MV Lyr (Pavlenko, 1998) and GK Per (Patterson, 1991). Kato et al. (1992) found “super-QPOs” with amplitudes of about 0.2 mag in SW UMa during superoutburst. Large

amplitude QPOs have also been observed in VW Hyi (Warner & Brickhill, 1978), EF Peg (Kato & Starkey, 2002a), TV Crv (VSNET 21 Feb 2001) and DW Cnc (Uemura et al., 2002). The nova-like variable TT Ari (Hollander & van Paradijs (1992) and Tremko et al. (1996)) also shows visible QPOs. For a full list of QPOs observed to date, see Table 1, and also Warner (2004).

As discussed in the previous chapter, some QPOs are clearly linked to DNOs, while others are not. We will refer to the former as DNO-related QPOs, and to the latter as unclassified QPOs. Figure 5 shows the number of QPOs as a function of QPO period, for all the QPOs in table 1. It can be seen that the majority of DNO-related QPOs have periods less than 750 s, with most falling below 500 s, while unclassified QPOs tend to have periods above 500 s. The CV with a DNO-related QPO at 5000 s is nova GK Per, which has the longest known CV orbital period of 47.76 h. Figure 6 shows the distribution of QPOs as a function of the ratio of the QPO period to the orbital period of the CV. While DNO-related QPOs are evenly distributed below $P_{QPO}/P_{orb} = 0.125$, unclassified QPOs increase in numbers until $P_{QPO}/P_{orb} = 0.125$, but over one third have $P_{QPO}/P_{orb} > 0.125$.

The profiles of QPOs vary from sinusoidal to extremely jagged. Large QPOs often drag the intensity well below the smooth lower envelope of the light-curve when at their minimum (Woudt & Warner, 2002a) (eg SW UMa (Kato et al., 1992)). In a run of VW Hyi, taken on 23 September 1984, the growth and decay of a QPO over five cycles was followed by the growth and decay over 4-6 cycles of another QPO of similar period, but phase-shifted by 0.4 cycles. Warner, Woudt & Pretorius (2003) note "It is commonly observed that a train of 5 to 10 coherent QPO cycles is visible, before changing in phase or disappearing altogether."

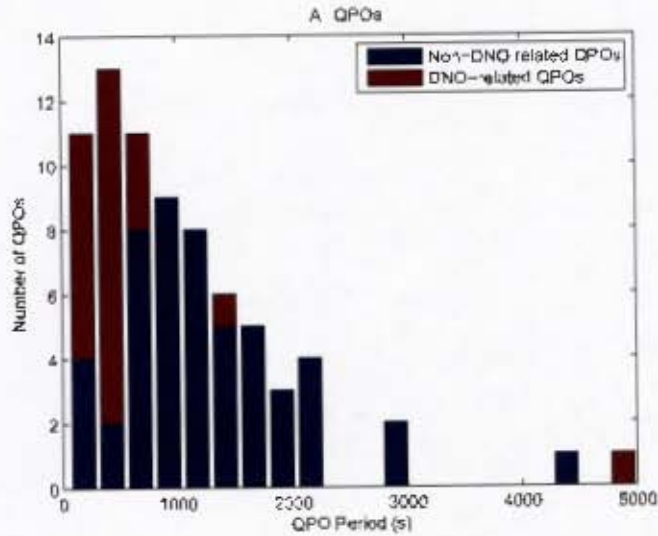


Figure 5: The period distribution of DNO-related and unclassified QPOs, across all CV types.

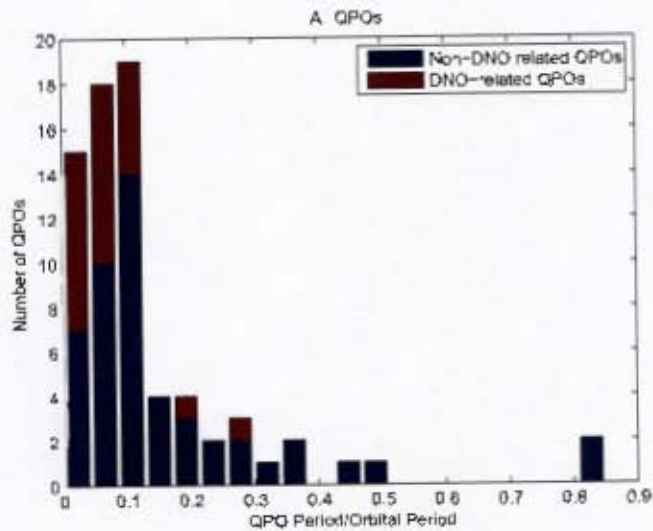


Figure 6: Number of DNO-related and unclassified QPOs, across all CV types, as a function of the ratio P_{QPO}/P_{orb} .

Woudt & Warner (2002a) noted that the DNOs visible in VW Hyi in some lightcurves were modulated by the QPOs, indicating a link between the two phenomena that was further strengthened by the discovery, discussed in Woudt & Warner (2002a), of a period-luminosity relationship for QPOs observed during outburst. P_{QPO} increased from 400 to 600 s while VW Hyi faded at a rate of 0.09 mag/h and RU Peg was observed to follow a large QPO period evolution over the 10 night decay to quiescence (Pretorius, Warner & Woudt, 2006). It has been shown (Woudt & Warner, 2002a) that for DNO-related QPOs, P_{QPO}/P_{DNO} is about 15. For a detailed overview of the interaction of QPOs and DNOs, see Woudt & Warner (2002a). QPOs are also observed to exhibit frequency doubling and tripling, as do DNOs. Woudt & Warner (2002a) noted that a 600 s QPO in VW Hyi halved in period near the end of a run. In Warner & Woudt (2006a), QPOs with first and second harmonics are observed, with the harmonics appearing at the same time as their DNO counterparts.

Pretorius, Warner & Woudt (2006) found an unclassified QPO in V1193 Ori that was relatively coherent over several nights, despite several changes of period; they note that the flickering previously observed in V1193 Ori (Bond et al. (1987), Warner & Nather (1988) and Papadaki et al. (2004)) may in fact be due to QPOs. Unclassified QPOs and DNO-related QPOs have been observed to co-exist; HX Peg during decline from outburst showed both a 347 s DNO-related QPO and a 746 s unclassified QPO (Warner, Woudt & Pretorius (2003) and Pretorius, Warner & Woudt (2006)).

In summary, then, QPOs occur in dwarf novae of all types, generally during normal or superoutburst, and very occasionally at quiescence. They also occur in nova-likes and old novae. QPO periods range from 50 to 5000 s,

and may change abruptly or gradually over the course of the observation. Two periods may be present simultaneously. Amplitudes range from 0.005 to 0.7 mag and may change during an observation. The phase may change abruptly, either with or without a change in period and amplitude. They may exhibit frequency doubling and tripling, and DNO-related QPOs and unclassified QPOs can coexist.

2.2 Detection through Periodogram Analysis

The periodogram (by which we mean the normalized squared Fourier transform of the time series, see section 4.3) is by far the most common analytical tool for detecting QPOs, perhaps since it is ubiquitous in the analysis of light curves. Tovmassian (1988) and Borisov (1992) detected QPOs using Deeming's method for calculating the Fourier transform (Deeming, 1975) and hence the periodogram. Semeniuk et al. (1987) computed periodograms using both Deeming's method and the maximum entropy method (MEM) proposed by Burg (1972); while the two techniques gave mutually consistent results, the QPOs were better seen in the MEM spectra.

QPOs may be visible in the periodogram, but not in the lightcurve, and vice versa. WW Cet, for example, has 263 s QPOs that are clearly visible in periodogram, but not obvious in the lightcurve (Pretorius, Warner & Woudt, 2006), while Robinson & Nather (1979) found 75 s QPOs in U Gem that were visible in the light curve but not in periodogram. Both BR Lup and AG Hya show QPOs clearly visible in both the lightcurve and the periodogram (Pretorius, Warner & Woudt, 2006).

While a full discussion of the problems associated with using periodograms to analyse a non-stationary signal is reserved for Chapter 4, we note that Skillman et al. (1995) found a quasi-period in a periodogram of MV Lyr which would have been ignored had Borisov (1992) not found the same period in previous observations. Skillman et al. (1995) argue that the real amplitude of the oscillation may be about two times greater than that estimated using the periodogram, since the periodogram technique assumes coherence.

A frequently used alternative to the periodogram is *phase dispersion minimisation* (PDM) (originally called periodogram analysis in the literature), proposed by Warner & Robinson (1972) and extended in Warner & Brickhill (1978). The mean and linear trends are subtracted from the data, and an epoch and trial period chosen. The phase of each datum point is calculated, and the entire phase diagram divided into overlapping bins. The points in each bin are added, and then divided by the number of points in the bin, giving a mean magnitude for each bin. The mean magnitudes from all the bins make up a mean light curve. Several methods have been suggested for deciding which mean light curve represents the best period. Warner & Robinson (1972) construct a discriminating function which reaches a maximum when the best period is found. Warner (1975) and Warner & Brickhill (1978) used this method on overlapping sections of the time series to show the evolution of QPOs in V436 Cen. Pezzuto et al. (1992) detected QPOs using a periodogram, but employed Warner and Robinson's method to improve on the period estimation. Ostriker & Hesser (1968) suggested choosing the value of the period which minimizes the dispersion about the mean. Stellingwerf (1978) further developed this method, giving it its current name. Hollander

& van Paradijs (1992) and Nogami et al. (1998) used PDM to detect QPOs successfully in TT Ari and SW UMa respectively. Stellingwerf (1978) notes that this method is particularly suited to non-sinusoidal time variations. However, Warner (1981) points out that phase dispersion minimisation will not give reliable results if the oscillation undergoes changes in phase and period, as many QPOs do.

2.3 Statistical Models

Robinson & Nather (1979) investigated QPOs in U Gem, SS Cyg, RU Peg, and KT Per. Power spectra of the light curves were calculated, which in some cases indicated the presence of QPOs. The autocorrelation function (ACF) was calculated by taking the lagged product of the pre-whitened light curve with itself. The QPOs present in the light curves were expected (and found) to give oscillations in the ACF with similar periods, evenly spaced. The ACFs were, however, affected by low-frequency trends and photon counting noise, scintillation and seeing which caused a spike at lags less than 30 s. From QPOs visible in the light curves, Robinson & Nather (1979) were able to deduce that the minimum amplitude necessary for detection by the ACF was 0.003 mag - considerably better than the sensitivity of periodogram.

Robinson & Nather (1979) continued the characterization of the QPOs by modelling the light curve as a second order autoregressive process (AR(2)). This equates the light curve with a driven damped oscillator in which the driving force is a series of random small displacements. If this model is used, the oscillations can be characterized by the mean angular frequency (or resonant frequency) ω , the mean amplitude variation σ and the damping (or

coherence) time γ . Robinson & Nather (1979) show that they can produce a light curve with many of the features observed in lightcurves exhibiting QPOs.

The periodogram of an AR(2) process can in principle be used to measure ω and γ , since it has a Lorentzian distribution peaking at ω , with a half-width at half-maximum of $\gamma/2\pi$. However, since the standard deviation of the power is equal to the power itself, the periodogram is very noisy, and often unsuitable for calculation. Robinson & Nather (1979) give an alternative method for finding ω and γ , using the ACF. For an AR(2) process the ACF is a cosine curve with exponentially decaying amplitude. The frequency of the cosine is the resonant frequency, while the decay time is the same as the decay time of the AR process. In practise it is necessary to have an ACF with prominent oscillations, and to remove the high-frequency spike and the low-frequency trends. Robinson & Nather (1979) note that the noise properties of the ACF are very difficult to evaluate, as noise at any lag is highly correlated with noise at nearby lags. This renders the ACF reliable only for lags less than 1/30 the length of the light curve.

Robinson & Nather (1979) found that in U Gem QPOs were not visible in any power spectra, and only in some ACFs, and only calculation of mean frequency was possible. For SS Cyg the mean frequency could be calculated from both periodogram and ACF (and the results agreed) but they could only use the periodogram for decay time.

Robinson & Nather (1979) concluded that QPOs are 'a stochastic process, not deterministic', since the randomly driven damped harmonic oscillator model successfully, they felt, characterized the QPOs. Robinson & Warner

(1984) used both the ACF and AR(2) modelling to find the coherence time of QPOs in VW Hyi, but Robinson et al. (1987) were not able to apply the same techniques as the periodograms were noisy and poorly defined. Pezzuto et al. (1992) could not extract the decay time from the ACF due to irregular behaviour of the function. Singh et al. (1993) used a periodogram of the ACF to detect QPOs, as it is usually less noisy than the periodogram of the original light curve.

The noise-excited damped harmonic oscillator model has also been used to model DNOs (Robinson & Nather, 1979; Hildebrand et al., 1981). As an alternative, Horne & Gomer (1980) and Cordova et al. (1984) model X-ray DNOs as a random walk in phase of a sinusoidal oscillator, but Jones & Watson (1992) note that this does not give simplified description of the data.

2.4 Other Methods

DNOs can be used as a signal for the possible presence of QPOs: a broad, low-amplitude 470 s signal in the FT of EC2117-54 would have been ignored if a double DNO with a beat frequency of 470 s had not been observed (Warner, Woudt & Pretorius, 2003).

Patterson et al. (1977) used an O-C (observed-minus-calculated) diagram of the times of observed maxima vs the times of periodic maxima to show that the light curve of RU Peg could not be modelled as a strictly periodic process with superposed random phase jitter. Large scatter and no pattern in the O-C also indicates a short coherence (Semeniuk et al., 1987). O-C diagrams

are frequently used in Woudt & Warner (2002a), Warner & Woudt (2002), Warner, Woudt & Pretorius (2003), Warner & Woudt (2006a), Pretorius, Warner & Woudt (2006) and Warner & Pretorius (2008). The choice of test frequency for the O-C diagram is important, and the 2π phase ambiguity of any point can lead to diagrams which appear chaotic (e.g Warner & O'Donoghue (1989)), but on reinterpretation have a clear structure (Warner & Pretorius, 2008).

Gnedin, Nagovitsyn & Natsvlishvili (1999) conducted an extensive study of QPOs in SS Cyg, using a formidable selection of analytic tools, including the Grassberger-Procaccia method for attractor dimension estimation, and wavelet decomposition followed by Hilbert analysis. They identify QPOs in three period ranges: 100-170 s, 2-5 min and 5-30 min. They also show the existence of a low-dimension strange attractor in the 1-10 min brightness oscillations of SS Cyg during quiescence, which they explain using an “impact oscillator” model.

2.5 Discussion

In order to study QPOs, we need to be able to detect them objectively and robustly. For possibly DNO-related QPOs, we also need to be able to detect the DNOs; Warner & Pretorius (2008) note that few DNOs are observed during the decline from outburst because “the DNOs often become too ‘incoherent’ to detect with Fourier transforms late in outburst.”

Visual identification is subjective, and cannot pick up low amplitude oscillations. Robinson & Warner (1972) point out that in some cases, QPOs

are barely visible in periodograms because the power is spread over a large range of frequencies and the periodogram is noisy, so the low-contrast QPO-induced hump gets lost in the noise.

On the other hand, Warner (1981) notes that “the combination of transitory oscillation of variable period can lead to an impressive periodogram (or power spectrum) containing well-resolved peaks whose widths are very little greater than that expected from the resolving power of the technique”, which can be very misleading. Periodograms are therefore also not reliable QPO detectors.

Robinson’s AR modelling relies on either the periodogram or ACF for calculation of characteristic time scales etc, and so is subject to the same restriction as both of these methods. It is also not clear that an auto-regressive process is the best model for QPOs.

We conclude that in order to detect QPOs, we need a method can identify relatively small numbers of cycles of oscillations buried in non-white noise, and allow us to track and measure the changes in the frequency and amplitude of the oscillations accurately. Once we are able to detect QPOs, the next step, if possible, is to find a mathematical model to describe them and shed light on the physics at work.

Table 1: Table of optical QPOs

CV	Type	Period (s)	Amp (mag)	Visual	Fourier	ACF	AR	O-C	Other	Ref
SS Aur	DN	1200-1800	0.7(?)	✓	✓	-	-	-	-	1
VW Hyi	DN(SU)	600, 833, 260?, 720?, 930-980	0.025	✓	✓	-	-	-	-	2
MV Lyr	NL	2256	0.5	✓	-	-	-	-	-	3
GK Per	N	380	0.03	✓	✓	-	-	-	-	4
Outburst										
RX And	DN(Z)	1000	-	✓?	-	-	-	-	-	5, 6
TT Ari	NL	900 - 2820	0.04 - 0.2	✓	✓	-	-	-	PDM	9, 10, 11, 12, 13, 14
WX Ari	NL	1180	-	-	-	-	-	-	-	15

Continued on next page

Table 1 – continued from previous page

CV	Type	Period (s)	Amp (mag)	Visual	Fourier	ACF	AR	O-C	Other	Ref
KR Aur	DN	400 - 900	-	✓	✓	✓	-	-	-	16, 17
CR Boo	AMC	300	0.02	✓	✓	-	-	✓	-	18
AM CVn	AMC	290,820	-	-	✓	-	-	-	-	19
OY Car	DN	320, 1500		×	✓	-	-	✓	-	18
V592 Cas	NL	1296	0.4, 0.05	✓	×, ✓	-	-	-	-	20
V436 Cen	DN	400 - 500	-	✓	-	-	-	-	-	21, 22
V442 Cen	DN	925	-	?	?	-	-	-	-	23
V842 Cen	N	750 - 1300	-	-	-	-	-	-	-	18
WW Cet	DN	263	-	✓	-	-	-	-	-	66
Z Cha	DN	585	-	?	✓	-	-	-	-	18
DW Cnc	NL	2220, 4400	0.3	✓	✓	-	-	-	-	24

Continued on next page

Table 1 – continued from previous page

CV	Type	Period (s)	Amp (mag)	Visual	Fourier	ACF	AR	O-C	Other	Ref
GO Com	DN(WZ)	1980	0.5	✓	✓	-	-	-	-	25
TV Crv	DN	600	-	✓	✓	-	-	-	PDM	26
SS Cyg	DN	100 - 170, 730	0.003	×	✓	✓	✓	×	Hilbert, Wavelet, Chaos	28, 29, 30
V751 Cyg	NL	1230	0.03 - 0.3	✓	✓	-	-	-		30
EC0528-58	VY Scl	900 - 1560	-	✓	✓	-	-	-	-	31
EC2117	NL	500	0.06	✓	✓	-	-	-	-	18
AQ Eri	DN(SU)	280	?	?	✓	-	-	✓	-	18
U Gem	DN	77	0.008	✓	×	✓	-	-	-	32, 27
V533 Her	N	1400	-	-	×		-	-	-	18

Continued on next page

Table 1 – continued from previous page

CV	Type	Period (s)	Amp (mag)	Visual	Fourier	ACF	AR	O-C	Other	Ref
V795 Her	NL	1150	0.015	✓	✓	-	-	-	-	33, 34, 35,36, 37, 38, 15
AG Hya	DN(UG)	983	-	✓	-	-	-	-	-	38
VW Hyi	DN(SU)	400 - 600	0.001 - 0.02	✓	✓	✓	✓	✓	-	32, 39, 27, 40, 2, 22, 18
WX Hyi	DN(SU)	190, 1140 , 1560	0.002	✓	-	-	-	-	-	41, 18
HP Lib	AMC	290	-	-	-	-	-	-	-	42
BH Lyn	NL	1030	-	-	-	-	-	-	-	38
MV Lyr	NL	2820	0.03	-	✓	-	-	-	-	43, 44, 45,46

Continued on next page

Table 1 – continued from previous page

CV	Type	Period (s)	Amp (mag)	Visual	Fourier	ACF	AR	O-C	Other	Ref
BR Lup	SU	650	-	✓	✓	-	-	-	-	66
AH Men	NL	1100	0.1	✓	✓	-	-	-	-	47
TU Men	DN	313	-	-	-	-	-	-	-	48
BT Mon	N	1800	-	-	-	-	-	-	-	18, 49
CW Mon	DN	2200	-	-	-	-	-	-	-	50
NSV 10934	DN	1300	-	-	-	-	-	-	-	51
V2051 Oph	DN	486, 1800	-	✓	✓	-	-	-	-	52
V426 Oph	DN(Z)	1680	-	×	✓	-	-	-	-	53
V442 Oph	NL	1000	-	✓	✓	-	-	-	-	38
V1193 Ori	NL	649, 1720	-	✓	✓	-	-	-	-	66
EF Peg	DN(SU)	360, 1080	0.02	✓	✓	-	-	✓	-	20

Continued on next page

Table 1 – continued from previous page

CV	Type	Period (s)	Amp (mag)	Visual	Fourier	ACF	AR	O-C	Other	Ref
HX Peg	DN	347, 746, 1400 - 1900	-	✓	×	-	-	-	-	18, 66
LS Peg	NL	1240	0.04	×	✓	-	-	-	-	54, 55,56
RU Peg	DN	55, 160, 250, 360	0.02, 0.035, 0.15	✓	✓	✓	✓	✓	-	57, 27, 66
GK Per	N	5000	-	✓		-	-	-	-	58
TY PsA	DN(SU)	245	-	✓	✓	-	-	-	-	59
VZ Pyx	DN	390, 3000	0.005	?	✓	-	-	-	-	18
RX1643+34	NL	1000	0.03 - 0.3	✓	✓	-	-	-	-	38
V893 Sco	DN	350	-							18, 60

Continued on next page

Table 1 – continued from previous page

CV	Type	Period (s)	Amp (mag)	Visual	Fourier	ACF	AR	O-C	Other	Ref
RW Sex	NL	620, 1280	0.005	✓	✓	-	-	-	-	62
WZ Sge	DN(SU)	742	-	✓	✓	-	-	-	-	22
SU UMa	DN(SU)	2280	-	-	-	-	-	-	-	-
SW UMa	DN(SU)	280 - 370	0.05 - 0.2	✓	✓	-	✓	-	Monte- Carlo	61, 63
UX UMa	NL	650		-	-	-	-	-	-	64, 2
IX Vel	NL?	500	-	✓	×	-	-	✓	-	65
TU Vir	DN	1000	-	-	-	-	-	-	-	48

References for Table 1:

1 Tovmassian (1988), 2 Woudt & Warner (2002a), 3 Pavlenko (1998), 4 Patterson (1991), 5 Szkody (1976), 6 Patterson (1981a), 7 Patterson (1979), 8 Patterson et al. (1980), 9 Semeniuk et al. (1987), 10 Hollander & van Paradijs (1992), 11 Suleimanov et al. (1996), 12 Tremko et al. (1996), 13 Kraicheva et al. (1999), 14 Andronov et al. (1999), 15 Patterson et al. (2002a), 16 Singh et al. (1993), 17 Kato (2002b), 18 Warner, Woudt & Pretorius (2003), 19 Provencal et al. (1995), 20 Kato & Starkey (2002a), 21 Warner (1975), 22 Warner & Woudt (2002), 23 Marino & Walker (1984), 24 Uemura et al. (2002), 25 Kato & Hirata (1990), 26 VSNET 21 Feb 2001, 27 Robinson & Nather (1979), 28 Gnedin, Nagovitsyn & Natsvlishvili (1999), 29 Mauche (2002), 30 Patterson et al. (2001), 31 Chen (1994), 32 Haefner et al. (1977), 33 Zhang et al. (1991), 34 Zwitter et al. (1994), 35 Patterson & Skillman (1994), 36 Rosen et al. (1995), 37 Rodríguez-Gil et al. (2001), 38 Patterson et al. (2002a), 39 Warner & Brickhill (1978), 40 Robinson & Warner (1984), 41 Kuulkers et al. (1991), 42 Patterson et al. (2002c), 43 Borisov (1992), 44 Skillman et al. (1995), 45 Kraicheva et al. (1999), 46 Pavlenko & Shugarov (1999), 47 Patterson (1995), 48 Pretorius (Unpublished), 49 Smith, Dhillon & Marsh (1998), 50 Kato et al. (2003), 51 Kato et al. (2004), 52 Warner & O'Donoghue (1987), 53 Szkody (1990), 54 Garnavich & Szkody (1992), 55 Taylor et al. (1999), 56 Rodríguez-Gil et al. (2002), 57 Patterson et al. (1977), 58 Nogami, Kato & Baba (2002), 59 Warner & O'Donoghue (1989), 60 Bruch, Steiner & Gneiding (2002), 61 Robinson et al. (1987), 62 Hesser et al. (1972), 63 Nogami et al. (1998), 64 Nather & Robinson (1974), 65 Williams & Hiltner (1984), 66 Pretorius, Warner & Woudt (2006),

3 Time-frequency Analysis Literature Review

In our quest to find suitable methods for detecting and characterizing QPOs and DNOs, it is helpful to review the non-stationary time-series analysis used in areas of astronomy other than CV research. In this chapter we do not focus on the specifics of each technique, but rather on how it is used, and the type of data analysed. Where available, we included details about the spacing of the data, the noise level, the number of frequencies present simultaneously, the duration of the oscillation(s) and the nature of any non-stationarity.

Time-frequency representations, which allow a signal to be visualized in the time and frequency domains simultaneously, are an intuitive and instructive way to begin the analysis of any non-stationary signal. We begin, therefore, in section 3.1, with a review of applications of the wavelet transform, one of the most commonly used time-frequency representations in astronomy. Several other time-frequency representations have been known to be more effective than the wavelet transform for some types of data: we discuss several examples of these in section 3.2.

As discussed in chapter 1, a variety of mathematical models have been suggested for QPOs. Section 3.3 gives a brief discussion of several possible alternative models which have been successfully used to characterize non-stationary oscillations, both in astronomy and other disciplines. In section 3.4, we discuss differences between our data and the data covered in the literature (such as signal-to-noise ratio and frequency range), and the implications these differences have for our choice of analytical tools.

3.1 Wavelet Analysis

Two different wavelet techniques, the discrete wavelet transform (DWT) and the continuous wavelet transform (CWT) are both used in astronomy. Both techniques decompose a signal into short, damped oscillations of varying length and amplitude, called wavelets; the DWT decomposes the signal into orthogonal components, while the CWT allows a continuous decomposition. A variety of differently shaped analyzing wavelets is available, allowing the decomposition to be tailored to the data.

Wavelet analysis is used in different forms in several areas of astronomy, including image analysis, spectroscopy and time series analysis. We give only a brief overview of wavelet use in the first two areas, since most of the developments are not relevant for time series data analysis as they do focus on detection of periodic phenomena; for a more detailed introduction, and many useful references, see Scargle (1997).

3.1.1 Wavelets for Image analysis

Image analysis requires the use of 2-D wavelets. Starck et al. (1994) use wavelets for image reconstruction of two evolved stars from infrared speckle interferometry, while Damiani et al. (1997) develop a method for X-ray image analysis, used by Pizzolato, Maggio & Sciortino (2000), amongst others. Brosch & Hoffman (1999) resolve the shape of the LoTr 5 planetary nebula while Rosati et al. (1995) and Lazzati et al. (1999) independently develop fast, efficient, automatic image analysis for ROSAT data, used by many subsequent authors (for example Donahue et al. (2002)). Da Rocha & Mendes

de Oliveira (2005) and Adami et al. (2005) use the *à trous* wavelet transform to study intra-group diffuse light in compact groups of galaxies and intracluster light in the Coma cluster respectively. More recently, Masciadri & Riga (2004) use the (2-D) Mexican hat wavelet to identify exoplanets in direct imaging observations.

3.1.2 Wavelets for Spectroscopy

Moffat et al. (1994) investigate stochastic substructure in emission lines of Wolf-Rayet stars, using the Mexican Hat DWT, following the method of Gill & Henriksen (1990), who carried out wavelet analysis to detect hierarchical structure in the spectra of molecular clouds. An alternative technique for analyzing discrete stochastic components is given by Lépine & Moffat (1996), involving both the DWT and CWT with the Mexican Hat wavelet, and is used by many subsequent authors (e.g. Dessart & Owocki (2002) and Lefèvre et al. (2005)). Starck, Siebenmorgen & Gredel (1997) use the *à trous* WT to extract shallow emission or absorption bands from low signal-to-noise spectra. The method of Szatmary, Vinko, & Gál (1994) is used by Townsend (1999) to analyse periodicities in absorption-line spectra of early-type stars. Meiksin (2000) and Theuns & Zaroubi (2000) independently use the DWT with the high order Daubechies wavelet (most appropriate for smooth data) to analyse absorption properties of the Ly α forest. Madgwick et al. (2002) detect quasars in 2dF galaxy Redshift Survey via an automated search algorithm that localizes the broad emission line signature of quasars. Riera et al. (2003) use wavelet analysis of Fabry-Perót observations to study H α observations of HH100 jet; this same method is used by Hiriart, Salas & Cruz-González (2004) to study molecular hydrogen

kinematics in Cepheus A.

3.1.3 Early Wavelet use in Time Series Work

One of the earliest papers using wavelet analysis for astronomical time-series data is Goupil, Auvergne & Baglin (1990), which investigates brightness fluctuations in non-radially pulsating white dwarfs. They investigate both the phase and amplitude spectra of the CWT. In a companion paper (Goupil, Auvergne & Baglin, 1991), 3-D plots and contour diagrams are used to represent the scalogram (squared CWT) graphically. To investigate the effect of noise on the CWT, tests are carried out using white and non-white noise and a constant star. In all the stars analysed, the periodicities examined were clearly visible in the periodogram.

Scargle et al. (1993) use Haar wavelets in the form of a scalegram (the time average of the CWT) to investigate QPOs in high time resolution EXOSAT observations of the low-mass X-ray binary Sco X1. An extension of this method is detailed in Andronov (2003) and applied to the symbiotic star CH Cyg in Andronov & Chinarova (2003). The scalegram does not give time resolution of the data, rather, it is an alternative to the periodogram.

Szatmary, Vinko, & Gál (1994) investigate the CWT (with Morlet wavelet) of synthetic signals created to mimic features in long-period pulsating stars such as Mira and SR variables. In particular, they consider multicomponent signals and pseudo-periodic signals, concluding that the wavelet scalogram is an appropriate tool for investigating light-curves in the time-frequency plane. In Szatmary, Gál & Kiss (1996) and Szatmary & Vinko (1992), the light-curves of the SRa-type variables V Boo and Y Lyn respectively are

studied, using the techniques discussed in Szatmary, Vinko, & Gál (1994). They point out that the unevenly spaced data sets typical of long period variable light curves often result in a CWT containing spurious, confusing signals.

3.1.4 The WWZ

This problem is largely addressed by Foster (1996), who develops a modification of the CWT called the weighted wavelet Z-transform (WWZ) which has been used in many subsequent papers. The WWZ uses a modified Morlet wavelet, and includes a parameter that controls the trade-off between time and frequency resolution. The WWZ tends to have poorer time-frequency resolution than the CWT, since removing the effect of uneven sampling involves smoothing the data. Foster (1996) also introduces the weighted wavelet amplitude (WWA), and recommends using the WWZ to locate a signal frequency, and the WWA to estimate the amplitude of the signal. The calculation of formal errors for the WWZ is extremely complicated, and while Foster (1996) gives an estimate of the formal errors for the period and amplitude, he notes that the most trustworthy method is Monte Carlo simulations.

Although there are exceptions, the WWZ is generally used to analyse long-period variables such as SR- and Miras-types, whose lightcurves 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” (Foster, 1996).

In Andronov (1998), the WWZ is extended to obtain the time-averaged

frequency-dependent WWZ periodogram; this method is applied to semi-regular and Mira-type variables to investigate period switching, biperiodic pulsations and long-term phase and amplitude modulation of the pulsations.

Kiss & Szatmáry (1999) and Kiss et al. (1999) use the WWZ to provide a detailed analysis of some 93 red semiregular variables, monitored over several years by various amateur groups, showing that 44 have two significant periods, and 12 show three periods. Wavelet analysis was useful since the instability of the periods resulted in decreased amplitudes in the periodograms. Kiss et al. (2000) investigate a sample of bright semiregular variables exhibiting repetitive mode changes, significant amplitude modulations and multiple periods, monitored over 70-90 years. The amplitudes of the periodicities discussed varies between about 0.1 and 1 mag, with a noise level of about 0.05 mag (which we estimate from the periodograms in the paper.)

Arentoft & Sterken (2000) use the WWZ and WWA to confirm their periodogram results for XX Pyx, but do not use the WWZ to resolve the temporal structure of the oscillations. Percy & Kastrukoff (2001) also use both the WWZ and WWA, supplemented by Fourier analysis and autocorrelation analysis, to study period and amplitude changes in small-amplitude pulsating red giants, concluding that gaps in the data can enhance the apparent reality of an alias period.

McSaveney et al. (2002) use the WWZ to analyse the RV Tauri star U Mon, which has an oscillation amplitude of about 1 mag and a noise level of about 0.5 mag, while Zijlstra, Bedding & Mattei (2002) use 340 years of observations to investigate the period changes in Mira-type R Hyd, which

oscillates with an amplitude of about 2.5 mag and has an estimated noise level of about 0.3 mag. Using WWZ transforms of long-term data from the low amplitude δ Scuti star HD 52788, Zijlstra et al. (2004) show that the star exhibits highly unstable pulsations.

Zhou & Qiu (2002) use the WWZ to investigate the evolution of superhumps in DN 1XS J23939+0628; they note that both the superhumps and orbital humps are clearly visible in both the periodogram and the WWZ.

Zijlstra et al. (2004) use the WWZ to show that the evolving variable BH Cru increased its period by 25 per cent over 25 years, with the semi-amplitude increasing simultaneously, to a value of 1.25 mag (with noise level ~ 0.3 mag). In the same paper, new periods for several other SC/CS stars are determined, some showing period evolution in the WWZ and WWA. Templeton, Mattei & Willson (2005) study the period change in 547 Mira variables using AAVSO data, deducing average rates of period change.

Sabin & Zijlstra (2006) also use the WWZ to investigate the period evolution of semiregular variables. They find that the proximity of the different modes, which create overlapping and aliasing frequencies, can make analysis difficult. They also show that three main trends occur in the periods of stars on the AGB: sudden change, meandering behaviour and continuous change. Wright et al. (2006) analyse photometric data from V2067 Oph, and show that the intermittent periodicity could in fact be the result of random variations.

Frick et al. (1997) developed an alternative modified (Morlet) wavelet transform to the WWZ of Foster (1996), again to overcome the effect of irregu-

larly spaced data. They studied chromospheric variations in several lower main-sequence stars.

3.1.5 Other Wavelet Methods

The WWZ has clearly been an extremely useful tool giving considerable insight into non-stationary time-series, particularly for long period variables. For regularly spaced data, the CWT, which allows a choice of analysing wavelet, has also been used to good effect.

Barthès & Mattei (1997) use the Morlet wavelet to investigate the pulsation modes of the LPV star σ Cet, from long-term AAVSO visual observations. They observe two main components, both clearly visible in light curve and periodogram, with amplitudes ~ 3 mag and a noise level of about 0.4 mag. Chinarova & Andronov (2000) use periodogram analysis, wavelet analysis and “running parabola” scalegram analysis to tabulate the characteristics of the brightness variations of 173 semi-regular stars. Kiss, Szabó & Bedding (2006) use the CWT to investigate the brightness fluctuations of 48 red semiregular or irregular stars.

Fritz & Bruch (1998) analyse some 776 photometric data sets of CVs showing flickering, using Haar, Daubechies and Coiflets wavelets. Their choice of analysing wavelet reflects the characteristic shape of flickering flairs in CVs. Halevin et al. (2002) use the CWT to investigate the origin of QPOs and flares in several magnetic cataclysmic variables.

Otazu et al. (2002) detect two superimposed periodicities in X-ray data using the *a trous* wavelet algorithm of Starck & Murtagh (1994), reaching

a lower signal-to-noise values (~ 0.7) than standard periodogram techniques such as PDM and CLEAN. This technique is extended to include unevenly sampled data in Otazu et al. (2004). This method is aimed at detecting periodic components with non-sinusoidal profiles, rather than quasi-periodic components; once the components have been identified using the *a trous* algorithm, PDM is used to identify the period correctly.

González Pérez et al. (2003) use the Pauli and Morlet wavelets to analyse the pulsating planetary-nebula nucleus NGC 246; they note that while the Morlet wavelet gives the best frequency resolution, the Pauli wavelet gives the best temporal resolution.

Torrence & Compo (1998) provide one of the first discussions of significance testing for the CWT; their work is used in many disciplines, although it has been slow to find favour in astronomy outside of solar astrophysics, where it has been used to analyse magnetic field measurements (Boberg & Lundstedt, 2002), coronal variability (Marsh, Walsh & Bromage, 2002), high-frequency oscillations in coronal loops (Katsiyannis et al., 2002) and chromospheric UV oscillations (McAteer et al., 2004). Mathioudakis et al. (2003) investigate intensity oscillations in a white-light flare on II Peg, finding a period of 220 s. They introduce the criterion that at least three cycles are necessary to identify an oscillation, rather than a random, pulse-like variation. Lachowicz & Czerny (2005) investigate millisecond variability in Cyg X-1, while Lachowicz, Czerny & Abramowicz (2006) detect a possible 2:1 or 3:2 period relationship in QPOs seen in two AGN. Mitra-Kraev et al. (2005) use the Morlet wavelet with this method to analyse the first X-ray observation of an oscillation during a stellar flare.

Régulo et al. (2007) develop a method for the rapid detection of planetary transits, involving de-noising using the DWT, followed by detection using the Torrence & Compo (1998) method, to select the part of signal where the transit search was to be done. They use the Paul wavelet, whose shape is similar to that of the feature to be detected. Husser et al. (2005) and Jenkins (2002) also suggested the use of wavelets for planet detection.

3.2 Alternative Time-frequency Representations

The wavelet transform is not the only time-frequency representation possible, and several other methods have been known to give better time and frequency resolution for some types of data. For example, Kiss & Szatmáry (2002) use the wavelet transform and the Choi-Williams distribution to investigate the behaviour of R Cyg, and show that the Choi-Williams gives higher time resolution than the wavelet. The wavelet, Choi-Williams, and Zhao-Atlas-Marks distributions are used by Szatmary, Kiss & Bebesi (2003) to probe the period change in T UMi; they conclude that the Choi-Williams and Zhao-Atlas-Marks distributions give much clearer time-frequency representations than the wavelet.

Boyd et al. (1995) investigate the Gabor transform, and apply it to synthetic signals with multiple periods and time-varying frequencies, and photometric data from a roAp star. They show that what had previously been thought to be a modulation of the major frequency due to beating was in fact an abrupt shift to a lower power level. Belloni, Parolin & Casella (2004) use the Gabor transform to analyse a 7 Hz QPO in the atoll source 4U 1820-30, while O'Brien et al. (2001) use the Gabor transform to analyse Keck II

spectroscopy of mHz QPOs in Her X-1.

Howarth & Greaves (2001) use a moving window analysis similar to the Gabor transform to investigate the phase and amplitude variation over time for the Mira-type variable T Cas; this method, developed in Howarth (1991), relies on first identifying the periodicity to be analysed in the periodogram.

Marchenko & Moffat (1998) use an adaptive form of the Choi-Williams distribution to analyse irregularly spaced data from three Wolf-Rayet stars. Buchler & Kolláth (2001) investigate both the Choi-Williams and Gabor distributions of a synthetic signal mimicking R Sct, as well as decomposing the signal into its constituent components, and using the analytic version of the signal to find the instantaneous amplitudes and periods.

3.3 Other Non-stationary Analysis Techniques

Time-frequency representations have clearly proved a good starting point for the analysis of non-stationary time series data. However, it is often desirable to be able to model the nature of the non-stationarity. Analysis of lightcurves in phase-space is one complementary technique to time-frequency analysis; of particular interest is the detection and characterization of a low dimensional attractor, which could imply a chaotic process. Several methods of chaos analysis have been developed for univariate time-series.

The correlation sum method of Grassberger & Procaccia (1983) has been used to measure attractor dimension; as mentioned in the previous chapter, Gnedin, Nagovitsyn & Natsvlishvili (1999) find that the QPOs of the DN SS Cyg can be modelled by a strange attractor with characteristic time-scale of

one to ten minutes. Misra et al. (2006) show that the low-frequency QPOs of the black hole system GRS 1915+105 exhibit behaviour which may be consistent with a chaotic attractor. Cannizzo et al. (1990) find no evidence for deterministic chaos in the light curves of Mira, R Leo and V Boo. The Grassberger-Procaccia method has however fallen out of favour: Buchler, Kolláth & Serre (1996) show that typical data sets are too short to provide reliable estimates.

The global flow reconstruction method of Serre, Kolláth & Buchler (1996) has been applied to a wide variety of pulsations. Buchler, Kolláth & Serre (1996) show that the pulsations of R Scuti, an RV Tau-type star, can be described by an attractor of dimension 3.1, while Buchler, Kolláth & Serre (1998) and Kolláth et al. (1998) show that the lightcurve of AC Her is not consistent with multi-periodic behaviour, but is described well by an attractor of dimension <2.45 . Kiss & Szatmáry (2002) show that the period-doubling of R Cyg can be seen as the switching between a periodic orbit and chaotic state. Buchler, Kolláth & Cadmus (2004) investigate five irregularly pulsating stars, finding evidence that the light curves of all five are generated by low-dimensional chaotic pulsation dynamics.

Scargle et al. (1993) use the scalegram to detect self-similarity in Sco X-1, which is then modelled by a 'dripping handrail' model. This is an example of transient spatial chaos, rather than low-dimensional deterministic chaos. Bedding et al. (1998) use the WWZ and WWA to investigate mode-switching in R Dor, which they suggest may be due to weak chaotic effects.

Empirical Mode Decomposition (Huang et al., 1998), combined with Hilbert analysis, has been used to study, for example, rotation residuals of the solar

convection zone (Komm, Hill & Howe, 2001). Looking beyond astronomy, West, Prado & Krystal (1999) use time-varying autoregressive decomposition methods to investigate structure in nonstationary electroencephalographic traces, while complex demodulation is used to detect a 90-min cardiac biorhythm by Orr & Hoffman (1974). Econometric and control theory methods, such as Kalman filtering (Kalman, 1960), and regime change detection (Nyblom, 1989), may also prove useful for modeling QPOs.

3.4 Discussion

We have introduced a variety of techniques available for the analysis of non-stationary data, and discussed their use in astronomy to date. Time-frequency representations in particular have been used to study period evolution in a huge variety of astronomical sources; they would thus seem a natural place to start in our search for a detection tool for QPOs.

It is not clear *a priori*, however, which time-frequency representation is the most appropriate for our data. While wavelet analysis is the most common time-frequency representation, we have discussed cases where alternative representations are more effective. There are also significant differences between our data and the data analysed in the literature, which need to be considered when choosing an analysis technique.

In all the examples given above (and, indeed, in all the astronomical cases we have found that use time-frequency representations) the periodicities under discussion are clearly visible in the periodogram (and, in most cases, the light curve). Thus the time-frequency representation is used as an adjunct to the

periodogram, rather than as a detection tool in its own right. If we are to use the time-frequency representation to detect QPOs, we need to have a means of assessing the validity of our detection; apart from Boyd et al. (1995), only those papers using the Torrence & Compo (1998) method give any formal significance test of their results. The requirement of significance testing means that the statistics of the time-frequency representation must be well understood and amenable to computation.

The time-frequency representation must also be amenable to signal detection in low signal-to-noise data: our lightcurves often have signal-to-noise less than one in the QPO frequency range, whereas the examples in the literature usually have a signal-to-noise ratios (significantly) greater than one.

We also require that the time-frequency representation gives good frequency resolution over a broad range of frequencies, since we must be able to detect both DNOs and QPOs. Typically the frequency range must span periods from tens to several thousands of seconds, which is a broader range than any required in the examples discussed above. Ideally, the representation must also have good time resolution, since both DNOs and QPOs are more intermittent than the majority of oscillations discussed in this chapter, and understanding the temporal nature of their occurrence is an important aspect of this study.

Our data thus present a combination of characteristics not previously encountered in astronomy, making an exploration of the different time-frequency representations necessary before a choice of the best technique can be made.

4 Time-frequency Representations

In this chapter we give an introduction to the theory of the time-frequency representations mentioned in the previous chapter. We begin with an overview of stationarity, the discrete Fourier transform and instantaneous frequency. After a discussion of some desirable properties for time-frequency representations (TFRs) we introduce the archetypal TFR: the Wigner-Ville distribution, and its Fourier transform dual, the Ambiguity Function. Although the Wigner-Ville distribution has excellent time-frequency resolution, it suffers from interference terms which make it all but useless for analysing multi-component signals; the bulk of the chapter focuses on an analysis of time-frequency representations developed to overcome this problem, using synthetic signals which mimic the most important features of our light curves. We conclude with a discussion of our findings, a more detailed analysis of some the applications of TFRs in astronomical time series mentioned in the previous chapter, and an overview of the available software.

4.1 Definitions

Vectors and matrices will be shown in boldface, with ' indicating transpose and * indicating complex conjugate. Unless otherwise specified, f refers to frequency (in Herz), and ω to angular frequency (in radians/second). E stands for the expectation operator.

We follow the notation of Chatfield (1975) by denoting a stochastic process, which may be defined as a collection of random variables, as $\{X(t) : t \in T\}$ where T denotes the set of times at which the process is defined. For

continuous T , we will denote the random variable at time t by $X(t)$, while for discrete T we will use X_t . Similarly, a finite observed time series that is a single realisation of stochastic process $\{X(t)\}$ will be denoted $x(t)$, for $0 \leq t \leq T$, if the observations are continuous, and by x_t , for $t_i = t_0, t_0 + \Delta t, \dots, t_0 + (N - 1)\Delta t$, if, as is the case for our astronomical data sets, the observations are discrete.

4.2 Stationarity

We confine our discussion initially to stochastic signals which are stationary, by which we mean that there is no systematic change in the mean or variance in the time-series, and all strictly periodic variations have been removed (Chatfield, 1975).

More precisely, a time series is said to be *weakly stationary* if all realizations $\{x_t\}$ of the process $\{X_t\}$ have the same mean

$$\mu = E(X_t), \quad (1)$$

the (auto)covariance, $r(\tau)$, between any 2 samples from the same realization, x_t and $x_{t+\tau}$, is a function of the lag τ only:

$$r(t, t + \tau) = E\{(x_t - \mu)(x_{t+\tau} - \mu)\} = r(\tau). \quad (2)$$

and the variance,

$$\sigma^2 = E(X_t - \mu)^2 = r(0), \quad (3)$$

is finite ($\sigma^2 < \infty$).

The autocorrelation function $\rho(\tau)$ is defined as:

$$\rho(\tau) = \frac{r(\tau)}{\sigma^2} \quad (4)$$

and is one way of characterizing the time-dependent properties a stochastic signal.

Since astronomical data are seldom weakly stationary, it is often necessary to induce stationarity in the data before any analysis relying on stationarity can be undertaken; we give a brief discussion of the chief methods used.

4.2.1 Stationarity of the Mean

There are many different ways of removing the trend, including smoothing, filtering and differencing, all of which equate to applying a high-pass filter. Visual inspection will usually show any long term trend, which may be modelled by smoothing the data using a moving average filter or by fitting a low-order polynomial. The fitted trend is then subtracted from the original time series to give mean-stationary data.

Differencing the data, a technique favoured by statisticians, may introduce

artifacts if the trend is not described by a polynomial. It also reduces the amplitude of periodic components and introduces a phase-shift (Harvey, 1993). Since this phase-shift depends on the frequency of the oscillation, DNOs and QPOs which may be related in phase but have different frequencies will be shifted so that the phase correlation might be lost.

Since most of our lightcurves show a monotonic but non-linear trend, we have usually chosen to fit a second-order polynomial, and to subtract the fitted polynomial from the data. An alternative method for trend removal, using Empirical Mode Decomposition, is discussed in section 6.

4.2.2 Removal of Periodicities

Periodicities in the data due to orbital modulations can be identified in the periodogram (defined in section 4.3) and may be removed by fitting sinusoids of the correct frequency, and subtracting the fitted sinusoids from the lightcurve.

Data which are dominated by orbital modulations were prewhitened by removing the fundamental, first and second harmonics of the orbital period (where appropriate) using EAGLE, a FORTRAN programme written by Darragh O'Donoghue.

4.2.3 Stationarity of the Variance

Once the trend and any strict periodicities have been removed, it is necessary to check that the variance remains constant over the entire length of data.

The cumulative sum (cusum) of a series y_t , defined as

$$c_t = \sum_{j=1}^t (y_j - \mu) \quad (5)$$

is a useful tool for finding changes in the mean of the series. By plotting the cusum of the squared residuals of the de-trended data, changes in the variance can be found. For homoskedastic data (i.e. which show no change in variance) the cusum should oscillate around zero, showing no clear long term trends.

Figure 7 shows the prewhitened data from RU Peg run s1674, and the resulting (variance) cusum. The variance of the light curve increases suddenly at about 6000 s; this is clearly seen in the cusum; for times prior to this the slope of the cusum is negative, indicating that the variance is below the mean variance, while for times after the increase in variance the slope of the cusum is positive.

In theory, if the data are found to be heteroskedastic (i.e. have time-varying variance), they should be transformed to homoskedasticity if possible, or analysed in homoskedastic sections. In practise, if an oscillation such as a DNO or QPO begins or ends during a run, then the variance of the data is likely to change. It is not practical or, in many cases possible, to transform the data to homoskedasticity, so we have simply used the cusum as a indicator for the presence of possible heteroskedasticity.




Figure 7: Variance and cusum of RU Peg run s1674.

4.3 The Power Spectrum and the Periodogram

4.4 Theoretical Spectra

In developing the equations for the spectra of the different kinds of signals, we follow the approach of Djurić & Kay (1998). For alternative derivations, see Deeming (1975) or Blinichikoff & Zverev (1976).

4.4.1 Spectra of Periodic signals

Most functions can be expressed as a superposition of sinusoids of different frequencies. For periodic functions, this superposition takes the form of an infinite sum, which on discretization in time reduces to a finite sum. A

periodic, discrete-time signal $x_p(n)$, $n \in \mathbb{I}$, with period N , can be expressed in the frequency domain by the discrete-time Fourier series:

$$\hat{x}_p(f_k) = \sum_{n=0}^{N-1} x_p(n) e^{-i2\pi f_k n}, \quad f_k = \frac{k}{N}, \quad k \in \{0, 1, \dots, N-1\}. \quad (6)$$

The complex function $\hat{x}_p(f_k)$, called the discrete time Fourier series of $x_p(n)$, can also be expressed in the form

$$\hat{x}_p(f_k) = |\hat{x}_p(f_k)| e^{i\phi(f_k)}$$

where $|\hat{x}_p(f_k)|$ is called the amplitude spectrum of $x_p(t_n)$, and $\phi_p(f_k)$ is the phase spectrum of $x_p(t_n)$. Defining the average power of the signal to be

$$P = \frac{1}{N} \sum_{n=0}^{N-1} |x_p(n)|^2,$$

we have (from Parseval's theorem),

$$P = \frac{1}{N^2} \sum_{k=0}^{N-1} |\hat{x}_p(f_k)|^2.$$

Then $|\hat{x}_p(f_k)|^2/N^2$ represents the power contributed to the total power by the component of $x_p(n)$ at frequency f_k , and hence represents the power spectrum density of $x_p(n)$ (Djurić & Kay, 1998).

4.4.2 Spectra of Aperiodic signals

If the discrete-time signal $x_a(n)$ is aperiodic, and obeys the finite-energy criterion

$$\sum_{n=-\infty}^{\infty} |x_a(n)|^2 < \infty,$$

then we can extend the discussion in the previous section to include $x_a(n)$ by allowing $N \rightarrow \infty$ (Goswami & Chan, 1999). Then the discrete-time Fourier series (6) becomes the discrete-time Fourier transform (DTFT):

$$\hat{x}_a(f) = \sum_{n=-\infty}^{\infty} x_a(n) e^{-i2\pi f t_n} \quad (7)$$

The total energy of the signal is given by

$$E = \sum_{n=-\infty}^{\infty} |x_a(n)|^2 = \int_0^1 |\hat{x}_a(f)|^2 df,$$

the last equality again following from Parseval's theorem. Since $|\hat{x}_a(f)|^2 df$ is the contribution from the frequency band $(f, f + df)$ to the total energy of the signal, $|\hat{x}_a(f)|^2$ represents the energy density spectrum of the signal $x_a(n)$.

4.4.3 Spectra of Stochastic Signals

The assumption of weak stationarity enables us to use a single realization of a stochastic process to make inferences about the underlying process. In particular, we can use a single realization of the process to estimate the power spectrum $S(f)$ of the underlying process (Kendall, Stuart & Ord (1983), Timmer & Koenig (1995)). The discrete-time Fourier transform of

the autocorrelation function

$$S(f) = \sum_{\tau=-\infty}^{\infty} \rho(\tau) e^{-i2\pi f\tau}$$

is defined as the power spectral density (PSD) and can be shown, by Wold's theorem, to be equal to

$$S(f) = \lim_{N \rightarrow \infty} E \left\{ \frac{|\hat{x}(f)|^2}{2N + 1} \right\} \quad (8)$$

(Djurić & Kay, 1998), where $\int_{f_1}^{f_2} S(f) df$ represents the average contribution to the total power of the components of $\{x(n)\}$ with frequencies in the range $(f, f + df)$ (Hamilton, 1994; Djurić & Kay, 1998). Note that the auto-correlation function and the power spectrum form a Fourier transform pair; they both contain the same information, but while the auto-correlation function occurs in correlation space (sometimes referred to as time-delay space), the power spectrum occurs in frequency space.

We have shown that the spectra of deterministic periodic signals and weakly stationary signals are power densities, the former defined only on the discrete set of frequencies $\{0, 1/N, 2/N, \dots, (N - 1)/N\}$. Deterministic aperiodic signals have spectra which are *energy* densities, defined on the continuous interval $[0, fs/2)$ where $fs = \frac{1}{\Delta t}$ is the sampling frequency for integration time Δt .

4.5 Estimating Power Spectra

The most common estimator for the discrete Fourier transform of an infinite continuous-time signal $x(t)$ that has been discretely sampled at times $t_i = t_0, t_0 + \Delta t, \dots, t_0 + (N - 1)\Delta t$ is

$$\hat{x}(f) = \sum_{i=1}^N x_i e^{-i2\pi f t_i} \quad (9)$$

where the frequencies $f = n/N; n = 1, 2, \dots, [N/2]^1$ at which $\hat{x}(f)$ are calculated are called the Fourier frequencies (Frescura, F. et al., 2007). The periodogram $P(f)$, which is the most common estimator for the power spectrum $S(f)$, is then given by

$$P(f) = \frac{2\Delta t}{\sigma^2 N} |\hat{x}(f)|^2 \quad (10)$$

where we have used the normalization of Vaughan (2005). Calculating $P(f)$ at the Fourier frequencies ensures that the $P(f)$ are statistically independent (Koen, 1990).

Unfortunately, the periodogram of a single realization of the underlying process is not a consistent estimator for $S(f)$: if the number of data points is increased, $P(f)$ does not approach $S(f)$. Specifically, for an independently, identically, normally distributed signal, the periodogram at a given frequency, $P(f)$, is scattered around the true power spectrum, $S(f)$, following a χ^2 distribution with two degrees of freedom:

¹ $[N/2] = (N - 2)/2$ if N is even or $[N/2] = (N - 1)/2$ if N is odd

$$P(f) \sim S(f)\chi_2^2/2 \quad (11)$$

(Vaughan, 2005). The standard deviation of the periodogram is equal to the mean, resulting in a periodogram which fluctuates wildly around the true value (Timmer & Koenig, 1995).

Thus, while periodic signals in the data appear as peaks in the periodogram, noise in the data can produce large spurious peaks in the periodogram too. For data with a red noise structure, for example, which has more power at low frequencies, the result is a periodogram with many noise peaks at low frequencies, which can obscure peaks due to true periodicities. It is therefore necessary to be able to estimate the likelihood that a given peak is due to a periodicity in the data, rather than a random noise fluctuation. This is particularly important for QPOs, which occur in a frequency range which usually has a high level of noise.

4.5.1 Confidence Levels

If the population variance is unknown, and is estimated by $s^2 = \frac{\sum x_i^2}{N-1}$, Koen (1990) shows that the periodogram follows an F-distribution ($F_{2,N-1}$). For large N , $F_{2,N-1} \rightarrow \chi_2^2$ so in practice, the chi-squared distribution is often used even when the variance is not known.

If the underlying noise continuum is a white noise process, so that the power spectrum is constant, either the χ^2 or F-distribution can be used to estimate the likelihood that the highest peak in the periodogram is due random fluc-

tuation in the noise spectrum, rather than a periodic signal (Koen, 1990).

For non-white noise, the power spectrum of the noise continuum is not constant, and it is necessary to fit a model for the noise before confidence levels can be calculated. Vaughan (2005) notes that one of the problems with fitting a continuum spectrum to non-white noise is that the variance of a non-white process is heteroskedastic, so ordinary least squares fitting does not provide an efficient estimator of the spectrum. However, the variance of the periodogram is homoskedastic in log-space:

$$\log[P(f)] = \log[S(f)] + \log[\chi_2^2/2] \quad (12)$$

The logarithm of the periodogram ordinate (corrected for a constant bias introduced by the chi-squared term), is an unbiased estimator for the logarithm of the spectrum, and is distributed independently and identically (about the underlying spectrum) at each frequency (Vaughan, 2005).

Vaughan (2005) fits a linear function to the log-periodogram, since his data show a red-noise power spectrum well modelled by power law (i.e. linear in log-space). Flickering (discussed in section 1.4.1) also gives our data a red noise spectrum, but CV flickering has less power at very low frequencies than the X-ray data of Vaughan (2005); we therefore fit a low-order polynomial (generally between third and fifth-order) in log-space, using an OLS fitting programme.

The Kolmogorov-Smirnov test can be used as a goodness-of-fit-test for the fitted continuum model $\tilde{S}(f)$ (Vaughan, 2005); we use this to choose the

order of the polynomial which best fits the data. The Kolmogorov-Smirnov statistic, D , compares the cumulative probability distribution of a data set to the assumed theoretical cumulative distribution function of the underlying series:

$$D = \max_{-\infty < x < \infty} |\text{cdf}(x_t) - \text{cdf}(X_t)|$$

The significance of any observed nonzero value of D (as a disproof of the null hypothesis that the distributions are the same) is given approximately (becoming asymptotically accurate as N becomes large) by the formula

$$\Pr(D > \text{observed}) = Q_{KS}(\sqrt{ND})$$

where

$$Q_{KS}(\lambda) = 2 \sum_{j=1}^{\infty} (-1)^{j-1} e^{-2j^2\lambda^2}$$

(Press et al., 1988).

Thus, by comparing the ratio $2\tilde{S}(f)/P(f)$ with the χ_2^2 PDF, we can estimate the likelihood that a given peak is due to a real periodicity in the data, rather than a fluctuation in the noise (Press et al., 1988). Confidence intervals can also be calculated using the appropriate values from the χ_2^2 distribution, scaled correctly for the frequency at which the level is to be calculated. Figure 8 shows the fitted log-periodogram of VW Hyi run s0484, and the resulting periodogram fit with 95% and 99% confidence levels calculated assuming a χ_2^2 -distribution for the noise spectrum.

Vaughan (2005) points out that comparing the ratio $2\tilde{S}(f)/P(f)$ with the




Figure 8: Log-periodogram (top) and periodogram of VW Hyi run s0484.

χ_2^2 PDF assumes that the power spectrum is known exactly, whereas we are in fact using \tilde{S} , which is an estimated model. This alters the probability distribution of the ratio, but, for large numbers of Fourier frequencies (i.e. for large N , such as we have in our data) the probability distribution using the estimated power spectrum tends to χ_2^2 , so we have used χ_2^2 .

The Fourier frequencies at which the periodogram $P(f)$ is calculated ensure that for each frequency f , $P(f)$ is statistically independent of the periodogram ordinate at any other frequency. However, it is often convenient to calculate the periodogram at frequencies other than the Fourier frequencies. In particular, oversampling increases the sensitivity of the periodogram to weak signals. Oversampling can lead to an overestimation of the significance of peaks in the periodogram (Vaughan, 2005), since the confidence levels

discussed earlier apply only to the periodogram calculated at the Fourier frequencies.

4.6 Instantaneous Frequency

The Fourier coefficients $\hat{x}(f)$ are obtained by correlating the time series $x(t)$ with a sinusoidal wave $e^{i\omega t}$. Support of $e^{i\omega t}$ covers the whole real line, so $\hat{x}(f)$ depends on values of $x(t)$ for all $t \in \mathbb{R}$ (Mallat, 1998).

Thus if the frequency content of a signal is changing over time, then the power spectrum is of limited use, since it considers the average power at each frequency over the entire data set. We would ideally like a time-dependent analogue to the power spectrum, which enables us to see how the frequency content of the signal evolves with time. This requires the definition of a time-dependent or instantaneous frequency. There are several different definitions in use; we have chosen to use that of Mallat (1998).

The angular frequency ω_0 of a cosine modulation $x(t) = a \cos(\omega_0 t + \phi_0) = a \cos \phi(t)$ is the derivative of the phase $\phi(t) = \omega_0 t + \phi_0$. Generalizing this, a real signal $x(t)$ may be written as a time varying amplitude $a(t)$ modulated by a time varying phase ϕ :

$$x(t) = a(t) \cos[\phi(t)]$$

with $a(t) \geq 0$. The instantaneous (angular) frequency is then defined as the positive derivative of the phase:

$$\omega(t) = \phi'(t) \geq 0$$

(Gonçalves et al., 1997). Since there are many possible choices of $a(t)$ and $\phi(t)$ for any given signal, $\omega(t)$ is not uniquely defined relative to f . However, if we define the analytic part $x_a(t)$ of a signal $x(t)$ as the inverse Fourier transform of

$$\hat{x}_a(f) = \begin{cases} 2\hat{x}(f) & \text{if } f \geq 0 \\ 0 & \text{if } f < 0 \end{cases}$$

then $x_a(t) = a(t) \exp[i\phi(t)]$ and, since $x(t) = \text{Re}[x_a(t)]$, $x(t) = a(t) \cos \phi(t)$. We then call $a(t)$ the analytic amplitude of $f(t)$ and $\phi'(t)$ its instantaneous frequency (Mallat, 1998).

Whether this instantaneous frequency has any physical meaning, however, depends on a number of factors. Since the instantaneous frequency is defined as a single value function of time, a signal consisting of multiple modes cannot be analysed directly. It must first be decomposed into its constituent 'monocomponent' signals (Cohen, 1995). The question of what exactly a 'monocomponent' signal is, however, remains largely unanswered. Initially, bandwidth limited functions were suggested, but the bandwidth limitations necessary to give a meaningful instantaneous frequency have never been established (Huang et al., 1998). Boashash (1992) showed that if the real part of the Fourier transform of a signal contains only positive frequencies, then the instantaneous frequency is physically meaningful.

4.7 Time-frequency Representations

Ideally, we would like some way of finding the instantaneous frequency for each frequency component in our light curves. The top panel of 9

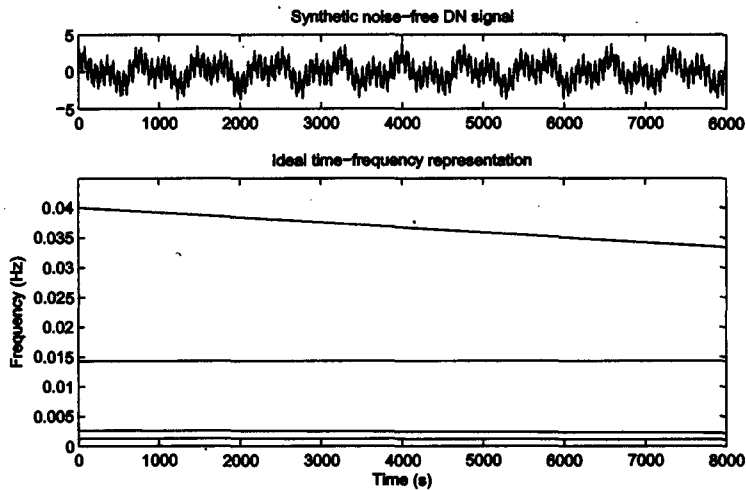


Figure 9: Top panel:

shows a synthetic light curve we have created which mimics a noise-free DN lightcurve with a first harmonic DNO changing linearly in period from 25 s to 30 s, fundamental and first harmonic QPOs starting with periods 750 s and 375 s respectively and tracking the period change of the DNO, and an lpDNO with a constant 70 s period. The light curve we have based this synthetic signal on is run s0484 of VW Hyi, discussed in section 7.4. For simplicity, the amplitudes of the constituent periodicities are all constant and equal. The integration time was chosen as 4 s, since this is the most common value for our data. The bottom panel of figure 9 shows an idealized time-frequency representation (TFR) of the signal: the top line shows the frequency evolution of the DNO, while the bottom two lines track the frequency change of the QPOs. The 70 s lpDNO is an horizontal line at 0.014 Hz.

4.7.1 Important Properties

While an ideal TFR such as the one in figure 9 is not realizable in practice, many effective approaches have been developed since the pioneering work of Gabor (1946), Ville (1948) and Page (1952).

Before considering specific TFRs, however, it is useful to consider the properties which are desirable in a TFR, with particular reference to the types of signal we wish to analyse. Hlawatsch & Boudreaux-Bartels (1992) and Auger et al. (1996), amongst others, give a detailed list of such properties; we follow here the formulation of the latter, and include only those that are relevant to our analysis. Williams (1998) gives a good discussion of the properties most relevant to analysis of noisy signals.

1. *Energy Conservation:* Since we are, in some sense at least, looking for a time-dependent energy spectral density, it seems reasonable to require that the TFR $T_x(t, f)$, when integrated over time and frequency, gives the total energy of the signal:

$$\begin{aligned} E_x &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} T_x(t, f) dt df \\ &= \int_{-\infty}^{\infty} |x(t)|^2 dt \\ &= \int_{-\infty}^{\infty} |\hat{x}(\omega)|^2 d\omega \end{aligned}$$

Note that we only consider finite energy signals, i.e. signals for which $E_x < \infty$.

2. *Marginal properties:* Following on from property 1, we would like $T_x(t, f)$, when integrated along frequency or time, to correspond to

the instantaneous power or spectral energy density respectively:

$$\int_{-\infty}^{\infty} T_x(t, f) dt = |\hat{x}(f)|^2$$

$$\int_{-\infty}^{\infty} T_x(t, f) df = |x(t)|^2$$

This requirement ensures that, if the signal is stationary, the time integral of $T_x(t, f)$ gives the (smoothed) power spectrum.

3. *Real-valued*: If we are using the TFR as a measure of instantaneous energy, then it should be real:

$$T_x(t, f) = T_x^*(t, f)$$

4. *Time covariance*: If we shift the signal $x(t)$ in time, then $T_x(t, f)$ should shift in time by the same amount, but otherwise remain the same:

$$y(t) = x(t - t_0) \Rightarrow T_y(t, f) = T_x(t - t_0, f)$$

5. *Frequency covariance*: If we shift the signal $x(t)$ in frequency, then T_x should shift in frequency by the same amount, but otherwise remain the same:

$$y(t) = x(t)e^{i2\pi f_0 t} \Rightarrow T_y(t, f) = T_x(t, f - f_0)$$

6. *Dilation covariance*: If we dilate the signal $x(t)$, then $T_x(t, f)$ should be dilated by the same amount, but otherwise remain the same:

$$y(t) = \sqrt{a}x(at) \Rightarrow T_y(t, f) = T_x(at, \frac{f}{a})$$

7. *Wide-sense support conservation*: if our signal is limited in time or frequency, then $T_x(t, f)$ should have the same limits in time and frequency:

$$\text{a) } x(t) = 0, |t| > T \Rightarrow T_x(t, f) = 0, |t| > T \quad (\text{time support})$$

$$\text{b) } \hat{x}(f) = 0, |f| > B \Rightarrow T_x(t, f) = 0, |f| > B \quad (\text{frequency support})$$

8. *Instantaneous frequency*: we would like the instantaneous frequency $f_x(t)$ of our signal $x(t)$ to be the average frequency at a given time of $T_x(t, f)$:

$$\frac{\int_f f T_x(t, f) df}{\int_f T_x(t, f) df} = f_x(t) = \frac{1}{2\pi} \omega(t)$$

9. *Unitarity*

To ensure energy conservation, we require that the scalar product is conserved from the time domain to the time-frequency domain (Moyal's formula):

$$\left| \int_{-\infty}^{+\infty} x(t) x^*(t) dt \right|^2 = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} T_x(t, f) T_x^*(t, f) dt df.$$

10. *Positivity*: If we are using the TFR as a measure of instantaneous energy, then it should be always positive. However, this requirement turns out to be incompatible with Properties 2, 4, 5, 6, 7 and 9 (Mecklenbräuker & Hlawatsch, 1997). Generally, a judgement has to be made about which properties are most important, and to what extent they can be relaxed. We discuss this in further detail for each TFR.

11. *Well defined stochastic distribution*: The probability distribution of $T_x(t, f)$ for a stochastic signal $x(t)$ must be well defined and computable, to enable significance testing of results. While this property

is seldom included in the standard texts on TFRs, for our data it is essential; we must be able to assess the significance of any detections.

12. *Good time and frequency resolution:*

By time resolution, we mean the spread in the time direction of $T_x(t, f)$ for a time-dirac signal $x(t) = \delta(t - t_0)$, while frequency resolution will refer to the spread in the frequency direction of $T_x(t, f)$ for a frequency-dirac signal $x(t) = e^{i\omega t}$. Ideally we would like $T_x(t, f) \rightarrow \delta(t - t_0)$ for a time-dirac signal centered at t_0 , and $T_x(t, f) \rightarrow \delta(f - f_0)$ for a frequency dirac signal centered at f_0 .

The uncertainty principle expresses the fundamental relationship between the standard deviation of a function and the standard deviation of its Fourier transform: they cannot simultaneously be made arbitrarily narrow. A narrow waveform will result in a broad spectrum, and vice versa (Cohen, 1989). For TFRs, the implication of the uncertainty principle is that they cannot be considered as a 'point' energy representations. Instead, we need to view the TFR as quantized into blocks of area $\sigma_t^2 \sigma_\omega^2$, where

$$\begin{aligned} \sigma_t^2 &= \frac{4\pi}{E_x} \int_{-\infty}^{\infty} (t - \bar{t})^2 |x(t)|^2 dt && \text{time spread} \\ \sigma_\omega^2 &= \frac{1}{2\pi E_x} \int_{-\infty}^{\infty} (\omega - \bar{\omega})^2 |\hat{x}(\omega)|^2 d\omega && \text{frequency spread} \\ \bar{t} &= \frac{1}{E_x} \int_{-\infty}^{\infty} t |x(t)|^2 dt && \text{average time} \\ \bar{\omega} &= \frac{1}{E_x} \int_{-\infty}^{\infty} \omega |\hat{x}(\omega)|^2 d\omega && \text{average frequency} \end{aligned}$$

Then the Heisenberg-Gabor inequality gives a lower bound on the time-bandwidth $\sigma_t \sigma_\omega$ product:

$$\sigma_t \sigma_\omega \geq \frac{1}{2},$$




Figure 10: Four different tilings of the TF plane. From Auger et al. (1996)

which holds with equality only for Gaussian functions (Mallat, 1998). Many different quantizations (or tilings) of the time-frequency plane are possible. Figure 10 shows four different TFRs, which result in different tilings of the TF plane, and allow different analyses.

Within the limits of the uncertainty principle, however, we would like to find TFRs that have the best possible time and frequency resolution.

4.8 The Wigner-Ville Distribution

With these basic properties in mind, we introduce our first TFR: the Wigner-Ville distribution (WVD). The Wigner-Ville distribution $W_x(t, f)$ of a signal $x(t)$ is given by

$$W_x(t, f) = 2 \int_{-\infty}^{+\infty} x(t + \tau)x^*(t - \tau)e^{4\pi f\tau i} d\tau \quad (13)$$

(Ville, 1948). It satisfies properties 1-9, and has the best time-frequency resolution of all TFRs: $T_x(t, f) = \delta(t - t_0)$ for a time-dirac signal centered at t_0 , and $T_x(t, f) = \delta(f - f_0)$ for a frequency dirac signal centered at f_0 (Mallat, 1998).

The WVD of a realization of a stochastic signal is an estimator for the theoretical WVD of the generating process (Mecklenbräuker & Hlawatsch, 1997). For a weakly stationary noise process with power spectrum $S(f)$, the WVD at each time is given by the power spectrum of the process: $E[W_x(t, f)] = S(f)$ (Qiu, 1993). For a deterministic signal corrupted with independent, identically distributed (i.i.d) gaussian noise, $E[W_x(t, f)]$ is the sum of the WVD of the deterministic signal, W_{signal} , and the power spectrum of the noise, $S_{noise}(f)$:

$$E[W_x(t, f)] = W_{signal} + S_{noise}(f) \quad (14)$$

(Flandrin & Martin, 1984). This means that the WVD is a biased estimator of the signal, and the larger the power of the noise, the worse the bias. The

WVD of the noisy signal, since it is determined by the power spectrum of the noise, is also not consistent:

$$W_x(t, f) \sim S_{noise}(f) \frac{\chi_2^2}{2} \quad (15)$$

Thus the WVD suffers from the same estimation problems as the periodogram. However, by estimating the power spectrum of the noise, as discussed in section 4.5.1, confidence contours (the 2-D equivalent of confidence levels) can be used to determine regions in the time-frequency plane at which the signal has power significantly above that expected for pure noise.

4.8.1 The Pseudo Wigner-Ville Distribution

In practise signals have finite duration, which is equivalent to multiplying the infinite signal in equation 13 by a rectangular window $h(t)$ that is zero outside the time range. The WVD of the truncated signal is called the pseudo Wigner-Ville distribution (PWVD), and is given by:

$$W_x(t, f) = 2 \int_{-\infty}^{+\infty} h(\tau) x(t + \tau) x^*(t - \tau) e^{4\pi f \tau i} d\tau. \quad (16)$$

The PWVD no longer satisfies the marginal property (2), the compact frequency support property (7b), or Moyal's formula (9), and loses the ideal frequency resolution of the WVD to some extent (depending on the length of the data set).

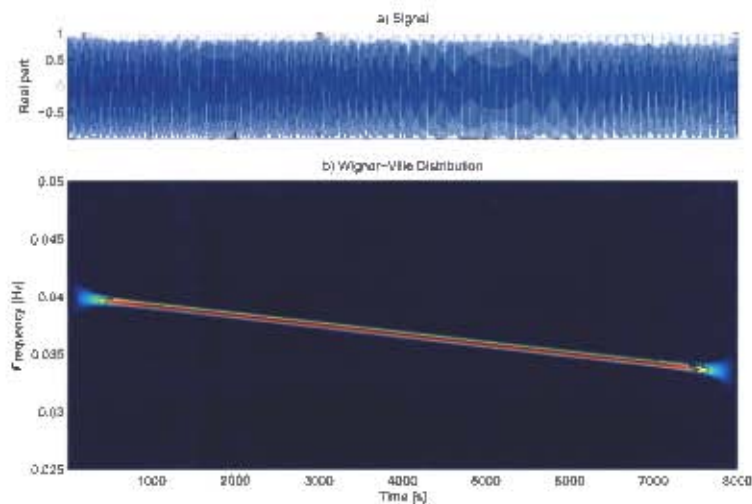


Figure 11: Top panel: A sinusoidal signal with period varying linearly from 25s to 30s. Lower panel: The Pseudo Wigner-Ville distribution of the signal shown above.

If $x(t)$ is sampled at times $t = 0, \Delta t, 2\Delta t, \dots$ then, unless the signal is over-sampled by a factor of at least 2, or the analytic signal is used, the WVD will suffer from spectral aliasing (Auger et al., 1996). From now on the signal $x(t)$ will be assumed to be its analytic version, as defined in section 4.6.

Figure 11 shows the (discretized) PWVD of an 8000 s long sinusoidal signal with period varying linearly from 25 s (0.04 Hz) to 30 s (0.033 Hz). The evolution of the frequency of the signal is clearly visible in the PWVD, although there is attenuation at the ends due to the finite length of the signal.

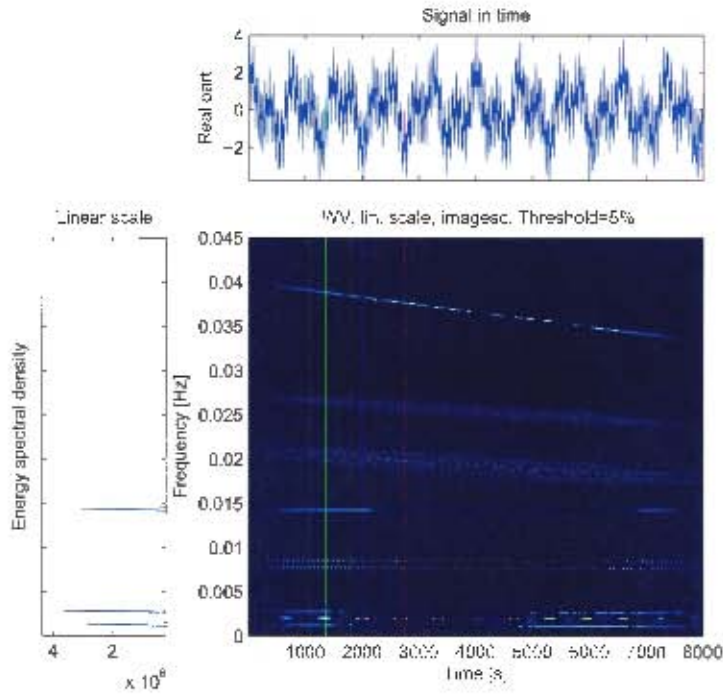


Figure 12: Signal, energy spectral density and PWVD of a synthetic DN lightcurve.

4.8.2 Interference Terms

The PWVD of our single sinusoid is an excellent TFR which allows us to see the evolution of the period of the sinusoid over time. Unfortunately, as soon as a signal has more than one frequency component, things become more complicated. The main panel of figure 12 shows the PWVD of the synthetic DN lightcurve we introduced in figure 9. While the changing DNO and QPO periodicities are clearly visible, as is the constant period lpDNO, there are also five sets of spurious signals, called interference terms.

Following the argument of Mecklenbräuker & Hlawatsch (1997), let us con-

sider a multi-component signal $x(t)$ consisting of N components x_k , each localised around frequency f_k and time t_k :

$$x(t) = \sum_{k=1}^N x_k(t). \quad (17)$$

Then the WVD of $x(t)$ can be expressed as

$$W_x(t, f) = \sum_{k=1}^N W_k(t, f) + \sum_{k=1}^N \sum_{l=1}^N W_{kl}(t, f) \quad (18)$$

where the k -th term in the first sum is the WVD of x_k , and each term in the second sum is an interference term (IT) caused by the interaction of the k -th and l -th signal components. Thus the WVD of an N -component signal consists of N signal terms and $\frac{N(N-1)}{2}$ ITs (Mecklenbräuker & Hlawatsch, 1997). In figure 12 there should be six sets of ITs, but only five are apparent: the broadest set of ITs (at around 0.02 Hz) is in fact the sum of the interference terms from the DNO with both QPOs. In figure 12 we can easily identify the sets of ITs, but this would not be possible if the true signal frequencies were unknown, or the periodogram could not be used to differentiate between signal and interference terms, as is the case with our data.

The IT corresponding to the interaction between the k -th and l -th signal components occurs at point $(t_{kl} = \frac{t_k+t_l}{2}, f_{kl} = \frac{f_k+f_l}{2})$ in the time-frequency plane (i.e midway between the signal terms), and oscillates in the time and frequency directions with 'frequencies' $\nu_{kl} = f_k - f_l$ and $\tau_{kl} = t_1 - t_2$ respectively (Hlawatsch & Urbanke, 1994). Hlawatsch & Boudreaux-Bartels

(1992) note that this implies that the closer together two components are in frequency, the slower their ITs will oscillate in the time direction. This is clear in figure 12: the ITs due to the interaction between the two QPO terms oscillate much more slowly in the time direction than the ITs due to the other interactions, whose generating signal terms are further apart in frequency (specifically, the ITs due to the QPO interaction oscillate in the time direction with a period of 900 s). In a noisy signal, these slowly oscillating ITs could easily be mistaken for an intermittent QPO, if no additional identification strategy is possible.

The presence of ITs means that the WVD is rarely used as a TFR for signal detection. However, since the ITs are oscillatory, they can be reduced by smoothing the WVD in the time and frequency directions, which in addition reduces the violation of the positivity property which the ITs cause. This is achieved by convolving the WVD with a 2-D smoothing kernel (Hlawatsch & Boudreaux-Bartels, 1992).

4.9 The Ambiguity Function & the Correlation Domain

Since the convolution of two functions is also the product of their Fourier transforms, and products are somewhat easier to visualize and work with than convolutions, we will analyse smoothing kernels in the the Fourier transform dual of the time-frequency domain: the correlation domain.

Just as the power spectrum and the autocorrelation function form a Fourier transform pair, the WVD and the ambiguity function (AF) form a Fourier transform pair. The AF occurs in the correlation domain, rather than the

time-frequency domain of the WVD and it can be interpreted as a joint time-frequency correlation function (Hlawatsch & Boudreaux-Bartels, 1992). The ambiguity function, $A_x(\tau, \nu)$, of a signal $x(t)$ is defined as:

$$\begin{aligned} A_x(\tau, \nu) &= 2 \int_t x(t + \tau) x^*(t - \tau) e^{-i4\pi\nu t} dt \\ &= 2 \int_f \hat{x}(f + \nu) \hat{x}^*(f - \nu) e^{i4\pi\tau f} df \end{aligned} \quad (19)$$

where τ is the time-delay, and ν is the Doppler value, which measures the difference in frequencies between points at a given time-delay (Hlawatsch et al., 1995). As shown in figure 13, for a signal localised in time over an interval T and frequency on an interval F , the AF is localised at the origin, with a spread in the time-delay of $[-\frac{T}{2}, \frac{T}{2}]$ and Doppler domain spread of $[-\frac{f}{2}, \frac{f}{2}]$ (Mecklenbräuker & Hlawatsch, 1997).

The top panel of figure 14 shows the AF of each of the separate components of our synthetic DN lightcurve. The DNO component appears as a diagonal from top left to bottom right: points which are separated in time by 2000 s, for example, have a difference in frequency of -0.002 Hz. The lpDNO component is horizontal along the τ -axis, since all points occur at the same frequency. The two QPO components have a much smaller overall change in frequency over the course of the light curve than the DNO, so they are almost parallel to the τ -axis.

The bottom panel of figure 14 shows the AF of the synthetic DN lightcurve. The interference terms appear as bands running parallel to the τ -axis. For a general multi-component signal as discussed in section 4.8.2, the ITs appear in the AF at (τ_{kl}, ν_{kl}) and $(-\tau_{kl}, -\nu_{kl})$ and oscillate in the τ direction with frequency f_{kl} and in the Doppler direction with ‘frequency’ t_{kl} .

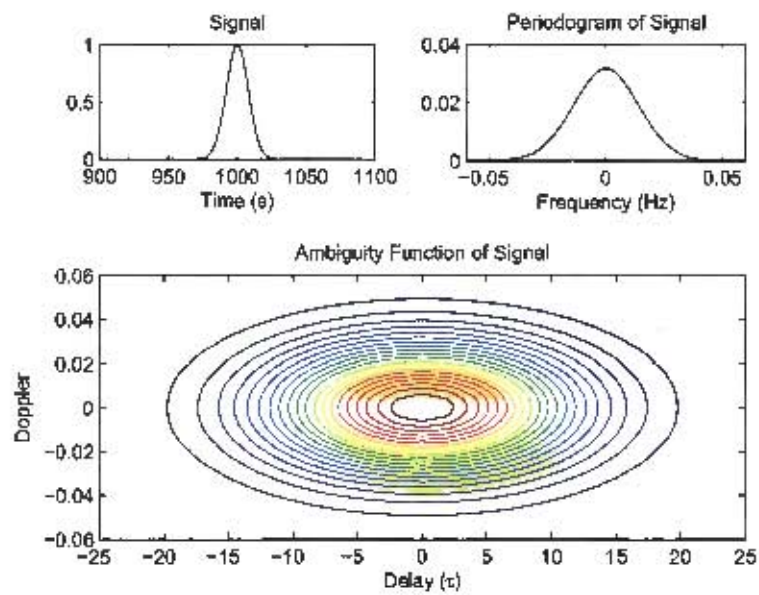


Figure 13: A Gaussian signal with time duration $T=40$ s, frequency spread $F=0.1$ Hz, its Fourier transform, and its ambiguity function.

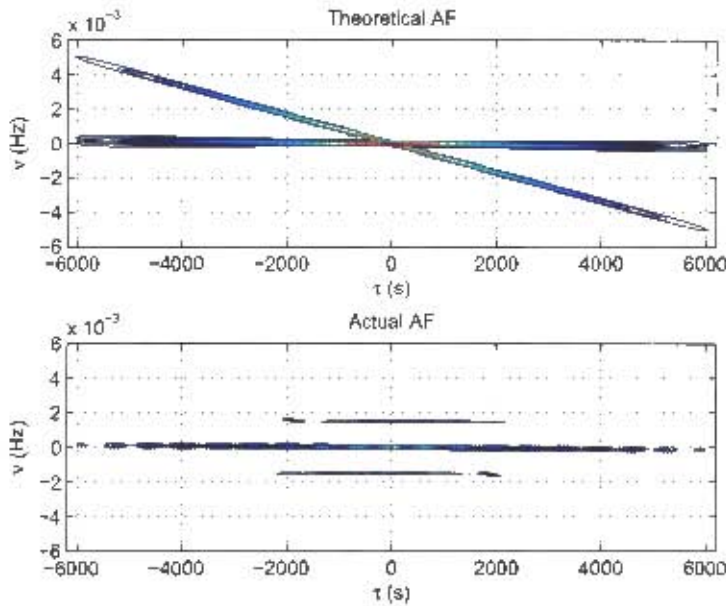


Figure 14: Top panel: Superposed AFs of the individual components of the synthetic DN lightcurve. Lower panel: AF of the synthetic DN lightcurve.

The advantage to viewing a signal in the correlative domain is that ITs are located away from the origin, while signal terms pass through the origin: if we can find a 2-D smoothing kernel that is localised at the origin in the correlation domain, multiplying the AF by it (or convolving the WVD with its FT) will produce a TFR with reduced or removed ITs, but signal terms intact.

Considerable work has gone into developing kernels which allow the resulting TFR to still fulfil at least some of the properties mentioned in section 4.7.1 while sufficiently attenuating the ITs. Smoothing comes at a cost, however: the more smoothing in time and/or frequency, the poorer the resolution in time and/or frequency (Cohen, 1989). Since, as we have seen, the geometry of the ITs depends on the components of the signal, different kernels have

been developed for different kinds of signals: there is no single kernel which is optimal for all signals. Many papers have been written comparing the effect of different kernels on various signal types (e.g. Carothers & Swenson (1992), Li, Jia & Zhao (1999), Hlawatsch et al. (1995), Buchler & Kolláth (2001), Kiss & Szatmáry (2002), Thayaparan & Kennedy (2003), Szatmary, Kiss & Bebesi (2003) or Figueiredo et al. (2004)). We follow a similar analysis to that of Hlawatsch et al. (1995) - while our synthetic signals mimic the nature of our data, we use their method of ambiguity domain analysis to establish the effect of the kernel under discussion.

We divide our discussion into two sections: kernels which allow variable degrees of IT attenuation, and kernels which, in most situations, have no ITs. Most kernels belong to the Cohen class (Cohen, 1989) or the affine class (Rioul & Flandrin, 1992) or both. Table 2 gives a summary of the kernels discussed and the properties that they satisfy. In our analysis, we will focus on the IT attenuation, noise behaviour and time-frequency resolution of the kernels, all three aspects of which are integral to the analysis of our data.

4.10 TFRs with Variable IT Attenuation

4.10.1 The Cohen Class

The Cohen class is defined as the set of all quadratic TFRs that satisfy property 4, i.e. are covariant with respect to time-shift and frequency shift operations (Cohen, 1989). The Cohen TFRs, $C_x(t, f)$, are derived from a time-frequency convolution of the WVD with a smoothing kernel $\psi(t, f)$:

Table 2: Kernels and the properties they satisfy.

Name	Kernel $\Psi(\tau, \nu)$	Properties
Wigner-Ville (WVD)	1	1, 2, 3, 4, 5, 6, 7, 8, 9.
Pseudo Wigner-Ville (PWVD)	$h(\frac{\tau}{2})h^*(-\frac{\tau}{2})$	1, 3, 4, 5, 6, 7a, 8.
Smoothed Pseudo Wigner-Ville (SPWVD)	$g(\frac{\tau}{2})g^*(-\frac{\tau}{2})\hat{h}(\nu)$	1, 3, 4, 5.
Choi-Williams (CWD)	$e^{-\frac{(2\pi\tau\nu)^2}{\sigma}}$	1, 2, 3, 4, 5, 6, 8.
Born-Jordan (BJD)	$\frac{\sin(\pi\tau\nu)}{\pi\tau\nu}$	1, 2, 3, 4, 5, 6, 8.
Cone-Kernel (CKR)	$g(\tau) \tau \frac{\sin(\pi\tau\nu)}{\pi\tau\nu}$	3, 4, 5, 7.
Affine Smoothed Pseudo Wigner-Ville (ASPWVD)	$h(\frac{\tau}{a})g(\frac{s-\tau}{a})$	1, 3.
Gabor Spectrogram (SPEC)	$W_h(-\tau, -\nu)$	1, 3, 4, 5, 10.
Scalogram (SCALO)	$W_\psi(\tau, \nu)$	1, 3, 10.

$$C_x(t, f) = \int_{t'} \int_{f'} \psi(t - t', f - f') W_x(t', f') dt' df' \quad (20)$$

or, equivalently, in the correlative domain, as

$$C_x(t, f) = \int_{\tau} \int_{\nu} \Psi(\tau, \nu) A_x(\tau, \nu) e^{i2\pi(t\nu - f\tau)} d\tau d\nu \quad (21)$$

where $\Psi(\tau, \nu)$ is the Fourier transform of $\psi(t, f)$.

Each member of Cohen's class is thus associated with a unique, signal independent kernel function $\Psi(\tau, \nu)$. Properties 1-10 can be expressed as constraints on the kernel, and hence used to construct kernels which satisfy particular sets of properties; for an excellent review of this method of kernel construction, see Loughlin, Pitton & Atlas (1993).

For any kernel satisfying Property 2 (existence of marginals), $C_x(t, f)$ of a deterministic signal corrupted with i.i.d gaussian noise is the sum of the noise-free Cohen TFR and the power spectrum of the noise, as was the case for the WVD. Thus these TFR estimators are also biased, with the amount of the bias depending on the noise level. We also have $C_x(t, f) \sim S_{noise}(f) \frac{\chi^2_2}{2}$ (Posch, 1990). In general, the variance of the TFR will depend on the amount of smoothing by the kernel: more smoothing results in a lower variance (Williams, 1998).

In the correlative domain the regions where the kernel $\Psi(\tau, \nu)$ is small or zero show where ITs will be attenuated, thus by considering the shape of the kernel in the correlation domain, it is relatively clear which kinds of signal will benefit from it.

4.10.2 Smoothed Pseudo WVD

One method of smoothing the WVD is to use the separable kernel

$$\Psi_{SPWVD}(\tau, \nu) = g\left(\frac{\tau}{2}\right)g^*\left(-\frac{\tau}{2}\right)\hat{h}(-\nu) \quad (22)$$

where $g(t)$ is a time-smoothing window and $\hat{h}(\nu)$ (the Fourier transform of a time-smoothing window $h(t)$) is a frequency-smoothing window (Auger et al., 1996). The lengths of the windows $g(t)$ and $h(t)$ control the amount of smoothing in the time and frequency domains respectively: a longer $g(t)$ gives more time smoothing (and hence worse time resolution), and a longer $h(t)$ gives less frequency smoothing (and hence better frequency resolution) (Hlawatsch et al., 1995). The SPWVD satisfies properties 1, 3, 4 and 5.

If the length of the window $g(t)$ is M , then for a deterministic signal corrupted by i.i.d gaussian noise,

$$C_{SPWVD}(t, f) \sim \frac{2}{2M-1} S_{noise}(f) \frac{\chi_2^2}{2}$$

(Flandrin & Martin, 1984). Thus the variance of the SPWVD decreases as the length of the time-smoothing window increases.

Figures 15 and 16 show the SPWVD of the synthetic DN lightcurve, with $L_g = 50$ (200 s) in both figures, but $L_h = 50$ in the first, and $L_h = 300$ in the second. The effect of lengthening $h(t)$ is marked: in the first figure,

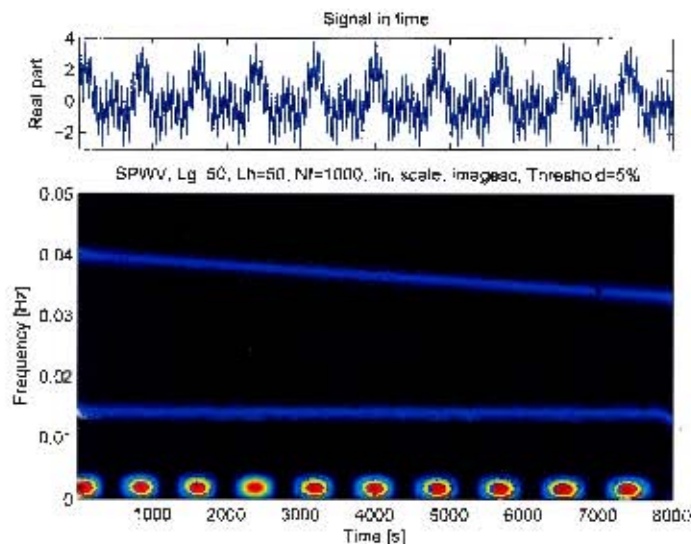


Figure 15: Smoothed Pseudo Wigner-Ville distribution of synthetic DN lightcurve, with $L_g = 50$ (200s), $L_h = 50$.

we cannot differentiate the two long period QPO components, but in the second they are clearly visible. Figures 17 and 18 show the ambiguity function of the synthetic signal and the ambiguity domain representation of the PSWVD kernels corresponding to those used figures 15 and 16. Since, in the ambiguity domain, the SPWVD is simply the product of the AF of the signal and the ambiguity domain version of the kernel, signal or interference terms which are present where the kernel is zero are completely removed. The ambiguity domain shows clearly the effect of lengthening $h(t)$: the SPWVD kernel becomes wider in the τ -dimension, resulting in less smoothing in the τ -dimension, which corresponds to less smoothing in the frequency dimension in the TFR.

Interference terms between the two long-period components in figure 16 are prominent. By lengthening the time-smoothing window $g(t)$ to 250 s, as in

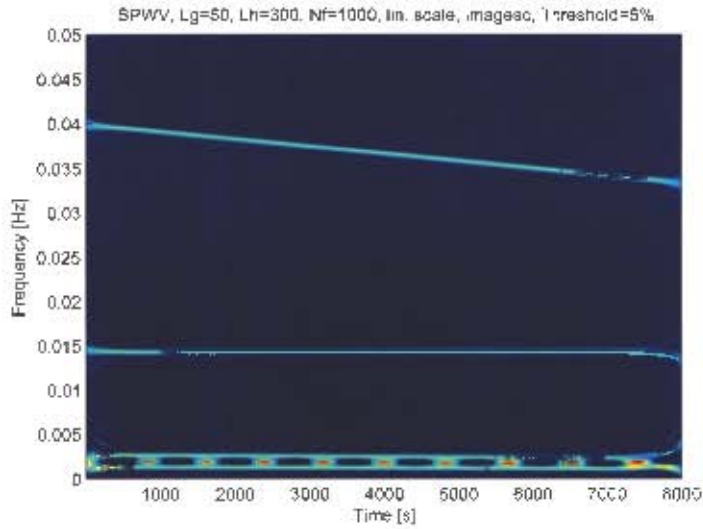


Figure 16: Smoothed Pseudo Wigner-Ville distribution of synthetic DN lightcurve, with $L_g = 50$ (200s), $L_h = 300$.




Figure 17: Ambiguity function of synthetic DN lightcurve, with SPWVD kernel superimposed, with $L_g = 50$ (200s), $L_h = 50$.




Figure 18: Ambiguity function of synthetic DN lightcurve, with SPWVD kernel superimposed, with $L_g = 50$ (200s), $L_h = 300$.

figure 19, the interference terms are completely attenuated. In the ambiguity domain (figure 20), we see that increasing the length of $g(t)$ corresponds to decreasing the height of the kernel in the ν -dimension. ITs which were previously not attenuated by the kernel now occur in a region where the kernel is zero, and are hence completely attenuated. Smoothing in the ν dimension in the ambiguity domain corresponds to smoothing in the time dimension in the TFR.

The time-smoothing introduced by lengthening $g(t)$ comes at a price: the ability of the SPWVD to clearly resolve temporal discontinuities in the signal. Figures 21 and 22 show the SPWVD and ambiguity domain analysis respectively of an intermittent sinusoid of period 100 s. The smoothing in the time domain results in a SPWVD that does not clearly show the beginning or end of each sinusoidal burst. Indeed, it gives the impression that

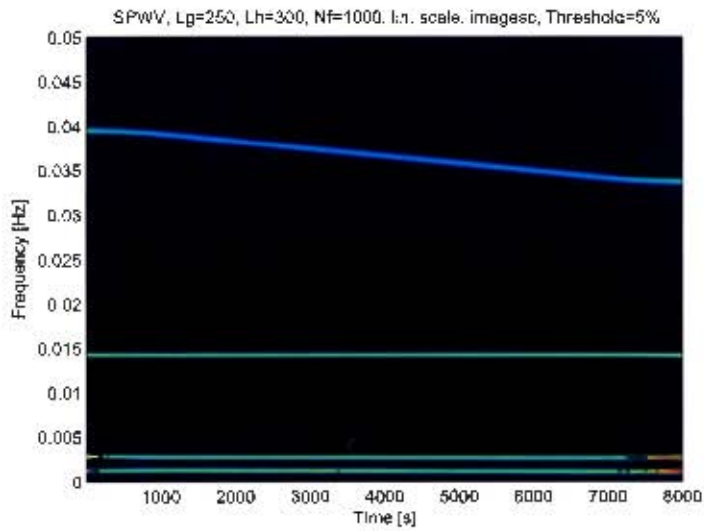


Figure 19: Smoothed Pseudo Wigner-Ville distribution of synthetic DN lightcurve, with $L_g = 250$ (1000s), $L_h = 300$.




Figure 20: Ambiguity function of synthetic DN lightcurve, with SPWVD kernel superimposed, with L_g 250 (100s), L_h 300.

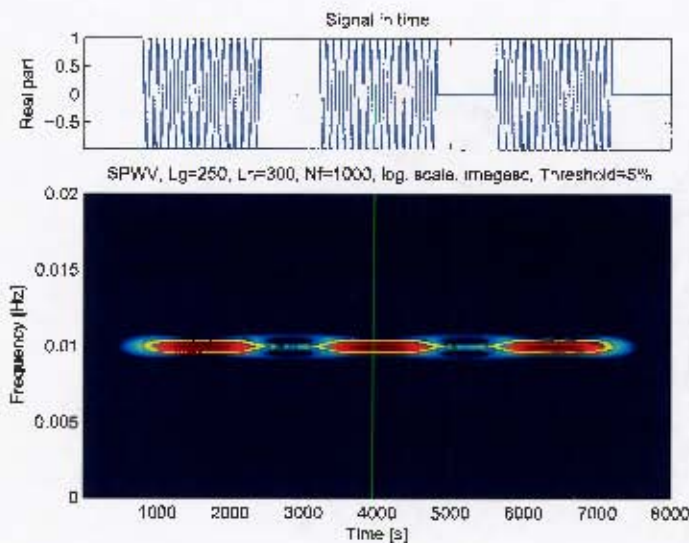


Figure 21: Smoothed Pseudo Wigner-Ville distribution of intermittent sinusoid, with L_g 250 (1000s), L_r 300.

the sinusoid is perhaps present for the duration of the signal, at varying amplitude. There are two reasons for this: signal terms, which should not be smoothed, are smeared in the time dimension by the narrow kernel, and ITs which oscillate in the frequency direction between the signal terms are not removed because we have set the length of the frequency smoothing window to ensure that we have optimal frequency resolution (i.e. minimal smoothing).

We are thus presented with the central issue that will dog many of the kernels we investigate: in order to ensure that ITs are attenuated, which is crucial for our analysis, as we cannot afford to confuse ITs and QPOs, we need to smooth in time (and minimally in frequency, if we are to have sufficient frequency resolution), but this results in poor time resolution, particularly at high frequencies.

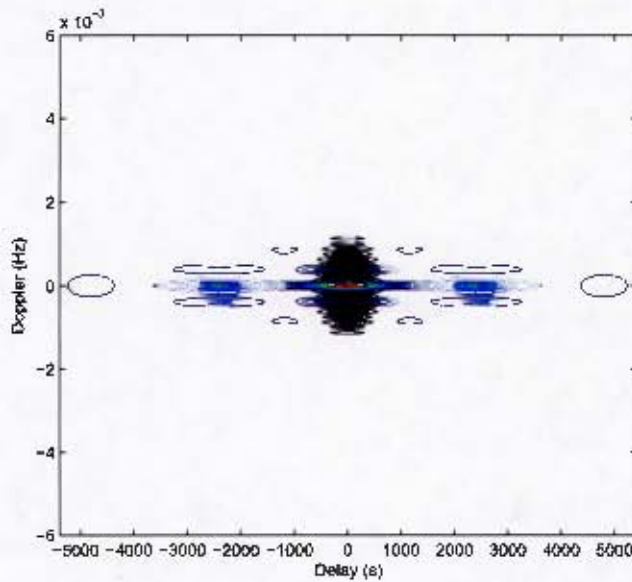


Figure 22: Ambiguity function of intermittent sinusoid, with SPWVD kernel superimposed, with L_g 250 (100s), L_h 300.

4.10.3 Choi-Williams Distribution

Instead of using a separable smoothing kernel, several kernels have been developed that depend on the product of τ and ν . These have the general form

$$\Psi(\tau, \nu) = f(\tau\nu). \quad (23)$$

This parametrization ensures that the marginal properties are respected, and, since the ITs are attenuated, the resulting TFRs are sometimes called Reduced Interference Distributions (Auger et al., 1996). One of the earliest examples these joint-smoothings of the WVD is the Choi-Williams distribu-

tion (CWD) (Choi & Williams, 1989), which is characterized by the kernel

$$\Psi_{CWD}(\tau, \nu) = \exp\left[-\frac{(2\pi\tau\nu)^2}{\sigma}\right], \quad (24)$$

where σ is a positive parameter that jointly controls the smoothing in time and frequency (Hlawatsch et al., 1995). A larger value for σ gives less smoothing, as it results in a broader $\Psi_{CWD}(\tau, \nu)$. In fact, as $\sigma \rightarrow +\infty$, the CWD tends to the WVD.

The CWD satisfies properties 1, 2, 3, 4, 5, 6 and 8. In particular, the time and frequency support properties are not satisfied. Pure sinusoids have the ideal time-frequency resolution of the WVD, but all other signals are broadened when compared with the WVD (Hlawatsch & Urbanke, 1994). Since the CWD satisfies property 2, the CWD of a noisy signal is the sum of the noise-free CWD and the power spectrum of the noise, as discussed in section 4.10.1.

The top panel of figure 24 shows the synthetic signal that we have analysed with the CWD; it consists of a two sinusoids, one with period 70s and one with period 100s, that overlap in the middle portion of the data set for 4000 s. The lower panel shows the CWD of the synthetic signal, with $\sigma = 100$. Reducing σ to 0.1, as shown in figure 24, which increases the amount of smoothing, does not fully eliminate the ITs; further reductions in σ have no noticeable affect.

If we view the signal's AF, with the kernel superimposed as in figure 25, then we can understand the reason for the lack of complete attenuation of




Figure 23: CWD of two sinusoids (periods 70 s and 100 s), overlapping in time from 2000 s to 6000 s, $\sigma = 100$.

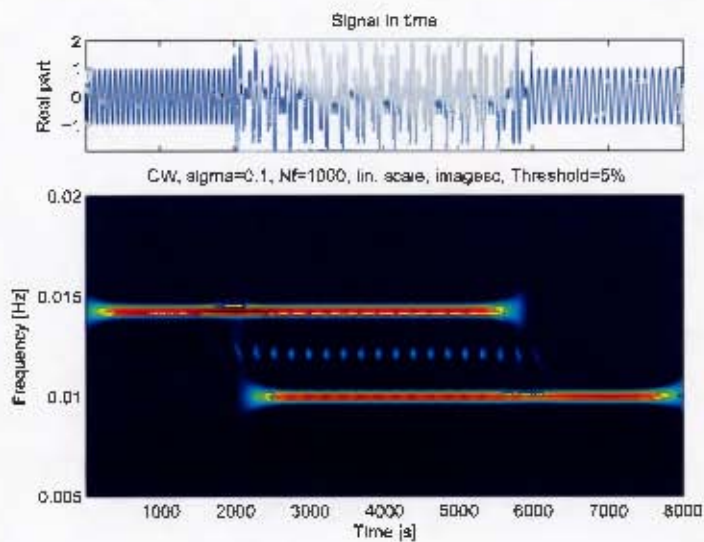


Figure 24: CWD of two sinusoids (periods 70 s and 100 s), overlapping in time from 2000 s to 6000 s, $\sigma = 0.1$.

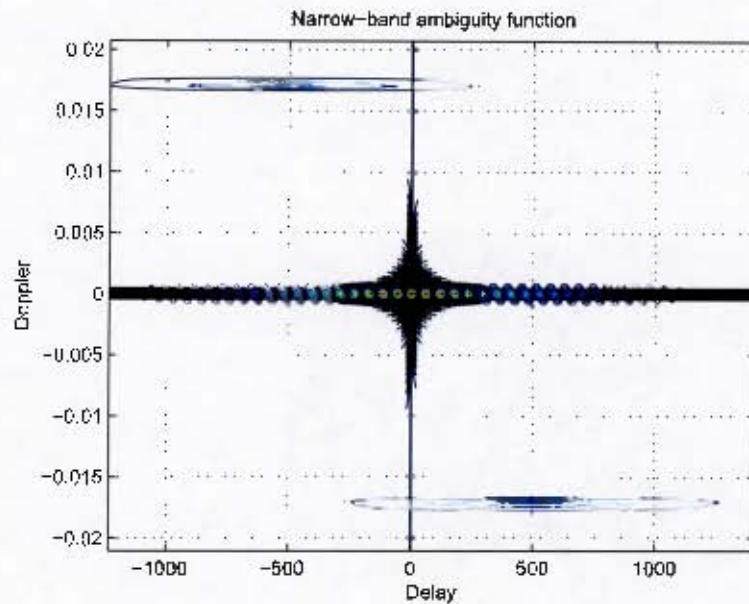


Figure 25: Ambiguity function of CWD kernel with $\sigma = 0.1$, superimposed on ambiguity function of overlapping sinusoids signal.

ITs: ITs occurring on the ν - or τ -axis cannot be completely attenuated. These ITs are due to either different frequency components occurring at the same time (i.e. for $t = 2000 - 6000$ s in our signal) or components at the same frequency occurring at different times; the former result in ITs on the ν -axis, the latter ITs on the τ -axis. Thus for signals such as ours, which frequently have different frequency components occurring at the same time, we will not be able to attenuate the ITs completely if we use the CWD.

4.10.4 Born-Jordan Distribution and Cone-shaped Kernel Representation

If we require that the TFRs resulting from the kernels described by equation 23 preserve finite time- and frequency-support (property 7), then the simplest choice for $\Psi(\tau, \nu)$ is

$$\Psi_{BJ}(\tau, \nu) = \frac{\sin(\pi\tau\nu)}{\pi\tau\nu} \quad (25)$$

which defines the Born-Jordan distribution (BJD) that satisfies properties 1 to 8. The BJD is the minimum variance distribution for white noise (Hearon & Amin, 1995), but does not allow smoothing of ITs in the frequency direction. Allowing smoothing of the BJD in frequency gives the Zhao-Atlas-Marks representation (Zhao et al., 1990), also known as the cone-shaped kernel representation (CKR), with the kernel

$$\Psi_{CKR}(\tau, \nu) = g(\tau)|\tau| \frac{\sin(\pi\tau\nu)}{\pi\tau\nu} \quad (26)$$

where $g(\tau)$ is a window function (Hlawatsch et al., 1995). The CKR satisfies properties 3, 4, 5 and 7. Most importantly, it does not satisfy property 1, since $\Psi_{CKR}(\tau, \nu)$ is zero on the entire ν -axis, and hence the CKR's integrals with respect to frequency, and over the entire time-frequency plane, are zero. It is thus not an energy distribution (in the statistical sense), which is why we refer to it as a representation, rather than a distribution (Hlawatsch et al., 1995).

The CKR of a deterministic signal corrupted by i.i.d gaussian noise is the sum of the CKR of the noise-free signal and $D(t, f; g)$, where

$$D(t, f; g) = \int_{-\infty}^{\infty} g(\tau) |r(\tau)| r(\tau) e^{-i2\pi f\tau} d\tau \quad (27)$$

and $r(\tau)$ is the auto-covariance of the noise (defined in section 4.1). This means that for white noise, $D(t, f; g) = 0$, and the CKR is an unbiased estimator (Oh & Marks, 1992). Hearon & Amin (1995) show that of all the Cohen TFRs the CKR has the highest variance, and hence the least desirable performance in the presence of noise. In addition, confidence contours cannot easily be constructed for the CKR since the stochastic distribution of the CKR is not calculable from the power spectrum as is the case for the other TFRs that obey the energy conservation property. While these issues mean that the CKR is not ideal for our data, we nonetheless include an analysis of the behaviour of the CKR, for completeness.

IT attenuation in the CKR is more complicated than in the kernels discussed so far because the kernel does not reach its maximum at the origin, but rather at $(\pm\tau_0, 0)$, where τ_0 is determined by $g(\tau)$. Figure 26 shows the kernel of the CKR, for window lengths 25, 75, 100 and 200, where we have used a Gaussian window. The asymmetry evident in figures 26b) and c) is due to asymmetric sampling in the τ domain resulting in asymmetric contours, rather than any inherent asymmetry in the kernel. The position of the maximum shifts out along the τ -axis as the window length increases.

Since $\Psi_{ZAM}(\tau, \nu)$ is zero on the entire ν -axis, any signal terms occurring on the ν -axis will be attenuated, resulting in a (misleading) narrowing of fre-

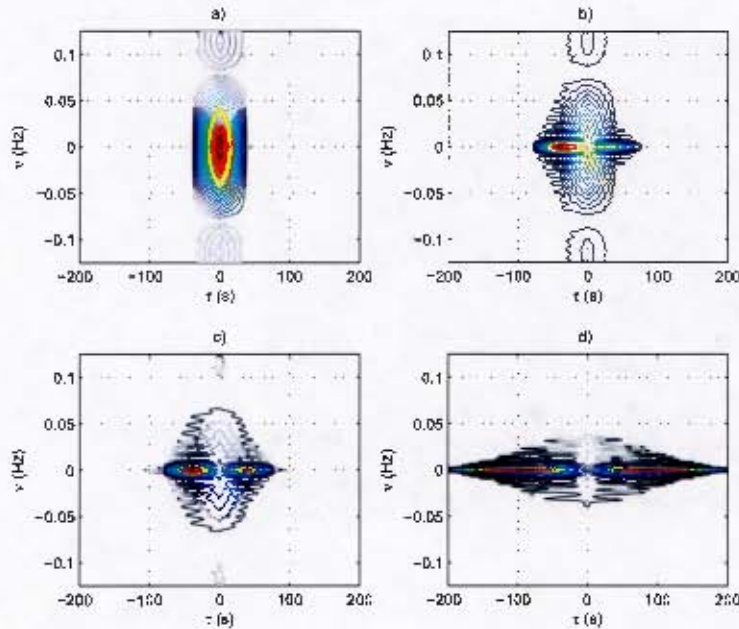


Figure 26: Ambiguity function of the CKR kernel, for different window lengths.

quency components in the TFR. For components which are close in frequency (i.e. $\nu \approx 0$), the attenuation will depend on the time distance between the components. If the time distance between components is $t = \pm\tau_0$, then the resulting IT will be amplified, rather than attenuated. However, if T_g is the length of the window, then all terms (signal and IT) with $|\tau| > \frac{1}{2}T_g$ will be completely suppressed. We are thus faced with a dilemma: increasing the length of the window improves frequency resolution, but also makes it more likely that ITs will be close to $(\pm\tau_0, 0)$, and hence enhanced (Hlawatsch et al., 1995).

Figure 27 shows the CKR of our synthetic DN lightcurve, where the window length, 250, has been chosen to be the minimum that attenuates all ITs.

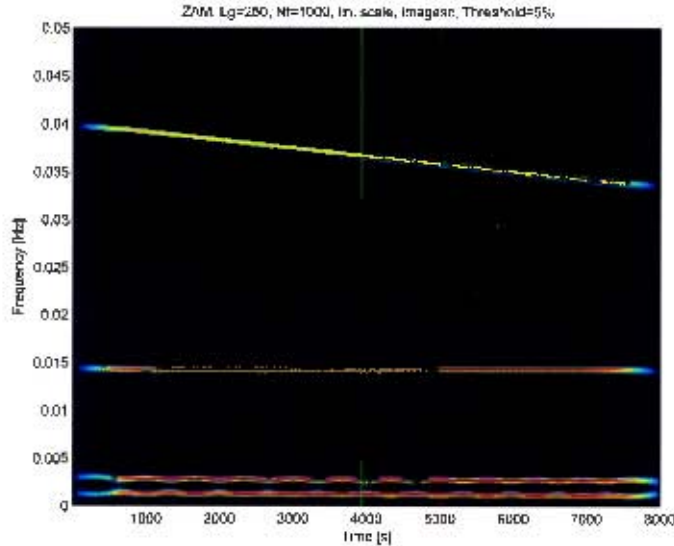


Figure 27: CKR of synthetic DN lightcurve, $L_g \approx 250$ (1000 s).

Using this window length for our intermittent sinusoid (Figure 28) gives a better result than the SPWVD; the signal terms are well-defined, and the ITs are more attenuated.

Figure 29 shows the effect of this kernel on a short-period intermittent, amplitude modulated signal, mimicking, for example, a DNO of changing amplitude. The amount of smoothing in time required to remove the low-frequency ITs in the synthetic DN lightcurve result in the signal terms of the intermittent DNO being completely smoothed. In addition, the ITs are amplified and smoothed (because they appear near where the kernel is maximum in the correlation domain), resulting in the appearance of five frequency bands. While this is a somewhat extreme example, it does show that the effects of the CKR must be carefully monitored.

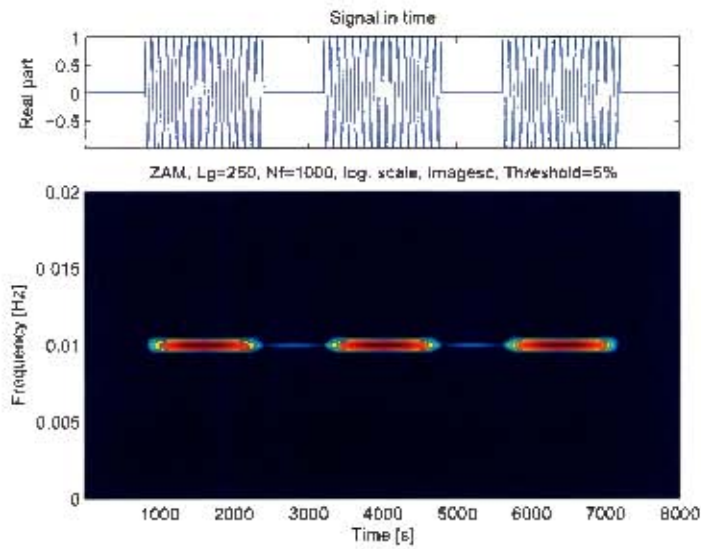


Figure 28: CKR of intermittent sinusoid, $L_g=250$ (1000 s).




Figure 29: CKR of a intermittent, amplitude modulated sinusoid of period 20 s, $L_g=250$ (1000 s).

4.10.5 The Affine Class

Affine TFRs are covariant in both translation and dilation, rather than translation and frequency shift as is the case for the Cohen class. A general affine TFR ($A_x(t, a; \Pi)$) of an analytic signal $x(t)$ is

$$A_x(t, a; \Pi) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Pi\left(\frac{s-t}{a}, a\xi\right) W_x(s, \xi) ds d\xi \quad (28)$$

where $\Pi(t, \nu)$ is a smoothing kernel. Equivalently, using the AF instead of the WVD, we have

$$A_x(t, a; \Psi) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Psi(a\xi, \tau/a) A_x(\xi, \tau) e^{-i2\pi\xi t} d\xi d\tau \quad (29)$$

where $\Psi(\xi, \tau)$ is the ambiguity domain version of $\Pi(t, \nu)$. Strictly speaking, affine TFRs are *time-scale* representations; a is called the dilation or scale parameter, and acts as a measure of the frequency since by reducing s the support of $\Pi_{t,s}$ is reduced in time and hence covers a larger frequency range (Goswami & Chan, 1999). If we let $a = \nu_0/\nu$, where ν_0 is the central frequency of the analyzing kernel, then an affine TFR can be displayed as a time-frequency representation (Rioul & Flandrin, 1992). Since the WVD satisfies the scale-invariant property, it is a member of the affine class. The Born-Jordan and Choi-Williams kernels can also be used for affine TFRs (Hlawatsch & Boudreaux-Bartels, 1992).

The important difference between affine and Cohen TFRs is that the affine kernel is scaled depending on the analysing frequency, while Cohen kernels

are simply shifted or translated, but retain the same 'size' at all frequencies. This means that for affine TFRs the time and frequency resolutions depend on the scale (or frequency) at which the signal is being analysed, while for Cohen TFRs the time and frequency resolutions are the same at all frequencies.

Figures 30 and 31 show the time-frequency plane tiling for a Cohen TFR (the Gabor spectrogram, which will be discussed in section 4.11.1) and an affine TFR (the wavelet scalogram, which will be discussed in section 4.11.2). For affine TFRs, the frequency resolution will typically be better at low frequencies (since the width of the time-frequency box in the frequency direction is longer at lower frequencies), while the time resolution will be better at higher frequencies (the height of the time-frequency box in the time direction decreases as the scale or frequency decreases).

The majority of the affine TFRs, such as the Bertrand (Bertrand & Bertrand, 1992), D-Flandrin (Flandrin, 1999) and Unterberger (Unterberger, 1983) distributions, were developed to analyse wideband monocomponent signals (Gonçalves & Baraniuk, 1998), and as such are not appropriate for our data, which are essentially composed of several components each of which is relatively narrowband. There are however two affine TFRs that are potentially useful: the affine smoothed pseudo Wigner-Ville distribution, which allows variable IT attenuation, and the wavelet scalogram, which does not have ITs. Since the kernel depends on the frequency, a correlation domain analysis is less helpful for affine TFRs.




Figure 30: Time-frequency plane tiling for the Cohen class Gabor spectrogram, from Mallat (1998).




Figure 31: Time-frequency plane tiling for the affine class wavelet scalogram, from Mallat (1998).

4.10.6 Affine Smoothed Pseudo WVD

The affine smoothed pseudo WVD (ASPWVD) is the affine counterpart of the SPWVD, and is defined as

$$A_{ASPWVD} = \frac{1}{a} \int \int_{-\infty}^{\infty} h\left(\frac{\tau}{a}\right) g\left(\frac{s-t}{a}\right) x\left(s + \frac{\tau}{2}\right) x^*\left(s - \frac{\tau}{2}\right) ds d\tau. \quad (30)$$

where $g(t)$ and $h(t)$ are smoothing windows as in the Cohen case. The ASPWVD only satisfies properties 1 and 3. $g(t)$ and $h(t)$ give a flexible choice of time and scale (or frequency) resolution.

If the length of the window $g(t)$ at scale s is M_s , then for a deterministic signal corrupted by i.i.d gaussian noise,

$$A_{ASPWVD}(t, s) \sim \frac{2}{2M_s - 1} S_{noise}(s) \frac{\chi_2^2}{2}$$

(Flandrin & Martin, 1984). Thus the variance of the ASPWVD decreases as the window length increases, as in the case of the SPWVD, but also depends on the analysing scale s .

Figures 32 and 33 show the ASPWVD of the synthetic DN lightcurve, with the length of $g(t)$ set to 0 in the first case (no time smoothing) and 80 (320 s) in the second (these are the lengths of the window at the largest scale). The length of $h(t)$ is 80 in both cases, which was the minimum length needed to resolve both QPOs. What is apparent in figure 33, when compared with the SPWVD shown in figure 19 where the length of $g(t)$ is 250, is that less




Figure 32: ASPWVD of synthetic DN lightcurve, with $L_g = 0$ and $L_h = 80$.

time smoothing is needed to remove the IT in the ASPWVD. However, the frequency resolution decreases in the APSWVD as the frequency increases - the DNO in the APSWVD is broader than the DNO in the PSWVD.

If we now compare the effect of the PSWVD and APSWVD (with kernels chosen to give good time-frequency resolution for our synthetic DN signal) on an intermittent sinusoid of period 100 s (figures 21 and 34), we see that in the frequency band of interest, the PSWVD maintains good frequency resolution, but has poor time resolution, while the APSWVD has good time resolution (since for affine distributions the time resolution improves as the frequency increases), but has poorer frequency resolution than the PSWVD.

Finally, we compare the effect of the CKR and APSWVD on an intermittent, amplitude modulated sinusoidal signal of period 20 s (figures 29 and 35). The time resolution of the APSWVD is somewhat better than the CKR:




Figure 33: ASPWVD of synthetic DN lightcurve, with $T_g = 80$ and $T_h = 80$.

individual groups of oscillations are visible, but the signal is still distorted by ITs. The frequency resolution is about the same for both kernels at this frequency.

4.11 TFRs with No ITs

There are two TFRs which do not produce ITs if their signal components do not overlap in time or frequency (Hlawatsch, 1991): the Gabor spectrogram, which is a member of the Cohen class, and its affine counterpart the wavelet scalogram. The price paid for the lack of ITs is resolution: both these TFRs have the poorest time-frequency resolution of all the TFRs discussed.




Figure 34: ASPWVD of an intermittent sinusoid, with $L_g = 80$ and $L_h = 80$.




Figure 35: ASPWVD of an intermittent, amplitude modulated sinusoid of period 20 s, with $L_g = 80$ and $L_h = 80$.

4.11.1 Gabor Spectrogram

The spectrogram is traditionally defined as

$$C_{SPEC}(t, f) = \left| \int_{t'} x(t') h^*(t' - t) e^{-i2\pi f t'} dt' \right|^2 \quad (31)$$

where $h(t)$ is a short time analysis window, localised around $t = 0$ and $f = 0$ (Hlawatsch et al., 1995). The integral inside the absolute value is called the Short-time Fourier transform (STFT). Auger et al. (1996) note that, using Moyal's theorem, the spectrogram can be expressed as a smoothing of the WVD:

$$C_{SPEC}(t, f) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} W_h(s - t, \xi - \nu) W_x(s, \xi) ds d\xi. \quad (32)$$

so the Cohen kernel for the spectrogram is $\Psi_{SPEC} = W_h(-\tau, -\nu)$, which is the WVD of the spectrogram's analysis window. This means that the time-frequency resolution of the spectrogram is dictated by the resolution of the analysing window, rather than the analysed signal. Smoothing in the spectrogram is thus extensive, resulting in virtually complete interference attenuation, but also poor time-frequency resolution (Hlawatsch & Boudreaux-Bartels, 1992). The choice of $h(t)$ determines the trade-off between the time spread and the frequency spread of the smoothing function: time resolution is proportional to the time duration of $h(t)$, while frequency resolution is proportional to the bandwidth of $h(t)$. The Gabor spectrogram satisfies properties 1, 3, 4 and 5, and, for a noisy signal, is the sum of the noise free Gabor spectrogram and the power spectrum of the noise, as for the WVD.

Williams (1998) shows that the spectrogram, when a rectangular window is used, gives the best performance of all Cohen TFRs in the presence of non-white noise, as it has the smallest variance.

Figure 36 shows the Gabor spectrogram of our synthetic DN signal, using a Gaussian window of length 500, which allows resolution of the QPO components. Each component is clearly visible, and there are no interference terms, although some attenuation and smoothing is apparent near the ends.

Figure 37 shows the spectrogram of an intermittent sinusoid, using the same window as in the previous example. We see that requiring good frequency resolution has again resulted in poor time resolution. Indeed, as Cohen (1989) and Zhao et al. (1990) point out, if one wishes to have accurate time and frequency measurements using the spectrogram, two separate analyses must be undertaken, using windows of different length.

4.11.2 The Wavelet Scalogram

We define the (continuous) wavelet transform of a signal $x(t)$ as

$$T_x(t, a; \Psi) = \int_{-\infty}^{\infty} x(s) \Psi_{t,a}^*(s) ds \quad (33)$$

where $\psi_{t,a}(s) = \frac{1}{\sqrt{a}} \psi_0\left(\frac{s-t}{a}\right)$ and ψ_0 is a fixed function, called the “mother wavelet”, which has zero mean and is well localized in time and frequency (Cohen, 1995; Mallat, 1998). There are clearly an infinite variety of functions fulfilling these criteria; one commonly used wavelet is the Morlet wavelet, given by




Figure 36: Gabor spectrogram of the synthetic DN lightcurve, with window length 500.




Figure 37: Gabor spectrogram of an intermittent sinusoid, with window length 500.

$$\psi(\eta) = \pi^{\frac{1}{4}} e^{-\frac{1}{2}\eta^2} e^{i6\eta}; \quad (34)$$

which is the wavelet we will use in this section as it has excellent frequency resolution. Chapter 5 will give a more detailed introduction to the theory of wavelets, and introduce further mother wavelets.

The wavelet scalogram, A_{SCALO} , is defined as

$$A_{SCALO}(t, a) = |T_x(t, a; \Psi)|^2 \quad (35)$$

which can be written as

$$A_{SCALO}(t, a) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} W_x(s, \xi) W_\psi\left(\frac{s-t}{a}, a\xi\right) ds d\xi \quad (36)$$

(Auger et al., 1996), showing that the kernel for the scalogram is $\Psi_{SCALO}(\tau, \nu) = W_\psi(\tau, \nu)$; the wavelet scalogram is the affine counterpart of the spectrogram. The scalogram validates properties 1 and 3, and is always positive. Percival (1995) shows that the time integral of the wavelet scalogram can be used as an unbiased, consistent estimator for the power spectrum, thus for practical purposes the wavelet scalogram satisfies the marginal properties. It is also the minimum affine variance estimator, and we have $A_{SCALO} \sim \frac{1}{2} S_{noise} \chi_2^2$.

By varying the width of the mother wavelet, the time-frequency resolution at the largest scale can be varied, which in turn determines how the resolution




Figure 38: Wavelet scalogram of synthetic DN lightcurve, with wavelet length 80

changes as the frequency increases at shorter scales. Figure 38 shows the wavelet scalogram of the synthetic DN lightcurve, when the mother wavelet is a Morlet wavelet of length 80 (320 s). All components are clearly visible. With this mother wavelet, the intermittent sinusoid with period 100s is resolved well in frequency but not time, as figure 39 shows. If we reduce the width of the mother wavelet to 35 (140 s), then the intermittent sinusoid with period 100 is well resolved in time (figure 40), but the frequency resolution is poor. The intermittent, amplitude-modulated sinusoid with period 20 s is well resolved in time - much better than with the CKR or ASPWVD, but again the frequency resolution is poor. With the mother wavelet of width 35 (140 s), the QPOs in the synthetic DN lightcurve are still well resolved in frequency (figure 42), but the lpDNO and especially the DNO are not.

The wavelet scalogram is thus able to give us the simultaneous high fre-




Figure 39: Wavelet scalogram of intermittent sinusoid, with wavelet length 80.




Figure 40: Wavelet scalogram of intermittent sinusoid, with wavelet length 35.




Figure 41: Wavelet scalogram of high-frequency intermittent sinusoid, with wavelet length 35.




Figure 42: Wavelet scalogram of synthetic DN lightcurve, with wavelet length 35.

quency resolution we require at low frequencies to detect QPOs, and the high time resolution at high frequencies needed to study DNOs, which none of the other TFRs have been able to do, but we lose the good frequency resolution at high frequencies, which is not ideal. However, Delprat et al. (1992) show that the wavelet scalogram reaches a maximum at the instantaneous frequency $\phi'(t)$ defined in section 4.6. Thus by finding all local maxima in the frequency direction (which we will call maximal ridges), we can detect the instantaneous frequency, even at high frequencies where the frequency resolution of the wavelet scalogram is poor. Figure 43 shows the synthetic DN lightcurve, analysed with a Morlet mother wavelet of length 35 as in figure 42, but with the maximal ridges of the shown. Thus the wavelet scalogram, with the addition of maximal ridges, provides a good TFR for data such as ours, which consist of components covering a wide frequency range, and requiring high time resolution at high frequencies.

4.12 Discussion

We have investigated a number of TFRs, with the aim of finding a TFR with good time and frequency resolution which enables us to detect and analyse DNOs and QPOs. There are, of course, many quadratic TFRs that we have not investigated. We have not discussed the Page distribution (Page, 1952), since the amplitude of frequency components in this distribution depends on the duration of the oscillation, which makes interpretation difficult, and significance testing virtually impossible. ITs in the Margenau-Hill distribution are shifted to appear at all signal frequencies for all times where the signal is non-zero. Since these ITs are unattenuated, they can appear as strong as signal terms, and, since they are superimposed on the signal terms, can give




Figure 43: Wavelet scalogram of synthetic DN lightcurve, with wavelet length 35, and maximum ridges.

a misleading appearance to the signal terms.

TFRs other than the quadratic forms mentioned are also available, such as the local ambiguity function of O'Neill & Flandrin (2000), which combines elements of the WVD and the AF. However, the ITs are even more complicated than those of the quadratic TFRs, and we have not pursued these methods. Reassignment methods, which involve post-processing of the spectrogram to improve resolution, are explored in, for example, Kodera, Villedary & Gendrin (1976). However, these representations do not allow an easy statistical interpretation of the distribution of the representation and so were not investigated further; for a good introduction, including examples, see Thayaparan & Kennedy (2003). Oung & Forsberg (1998) introduced adaptive affine distributions, but since these rely on a *priori* knowledge

of the instantaneous frequency and bandwidth of the signal to smooth the signal, and are best suited to monocomponent signals, they were also not pursued.

While we have focused our analysis of kernels on ambiguity domain analysis to understand IT attenuation and time-frequency resolution, alternative analyses are possible, often relying on the development of appropriate measures for time-frequency resolution to choose the optimal kernel analytically. We briefly discussed one such measure (the time-bandwidth product); for alternative measures, see Stanković (2001) and references therein.

With the theory of TFRs behind us, it is insightful to re-examine in greater detail the use of different TFRs introduced in chapter 3.

Boyd et al. (1995) use the Gabor transform to investigate HD 60435 (see fig. 7 of Boyd et al. (1995)). Because they are analysing a specific oscillation (11.64 min) they are able to choose an analysis window width which gives the required time-frequency resolution in this narrow frequency band. Belloni, Parolin & Casella (2004) also consider a relatively narrow frequency band, and are hence able to achieve an appropriate time-frequency resolution balance. Dolan et al. (1998) gives one of the the few examples of the analysis of a transient signal with the Gabor spectrogram, but since the frequency range is relatively narrow, they are able to choose a window length which gives the optimal time-frequency resolution for analyzing the signal. Figures 1.10 and 1.11 of Buchler & Kolláth (2001) compare the CWD and Gabor spectrogram of a synthetic R Sct light-curve; the higher time and frequency resolution of the CWD enable more detailed analysis of the light curve than the Gabor; the ITs do not interfere with interpretation since they

can readily be identified as such with the help of the periodogram, in which the true periodicities are clear.

Kolláth & Buchler (1997) compare the (Morlet) wavelet scalogram, Gabor spectrogram and CWD for long-term light curves of T UMi and R Sct (Figures 1 and 2 of Kolláth & Buchler (1997)). It can be seen that at higher frequencies, the wavelet has better time resolution than the Gabor spectrogram or CWD, while overall the CWD has better frequency resolution than either of the other two, but is somewhat distorted by ITs. Buchler & Kolláth (2001) compare the CWD and Gabor spectrogram of a synthetic light curve of the long-period variable R Sct, giving a good example of how the Gabor spectrogram, though lacking ITs, has poorer time-frequency resolution than the CWD.

Figure 5 of Szatmáry, Kiss & Bebesi (2003) shows the wavelet, CWD and CKR of T UMi. The CWD and CKR give far better frequency resolution than the wavelet scalogram, and allow harmonics not visible in the wavelet scalogram to be seen. Because the harmonics are relatively far apart, the interference terms, especially noticeable in the CKR, oscillate quickly in the time direction; they do not present an interpretational challenge. In Kiss & Szatmáry (2002), the better time resolution of the CWD, when compared with the wavelet map (Fig. 8 of Kiss & Szatmáry (2002)), enables the increase in frequency of the main pulsation to be seen clearly. The strong interference terms caused by the interaction of the main component with its harmonics are not a problem, since they can be easily identified as spurious using the periodogram. The good time resolution also allows the first harmonic period evolution to be seen in detail, at higher frequency resolution than the wavelet. The wavelet does, however, show the second

harmonic, which is not visible in the CWD. In Figures 2 to 4 of Buchler, Kolláth & Cadmus (2004), the CKR is seen to give sharp images for the five large-amplitude irregular pulsators under discussion. The fundamental period and harmonics are visible; since these are relatively widely spaced in frequency, the ITs are easily identifiable.

The examples above show that if ITs can be identified as such (either by eye or with the use of the periodogram, for example) then the CWD and CKR are superior to the wavelet scalogram. For our data, however, we cannot rely on the periodogram to detect ITs, and so have to ensure by our choice of kernel parameters that they are attenuated as far as possible, resulting in poor time resolution for the CWD and CKR. This is particularly acute if we have to attenuate slowly-oscillating ITs resulting from components close in frequency, such as QPO harmonics, since this requires broad smoothing in the time-domain.

Apart from Dolan et al. (1998), the oscillations in the examples discussed above are long lived, and time resolution is of lesser importance. For transient oscillations, such as DNOs and QPOs, we have to have good time resolution to analyse them in a meaningful way; this makes the CWD and CKR, with kernels chosen to attenuate ITs, less effective for our data.

As an alternative to the CWD and CKR we have shown that the wavelet scalogram, with maximal ridges to compensate for its poorer frequency resolution at shorter scales than the CWD and CKR, does not suffer from ITs and has the requisite high time resolution at high frequencies. In addition, it is a minimum variance distribution in the presence of non-white noise, with a statistical distribution easily computable from the power spectrum.

It is this combination of time-frequency resolution attributes and desirable properties in the presence of noise that makes the wavelet scalogram an appropriate TFR for further study as a tool for detecting DNOs and QPOs.

4.13 Software

A variety of software is available for time-frequency analysis. The figures of TFRs in this chapter were generated using the Matlab Time-Frequency Toolbox, developed by François Auger, Olivier Lemoine, Paulo Gonçalves and Patrick Flandrin. It is available at <http://tftb.nongnu.org/>, and is also compatible with GNU Octave (available at <http://www.gnu.org/software/octave/>). An Ansi C version is available at <http://www-lagis.univ-lille1.fr/~davy/-toolbox/Ctftbeng.html>.

Some other (open-source) time-frequency packages available on the web include:

- **WWZ**: Fortran and Windows VB.net versions of Foster (1996)'s WWZ, available from <http://www.aavso.org/data/software/>.
- **TiFrAn**: a time-frequency package developed by Z. Csubry and Z. Kolláth for the analysis of light-curves, available for both Windows and Linux systems, at <http://www.konkoly.hu/tifran/index.html>.
- **WaveLab**: a comprehensive wavelet toolbox for Matlab and GNU Octave, available at <http://www-stat.stanford.edu/~wavelab/>.
- **WMTSA**: an implementation of the wavelet-based techniques for the analysis of time series presented in Percival & Walden (2000), for Mat-

lab.

- Fortran, IDL and Matlab implementation of the continuous wavelet transform, including significance testing, based on Torrence & Compo (1998), available at <http://atoc.colorado.edu/research/wavelets>.

5 The Wavelet Scalogram

We introduced the wavelet scalogram briefly in the previous chapter as one method of smoothing the WVD to remove interference terms, and showed that, with the addition of maximal ridges, the scalogram provides an appropriate method for analysing our data, which require good time and frequency resolution over a broad frequency range, and well-defined behaviour in the presence of noise. In this chapter we give a more detailed description of wavelet theory, focusing on practical implementation issues such as the choice of analysing wavelet. We then analyse a number of synthetic light curves with features similar to those of our data, in order to explore the behaviour of the wavelet scalogram in detail.

5.0.1 Introduction

We begin with a more traditional introduction to wavelets than that given in the previous chapter, following the standard approach of, for example, Daubechies (1992), Goswami & Chan (1999) or Chui (1997).

First, we define a family of functions

$$\psi_{s,b}(t) = \frac{1}{\sqrt{s}} \psi_0 \left(\frac{t-b}{s} \right), \quad s > 0, b \in \mathbb{R}$$

where ψ_0 is a fixed function, called the “mother wavelet”, that is well localised in both time and frequency i.e.

$$|\psi_0(t)| \leq c(1 + |t|)^{-1-\epsilon}, \quad |\hat{\psi}_0(\omega)| \leq c(1 + |\omega|)^{-1-\epsilon} \quad (37)$$

for some $\epsilon > 0$ (Cohen & Kovačević, 1996).

If we form the scalar product of $\psi_{s,b}$ and $x(t) \in \mathbb{R}$, then, just as $\langle x, e^{-i\omega t} \rangle$ is the continuous Fourier transform (FT),

$$T_{wav}(s, b; \psi) = \langle x(t), \psi_{s,b} \rangle = \frac{1}{|s|^{1/2}} \int x(t) \psi_0^*\left(\frac{t-b}{s}\right) dt, \quad (38)$$

gives the continuous wavelet transform (CWT). The main difference between the FT and the CWT is that $T_{wav}(s, b; \psi)$ is a 2-dimensional function of both time and frequency, while $\hat{f}(\omega)$ is a 1-D function of frequency only. If the wavelet function $\psi_{s,u}(t)$ is complex then the wavelet transform will also be complex, so it is useful to define the wavelet scalogram, analogous to the Fourier power spectrum, as $A_{SCALO}(s, b; \psi) = |T_{wav}(s, b; \psi)|^2$. We will use the terms wavelet scalogram and wavelet spectrum interchangeably.

s is called the dilation parameter, or the scale of the wavelet, and acts as a measure of the frequency of the wavelet, since as s changes, $\psi^{s,0}(t) = \frac{\psi(\frac{t}{s})}{|s|^{1/2}}$ covers a different frequency range (Goswami & Chan, 1999). b is called the translation parameter as it indicates the location of the wavelet along the time axis.

The wavelet function at each scale s is normalized to have unit energy:

$$\hat{\psi}(s\omega_k) = \left(\frac{2\pi s}{\Delta t}\right)^{\frac{1}{2}} \hat{\psi}_0(s\omega_k)$$

where $\int_{-\infty}^{\infty} |\hat{\psi}_0(\omega')|^2 d\omega' = 1$, which enables comparisons between wavelet transforms at each scale s and to the transforms of other time series (Torrence & Compo, 1998).

There are several constraints that we can impose on our choice of suitable functions for wavelets, over and above the localisation constraints given in equation 37. Ideally, for example, we would like to be able to reconstruct $x(t)$ given the wavelet components $T_{wav}(a, b; \psi)$, just as a signal can be reconstructed from its Fourier transform components. Since the wavelet transform is the 2-D projection of a 1-D function it is redundant, so the inverse is not uniquely defined (Fritz & Bruch, 1998). If, however, we impose the finite energy constraint $C_\psi = \int_{-\infty}^{+\infty} \frac{|\hat{\psi}(\omega)|^2}{|\omega|} < +\infty$, then $f(t)$ can be reconstructed by

$$f(t) = C_\psi^{-1} \int_0^{+\infty} \frac{da}{a^2} \int_{-\infty}^{+\infty} [W_\psi f](a, b) \psi_{a,b}(t) : db$$

(Cohen & Kovačević, 1996). The finite energy constraint implies that both $\hat{\psi}(0) = 0$ and $\int_{-\infty}^{+\infty} \psi(t) dt = 0$, so only (time-localised) oscillatory $\psi_{a,b}(t)$ are permissible (Fritz & Bruch, 1998), hence the name wavelet.

It is helpful to be able to plot the wavelet transform as a function of frequency rather than scale. Since the peak in the Fourier transform of the wavelet, $\hat{\psi}(s\omega)$, does not necessarily occur at a frequency of s^{-1} , Meyers, Kelly & O'Brien (1993) show that for wavelets with a dominant frequency, a cosine wave of known frequency can be substituted into equation 38, and the scale s at which the wavelet spectrum reaches its maximum calculated.

Figure 44 shows how wavelets with different time and scale parameters analyse the signal at different parts of the time-frequency plane. Notice that, as discussed in the previous chapter, as the analysing frequency increases, times support decreases and frequency support increases.




Figure 44: The time-frequency boxes of a wavelet scalogram. From Mallat (1998).

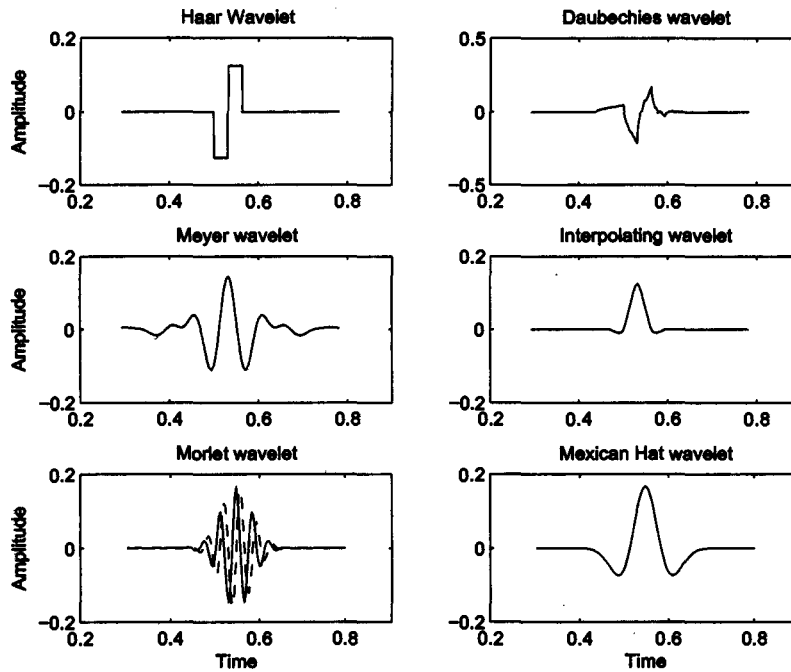


Figure 45: Examples of different mother wavelets.

5.1 Choice of Analysing Wavelets

There are an infinite variety of functions, both real and complex, which satisfy the restriction given in section 5.0.1. Figure 45 shows the Haar, Daubechies, Meyer, Interpolating, Morlet and Mexican Hat wavelets.

The choice of a particular wavelet depends on the signal to be analysed, and the type of analysis desired. If the shape of the signal to be detected is known, then the wavelet with the nearest matching shape should be chosen (Meyers, Kelly & O'Brien, 1993). Considering the often (amplitude modulated) sinusoidal shape of DNOs and QPOs, smoothly oscillating wavelets such as the Morlet, Mexican Hat and Paul wavelet are indicated, rather

than the abrupt step function of the Haar wavelet, or the spiky Daubechies wavelet.

We noted in section 4.7.1 that the uncertainty principle implies that a function cannot be arbitrarily well localised in both time and frequency; if σ_t and σ_ω are the time and frequency spread $\psi_{s,u}(t)$ respectively (defined in section 4.7.1),

$$\sigma_t \sigma_\omega \geq \frac{1}{2}. \quad (39)$$

If σ_t and σ_ω are finite, as is the case for any $\psi(t)$ obeying the constraints(37), then $\psi(t)$ is called a time-frequency window and $T_{wav}(s, b; \psi)$ contains the information of $x(t)$ in the time window

$$[s\bar{t} + b - s\sigma_t, s\bar{t} + b + s\sigma_t]$$

and the frequency window

$$\left[\frac{1}{s}(\bar{\omega} - \sigma_\omega), \frac{1}{s}(\bar{\omega} + \sigma_\omega)\right]$$

(Goswami & Chan, 1999).

Different wavelets make different compromises between σ_t and σ_ω , and hence between time and frequency resolution.

An additional consideration is that for complex wavelets, such as the Morlet, the phase and amplitude information of the data are separated: amplitude information about the signal is contained in $|W_n(s)|$, while phase information

is contained in $\tan^{-1} [\text{Im}\{W_n(s)\}/\text{Re}\{W_n(s)\}]$ (Torrence & Compo, 1998). This means that complex wavelets have good frequency resolution, but have constant power across the time duration of an oscillation, which does not enable extrema to be detected. For real wavelets, such as the Mexican Hat wavelet, phase and amplitude information is superposed in the time-frequency plane; they are useful for detecting extrema and discontinuities (Percival & Walden, 2000).

In order to explore fully the behaviour of wavelets we thus investigate two smooth wavelets with very different features: the complex Morlet wavelet and the real Mexican Hat wavelet.

5.1.1 The Morlet Wavelet

The Morlet wavelet is described by

$$\psi(\eta) = \pi^{-\frac{1}{4}} e^{-\frac{1}{2}\eta^2} e^{i\omega_0\eta} \quad (40)$$

where $\eta = t/s$ is a non-dimensional time parameter and ω_0 is the wavenumber (Torrence & Compo, 1998). The Morlet wavelet is a complex exponential of frequency $\frac{\omega_0}{2\pi}$, modulated in amplitude by a Gaussian window (Percival & Walden, 2000).

The equivalent Fourier period for a given scale s is given by

$$\text{Period} = \frac{4\pi s}{\omega_0 + \sqrt{2 + \omega_0^2}}$$

Figures 46 a) and b) show the Morlet wavelet for $\omega_0 = 3$ and $\omega_0 = 7$ respectively, from which it can be seen that increasing ω_0 increases the number of oscillations in the wavelet. Figures 46d) and e) give the periodograms of the wavelets in figures 46a) and b), showing that increasing ω_0 decreases the frequency spread of the wavelet, increasing the frequency resolution. Note that the Morlet wavelet is analytic ($\hat{\psi}_0(\omega) = 0 \forall \omega < 0$) and hence can be used for instantaneous frequency calculation, as will be discussed in section 5.1.3.

We have chosen to use $\omega_0 = 6$ in our analysis, which gives a Morlet wavelet which has 5 oscillations. This is the most commonly used value of ω_0 , but is also appropriate given the QPO detection criteria which will be discussed in chapter 7.

5.1.2 The Mexican Hat Wavelet

For real wavelets, the wavelet transform measures the variation of $x(t)$ in a neighbourhood of time u , whose size is proportional to scale s . Real wavelets have good resolution in the time domain because they only produce power at times when the signal has an extremum or a sharp discontinuity. They can be used to map signal extrema and detect discontinuities (Percival & Walden, 2000)

One of the most commonly used real wavelets is the Mexican Hat Wavelet (MHAT), with

$$\psi(\eta) = \frac{1}{\sqrt{\Gamma(2.5)}} \left(1 - \frac{1}{\eta^2}\right) e^{-\frac{\eta^2}{2}} \quad (41)$$

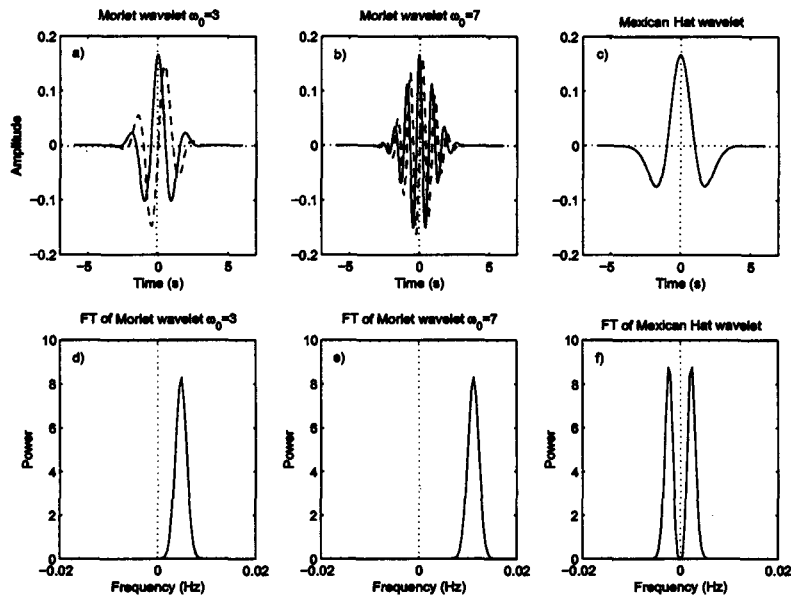


Figure 46: a) and b): Morlet wavelet for two different values of ω_0 . The real part is indicated by a solid line, the imaginary part by a dashed line. c) The Mexican Hat wavelet. d), e) and f): Fourier transforms of the wavelets shown in a), b) and c).

where again $\eta = t/s$. The equivalent Fourier period for a given scale s is

$$\text{Period} = \frac{2\pi s}{\sqrt{2.5}}$$

The MHAT wavelet is the normalized second derivative of a Gaussian function; figure 46c) shows the mother wavelet. When compared with the Morlet wavelet the MHAT wavelet has good resolution in the time domain, but is dispersed in the frequency domain, as can be seen in figure 46f).

5.1.3 Maximal Ridges and Instantaneous Frequency

Delprat et al. (1992) showed that, for analytic wavelets such as the Morlet wavelet, $|A_{SCALO}|^2$ reaches a maximum at the instantaneous frequency $\phi'(t)$ defined in section 4.6, for slowly changing low frequency signals. A detailed discussion is given in Mallat (1998), who extended the proof to show that the ridge algorithm of Delprat et al. (1992) also accurately measures instantaneous frequency for components changing rapidly at high frequency.

For a monocomponent signal containing a periodic or quasi-periodic signal, the wavelet maxima will form a horizontal line in the Morlet spectrum at the period of the oscillation, called a wavelet ridge. For a signal that has several components in different frequency bands, the local (frequency) maxima will give the instantaneous frequency in each band.

In the MHAT spectrum, the wavelet maxima present as points at the times of the maxima and minima of oscillations in the signal, at the period of the oscillation. While these points do not show the instantaneous frequency, they do show the time at which extrema occur, which we will use as part of our method for detecting QPOs, described in chapter 7.

5.1.4 Cone of Influence

Since the wavelet is localized in time and frequency, $T(s, b; \psi)$ is affected only by the signal around $t = b$, in a radius that depends on the width of the wavelet as determined by the scale s . More specifically, each wavelet has (compact) support on a finite time interval $[-C, C]$. The cone of influence (COI) of a point (t', s') on the time-frequency plane is the set of all points

(t'', s) for which $x(t')$ affects $T(s, t''; \psi)$. Mallat (1998) shows that the COI is defined by $|t'' - t'| \leq Cs$.

Figure 47, adapted from Mallat (1998), shows the COI of the point $(t', 0)$.

Finite data length means that for points near the beginning or end of the data set, the support of $T(s, b; \psi)$ thus effectively includes a number of points with zero amplitude (i.e. points before or after measurement started), and therefore has significantly reduced power. Torrence & Compo (1998) define the e-folding time for the wavelet as the time-radius at which the wavelet power first drops by a factor of e^{-2} at an edge. The time extent to which the decreased amplitude affects the wavelet transform depends on the scale of the wavelet and the type of wavelet; at shorter scales (shorter periods) the effect is minimal, while at very long scales, the effect is marked.

5.1.5 Confidence Contours

We have assumed that the noise structure of our data remains constant, and is i.i.d Gaussian, so that we can fit a stationary model to the noise continuum using the periodogram (as discussed in section 4.5.1). After fitting the appropriate noise spectrum to the periodogram, and choosing a particular confidence level for χ_2^2 , one can calculate the confidence level at each scale, since as discussed in 4.11.2, $|W_n(s)|^2 / \sigma^2 \sim \frac{1}{2} P_k \chi_2^2$ for complex wavelets. For real wavelet $|W_n(s)|^2 / \sigma^2 \sim \chi_1^2$ (Torrence & Compo, 1998).

To check that our assumption of a stationary noise structure for each data set was correct, we calculated the best-fit noise-model for sections of each run as well as for the entire run; in all cases we found that the models for




Figure 47: The cone of influence of the point $(t', 0)$. Adapted from Mallat (1998).

each section were very similar to the model for the complete time series.

5.2 Implementation

Torrence & Compo (1998) show that the wavelet transform can be calculated quickly and efficiently using

$$W_n(s) = \sum_{k=0}^{N-1} \hat{x}_k \hat{\psi}^*(s\omega_k) e^{i\omega_k n \Delta t}. \quad (42)$$

Rather than use linear scales, we write the scales as fractional powers of two:

$$s_j = s_0 2^{jd_j}, j = 0, 1, \dots, J$$

$$J = d_j^{-1} \log_2(N \Delta t / s_0)$$

where s_0 is the smallest resolvable scale and J determines the largest scale. In practise we are limited in the range of s (or equivalently frequency) by the integration time (Δt) and the finite length of the data. The highest frequency which we can use is the Nyquist frequency ($\frac{1}{2\Delta t}$), which gives an upper bound on s_0 . Smaller values of d_j give finer resolution; we typically use $d_j = 0.025$.

We have used the Matlab toolbox of Torrence & Compo (1998) as a starting point, but have added instantaneous frequency algorithms and confidence contour calculations appropriate to our data.

5.3 Wavelet Examples

In this section, we examine the wavelet scalograms of a variety of synthetic signals, following the example of Szatmary, Vinko, & Gál (1994), who studied wavelet scalograms of synthetic signals mimicking features of LPV lightcurves. We have created synthetic light curves that test wavelet spectrum behaviour for various aspects of DNO and QPO behaviour. All data are 8000 s long, and are sampled at 4 s intervals, which is the most common integration time for our data.

The colour of the wavelet spectrum at any point indicates the amplitude of the wavelet spectrum at that point. In the colour map that we have used, small amplitudes are blue, medium amplitudes are green or yellow, and large amplitudes are red; the deeper the red, the larger the amplitude. Since the wavelet scalogram is the square of the wavelet transform, negative minima are converted to (positive) maxima.

5.3.1 A Simple Sinusoid

Figure 48 shows the Morlet and MHAT spectra of a sinusoid with period 750 s. Note that while it appears in figure 48 that the Morlet scalogram gives the frequency of the oscillation as 725 s, this is an artefact of the logarithmic period scale, which makes it difficult to estimate periods by eye without magnifying the period range of interest: zooming in on the scalograms (figure 49) shows that the period is, in fact, 750 s.

The Morlet spectrum shows a narrow frequency band of high power centred at 750 s, but the amplitude of the signal is constant for the duration of the

signal; there is no indication of where the signal reaches local maxima or minima. In contrast, the MHAT spectrum of the same signal shows narrow (in the time domain) oval shapes centered at the times of the maxima and minima of the sinusoid, but dispersed over a wide frequency range.

The instantaneous frequency appears as a continuous (black) line in the Morlet scalogram at the frequency of the signal. There are several maximal ridges at longer periods, but these are relatively short, and are due to the finite length of the data. They do not occur in regions of high amplitude relative to the amplitude at 750 s.

Comparing the maxima of the MHAT scalogram (shown as black dots in the MHAT scalograms in figures 48 and 49) with the signal, we can see that the maximal points show the times (and period) at which the signal reaches an extremum.

The region below the black 'u' in both wavelet spectra in figure 48 shows the region in which the edge effects are significant; short periods are relatively unaffected by edge effects, but longer periods can be seriously affected. The effect is more pronounced in the Morlet spectrum, as the Morlet wavelet has larger time support than the MHAT at a given scale, and hence a broader COI.

Figure 48 highlights the different information available from the Morlet and MHAT spectra. Using the well-resolved frequency information from the Morlet scalogram, and the well-resolved time (or phase) information from the MHAT scalogram, we can gain a far more accurate picture of the original time series than by using either wavelet spectrum alone.




Figure 48: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period 750 s (top).

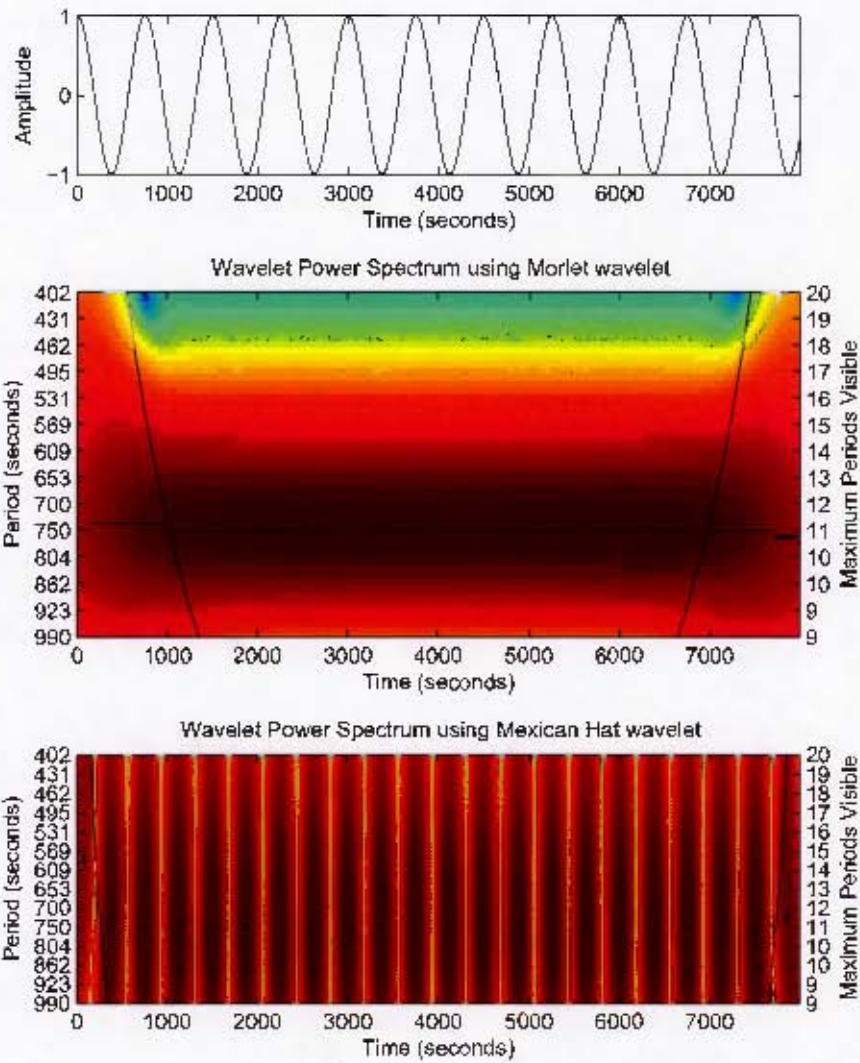


Figure 49: Detail of Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period 750 s (top).

5.3.2 Gaps

Our data are evenly sampled, but a few sets contain short gaps of at most a few hundred seconds. We linearly interpolate between the endpoints of the gap, rather than setting the missing points to zero. To investigate the effect of gaps, figures 50 and 51 show the Morlet and MHAT scalograms of a 100 s sinusoid sampled every 4 seconds; while figure 50 contains no gaps, figure 51 contains three gaps, of 16 s, 75 s and 100 s respectively.

We can again see that due to the larger time support of the Morlet wavelet, the Morlet spectrum is more affected by gaps than the MHAT. While at short periods, the effect of the gaps is relatively well localised in time in both wavelet spectra, at longer periods the effect is pronounced, giving much higher amplitudes than the ungappy data at the same period. In the Morlet spectrum, the gaps result in many more maximal ridges, at higher amplitudes.

5.3.3 Frequency Modulations

Since both DNOs and QPOs change frequency in a variety of ways, it is important to have an understanding of how well the wavelet spectra can detect these changes.

Figure 52 shows the wavelet spectra of a sinusoidal signal with period varying linearly from 200 s to 400 s. In both wavelet spectra the period of the sinusoid is accurately depicted by the maximal ridges/points.

The signal in figure 53 is a sinusoid with period changing abruptly from 200 s




Figure 50: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period 100 s (top).

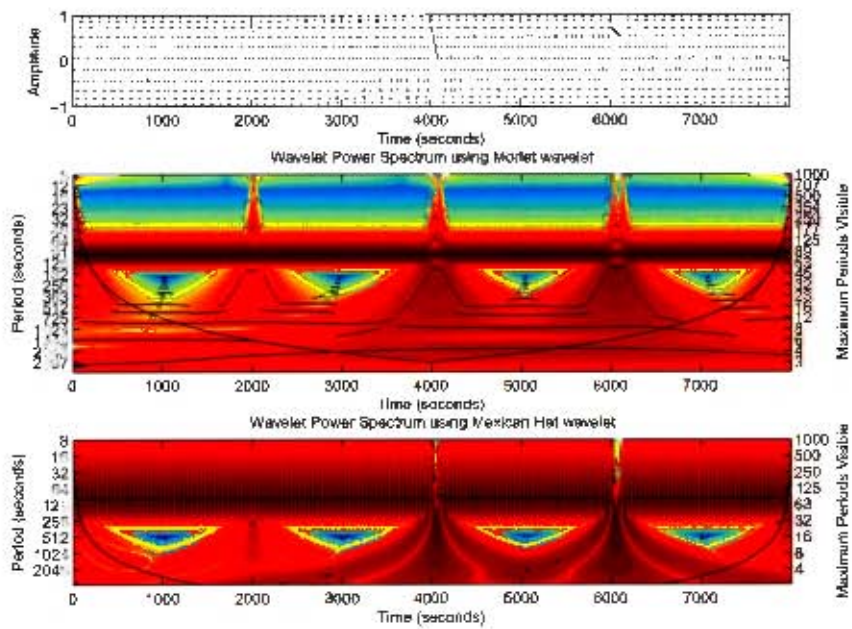


Figure 51: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period 100 s (top), with three gaps of lengths 20 s, 60 s and 100 s respectively.

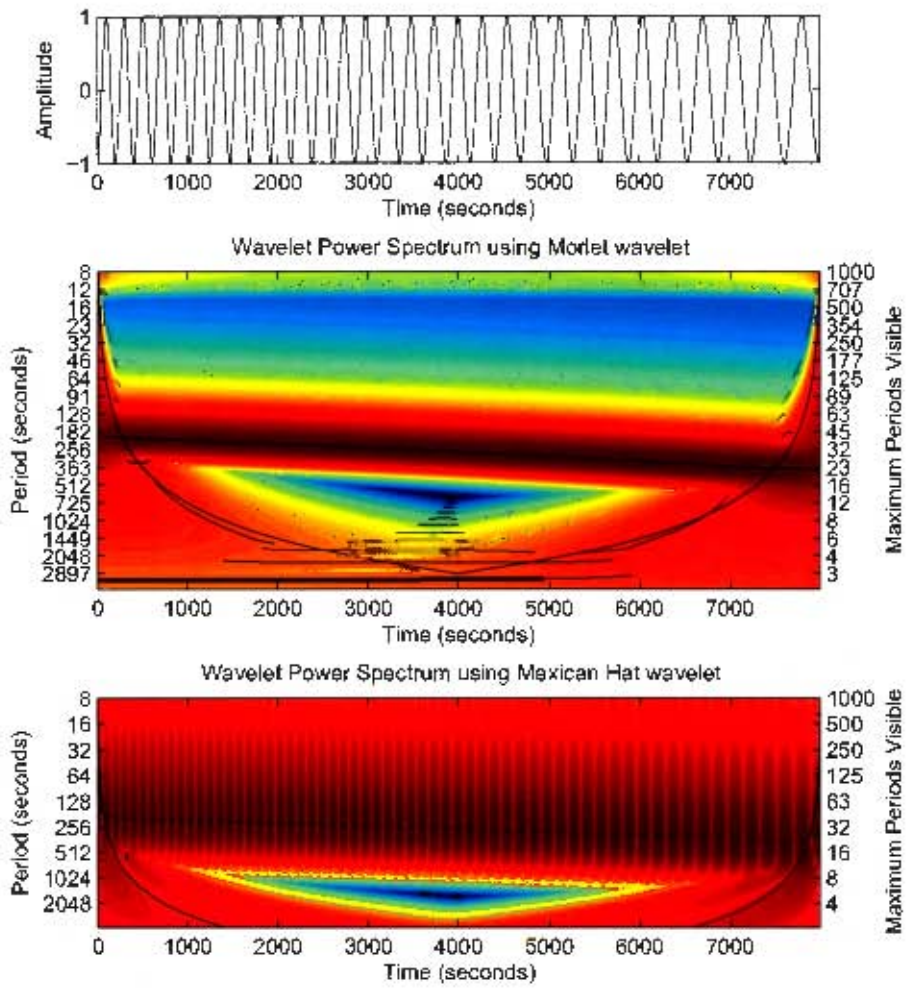


Figure 52: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period varying from 200 s to 400 s (top).

to 300 s at time 4000 s. Away from the change point, the Morlet spectrum accurately gives the period of the signal. However, the change in period is not abrupt in the Morlet spectrum; rather, it appears to change linearly over several hundred seconds. The Morlet spectrum also shows large amplitudes at shorter periods near the change point.

In contrast, the MHAT has only a single maximum between the two true frequencies; it accurately reflects the abrupt change in the period of the signal. Shorter periods do not appear to be affected by the change in period as was the case for the Morlet wavelet.

Figures 54, 55 and 56 show sinusoids with periods centred at 600 s but varying periodically (at a period of 2000 s); the amplitude of the period change is 10 s in figure 54, 20 s in figure 55 and 30 s in figure 56.

In all three figures, both wavelet spectra accurately track the period change. However, in the Morlet spectra spurious maximal ridges appear at a period of about 900 s when the period of the signal is at a minimum. There is also a maximal ridge at a period of about 1600 s in all three Morlet spectra.

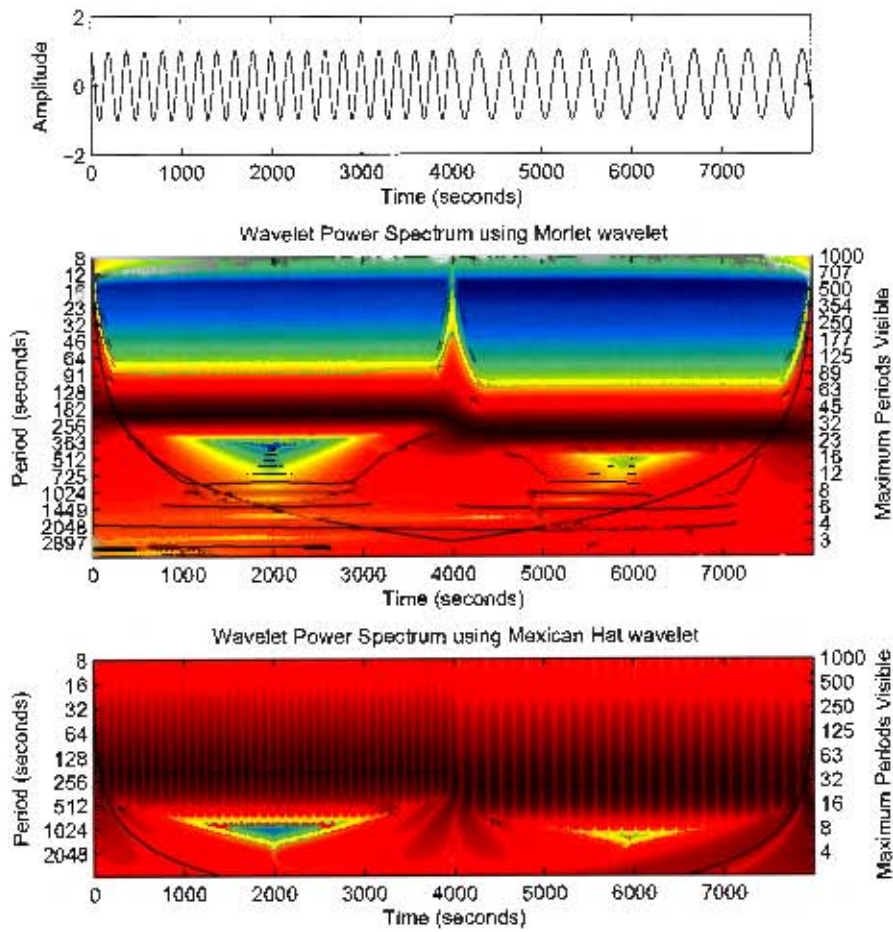


Figure 53: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period changing from 200 s to 300 s at 4000 s (top).

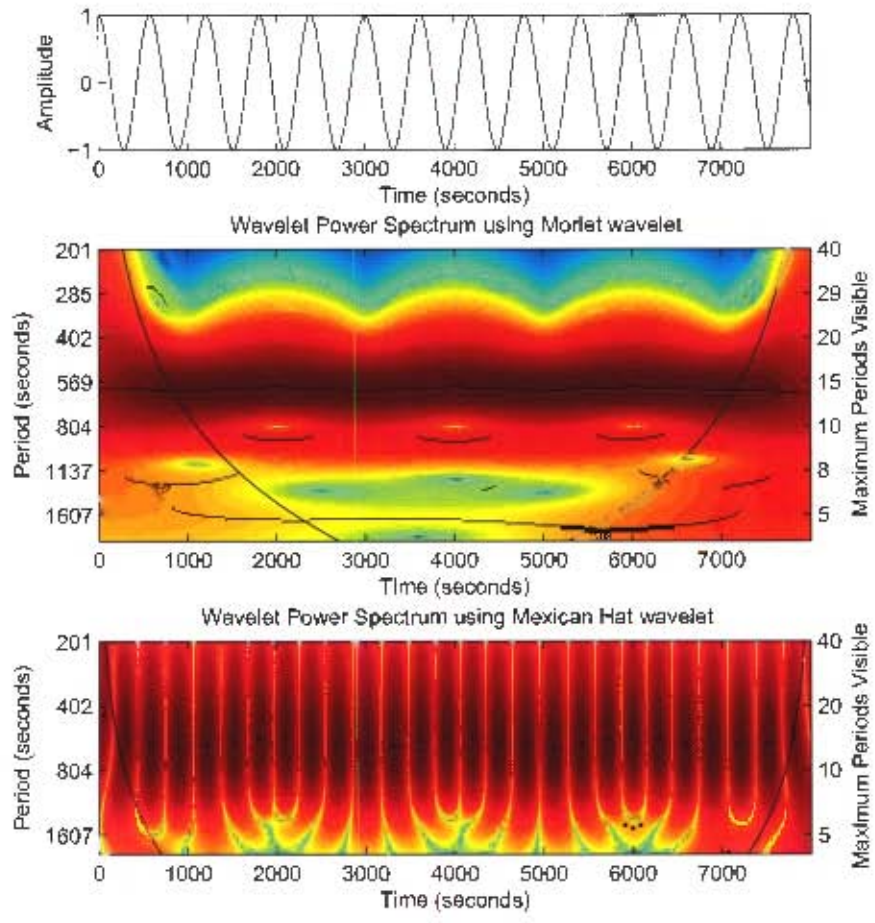


Figure 54: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period $600 + 10 \cos(\frac{2\pi t}{2000})$ (top).




Figure 55: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period 600 - $20 \cos\left(\frac{2\pi t}{2000}\right)$ (top).




Figure 56: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with period $600 + 30 \cos(\frac{2\pi t}{2000})$ (top).

5.3.4 Phase Jumps

Both DNOs and QPOs exhibit abrupt changes in phase, so in this section we investigate the effect of phase jumps on the wavelet spectrum.

Figure 57 shows the wavelet spectra of a sinusoidal signal with period 400 s, with a single phase change of 0.1 of a cycle (40 s) at 4000 s (marked by an arrow above the signal). Both Morlet and MHAT spectra are affected near to phase jump, the Morlet more than the MHAT. The phase change appears as a brief decrease in period in both spectra, as one would expect since the effect of the phase change is to decrease the length of the affected cycle. As with the abrupt period change in figure 53, the Morlet spectrum is affected at shorter periods near the phase jump.

Figure 58 shows the wavelet spectrum of a sinusoidal signal with period 400 s, with a single phase change of 0.4 of a cycle (160 s) at 4000 s (marked by an arrow above the signal); figure 59 shows the periodogram of this signal, which appears to contain two peaks of different amplitude on either side of the correct signal frequency (0.0025 Hz). This gives an excellent example of the use of the wavelet scalogram: in both the Morlet and MHAT spectra the true (single) period nature of the signal is visible, at the correct period, and the phase change can be seen as the time-localised event that it is. This larger phase change has a longer time effect on both wavelet spectra; about 1000 s for the Morlet wavelet and about 500 s for the MHAT wavelet.

In figure 60, the signal is again a 400 s period sinusoid, with 10 evenly spaced random phase jumps, with a maximum jump of 160 s. Here the effect of the phase jumps is far more pronounced than in the preceding figures; in




Figure 57: Morlet (middle) and Mexican Hat (bottom) scalograms of a 400 s sinusoid with a phase jump of 40 s at time 4000 s (top).

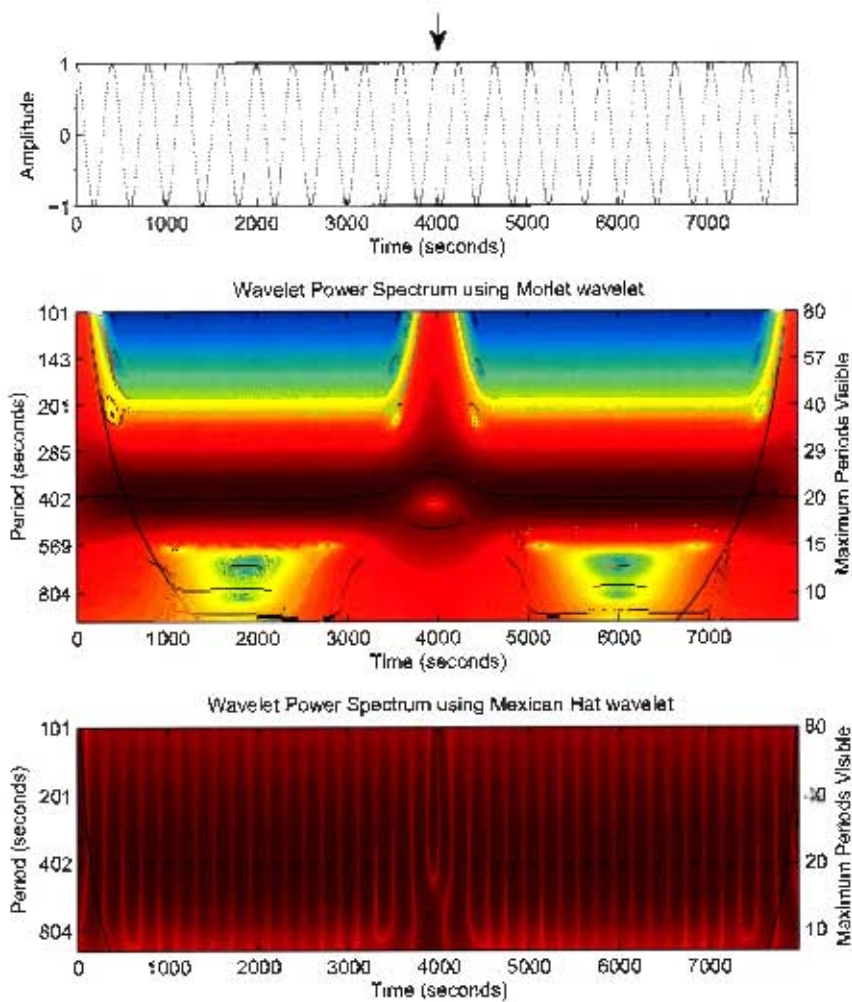


Figure 58: Morlet (middle) and Mexican Hat (bottom) scalograms of a 400 s sinusoid with a phase jump of 160 s at time 4000 s (top).




Figure 59: Periodogram of a 400 s sinusoid with a phase jump of 160 s at time 4000 s (top).




Figure 60: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with 10 random phase jumps (top).

both wavelet spectra the instantaneous frequency appears to vary around the true period, with the largest fluctuations near the largest phase jumps. The Morlet spectrum also shows several maximal ridges at various periods other than the true period.

5.3.5 Amplitude Modulations

Figures 61, 62 and 63 show a 400 s sinusoid, with amplitude modulated by a sinusoid of period 2000 s in the first two figures, and 3500 s in the third. In figure 61 the range of the amplitude change is 0.2, in figure 62 it is 0.5 and in figure 63 0.8. In all three wavelet spectra the instantaneous frequency is found accurately, although as the amplitude modulation increase there are additional maximal ridges in the Morlet wavelet at longer periods. In figures 61 and 62 the MHAT wavelet maximal ridges show the extrema accurately in time and period, but in figure 63 the maximal ridges are inaccurate in period when the amplitude of the signal is very small. In neither the Morlet nor the MHAT spectra is it easy to assess the magnitude of the amplitude fluctuation.

5.3.6 Multi-component Signals

Figure 64 shows the wavelet spectra of the synthetic DN lightcurve introduced in the chapter 4. While the instantaneous frequency of the Morlet spectrum shows the correct periods for the 4 components, the maximal ridges in the MHAT are less satisfactory, especially at the short DNO period (25-30 s) for which the maxima are correctly identified in time, but spread in the frequency direction. The MHAT maxima of longest period QPO (750-900 s) are the most accurate in period.

Figure 65 shows a signal consisting of two sinusoids of equal amplitude, and periods 400 s and 600 s. The maximal ridges in the Morlet spectrum show both periodicities, although there is a small oscillation in the estimated




Figure 61: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with amplitude $1 - 0.2 \sin(\frac{2\pi t}{2000})$ (top).




Figure 62: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with amplitude $1 + 0.5 \sin\left(\frac{2\pi t}{2000}\right)$ (top).




Figure 63: Morlet (middle) and Mexican Hat (bottom) scalograms of a sinusoid with amplitude $1 + 0.8 \sin(\frac{2\pi t}{2000})$ (top).

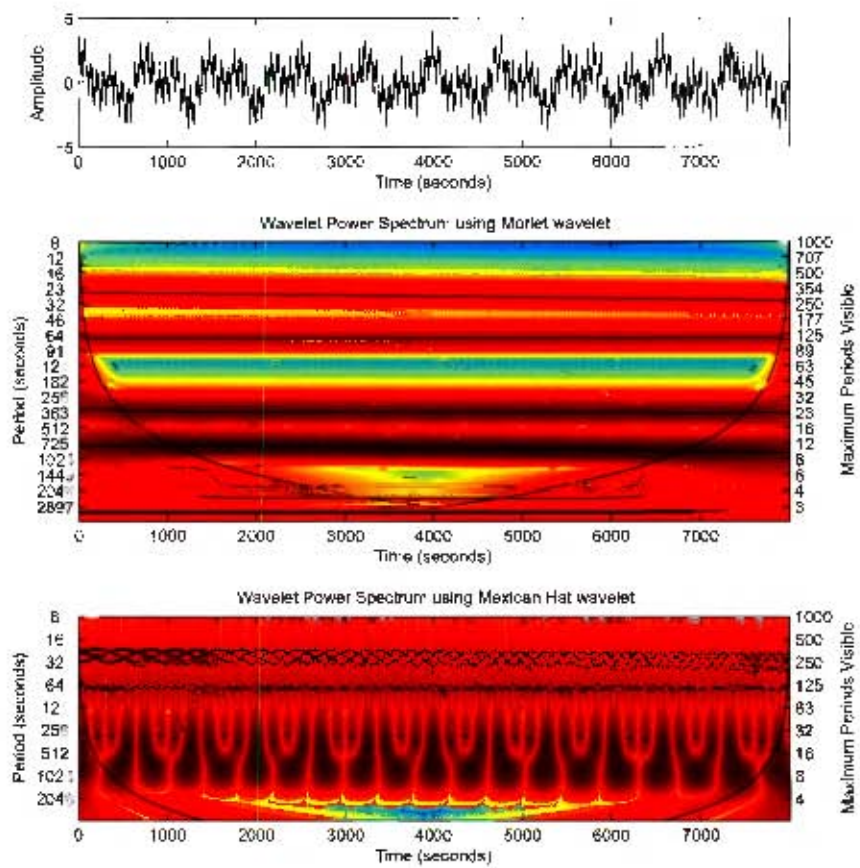


Figure 64: Morlet (middle) and Mexican Hat (bottom) scalograms of the synthetic DN signal (top).

period for both periodicities. The maximal ridges in the MHAT identify the maxima and minima in the signal at the correct times, but the periods are not correctly resolved; there is only a single line of maxima points, fluctuating around a period of about 500 s.

Figures 66 and 67 both show a signal consisting of the sum of an 800 s and a 1200 s sinusoid, with the phase difference between the sinusoids 0 cycles in figure 66 and 0.3 cycles in figure 67. While the periodicities are resolved in both Morlet spectra, there is again a slight oscillation in both periods. In addition, the Morlet spectrum in figure 67 is badly affected at the ends of the signal, where the double period is no longer correctly resolved.

In both figures the MHAT wavelet does not resolve the two periodicities, except at a few times in figure 66. As in the previous example, the dominant period appears to be the average of the two true periods.

Figure 68 again shows a signal that is the sum of two sinusoids, this time with periods 900 s and 1200 s. Here, the Morlet spectrum only correctly resolves the two periodicities at a few points, as does the MHAT spectrum.




Figure 65: Morlet (middle) and Mexican Hat (bottom) scalograms of a signal consisting of two sinusoids, with periods 400 s and 600 s (top).




Figure 66: Morlet (middle) and Mexican Hat (bottom) scalograms of a signal consisting of two sinusoids, with periods 800 s and 1200 sand phase different 0 cycles (top).




Figure 67: Morlet (middle) and Mexican Hat (bottom) scalograms of a signal consisting of two sinusoids, with periods 800 s and 1200 s and phase different 0.3 cycles (top).




Figure 68: Morlet (middle) and Mexican Hat (bottom) scalograms of a signal consisting of two sinusoids, with periods 900 s and 1200 s (top).

5.3.7 Noisy Signals

To investigate the effect of noise on the wavelet spectra, we have constructed a pure noise signal, with a red noise power spectrum similar to the noise of our data. We used the Time Frequency Toolbox function `noisecg`, with parameters $A_1 = 0.4$ and $A_2 = 0.2$, which gives complex colored gaussian noise obtained by filtering a white gaussian noise through a $\frac{\sqrt{1-A_1^2-A_2^2}}{(1-A_1z^{-1})-A_2z^{-2}}$ second order filter.

The noise signal is shown in the top panel of figure 69, with the periodogram in the lower panel. The top panel of figure 70 shows the log periodogram, with the third order polynomial fit, while the lower panel shows the periodogram, with the fitted periodogram and 95% and 99% confidence intervals.

Figure 71 shows the wavelet spectra of the pure noise signal. In the colour scheme we have used, the noise background changes from yellow (i.e. low amplitude) at the short periods to red (i.e. higher amplitude) at the longer periods, as expected for a red-noise continuum. The blue confidence contours indicate where the signal has power above the 95% significantly above the level expected from the noise at that period; as expected, even for a pure noise signal, there are some regions which are identified as having significant power, although they are generally of short duration. While the Morlet spectrum shows some maximal ridges at longer periods which appear to indicate a periodicity, they are not of significant power.

In figure 72 we have added the red noise signal to the synthetic DN lightcurve; the average amplitude of the noise is 1, while the amplitude of each of the components of the synthetic light curve is also 1. In the Morlet spectrum,




Figure 69: Top: red noise signal (see text for details). Bottom: periodogram of red noise signal.

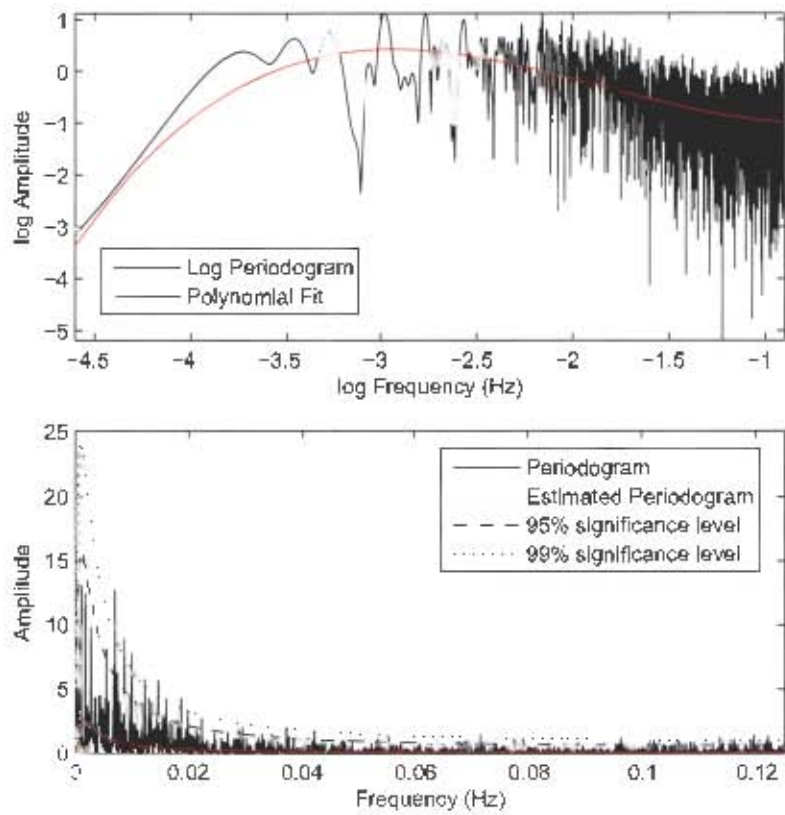


Figure 70: Top: Log periodogram and polynomial fit of red noise signal. Bottom: periodogram, fits and confidence levels of red noise signal.




Figure 71: Morlet (middle) and Mexican Hat (bottom) scalograms of red noise signal (top).

the signal components are all correctly detected in the 95% confidence intervals. While the instantaneous frequency of the components is no longer as straight as it was in the noise-free case, it is still possible to track the period change of the DNO and QPO components accurately.

The MHAT spectrum is most affected by the noise at shorter periods; the extrema of the QPOs are still readily resolved, and occur at the correct times, when compared with the noise-free MHAT spectrum. Figure 73 shows the periodogram of the noisy DN lightcurve; all signal components are clearly visible in the spectrum, although the DNO component is evident only as a broad hump between 0.03 and 0.04 Hz.

Figures 74 and 75 again show the wavelet spectra and periodograms of a noisy version of the synthetic DN lightcurve; the amplitude of the added noise has been doubled. The components are all visible in both the wavelet spectrum and the periodogram, although the 70 s lpDNO does not appear as a continuous oscillation of significant power in the wavelet spectrum. It is still possible to use the maximal ridges in the Morlet wavelet to find the period of the components. The maximal points in the MHAT spectrum are not of use at the lpDNO and DNO periods, where the period estimate is greatly affected by the noise. However, at the QPO periods, it is possible to use them to detect the extrema of the oscillations correctly.

Figures 76 and 77 show the wavelet spectra and periodograms of the synthetic DN lightcurve, but with noise of amplitude 3 added. Now detection of all components is affected by the noise and barely any confidence contours identify the 70 s lpDNO. While QPOs only register as above the confidence level over some time intervals, the maximal ridges in the Morlet




Figure 72: Morlet (middle) and Mexican Hat (bottom) scalograms of a synthetic DN signal with red noise of amplitude 1 added (top).




Figure 73: Periodogram of synthetic DN signal with red noise of amplitude 1 added.




Figure 71: Morlet (middle) and Mexican Hat (bottom) scalograms of a synthetic DN signal with red noise of amplitude 2 added (top).




Figure 75: Periodogram of synthetic DN signal with red noise of amplitude 2 added.

spectrum still identify the periodicities, and show an unbroken instantaneous frequency for each QPO. The DNO appears as a series of significant regions linked by maximal ridges; the instantaneous frequency fluctuates, and shows breaks.

In the time intervals where the QPOs are detected about the significance level, the MHAT spectrum can be used to identify the extrema, all though not as accurately in time as in the previous less noisy examples.

All components are visible in the periodogram, although there are now several peaks due to noise which are above the 95% confidence level, and one that is above the 99% confidence level (at 0.007 Hz (142 s)). However, since the 70 s lpDNO is not detected at this noise level, it would not be possible to use the Morlet spectrum to identify the 0.007 Hz peak as noise or signal, if the true signal periodicities were unknown.

Figure 78 shows a sinusoidal signal with a basic period of 600 s but modulated in period and amplitude, with period range 580 to 620 s and amplitude range of 1 to 1.7. Figure 79 shows the periodogram of the signal, which has a main peak at 0.00181 Hz (550 s), and a secondary peak at 0.00165 Hz (620 s). The maximal ridge in the Morlet spectrum shows a single continuous component fluctuating in period correctly. There are however several maximal ridges at longer periods, although not at large amplitudes. The MHAT spectrum maximal points identify only a single component, and show the period fluctuation correctly.

In figure 80, we have added red noise of amplitude 1.5 to the signal in the previous two figures. When the amplitude of the signal is large then the




Figure 76: Morlet (middle) and Mexican Hat (bottom) scalograms of a synthetic DN signal with red noise of amplitude 3 added (top).




Figure 77: Periodogram of synthetic DN signal with red noise of amplitude 3 added.




Figure 78: Morlet (middle) and Mexican Hat (bottom) scalograms of an amplitude and frequency modulated sinusoid, with a basic period of 600 s but modulated in period and amplitude, with period range 580 to 620 s and amplitude range of 1 to 1.7 (top).




Figure 79: Periodogram of an amplitude and frequency modulated sinusoid, with a basic period of 600 s but modulated in period and amplitude, with period range 580 to 620 s and amplitude range of 1 to 1.7.




Figure 80: Morlet (middle) and Mexican Hat (bottom) scalograms of an amplitude and frequency modulated sinusoid, with a basic period of 600 s but modulated in period and amplitude, with period range 580 to 620 s and amplitude range of 1 to 1.7, with red noise of amplitude 1.5 added (top).

signal is detected as above the confidence level, but during the low amplitude times it is not. However the maximal ridge in the Morlet spectrum identifies the periodicity correctly, although the fluctuations in period are not as clear as in the noise free case. When the signal is at large amplitude (above the confidence level in the Morlet spectrum) the MHAT maximal ridges can be used to identify the extrema at the correct period and time.

In figures 82 and 83 we have add the amplitude-frequency modulated sig-




Figure 81: Periodogram of an amplitude and frequency modulated sinusoid, with a basic period of 600 s but modulated in period and amplitude, with period range 580 to 620 s and amplitude range of 1 to 1.7, with red noise of amplitude 1.5 added.

nal to the noise signal for the first half of the time series only. The large amplitude times of the signal are again detected as significant in the Morlet spectrum, and during these times the MHAT wavelet correctly identifies the (changing) period of the extrema. In the periodogram, while the signal is still clearly visible as a peak, there are a number of noise spikes with amplitudes above 99% confidence level.

5.4 Discussion

In this chapter we have given a detailed introduction to wavelet theory, including a discussion of the features of different types of analysing wavelet and the factors that need to be considered when choosing an analysing wavelet. We concluded that smooth, sinusoidally varying wavelets are most appropriate for our data, and chose to investigate both the complex Morlet and real MHAT wavelets in greater detail, since the Morlet wavelet is ideal for frequency estimation, while the MHAT wavelet can be used to analyse extrema and discontinuities.

We have shown that the maximal ridge algorithm of Delprat et al. (1992) gives accurate values for the instantaneous frequency of a multicomponent or time-varying signal when the Morlet wavelet is used, and that the maximal points of the MHAT spectrum can be used to accurately measure the time at which extrema occur in a signal. The Morlet spectrum does sometimes show maximal ridges at frequencies other than the true signal frequency, but the amplitude at which these ridges occur is usually significantly lower than the the true signal amplitude. If confidence contours are used, then there is little danger of identifying these ridges as significant frequency components.




Figure 82: Morlet (middle) and Mexican Hat (bottom) scalograms of an amplitude and frequency modulated sinusoid, with a basic period of 600 s but modulated in period and amplitude, with period range 580 to 620 s and amplitude range of 1 to 1.7, lasting for 4000 s, with red noise of amplitude 1.5 added (top).




Figure 83: Periodogram of an amplitude and frequency modulated sinusoid, with a basic period of 600 s but modulated in period and amplitude, with period range 580 to 620 s and amplitude range of 1 to 1.7, lasting for 4000 s, with red noise of amplitude 1.5 added

Gaps in the data affect the Morlet spectrum more than the MHAT spectrum. However, since our data rarely show gaps, we have not focussed on addressing this issue. For data with significant gaps the WWZ, discussed in chapter 3, should be used.

We investigated the effect of phase jumps, where again the Morlet spectrum is seen to be more affected than the MHAT spectrum. However, the maximal ridges can still be used to estimate the instantaneous frequency, even when 10 phase jumps were included in an 8000 s test sinusoid. The maximal points in the MHAT wavelet are also still an accurate estimate of the extrema of the signal in the presence of phase jumps.

Amplitude modulations do not seem to affect the ability of either wavelet spectrum to detect the period of a signal significantly. We investigated the ability of the Morlet and MHAT spectra to resolve two close periodicities, and found that for long period components (>1000 s), the minimum period difference that can be resolved is about 300 s, although the relative phase of the components can affect this. For multicomponent signals, the MHAT spectrum is most useful for components with periods greater than about 100 s, for which it can be used to track the extrema.

We constructed a pure noise signal with a red noise periodogram similar to that seen in our data, which we added to various signals to test the ability of the Morlet spectrum, with confidence contours, to detect periodic components in the presence of noise. We showed that even when the signal-to-noise ratio is 0.5, we are still able to detect signal terms. Both the maximal ridges and maximal points can be used to measure the instantaneous frequency and extrema times within the confidence contours.

Overall, we have shown the Morlet and MHAT wavelet provide complementary ways of viewing non-stationary, multi-component data with signal-to-noise ratios less than 1. This builds on previous work in astronomy, discussed in chapter 3, which focus on signals that are less intermittent, have higher signal-to-noise ratios, and cover narrower frequency bands than ours.

6 Empirical Mode Decomposition

6.1 Introduction

In the previous chapter we showed that wavelet analysis, with the addition of confidence contours and wavelet ridges, can be used to detect non-stationary signals in noise. We would like to be able to analyse the amplitude behaviour of the detected signals in greater detail than the wavelet spectrum allows. An alternative decomposition of the signal into ‘meaningful’ frequency bands is given by Empirical Mode Decomposition (EMD), which is a useful adjunct to wavelet analysis as it allows detailed measurement of the amplitude of components.

In this chapter we introduce the theory of EMD, and investigate the EMD of the synthetic signals introduced in the previous chapter to gauge its usefulness for our data. We also carry out Monte Carlo simulations to investigate the effect of noise on the amplitude of the EMD components. We conclude with a discussion of possible uses for EMD.

EMD has been used in astronomy, primarily in solar astrophysics, where it has been used to analyse the sunspot cycle (Astley, Komm & Howe (2004), Li et al. (2007) and Xu et al. (2008)), solar coronal oscillations (Terradas, Oliver & Ballester, 2004) the heliospheric magnetic field (Smith, Zhou & Ruzmaikin, 2006) and rotation residuals of the solar convection zone (Komm, Hill & Howe, 2001). It has also been used in geophysics to investigate variations in stratospheric ozone and temperature (Ruzmaikin et al., 2007), the Earth’s magnetic field (Roberts, Yu & Russel, 2007) and coastal currents (Rossi & Salusti (1995) and Salon et al. (2003)).

6.2 Method

Huang et al. (1998) developed this method to decompose multi-component data into a set of functions for which a meaningful instantaneous frequency can be defined everywhere, called Intrinsic Mode Functions (IMFs). It is assumed that the data have at least two extrema (one maximum and one minimum) and that the characteristic time scale (of each IMF) is defined by the time lapse between extrema. This does require that the data are oversampled sufficiently that the extrema are well-defined (Rilling, Flandrin & Gonçalvès, 2003). If the data contain no extrema, only inflection points, then differencing the data once or more will reveal the extrema, and the final results can be obtained by summing the components.

To extract the IMFs from the original data, a sifting process is followed:

1. All local maxima are identified, and an upper envelope, $e_{max}(t)$ is constructed by interpolating between the local maxima with a cubic spline. The procedure is repeated for the local minima, forming a lower envelope, $e_{min}(t)$. Figure 84 a) shows $e_{max}(t)$ (red) and $e_{min}(t)$ (green) for part of the synthetic DN lightcurve discussed in section 4; the data is shown as a thin black line.

2. The mean $m_1(t)$ of the upper and lower envelopes is calculated:

$$m_1(t) = \frac{(e_{max}(t) - e_{min}(t))}{2}$$

In figure 84 a), $m_1(t)$ is indicated by a thick black line.

3. By subtracting $m_1(t)$ from $x(t)$ the first IMF, $h_1(t)$, is calculated:

$$h_1(t) = x(t) - m_1(t)$$

Figure 84 b) shows $h_1(t)$ for the synthetic DN data. $h_1(t)$ contains the finest scale or shortest period component of the data; in the case of the synthetic DN lightcurve, it consists of the 25 - 30 s DNO component.

4. The process is repeated, using $m_1(t)$ as the new data, until the extracted mean is monotonic. Figures 84 c), e) and g) show successive iterations of the the sifting processing, using the mean calculated in the previous iteration as the input for the next iteration. Figure 84 d), f) and h) show the extracted IMFs: they correspond to the 70 s lpDNO, and the 375 - 450 s and 750 - 900 s QPOs respectively.

The time series is thus decomposed into n IMFs and a residual, $m_n(t)$ which is either the mean trend or a constant:

$$x(t) = \sum_{j=1}^n h_j(t) + m_n(t) \quad (43)$$

6.2.1 Implementation

In order to find the extrema necessary to fit the splines, the second derivative of the data is needed, which requires second differencing of the data. This decreases the Nyquist frequency and the resolution at which the data can be analysed. To compensate for this, we interpolate between points so that after second differencing, the number of data points is the same as in the

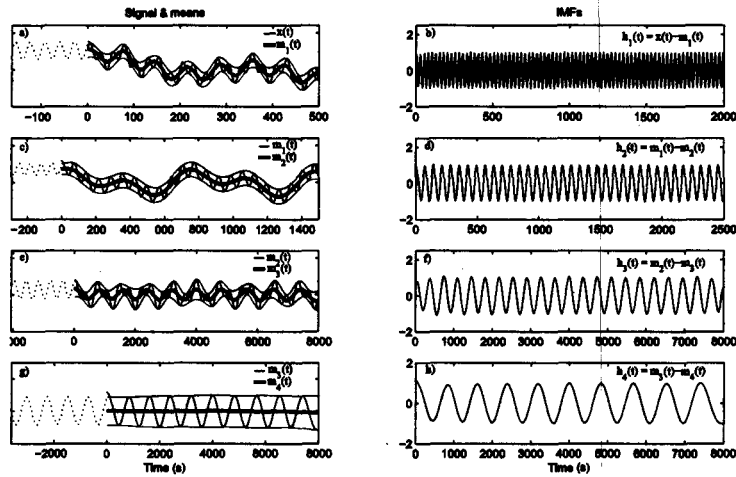


Figure 84: EMD of synthetic DN light curve. Thin black lines indicate the data, thick black lines indicate the fitted mean. Red lines indicate the fitted maximum enveloped, green lines indicate the fitted minimum envelope. Dotted lines indicate the oscillations added to constrain the spline curves at the beginning and end of the data.

original data set.

Unless the spline ends are correctly constrained, they have a tendency to propagate unwanted oscillations into the envelopes, affecting the calculated mean, and hence the decomposition. To constrain the spline ends accurately, we have added ~ 5 oscillations at each end of the time series, constructed to bring the splines gradually to zero. The oscillations mimic the behaviour of the data near the end, ensuring that the splines do not oscillate wildly. In Figure 84, the oscillations added at the beginning of each iteration are indicated by dotted black lines.

Ideally, $h_1(t)$ should be an IMF. However, a gentle slope may be amplified to become a local extremum in changing from rectilinear to curvilinear co-

ordinates, so the sifting process is repeated at each step using h_1 as the input until the extracted signal is an IMF. The usual criterion for stopping is when the number of extrema equals the number of zero crossings (Huang et al., 1999). An additional constraint sometimes used is that the mean of the upper and lower envelopes has to be 'close' to zero, for some threshold. Rilling, Flandrin & Gonçalves (2003) have extended this to include two thresholds: one criterion to ensure globally small fluctuations from the mean, and one to allow locally large deviations from the mean. However, we have not found this necessary for our data.

6.3 Testing

EMD is a signal-dependent adaptive technique which often requires numerical experiments to understand its properties, as theoretical analysis is seldom possible. Rilling, Flandrin & Gonçalves (2003) show that even for a pure tone, EMD may not extract the single mode correctly if the sampling period is insufficient. They also give an empirical relationship for the minimum distance necessary between frequencies for them to be correctly resolved. Flandrin, Rilling & Gonçalves (2004) show that the EMD of noisy signals can act as a dyadic filterbank, similar to wavelet analysis; based on this, Flandrin, Gonçalves & Rilling (2004) discuss the selection of true signal modes from noisy decompositions, and the construction of confidence intervals for the extracted modes. In this section we investigate the EMD of several of the synthetic signals introduced in the previous chapter, and carry out Monte Carlo testing of the amplitude of recovered modes from noisy signals.

6.3.1 Gaps

Figure 85 shows the EMD of a sinusoid of period 100 s, with three gaps of lengths 20 s, 60 s and 100 s respectively, which have been linearly interpolated over. h_1 , the first IMF, contains the signal; the amplitude is correct except near the second gap. The second and third IMFs, h_2 and h_3 , have significant amplitude only near the second gap. The final four IMFs contain very low amplitude long period oscillations; this is typical of the EMD of a finite length time series with non-zero mean. It is interesting to note that the third gap, which is the longest, is accurately represented in the first IMF, and does not affect other IMFs. We conclude that gaps in the data can result in apparent transient oscillations near the gap, and hence EMD of gappy data must be analysed with care.

6.3.2 Frequency Modulation

We wish to check that a frequency modulated signal is decomposed into a single signal IMF, rather than different IMFs covering different frequency bands. Figure 86 shows the EMD of a sinusoid with period changing linearly from 200 s to 400 s; the signal is correctly extracted in the first IMF, h_1 , while the remaining 5 IMFs contain only very low amplitude trend like behaviour.

Similarly, in figure 87, the first IMF contains the signal (a sinusoid with period changing abruptly from 200 s to 300 s.) at time 4000 s, while the remaining 5 IMFs contain low amplitude trend information. The EMD of a sinusoid with period with period given by $600 + 30 \cos(\frac{2\pi t}{2000})$ is shown in figure 88; again, the first IMF contains the signal, while the remaining three




Figure 85: Empirical Mode Decomposition of a sinusoid with period 100 s, with three gaps of lengths 20 s, 60 s and 100 s respectively.




Figure 86: Empirical Mode Decomposition of a sinusoid with period changing linearly from 200 s to 400 s.

IMFs contain low amplitude trend information.

6.3.3 Phase Jumps

Figure 89 shows the EMD of a sinusoid with period 100 s, with a phase jump of 160 s at time 4000 s. The first IMF contains the signal, while the remaining 4 IMFs contain trend information, and low amplitude oscillations near the time of the phase jump.




Figure 87: Empirical Mode Decomposition of a sinusoid with period changing abruptly from 200 s to 300 s at time 4000 s.




Figure 88: Empirical Mode Decomposition of a sinusoid with period with period $600 + 30 \cos(\frac{2\pi t}{2000})$.




Figure 89: Empirical Mode Decomposition of a 100 s sinusoid with a phase jump of 160 s at time 4000 s.

In figure 90, we show the EMD of a 100 s sinusoid with 10 random phase jumps. While the first IMF captures the majority of the signal, near the smallest phase jump, at time 7500 s, the signal is split over the first and second IMFs.




Figure 90: Empirical Mode Decomposition of a 100 s sinusoid with 10 random phase jumps.




Figure 91: Empirical Mode Decomposition of a 400 s sinusoid with amplitude $1 + 0.8 \sin(\frac{2\pi t}{2000})$.

6.3.4 Amplitude Modulation

Figure 91 shows the EMD of a 400 s sinusoid with amplitude modulated by $1 + 0.8 \sin(\frac{2\pi t}{2000})$. The first IMF contains the signal, with the correct amplitude; the remaining 5 IMFs contain trend information. It can be seen that even when the splines are constrained using characteristic oscillations added at the ends, the extracted mean can show fluctuations; this is particularly acute near the ends of the second and fourth IMFs.

6.3.5 Multiple Components

Figure 92 shows the full EMD of the synthetic DN signal, discussed in section 6.2. For a signal such as this, with no noise and components well spaced in frequency, EMD works very well: the first four IMFs contain the DNO, the lpDNO and the two QPOs respectively. However, for signals containing two components that are close in frequency, EMD does not always correctly decompose them into meaningful components. Figure 93 shows the EMD of a signal consisting of two sinusoids, with periods 400 s and 600 s respectively. The first IMF contains an amplitude modulated version of the true signal, while the second and third IMFs consist of low amplitude versions of the 400 s and 600 s components respectively.

6.3.6 Noisy Signals

So far we have considered only the decomposition of noise-free signals. However, since all our data are noisy, we need to gauge the effect of noise on the accuracy of the decomposition. To estimate the error in the EMD of their signal, Komm, Hill & Howe (2001) construct realizations of their time series which they use for Monte Carlo simulations. Each point in each realization is drawn from a distribution with mean given by the value of the true data point, and standard deviation given by the measurement error. They find that the realization and the original IMF components are in good agreement.

We use a different Monte Carlo simulation to estimate the effect of noise on the EMD; instead of using experimental data, we create 10 000 sinusoids




Figure 92: Empirical Mode Decomposition of a synthetic DN lightcurve.




Figure 93: Empirical Mode Decomposition of a signal consisting of two sinusoids, with periods 400 s and 600 s.




Figure 94: Average IMF (thick black line) and signal (red line) for the Monte Carlo simulation with noise of amplitude 0.3 added to the 400 s sinusoid. The standard deviation of each point is shown as a dotted line above and below the average curve.

of period 400 s, amplitude 1 and length 1000 s. To each sinusoid we add a realization of the red noise process discussed in section 5.3.7, and decompose the resulting noisy signal into IMFs. In order to find the IMF that contains the signal we find the highest peak in the periodogram of each IMF, and pick the IMF with the peak nearest the (known) signal frequency. We then average all the (signal containing) IMFs, to find an average IMF, and for each point on the average IMF, we find the standard deviation. We then compare the original signal with the average (signal containing) IMF. Figure 94 shows the (noise-free) signal (red line), average IMF (thick black line) and standard deviation (dotted lines) for Monte Carlo simulations with noise of amplitude 0.3. We repeated this process for different levels of added noise.

Table 3 summarizes our results: the first column gives the amplitude of the noise added, the second column gives the amplitude of the resulting average (signal containing) IMF and the third column gives the mean standard de-

Table 3: Average standard deviation for Monte Carlo simulations

Noise Amplitude	Average IMF Amplitude	Average std
0.1	0.96	0.14
0.2	0.91	0.19
0.3	0.87	0.26
0.5	0.82	0.27
0.8	0.75	0.30
1.0	0.72	0.34
1.2	0.70	0.36

viation of the average IMF. It is clear that as the noise level increases, the average IMF increasingly underestimates signal amplitude; for signals with a high noise level, we can expect that the IMF will give an underestimate of the signal amplitude.

6.4 Uses for EMD

While Huang et al. (1999) originally developed EMD as a means of decomposing signals prior to Hilbert analysis, we use the IMF's resulting from EMD directly, either for prewhitening or amplitude analysis. Here we give a description of its use for prewhitening; several detailed examples of amplitude analysis with EMD are given in chapter 7.

6.4.1 Prewhitening using EMD

As mentioned in section 4.2, it is often necessary to prewhiten astronomical data to ensure that they are stationary prior to further analysis. EMD provides an efficient way of prewhitening data which do not display a strictly linear or polynomial trend, or which contain a single eclipse that cannot be




Figure 95: RU Peg s1674: Raw and fitted data.

adequately removed using fitted sinusoids.

Figure 96 shows the empirical Mode decomposition of the raw light curve of RU Peg, run s1674. While the top 5 IMFs clearly contain oscillatory information, the bottom 5 IMFs contain the trend information. By summing the latter, a good fit to the trend can be achieved; figure 95 shows the raw data and the fitted trend. Subtracting this trend prewhitens the data, ensuring stationarity of the mean.

EMD can also be used to remove eclipses. Figure 98 shows the EMD of U Gem s1679. Once again, the lower IMFs contain information about the long-term change in the light-curve, this time due to the eclipse. Figure 97 shows the raw data and the fitted eclipse found by summing the six longest period IMFs. The prewhitened data show less residual influence from the eclipse than data prewhitened using fitted sinusoids.




Figure 96: Empirical Mode decomposition of RU Peg s1674.




Figure 97: U Gem s1676: Raw and fitted data.

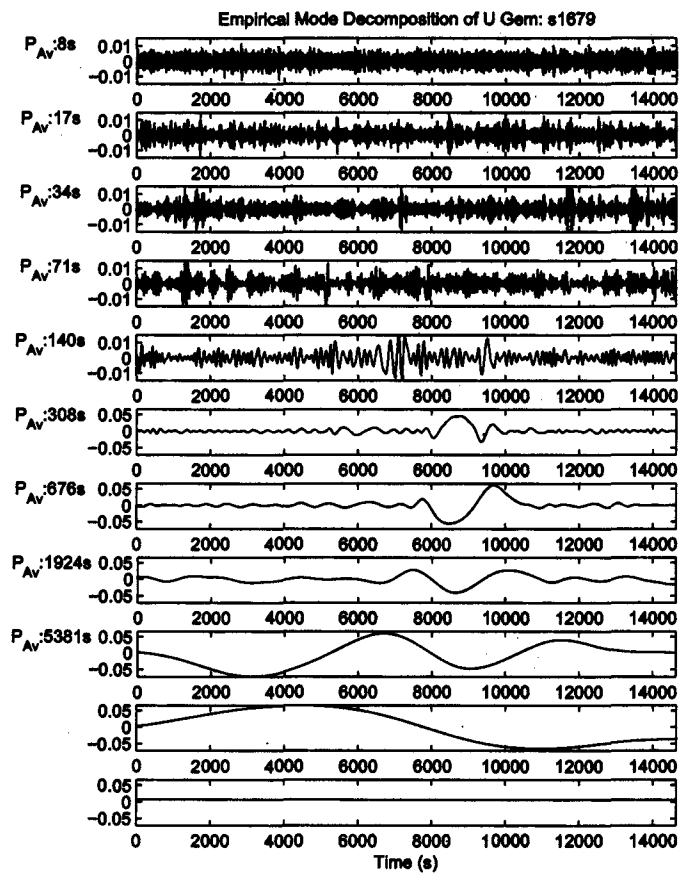


Figure 98: Empirical Mode decomposition of U Gem s1676.

6.5 Discussion

In this chapter we introduced Empirical Mode Decomposition as a method of adaptively decomposing a multi-component time-series into its constituent modes. We investigated the effect of decomposing signals with frequency and amplitude modulations, phase jumps and gaps and carried out Monte Carlo simulations to estimate the accuracy of the amplitude of the extracted signal IMF, for a sinusoidal signal buried in red noise. We found that as the noise level was increased, the amplitude of the extracted IMF was less than the amplitude of the true signal. While this provides an avenue for further research, to find a way to correct for the decreased amplitude, we have chose in the remainder of this work to use the amplitude given by the EMD without correction, baring in mind that the true amplitude may be somewhat larger. We also discussed the use of EMD as a means of prewhitening time series with eclipses and non-linear trends.

7 Application of Time-Frequency Techniques to Four Cataclysmic Variables

In this chapter we apply the wavelet and EMD techniques discussed in the previous two chapters to quiescent and outburst lightcurves from four CVs: RU Peg, U Gem, TY PsA and VW Hyi. Our purpose is threefold. Firstly, since all of these runs have been analysed in detail previously, we see how previously detected DNOs and QPOs appear in the wavelet scalogram, and hence build up a set of criteria for objectively detecting QPOs using the wavelet scalogram. We then use these criteria to see if there are QPOs in the VW Hyi data which have not been identified by the old methods and, finally, we investigate the non-stationary behaviour of all DNOs and QPOs detected, using both the wavelet scalogram and EMD.

This chapter is important for establishing the reliability of the new techniques (by checking that all previously identified DNOs and QPOs are identified by our techniques), and provides an opportunity to compare the information provided by the new techniques with that elicited by standard techniques such as the Fourier Transform and O-C curves.

7.1 RU Peg

RU Peg is a dwarf nova with an orbital period of 8.99 h. The four runs analysed here were taken on consecutive nights in November 1975 when RU Peg was at outburst maximum (Robinson & Nather, 1979). Details of the individual observing runs are given in Table 4; a preliminary report of the observations is given in Patterson et al. (1977) and details of the

Table 4: RU Peg Observation Details

Run Number	HJD Start +2440000	Int time (s)	Run Length (hrs)
s1674	2721.64	4	2.97
s1677	2722.56	4	3.82
s1680	2723.56	4	3.81
s1683	2724.56	4	3.77

QPOs detected are given in Robinson & Nather (1979). The data were obtained with the 2.1 m telescope at McDonald Observatory using the McDonald high-speed photometer (Nather, 1973) and were kindly provided by Ed Robinson. As we did not have data on a comparison star, the data are scaled arbitrarily.

7.1.1 s1674

Figure 99 shows the wavelet spectrum of run s1674 of RU Peg, the first light curve in which QPOs were identified, by Patterson et al. (1977). The ~ 50 s QPOs are clearly visible in the Morlet wavelet spectrum as little islands of high power (within 95% confidence contours) at around 50 s. They are most coherent from 3000 s to about 6000 s, when the DNOs at 12 s begin. Two IMFs of s1674 are shown in figure 100; the IMF with average period 12 s shows the DNOs, while the 52 s IMF shows the QPOs. The 12 s IMF confirms that the DNOs show a huge increase in amplitude from about 0.002 mag to 0.006 mag, starting from about 6000 s. It appears that the QPOs decrease in coherence from the start of the DNOs. The mean amplitude of the QPOs is about 0.004 mag, ranging from 0.002 to 0.01mag. Power spectra of times 0 to 3000 s, 3000 to 6000 s and 6000 to the end confirm that the QPOs, present as peaks at 0.016 Hz (62 s), 0.0175 Hz

(55 s) and 0.019 Hz (53 s) are most coherent from 3000 s to 6000 s, where they show well defined peaks significantly higher than the 99% confidence level, with the 0.016 Hz and 0.019 Hz being the most prominent. The MHAT spectrum is not useful for analysis of oscillations at these short periods.

7.1.2 s1677

The wavelet spectrum of run s1677 is shown in figure 102. Again it appears that when the DNOs are strongest, the QPOs are least coherent. The periodogram of the entire run is shown in figure 103, where the QPOs are seen as a broad hump of power between about 0.015 Hz and 0.022 Hz. Analyzing the run in sections does not resolve the peaks; it appears that the QPOs are changing frequency frequently; this is confirmed by the Morlet spectrum, which shows that the instantaneous frequency is changing constantly. The mean magnitude of these QPOs is about 0.004 mag.

7.1.3 s1680

The Morlet wavelet spectrum of s1680, shown in figure 104, shows several significant bands of power at around 750 s. However, the maximal points for these oscillations in the MHAT spectrum show that the period of these oscillations is changing dramatically, precluding their identification as QPOs. The 50 s QPOs are visible in patches, but are clearly very incoherent, which is supported by the periodogram (figure 105), which again shows a broad hump between 0.016 Hz to 0.022 Hz. The mean amplitude of the QPOs is 0.003 mag, ranging from invisible to 0.005 mag. The DNOs are only visible for a few stretches, and also appear to be less coherent than on previous

nights.

The periodogram in figure 105 shows a number of close, similarly-spaced significant peaks between 0.012 Hz and 0.21 Hz; the top panel of figure 106 shows a close up of this region of the periodogram, which is typical of a source showing stochastic flickering. In the lower three panels of figure 106 we show the periodograms for the first, second and third 1100 s of the run. Notice that, apart from the peaks due to the QPO between 0.017 Hz and 0.2 Hz, the significant peaks in the full periodogram are due to flickering at different times in the run; they do not represent reproducible periodicities. While peaks due to QPOs can recur from night to night in a narrow period range, peaks due to stochastic flickering do not.

7.1.4 s1683

The QPOs in run s1683 are again highly incoherent (figure 107), with a mean amplitude of 0.003 mag. It also appears that, apart from the ~ 50 s QPOs, there are oscillations at about 120 s. This is supported by the periodogram (figure 108), where there is a hump of power between 0.007 Hz and 0.01 Hz. This may be an example of period doubling. Again, there is power at about 750 s, but this is not resolvable into even a quasi-periodic oscillation.

Robinson & Nather (1979) noted that s1677 is the only run which showed clear power at the QPO period in the periodogram. We reproduce figure 2 in Robinson & Nather (1979) in 109; note that here the periodogram is plotted as $\log(\text{power})$ vs frequency. We note that in all four runs, using amplitude as the ordinate (rather than $\log(\text{power})$), with the confidence levels from the fit, the QPOs are detectable in the power spectrum.

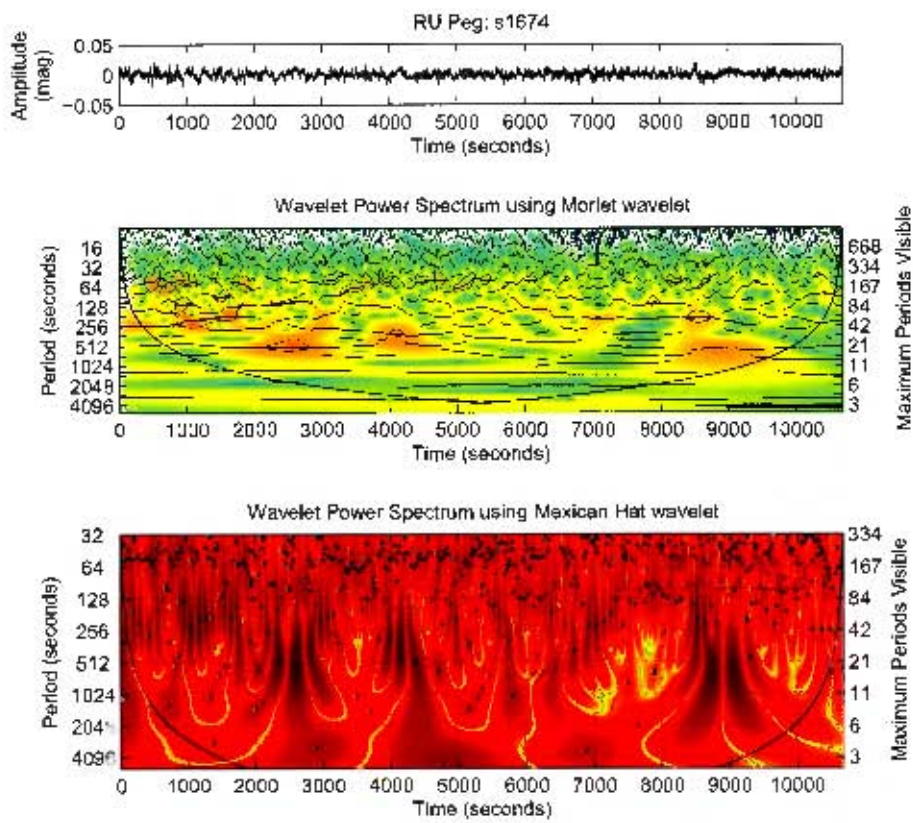


Figure 99: Wavelet analysis of s1674.




Figure 100: 1st and 3rd IMFs of s1674, showing the 12 s DNOs and 52 s QPOs respectively.

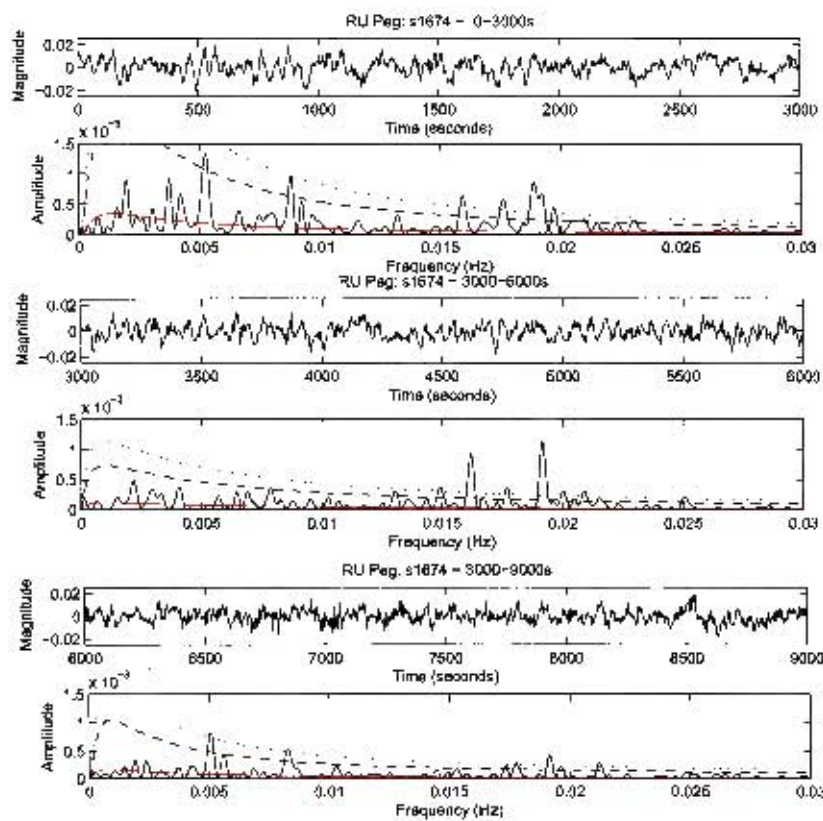


Figure 101: Lightcurve of three sections of RU Peg run s1674, with periodograms below each section. In the periodograms, the red line indicates the estimated periodogram, and the dashed and dotted lines the 95% and 99% confidence levels respectively.

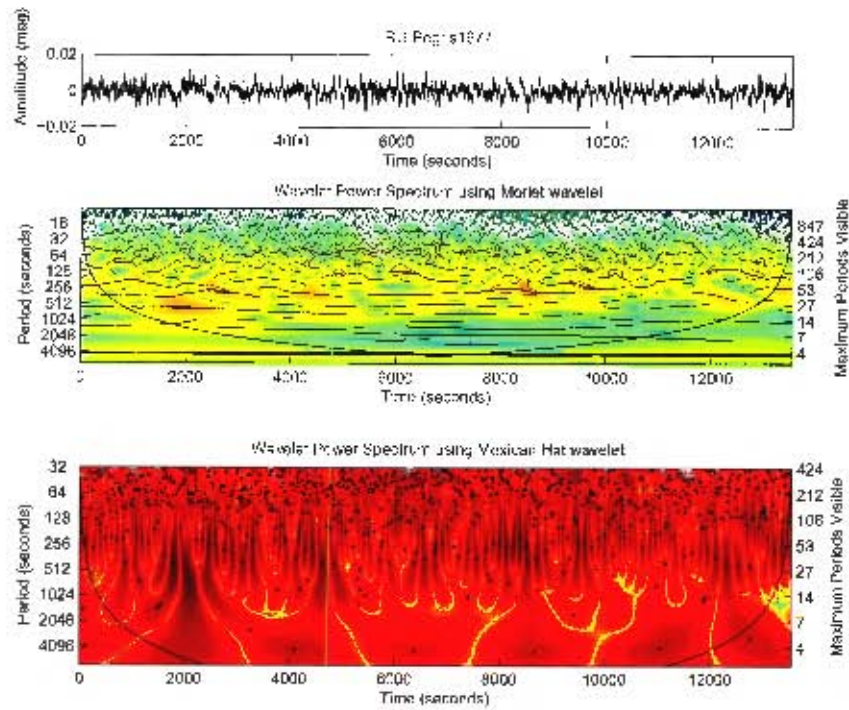


Figure 102: Wavelet analysis of s1677.




Figure 103: Periodogram of RU Peg s1677

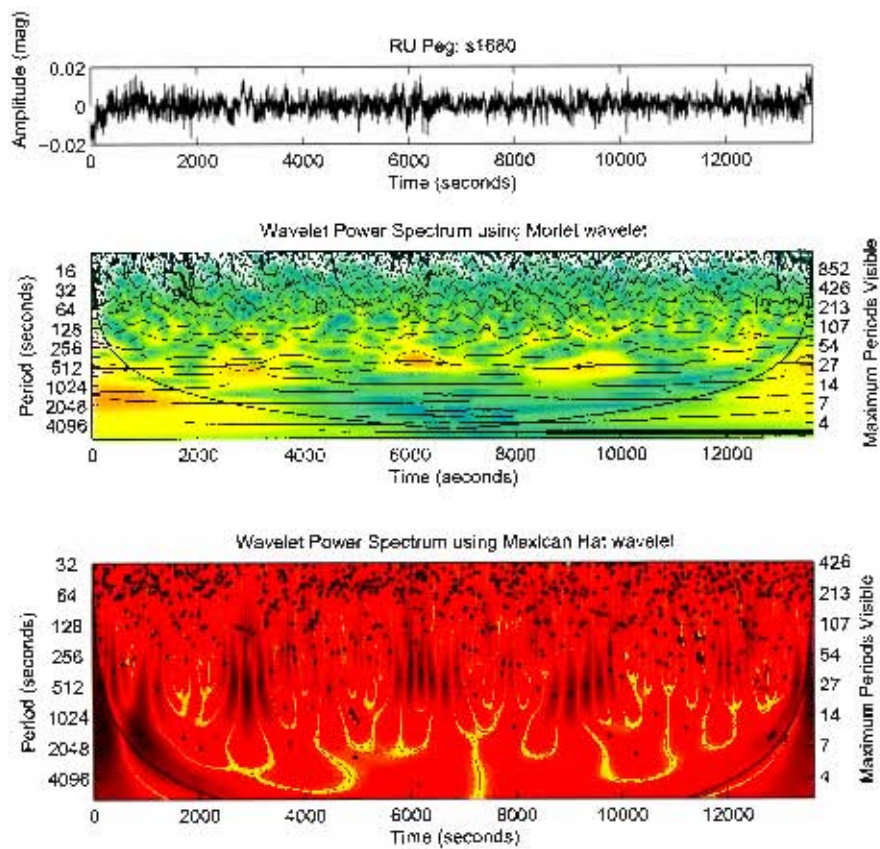


Figure 104: Wavelet analysis of s1680. The red line indicates the estimated periodogram, and the dashed and dotted lines the 95% and 99% confidence levels respectively.




Figure 105: Periodogram of RU Peg s1680

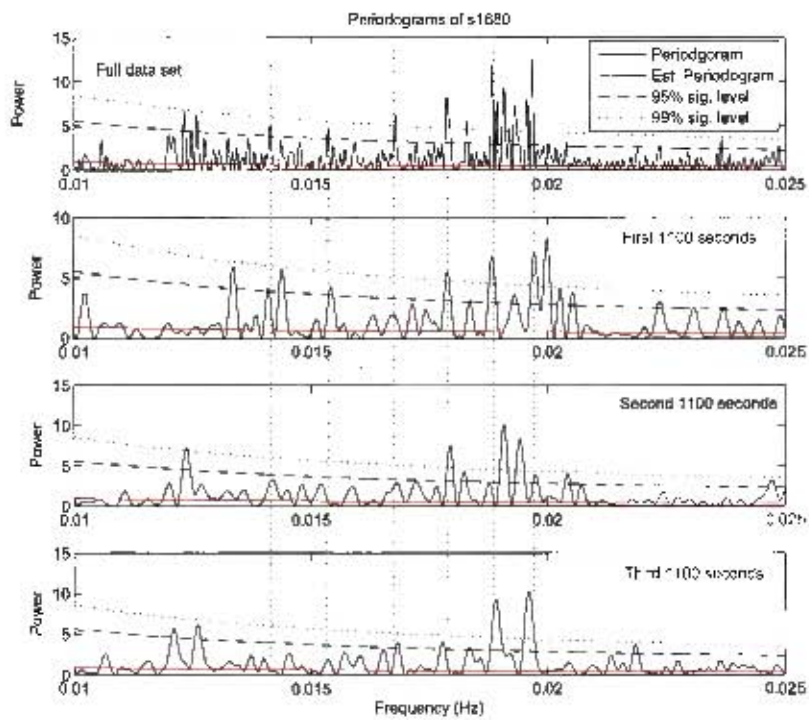


Figure 106: Periodogram of short sections of RU Peg s1680

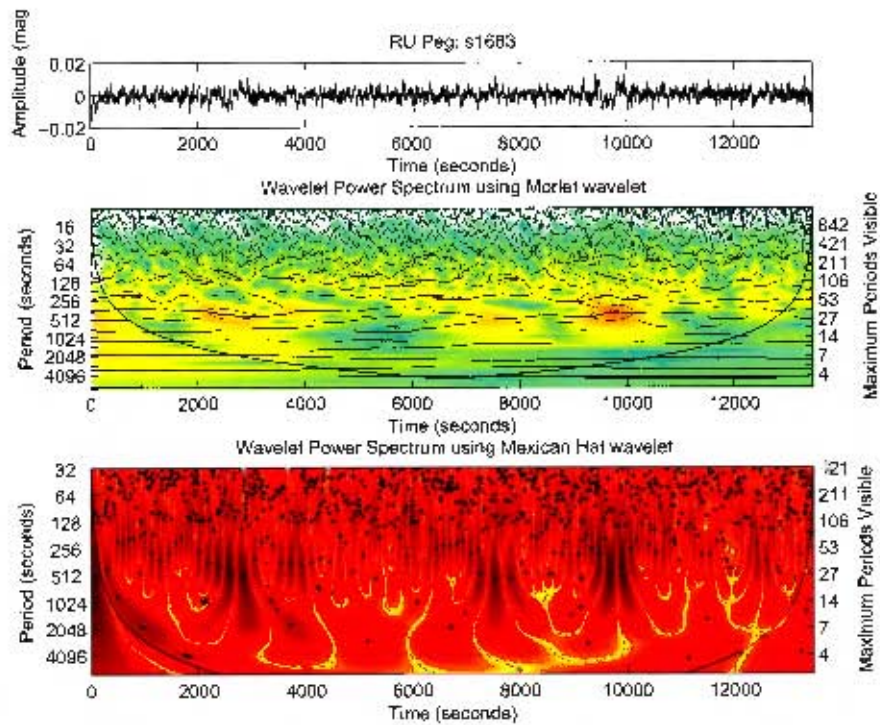


Figure 107: Wavelet analysis of s1683.




Figure 108: Periodogram of RU Peg s1683




Figure 109: Periodogram of RU Peg s1677, from Robinson & Nather (1979)

Table 5: U Gem Observation Details

Run Number	HJD Start +2440000	Int time (s)	Run Length (hrs)
s1676	2721.87	2	3.26
s1679	2722.83	3	4.12
s1682	2723.83	4	4.40
s1685	2724.84	4	2.64
s1691	2726.84	4	4.30

7.2 U Gem

U Gem is a dwarf nova with an orbital period of 4.25 h. The data analysed here were taken on consecutive nights in November 1975, 13-18 d after an outburst that reached a maximum on 1975 October 22, and were first analysed in Robinson & Nather (1979). Details of the individual observing runs are given in Table 5. The data were obtained with the 2.1 m telescope at McDonald Observatory using the McDonald high-speed photometer (Nather, 1973), and were kindly provided by Ed Robinson. As we did not have data on a comparison star, the data are scaled arbitrarily.

7.2.1 s1676

The ~ 73 s second oscillation seen in run s1676 is visible in the wavelet spectrum (figure 110); it appears that over the course of the run the period increases from about 64 s to 90 s, but the period often changes abruptly, as mentioned in Robinson & Nather (1979). The periodogram (figure 111) shows many peaks between 0.012 Hz and 0.016 Hz which fall above the 99% significance level. The 98 s IMF in figure 112 shows the IMF corresponding to the 70 s QPO; the average amplitude is about 0.004 mag, with a maximum of about 0.01 mag. In addition to the 70 s QPO the wavelet spectrum

shows some evidence of power at about 32 s, particularly 5000 s and 8000 s, although with very low coherence, as seen in the 49 s IMF in figure 112. There is also evidence of power of low coherence at about 160 s.

7.2.2 s1679

Figure 113 shows the wavelet spectrum of run s1679. It is similar to that of run s1676, with relatively incoherent oscillations occurring at about 90 s, with a mean amplitude of about 0.005 mag. Again, there is evidence of power at about half the 90 s period, at about 50 s, but this is very incoherent (although there is a significant peak in the power spectrum at this period - see figure 114). Note that while the wavelet spectra are of the pre-whitened data (eclipse removed) we have shown the light curve of the raw data, as we wish to investigate whether or not the oscillations change during the eclipse. The same plotting is used for runs s1682, s1685 and s1691. The 137 s IMF, shown in figure 115, corresponds to the 90 s QPO, and shows that the amplitude of the QPO appears to increase slightly during the eclipse occurring between 8000 s and 9500 s, while the period increases slightly as well. Once again there is some evidence of power of low coherence at about 160s, especially prior to the eclipse when the bright spot is in view.

7.2.3 s1682

The wavelet spectrum of run s1682 is shown in figure 116. There appear to be two time scales - about 80 s and about 160 s - at which power is present, but they show low coherence. The amplitude of the 80 s QPO (as identified by Robinson & Nather (1979)) is about 0.005 mag, increasing to 0.006 mag

just before eclipse. As the bright spot comes into view, from about 9000 s to the start of the eclipse at 13000 s, the oscillations become larger and even less coherent - this can be seen clearly in the 100 s IMF (see figure 118). The periodogram of the complete run, shown in figure 117, has a broad hump between 0.012 Hz and 0.016 Hz, with a prominent peak at 0.014 Hz, corresponding to the 80 s QPO. The power seen in the previous two runs at 160 s is again present, increasing in amplitude when the bright spot is in maximum view.

7.2.4 s1685 and s1691

Figure 119 shows the wavelet spectra of run s1685. While there is clearly power present in the 60 to 180 s range, it is seldom resolvable into a narrower frequency range. The amplitude appears to decrease during eclipse (see figure 121), however, there is another range of decreased amplitude, from 7000 to 8500 s, which does not occur during an eclipse. A similar pattern is seen in s1691. In neither run are QPOs of any coherence detected.

Robinson & Nather (1979) differentiate between the QPOs found in runs s1679 and s1682, which continue to be present during eclipse, and flickering from the bright spot, seen in runs s1685 and s1690, which disappears during eclipse. However, in both of the latter cases, the amplitude of the flickering during eclipse appears to return to the level prior to the bright spot coming in to view. It is therefore possible that the flickering due to the bright spot is occurring on a similar time scale to the relatively coherent QPOs, which are obscured during bright spot maximum.

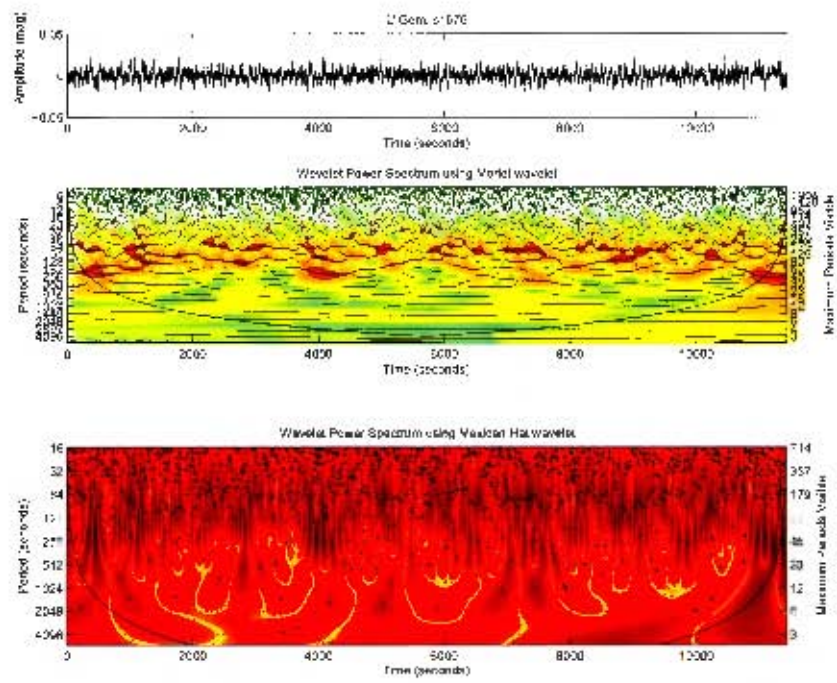


Figure 110: Wavelet analysis of s1676.




Figure 111: Periodogram of U Gem s1676




Figure 112: IMF of U Gem s1676

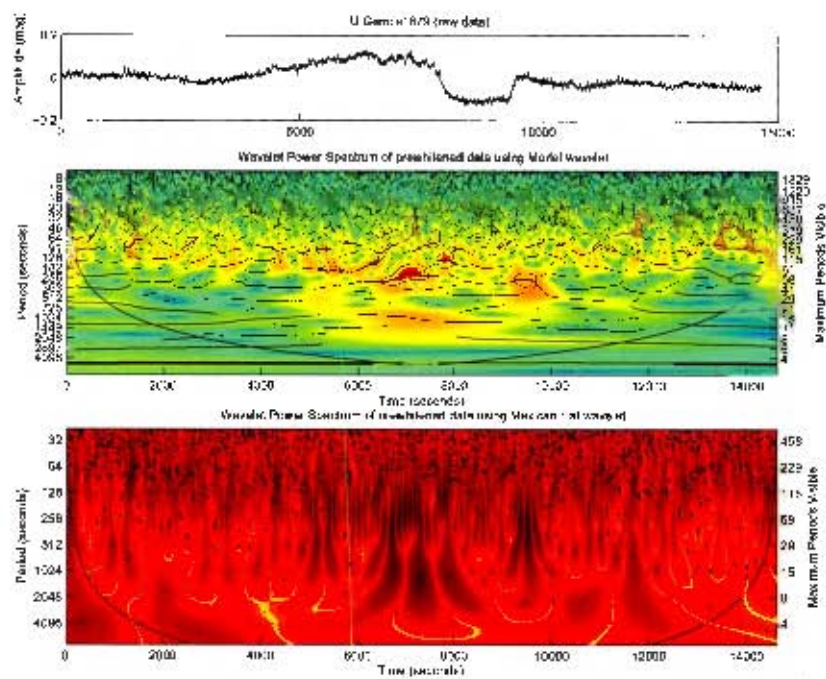


Figure 113: Wavelet analysis of s1679. Note that the wavelet analysis is of the prewhitened data, but the raw data are shown at the top.




Figure 114: Periodogram of U Gem s1679




Figure 115: IMF of U Gem s1679




Figure 116: Wavelet analysis of s1682. Note that the wavelet analysis is of the prewhitened data, but the raw data are shown at the top.




Figure 117: Periodogram of U Gem s1682




Figure 118: IMF of U Gem s1682




Figure 119: Wavelet analysis of s1685. Note that the wavelet analysis is of the prewhitened data, but the raw data are shown at the top.




Figure 120: Periodogram of U Gem s1685




Figure 121: IMF of U Gem s1685




Figure 122: Wavelet analysis of s1691. Note that the wavelet analysis is of the prewhitened data, but the raw data are shown at the top.




Figure 123: Periodogram of U Gem s1691




Figure 124: IMF of U Gem s1691

7.3 TY PsA: run s3412

This light curve was taken in September 1984, using the UCT photometer on the 40" telescope at the Sutherland site of the South African Astronomical Observatory, when the dwarf nova TY PsA was in normal outburst. TY PsA showed DNOs large enough to be seen directly in the lightcurve, which was first analysed by Warner & O'Donoghue (1989), who also found a ~ 245 s QPO. Warner & Pretorius (2008) re-analysed the run, and found a DNO with period roughly 38 s, but with many rapid changes in period and phase. We use wavelet analysis and EMD to attempt to gain details on the period and amplitude changes of the DNO, to test the efficiency of our methods at relatively short periods. Since there was a gap in the data, we analyse them in two separate sections; figures 125 and 126 show the wavelet spectrum of each section. While the first section shows only hints of the possible DNO, it is clearer in the second section, together with the ~ 245 s QPO which is clearly visible, especially in the second half, and a ~ 500 s modulation. It is clear that the DNO has very low coherence, and changes amplitude frequently. The MHAT spectrum shows the extrema of the 245 s QPOs; 5.5 cycles are present.

Figures 127 and 128 show the IMFs capturing the DNOs. The periods shown are the average period of the DNO for the section under consideration; table 6 shows all the DNO periods found.




Figure 125: Wavelet analysis of first part of s3412




Figure 126: Wavelet analysis of second part of s3412.




Figure 127: IMF showing DNOs in first part of s3412.

Table 6: DNO Periods measured in TY PsA run s3412.

Time since start (s)	Period (s)	Amplitude (mag)
First section, starting HJD 7.25149		
50 - 300	38.5	0.016
650 - 850	29.2	0.009
1250 - 1500	27.0	0.01
Second section, starting HJD 7.280101		
100 - 300	41.7	0.01
700 - 950	32.3	0.01
1350 - 1500	34.5	0.006
1700 - 1850	36.9	0.012
2250 - 2500	35.5	0.011
2700 - 2900	38.5	0.012
3300 - 3500	28.4	0.005
3500 - 3700	31.3	0.01
4100 - 4400	45.1	0.007
4500 - 4650	35.7	0.01
4700 - 4800	30.0	0.006




Figure 128: IMF showing DNOs in second part of s3412.

7.4 VW Hydri

VW Hydri is an SU UMa dwarf nova with an outburst time scale of the order of 28 days (Warner, 1995a). Since it is circumpolar at the Sutherland site of the South African Astronomical Observatory, and is relatively bright (9–13 mag), VW Hyi presents an extended season for maximum time-resolution observation. In addition to quiescent observations, many observations have been made selectively towards the end of outburst as the star returns to quiescence specifically to study DNO and QPO behaviour (Warner & Woudt, 2006a), and several groups of runs taken on consecutive nights allow for the study of persistent or evolutionary phenomena across nights. Since VW Hydri exhibits many of the DNO and QPO phenomena observed in dwarf novae, it provides a handy single star from which much about the entire class can be learned.

Our VW Hyi data are from the UCT archive, and have been previously analysed in Woudt & Warner (2002a), Woudt & Warner (2002b), Warner, Woudt & Pretorius (2003), Warner & Woudt (2006a) Pretorius, Warner & Woudt (2006) and Warner & Pretorius (2008). Details of the individual observing runs for outburst and quiescent observations are given in tables 7 and 8 respectively. All observations were taken at the Sutherland site of the South African Astronomical Observatory, using the 20", 40" or 74" telescope as indicated. Data acquisition for observations on the 20" and 40" telescope up to and including run s5248 used the UCT photometer (Warner & Nather, 1971); subsequent runs on this telescope and the 74" were made using the University of Cape Town CCD (O'Donoghue, 1995) operating in frame-transfer mode. Most CCD observations were taken in white light;

observations made with a B filter are marked with an asterisk.

VW Hyi outbursts can be divided into four groups: superoutburst, long normal outburst, medium normal outbursts and short normal outbursts (Woudt & Warner, 2002a). The RASNZ VW Hyi data, spanning 1972-2004, have been used to construct average outburst profiles (figure 129), onto which each outburst observation can be placed (Woudt & Warner, 2002a). Each observation is thus assigned a phase (column 'Start' in tables 7 and 8), relative to the arbitrary $T = 0$ d (column 'T=0' in tables 7 and 8) of the template. We have assigned quiescent observations a phase relative to $T = 0$ d of the next outburst, as we have several quiescent observations occurring a few days before an outburst for which we also have data.

We first analysed each run and noted how the QPOs previously detected in the run appeared in the Morlet and MHAT spectra. We found that all previously detected QPOs were visible in the Morlet spectrum, with power above the 95% confidence level and a single continuous instantaneous frequency line. Using the MHAT spectrum, we counted the number of extrema (given by the maximal points) which occurred during the time span of the confidence contour, and found that while 3 cycles (6 maximal points) were sometimes sufficient to have been identified as a QPO in the periodogram (e.g. run s0030), usually 5 or more cycles were present. We thus defined a QPO as at least 5 cycles (10 maximal points in the MHAT) at power above the 95% confidence level (on the Morlet spectrum), or 6 significant cycles spread over two confidence contours at similar frequency. We only deviated from this definition when there was a clear triplet of QPOs with fundamental, first and second harmonic, and one of these contained less than 5 cycles, in which case it was identified as a QPO (e.g. run s1307). We used




Figure 129: Averaged outburst profiles for the four different types of outburst seen in VW Hyi. Solid line indicates superoutburst, dashed line indicates long normal outburst, dash-dot indicates medium normal and dotted indicates short normal outburst.

the instantaneous frequency from the Morlet spectrum to find the (average) period of the detected QPO.

While it would have been possible to combine the confidence contours and instantaneous frequencies from the Morlet spectrum with the maximal points from the MHAT spectrum into a single spectrum, we chose simply to work with both spectra. While the MHAT wavelet was less useful in our analysis of RU Peg, U Gem, and TY PsA, in our analysis of VW Hyi, as outlined above, it played a crucial role.

Table 9 summarizes the details of each oscillation detected in our analysis; we include the run number, the type of outburst, the start and end times of each oscillation (in seconds from the beginning of the run), the period of the oscillation at start and finish (in seconds), and the average amplitude (in magnitudes) as measured from the IMF. Oscillations showing non-monotonic period changes were analysed in monotonic sections. For oscillations with periods >100 s the number of cycles observed is included. DNOs are classified as fundamental (DF), first(D1) or second harmonic(D2) or lpDNOs (DL). DNOs with periods between 70 and 120 s may be either fundamental period DNOs, or lpDNOs; we have decided to categorize all DNOs in this period range as lpDNOs. DNO-related QPOs are classified as fundamental (QF), first(Q1) or second harmonic (Q2) or subharmonic(QS). QPOs that do not appear to be DNO-related are classified as QU, and QPOs from quiescent light curves are classified as QQ. The second last column, 'Prev?', indicates whether the oscillation had been detected in previous analyses ('y') or not ('n'), and the final column includes comments.

In this chapter we discuss in detail the wavelet and EMD analysis of runs

and groups of runs of particular interest; in the following chapter we use the results from our analysis to investigate the nature of DNOs and QPOs in VW Hyi.

7.4.1 Runs s0018, s0019, s0026 and s0030 (mid-fall to return to quiescence)

These four runs were taken on consecutive nights, and span a long normal outburst from $T = -0.23$ to $T = 3$ d.

In addition to the 20 s DNO and 75 s lpDNO seen by Woudt & Warner (2002a), Warner, Woudt & Pretorius (2003) and Warner & Woudt (2006a) in s0018, we find, in the Morlet spectrum (Figure 130) shows significant power at 346 s near the beginning of the run, which is resolved as 4 cycles by the MHAT spectrum. Later in the run three cycles at 286 s have significant power. Unfortunately the oscillations are divided between two IMFs (see figure 131), making further analysis difficult. This is the earliest observation of DNO-related QPOs in normal outburst. There is also evidence of power at about 128 s, but this is not resolvable into a coherent oscillation.

The 75 s lpDNO is difficult to see in the Morlet spectrum, and the IMF shows that the amplitude is changing rapidly. The 0.009 mag 20 s DNO in s0018 is only clearly visible in the Morlet spectrum for about a thousand seconds, in contrast to the 0.02 mag 30 s DNO in s0019 (see figure 132), which is clearly visible throughout; both DNOs, however, are clearly visible in the periodogram. The 514 s QPOs of s0019 (Warner & Woudt, 2006a), when viewed with the Morlet and MHAT spectra, resolves into three separate runs of between 3 and 7 cycles, with periods of 487 s, 400 s and 510 s respectively,

Table 7: VW Hyi: Outburst observations from the UCT archive. Runs marked with an asterisk were taken using a B filter.

Run	Telescope	Length (hrs)	Int time (s)	Outburst Type	HJD Start +2440000	$T = 0$ (d)	T Start (d)
s0018	20	2.04	2	L	1572.52	1572.75	-0.23
s0019	20	4.30	5	L	1573.44	1572.75	0.69
s0026	20	3.32	5	L	1574.50	1572.75	1.75
s0030	20	6.92	5	L	1575.47	1572.75	2.72
s0110	40	2.47	2	Super	1662.30	1677.20	-14.90
s0111	40	0.97	2	Super	1662.41	1677.20	-14.79
s0112	40	2.31	2	Super	1662.48	1677.20	-14.72
s0122	40	2.56	5	Super	1675.33	1677.20	-1.87
s0124	40	1.84	5	Super	1676.37	1677.20	-0.83
s0127	40	3.77	4	Super	1677.29	1677.20	0.09
s0129	40	1.86	5	S	1691.30	1690.20	1.10
s0480	40	1.73	5	Super	2017.38	2022.65	-5.27
s0484	40	3.91	4	Super	2023.31	2022.65	0.66
s1277	40	1.78	4	M	2352.39	2354.20	-1.81
s1307	40	1.33	4	M	2354.39	2354.20	0.19
s1322	40	4.91	4	M	2355.39	2354.20	1.19
s1571	40	1.99	4	Super	2402.32	2403.40	-1.08
s1616	40	2.27	3	Super	2404.31	2403.40	0.91
s2241	40	2.63	5	Super	2758.41	2770.00	-11.59
s2243	40	1.65	5	Super	2759.40	2770.00	-10.60
s2623	40	3.60	5	M	3515.30	3514.60	0.70
s2911	40	3.14	4	L	4934.43	4937.10	-2.67
s2915	40	0.73	4	L	4937.28	4937.10	0.18
s2917	40	2.18	4	L	4939.49	4937.10	2.39
s3078	40	1.95	4	Super	5323.34	5328.50	-5.16
s3416	40	1.19	2	M	5967.55	5967.00	0.55
s5248	40	4.66	5	M	8202.36	8201.75	0.61
s6138_I	40	2.09	4	M	11898.28	11897.85	0.43
s6138_II	40	3.90	5	M	11898.43	11897.85	0.58
s6184	40	1.68	5	L	11957.28	11957.25	0.03
s6316	40	5.72	5	L	12354.24	12353.95	0.29
s6528_I	74	5.81	4	L	12520.22	12519.65	0.57
s6528_II	74	4.18	5	L	12520.47	12519.65	0.82
s7222_I*	74	3.61	4	Super	13007.31	13006.40	0.91
s7222_II*	74	3.03	4	Super	13007.46	13006.40	1.06
s7301	40	2.94	5	M	13087.24	13086.10	1.14
s7311*	74	5.41	4	M	13139.19	13138.95	0.24
s7342	40	3.07	5	M	13155.19	13154.90	0.29
s7621	40	4.03	6	L	13463.23	13463.10	0.13

Table 8: VW Hyi: Quiescent observations from the UCT archive

Run	Telescope	Length (hr)	Int time (s)	HJD Start +2440000	$T = 0$ (d)	T Start (d)	HJD Next +2440000
s0073	40	2.83	5	1648.25	1613.00	-28.95	1677.2
s0077	40	4.11	5	1602.44	1580.00	-10.56	1613.00
s0085	40	1.67	5	1604.44	1580.00	-8.56	1613.00
s0093	40	3.33	5	1605.45	1580.00	-7.54	1613.00
s0102	40	1.31	5	1657.33	1613.00	-19.87	1677.2
s0105	40	1.80	2	1660.29	1613.00	-16.91	1677.2
s1414	40	2.89	4	2384.30	2370.00	-19.10	2403.4

and amplitudes of 0.05 mag.

Run s0026 (figure 133) shows a 0.04 mag 1042 s QPO persisting throughout the run, and s0030 (figure 135) shows a 1124 s QPO of similar magnitude, but a comparison of their IMFs (figures 134 and 136) shows that the QPO of run s0030 is far less coherent than the previous night's. These two QPOs, which are not related to a DNO, are discussed in further detail in chapter 8.

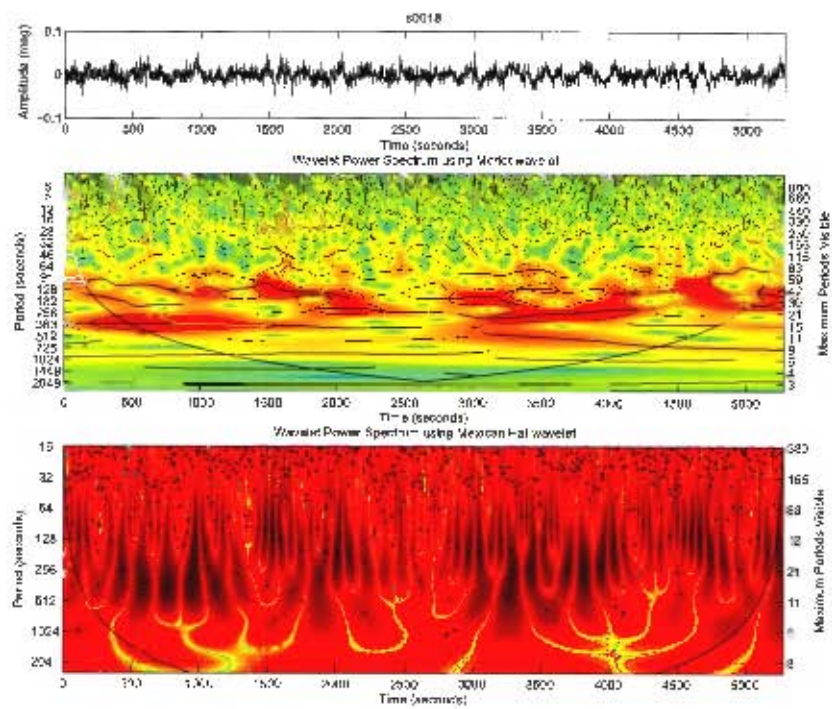


Figure 130: Wavelet analysis of s0018.




Figure 131: Fifth ($P_{av} = 118s$) and sixth ($P_{av} = 118s$) IMFs of s0018.

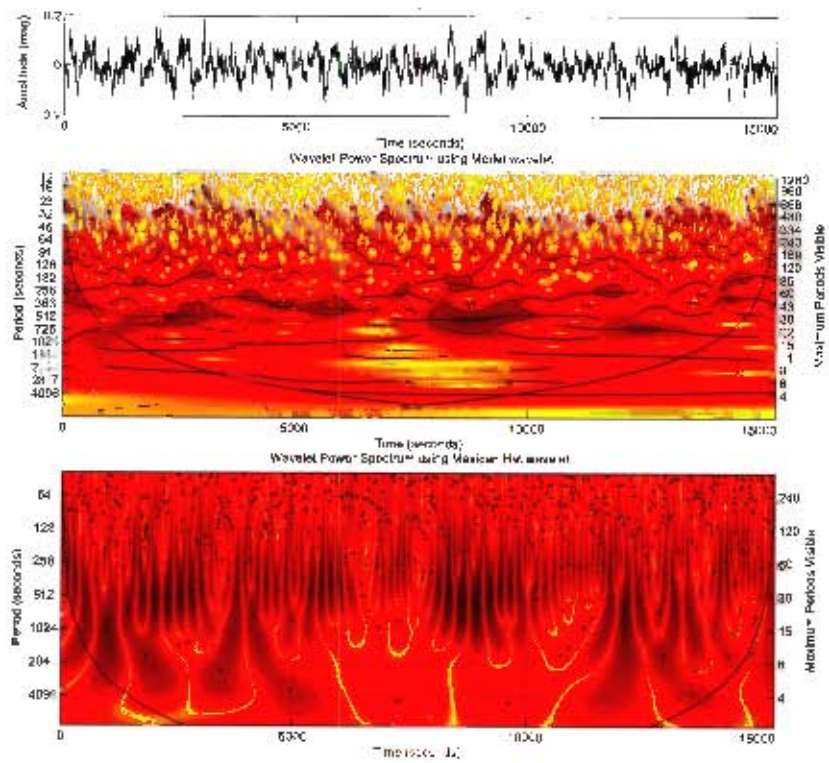


Figure 132: Wavelet analysis of s0019.




Figure 133: Wavelet analysis of s0026.




Figure 134: Periodogram and IMF analysis of s0026, 1170s unclassified QPO.




Figure 135: Wavelet analysis of s0030.

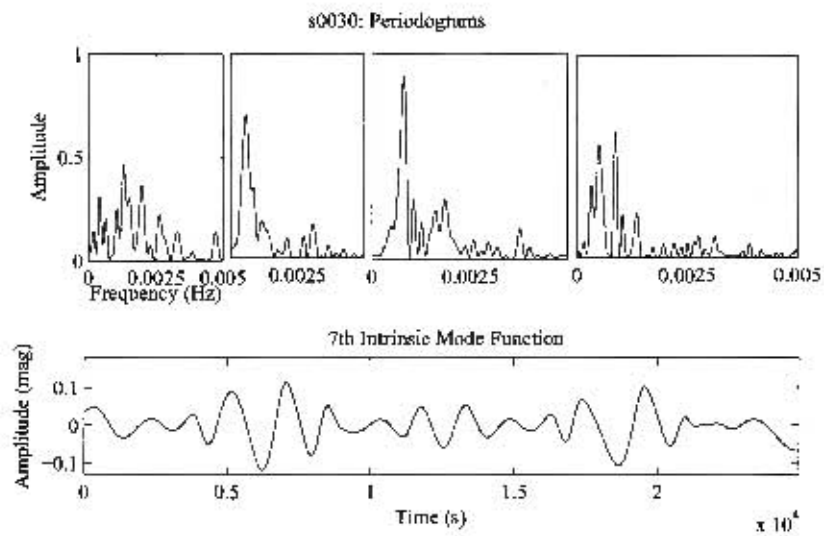


Figure 136: Periodogram and IMF analysis of s0030 1124s unclassified QPO.




Figure 137: Wavelet analysis of s0093.

7.4.2 Runs s0085 and s0093 (Quiescence)

Woudt & Warner (2002a) report a 719 s QPO in run s0085, which has an amplitude of 0.05 mag in our analysis. The following night, in run s0093, the Morlet spectrum shows runs of a 675 s oscillation, each lasting for 5 cycles (figure 137, MHAT spectrum), with amplitude 0.028 and 0.031 mag respectively (figure 138).




Figure 138- 6th IMF of s0093.

7.4.3 Runs s0102, s0105, s0110, s0111, s0112, s0122, s0124 and s0127 (a few days before and spanning superoutburst)

Runs s0102 and s0105 were taken during quiescence, some four to five days before the start of a superoutburst. In both, a strong oscillation with first harmonic is seen; in s0102 at 833 s with 1st harmonic at 411 s and in s0105 at 530 s with first harmonic at 260 s. The amplitude of the fundamental is 0.029 mag in s0102 and 0.026 mag in s0105, while the first harmonics have amplitudes of 0.035 mag and 0.037 mag respectively. In both cases, the oscillations are present for part of the run, and completely absent for part; the nature of the data changes quite significantly (see Figure 139). The usefulness of the time definition of the MHAT wavelet is apparent: the durations of the oscillations can be seen clearly, allowing FTs of only the relevant parts to be calculated.

Runs s0110, s0111 and s0112 were taken on the same night during the rise to supermaximum ($T = -14.90$ d), and essentially consist of a single long run with short gaps; the only oscillation is a 300 s oscillation visible in the first half of s0111, with amplitude 0.007 mag (figure 140).

Runs s0122, s0124 and s0127 are taken on consecutive nights during decline, from $T = -0.83$ to 0.09 d. In addition to the 410 s QPO seen in Woudt & Warner (2002a) in s0122 we also see a 909 s oscillation persisting throughout most of the run, especially prominent in the last half (figures 141 and 142). s0124 is oscillation-free, but s0127 exhibits a plethora of phenomena. The Morlet spectrum in figure 143 clearly shows the 27 s DNOs, increasing in period over the duration of the run to 34 s. There are also suggestions of the 15 s first harmonic for the first 3000 s, and again near the end of




Figure 139: Wavelet analysis of s0105.

the run. The 450 s QPO found by Warner & Woudt (2006a) is clearly visible in three groups with periods of 454 s, 498 s and 487 s and amplitudes 0.02 mag, 0.02 mag and 0.037 mag respectively (figure 144). In addition there is evidence of a 1250 s oscillation, changing to 1075 s after 8000 seconds, which we categorize as a subharmonic of the 480 s oscillation.

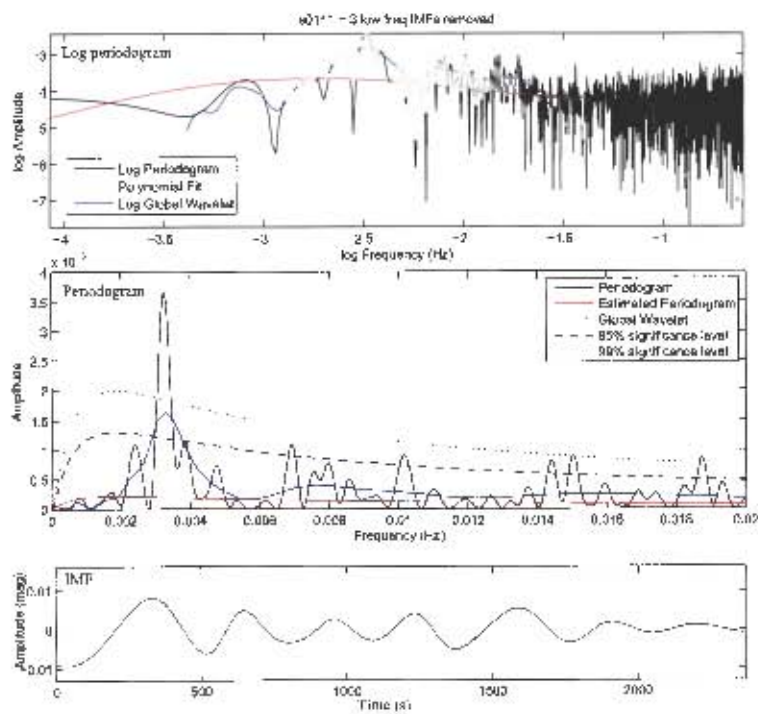


Figure 140: Periodogram and IMF analysis of the 304s uncharacterized QPO in s0111.




Figure 141: Wavelet analysis of s0122.

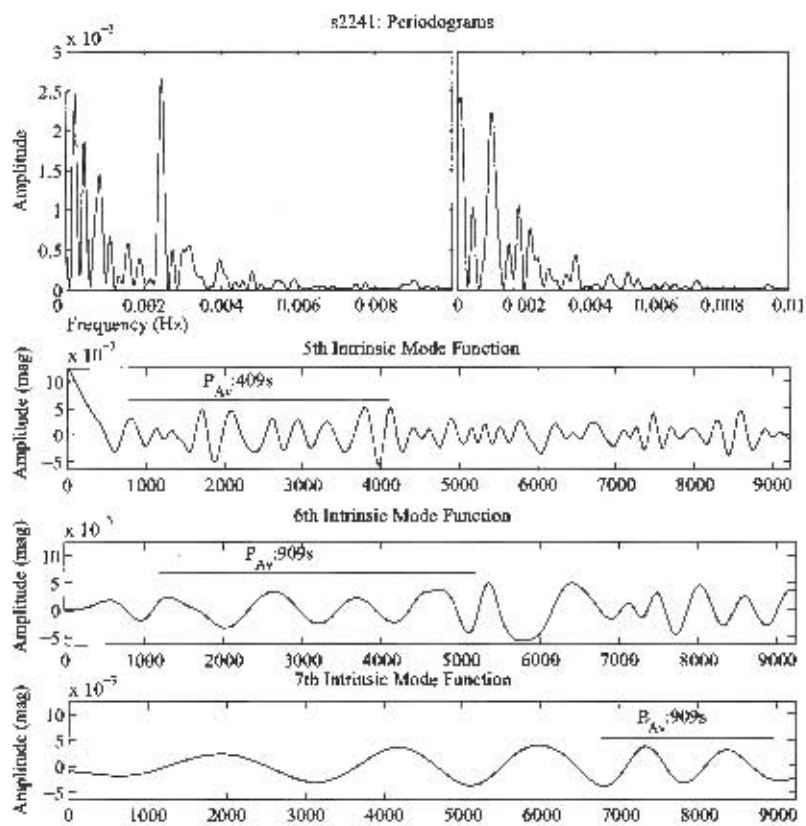


Figure 142: Periodogram and IMF analysis of s0122.




Figure 143: Wavelet analysis of s0127.




Figure 144: 5th and 6th IMFs of s0127.

7.4.4 Runs s1277, s1307 and s1322 (beginning, middle and end of medium normal outburst)

Runs s1277, s1307 and s1322 were taken at $T = -1.81, 0.19$ and 1.19 d respectively, spanning a medium normal outburst. s1277 shows no obvious oscillations.

In addition to the 18 s DNO found by Woudt & Warner (2002a) in s1307, there is some evidence in the Morlet spectrum of power at about 32 s (which would be the fundamental frequency of the DNO) - see figure 145. The fundamental (500 s), first (250 s) and second harmonic (145 s) QPO frequencies are also visible at various times (figure 146). The Morlet spectrum of run s1322 (figure 147) shows the 0.01 mag 28-32 s DNOs and 0.02-0.05 mag 420-476 s QPOs mentioned in Warner & Woudt (2006a), as well as possible QPO subharmonics at 1052 s with amplitude 0.02 mag (figure 148).




Figure 145: Wavelet analysis of s1307.




Figure 146: 5th and 6th IMFs of s1307




Figure 147: Wavelet analysis of s1322.




Figure 148: 6th and 7th IMFs of s1322.

7.4.5 Runs s1571 and s1616 (beginning and end of decline after superoutburst)

Runs s1571 and s1616 were taken two days apart, at the beginning and end of return to quiescence after a superoutburst, at $T = -1.081$ and 0.912 d respectively. Woudt & Warner (2002a) found a possible 1151 s QPO in s1571, which is clearly visible in the wavelet (figure 149) and two IMFs (figure 150), increasing in amplitude from 0.005 to 0.011 mag. A 694 s, 0.03 mag QPO appears twice in s1616, as well as a 1282 s 0.05 mag QPO for the first half, and 7 cycles of 333 s QPO at 0.02 mag, giving a triplet of QPO fundamental, first and second harmonic. There is also evidence of an 87 s 0.01 mag oscillation near the end of the run; this may be either an lpDNO or the DNO fundamental.

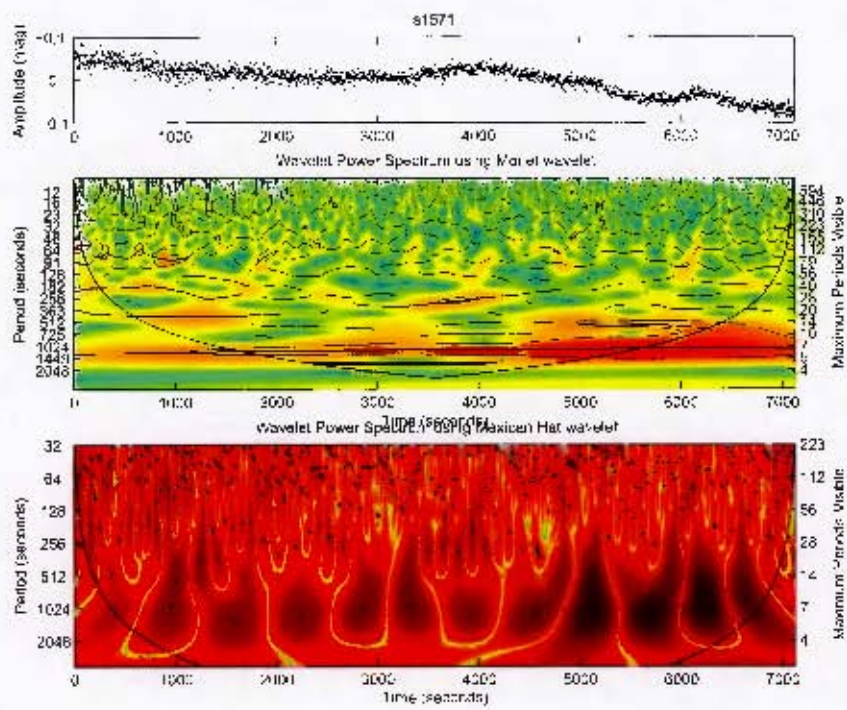


Figure 149: Wavelet analysis of `s1571`.




Figure 150: EMD analysis of s1571.

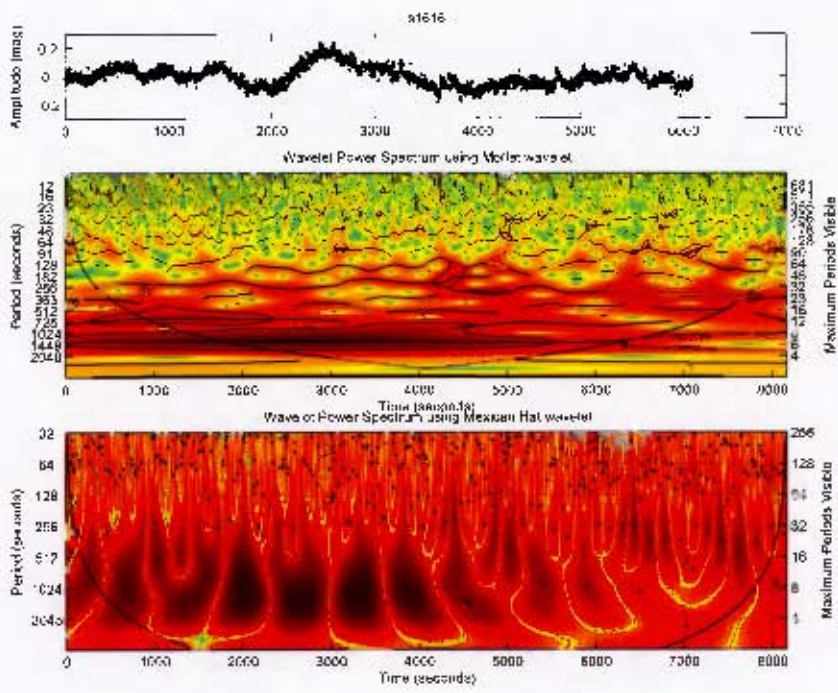


Figure 151: Wavelet analysis of s1616.

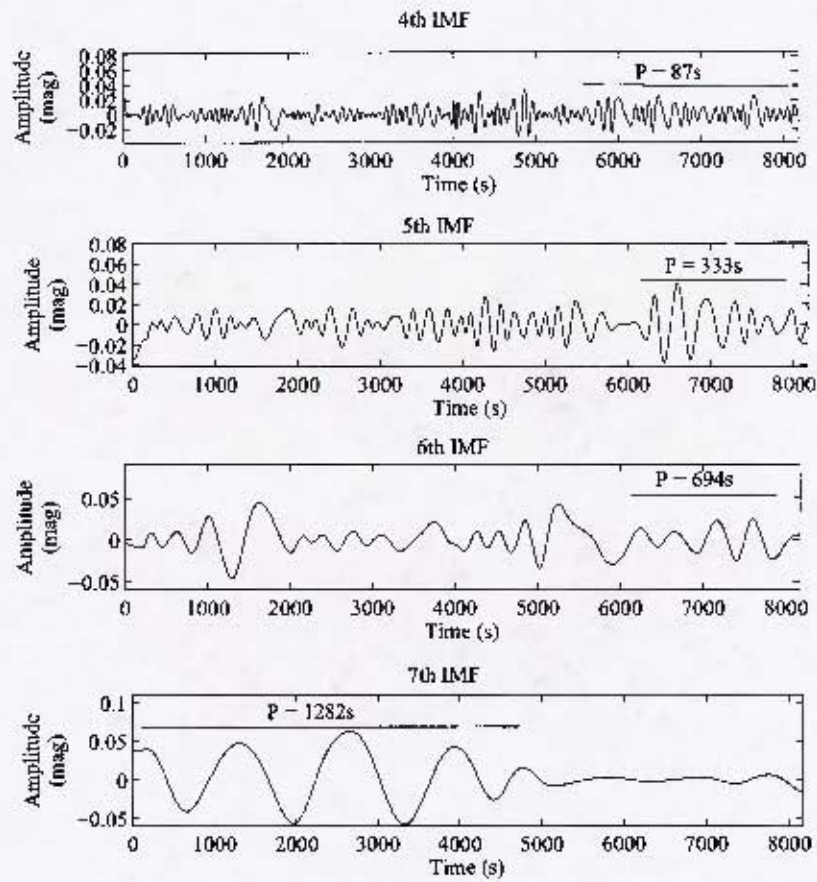


Figure 152: 4th, 5th, 6th and 7th IMFs of s1616.

7.4.6 s0129

s0129 is a good example of the wavelet spectrum's use in detecting oscillations of low coherence. In addition to the DNO (either DF or lpDNO) at 93 s identified by Woudt & Warner (2002a), there is evidence of power at about 33 s throughout the Morlet spectrum (see figures 153 and 154) as well as three runs of QPOs with periods 578 s, 658 s and 714 s, and amplitudes of 0.02 mag. The most coherent run lasts for only 5 cycles (MHAT spectrum, figure 153). Neither the 33 s nor the QPOs are visible in the power spectrum.

7.4.7 s2241

Run s2241 was taken during superoutburst, and shows unclassified QPOs (Woudt & Warner, 2002a) with a period of about 700 s that are only apparent during superhumps (see figure 155). In figure 156, we show the raw lightcurve of s2241, showing the superhumps, with the duration of the QPOs marked. The MHAT spectrum enables us to see clearly where the oscillations begin and end, and for how many cycles they last. There is some evidence for a first harmonic at about 350 s, but this may be due to the non-sinusoidal profile of the QPOs.

7.4.8 s2623

Figure 157 shows the wavelet spectra of run s2623. The ~ 24 s DNO can be seen clearly throughout in the Morlet spectrum, with discontinuities as well as slow changes in period. The QPO fundamental, which decreases in period

from 1000 to 833 s, tracks the decrease in DNO period change from 25 to 23 s, (from time 4000 to 9000 seconds) where the DNO has a discontinuity, jumping to 27 s, and the QPO fundamental disappears completely. The 400 s QPO first harmonic is present for several cycles before 9000 s, but also disappears abruptly when the DNO period jumps. The 250 s QPO second harmonic, which is present throughout, does not appear to track the DNO.

7.4.9 s6138

s6138 was analysed in two sections (figures 158 and 159), as the run contained a large gap and change in integration time. In both parts, in addition to the 25 s DNO first harmonic and ~ 350 s QPO first harmonic found by Warner & Woudt (2006a), we see evidence in both the Morlet and MHAT spectra of the QPO fundamental (at 667 s in the first half, and 769 s in the second half) and the QPO subharmonic (1030 s and 1450 s). Figure 160 shows this QPO triplet appearing in the second half of the run.

7.4.10 s6184

In the Morlet spectrum in figure 161, in addition to the DNO that changes from 24 s to 27 s over the course of the run, and the 80 s lpDNO, three sections of an ~ 400 s oscillation can be seen, although it is of low coherence. We interpret this as the fundamental frequency of a DNO-related QPO.

from 1000 to 833 s, tracks the decrease in DNO period change from 25 to 23 s, (from time 4000 to 9000 seconds) where the DNO has a discontinuity, jumping to 27 s, and the QPO fundamental disappears completely. The 400 s QPO first harmonic is present for several cycles before 9000 s, but also disappears abruptly when the DNO period jumps. The 250 s QPO second harmonic, which is present throughout, does not appear to track the DNO.

7.4.9 s6138

s6138 was analysed in two sections (figures 158 and 159), as the run contained a large gap and change in integration time. In both parts, in addition to the 25 s DNO first harmonic and ~ 350 s QPO first harmonic found by Warner & Woudt (2006a), we see evidence in both the Morlet and MHAT spectra of the QPO fundamental (at 667 s in the first half, and 769 s in the second half) and the QPO subharmonic (1030 s and 1450 s). Figure 160 shows this QPO triplet appearing in the second half of the run.

7.4.10 s6184

In the Morlet spectrum in figure 161, in addition to the DNO that changes from 24 s to 27 s over the course of the run, and the 80 s lpDNO, three sections of an ~ 400 s oscillation can be seen, although it is of low coherence. We interpret this as the fundamental frequency of a DNO-related QPO.




Figure 153: Wavelet analysis of s0129.




Figure 154: 2nd, 3rd and 6th IMFs of s0129.

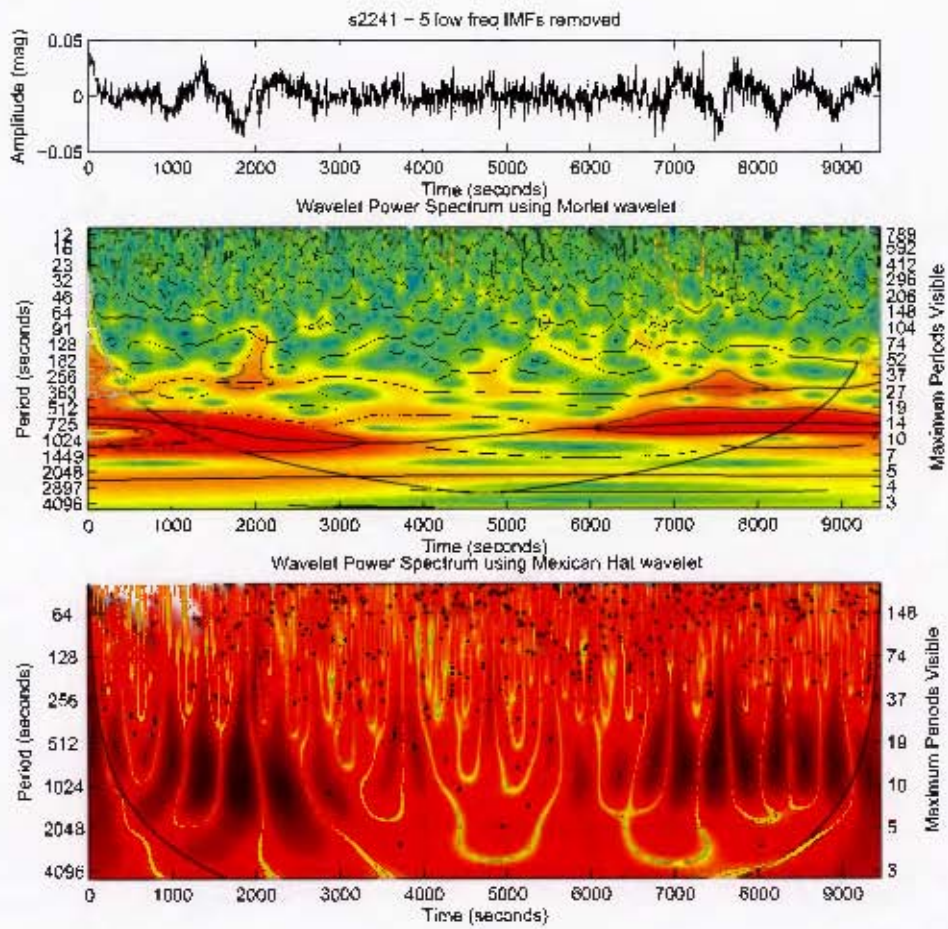


Figure 155: s2241 wavelet




Figure 156: s2241, showing superhumps.




Figure 157: Wavelet analysis of s2623.




Figure 158: Wavelet analysis of first half of s6138.




Figure 159: Wavelet analysis of second half of s6138.




Figure 160: 5th, 6th and 7th IMFs of the second half of s6138.

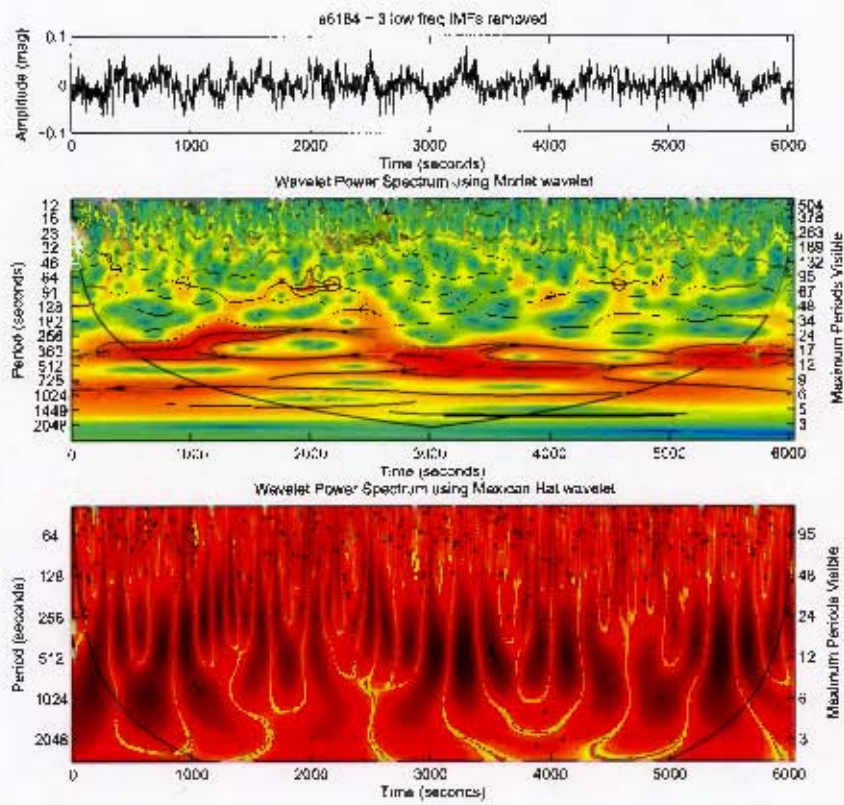


Figure 161: Wavelet analysis of s6184.

7.5 Conclusions

From the detailed discussion above, it is apparent that wavelet analysis and EMD are powerful tools for detecting both coherent and incoherent oscillations across a broad frequency range, for signals with a low signal-to-noise ratio.

The Morlet and MHAT spectra can be used in together to detect oscillations of only a few cycles (<5) (eg s1307, s6138 and s6184); it is this feature that we have used to identify QPOs objectively against criteria based on the appearance of QPOs identified in previous analyses, as discussed in section 7.4. We have added some 62 QPOs to the existing 44 of VW Hyi using this method.

For low amplitude coherent oscillations at short periods (ie < 50 s), it is no surprise that the periodogram is still the best means of detection. However, EMD can be useful in these situations for investigating possible amplitude changes in the oscillation. For oscillations in this range that have phase changes or gaps (e.g. s0129, s1307, s6528), the Morlet spectrum can be used to determine the structure of the oscillation, but only for amplitudes > 0.005 mag. The Morlet spectrum can also be used to detect short period oscillations of low coherence, which are not visible in the periodogram (e.g. s0129).

Coherent oscillations with periods greater than 800 s (eg s0127, s0484, s1322, s6138) fall in a region of the periodogram that often shows considerable noise in CV data, making differentiation between signal and noise difficult, especially if the amplitude of the oscillation is small. Fitting a model to

the noise continuum helps with the identification of possible QPOs in both the periodogram and the Morlet spectrum; even in the case of U Gem, where the QPOs are seen in the periodogram only as a broad hump of low power, correct modelling of the noise continuum enables accurate detection. EMD can then be used to provide details on the profile of the oscillation, as these are often non-sinusoidal. We find that the average amplitudes found from the IMFs are in good agreement with those published in the literature for relatively coherent QPOs, but tend to be somewhat larger than the published amplitudes for more incoherent QPOs. This is unsurprising, as the published amplitudes are calculated from the periodogram, which will underestimate the amplitude of an incoherent oscillation. Our amplitudes therefore give a new lower bound on the true amplitudes of QPOs, since, as was discussed in chapter 6, the presence of noise decreases the estimated amplitude in EMD.

The time resolution of the MHAT spectrum enables groups of spikes in the periodogram to be resolved, either as a single oscillation with discrete period changes (e.g. s0019), or (along with the Morlet spectrum) as multiple simultaneous oscillations of similar period (e.g. s2915, s2623, s3416 and s1322).

The wavelet spectra have also enabled us to develop far more detailed knowledge of the temporal behaviour of DNOs and QPOs. While the QPOs in s2241 had been observed previously, the wavelet spectra show that they only occur during superhumps, which was not previously known. We also observed further links between DNO and QPO behaviour, such as period tracking and simultaneous initiation/termination of oscillations (e.g. 2623).

Table 9: Table of QPOs detected in VW Hyi.

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Prev?	Comments
s0018	L	650	1900	4	346	346	0.008	QF	n	
s0018	L	1000	2000		75	75	0.007	DL	y	
s0018	L	1500	2500		20	20	0.009	DF	y	
s0018	L	3600	4600	3	286	286	0.007	QF	n	
s0019	L	0	0		30	32	0.02	D1	y	
s0019	L	1600	3200	3	487	487	0.05	Q1	y	
s0019	L	3500	6250	7	370	400	0.04	Q1	n	
s0019	L	5300	10000		20	20	0.02	D2	y	
s0019	L	5300	10000	4	510	510	0.05	Q1	y	
s0026	L	0	0	12	1042	1042	0.04	QU	y	
s0030	L	10000	0	3	1124	1124	0.04	QU	y	
s0111	Super	150	1300	4	307	307	0.007	QU	n	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s0122	Super	0	5000	12	409	409	0.005	QU	y	
s0122	Super	2000	8000	7	909	909	0.003	QU	n	
s0127	Super	0	2400		27.9	28.7	0.02	DF	y	
s0127	Super	1000	7000	5	1250	1250	0.03	QS	n	
s0127	Super	2000	5000	7	454	454	0.02	QF	y	
s0127	Super	3600	4800		29.3	29.3	0.02	DF	y	
s0127	Super	5100	8200		29.7	30.5	0.03	DF	y	
s0127	Super	6000	8200	5	498	498	0.02	QF	y	
s0127	Super	8000	11000	3	1075	1075	0.03	QS	n	
s0127	Super	9100	12000		32.2	32	0.015	DF	y	
s0127	Super	10000	12000	4	487	487	0.037	QF	y	
s0129	S	0	2500		33.5	33.5	0.02	D2	n	
s0129	S	0	2500	4	578	578	0.02	Q1	n	Implied first harmonic

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s0129	S	2000	3500		93	93	0.03	DL	y	DF in Warner - may be DL
s0129	S	3000	5000	3	658	658	0.02	Q1	n	Implied first harmonic
s0129	S	3500	5000		93	93	0.019	DL	y	DF in Warner - may be DL
s0129	S	5000	8000		33	33	0.02	D2	n	
s0129	S	5000	8000	4	714	714	0.02	Q1	n	Implied first harmonic
s0484	Super	0	4000		30	32	0.02	D1	y	
s0484	Super	0	3000	2.5	1282	1282	0.05	QF	y	
s0484	Super	0	12000	5	2857	2757	0.05	QS	n	
s0484	Super	2000	6000		68	68	0.009	DL	y	
s0484	Super	4000	6500		37	37	0.02	D1	y	
s0484	Super	4900	7500	6	488	488	0.03	Q1	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s0484	Super	5000	14000	7	1325	1325	0.05	QF	y	
s0484	Super	6500	12000		34	34	0.02	D1	y	
s0484	Super	8000	14000	11	530	530	0.03	Q1	y	
s1307	M	0	1000	2	500	500	0.01	QF	n	
s1307	M	750	3000		31	31	0.005	DF	n	
s1307	M	750	3000	5	250	250	0.005	Q1	n	
s1307	M	1500	3000		18	18	0.005	D1	y	
s1307	M	3500	4000		18	18	0.005	D1	y	
s1307	M	3500	4500		38	38	0.005	DF	n	
s1307	M	3500	4500	5	175	175	0.01	Q2	n	
s1322	M	0	4000		31	31	0.01	D2	y	
s1322	M	0	2500	11	357	357	0.05	Q2	n	
s1322	M	3000	6000	4	476	476	0.05	Q2	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s1322	M	6000	10000		29	30	0.01	D2	y	
s1322	M	6000	10000	8.5	465	465	0.05	Q2	y	
s1322	M	10000	12000		32	32	0.01	D2	y	
s1322	M	12000	16000		28	31	0.01	D2	y	
s1322	M	12500	15200	6	420	420	0.02	Q2	y	
s1571	Super	0	0	5.5	1170	1170	0.005	QU	y	
s1616	Super	0	2500	5	694	694	0.03	Q1	n	Implied fundamental
s1616	Super	0	5000	4	1282	1282	0.05	QF	n	Implied fundamental
s1616	Super	6000	8000		87	87	0.01	DL	n	
s1616	Super	6000	8000	7	333	333	0.02	Q2	n	Implied fundamental - low coherence DNOs visible in periodogram
s2241	Super	0	2000	3	714	714	0.02	QU	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s2241	Super	5500	9400	5	740	740	0.02	QU	y	
s2243	Super	0	0		75	75	0.002	DL	y	
s2623	M	0	2000		22	22	0.015	D2	y	
s2623	M	0	2000	5	408	408	0.02	Q1	y	
s2623	M	1000	2500	6	244	244	0.03	Q2	y	
s2623	M	2000	4000		25	25	0.014	D2	y	
s2623	M	4000	6000		24	24	0.014	D2	y	
s2623	M	4000	9000	3	1000	833	0.005	QF	y	Implied fundamental
s2623	M	5000	7000	7	263	263	0.036	Q2	y	
s2623	M	7000	10000	13	253	253	0.035	Q2	y	
s2623	M	7000	10000	6	500	500	0.02	Q1	n	
s2623	M	7500	10000		26	27	0.021	D2	y	
s2623	M	10000	12500		25	27	0.019	D2	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s2623	M	10000	12500	4	416	416	0.007	Q1	n	
s2915	L	700	1500	4	250	250	0.006	Q1	n	
s2915	L	1500	2500		20	20	0.011	D1	y	
s2915	L	1800	2300	2	250	250	0.006	Q1	n	
s3078	Super	0	7000		14	14	0.004	DF	y	
s3078	Super	400	6000		90	90	0.002	DL	n	
s3078	Super	2000	4000	4	183	183	0.004	QF	n	
s3078	Super	5500	6500	4	182	182	0.002	QF	n	
s3416	M	0	1000		25	25	0.017	D1	y	
s3416	M	0	1000	4	270	270	0.026	Q2	n	
s3416	M	1200	2500		27	27	0.015	D1	y	
s3416	M	1200	2600	4	350	350	0.06	Q1	y	
s3416	M	2600	4200		27	27	0.017	D1	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s5248	M	0	4000		36	36	0.02	D1	y	
s5248	M	0	0		96	96	0.02	DL	n	
s5248	M	500	4000	6	454	454	0.08	Q1	n	
s5248	M	4000	8000		40	40	0.005	D1	y	
s5248	M	4000	8000	5	694	694	0.07	QF	n	
s5248	M	5500	13000	5	1030	1030	0.05	QF	y	
s5248	M	5500	13000	3.5	2000	2000	0.06	QS	y	
s5248	M	8000	16000		40	40	0.005	D1	y	
s6138-I	M	0	3000		25	25	0.01	D1	y	
s6138-I	M	0	2200		25	26	0.01	D1	y	
s6138-I	M	0	2200	7	317	317	0.02	Q1	n	
s6138-I	M	0	3000	3	1030	1030	0.03	QS	n	
s6138-I	M	3200	5000	3	667	667	0.05	QF	n	

Continued on next page

Table 9 -- continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s6138-I	M	3300	5000		26	28	0.01	D1	y	
s6138-I	M	3300	5000	3	370	370	0.06	Q1	n	
s6138-I	M	5000	7300		28	28	0.005	D1	y	
s6138-II	M	0	3000		28	30	0.005	D1	y	
s6138-II	M	0	5000		70	70	0.02	DL	y	
s6138-II	M	5000	14000		30.6	30.6	0.007	D1	y	
s6138-II	M	5000	14000	13	344	392	0.04	Q1	y	
s6138-II	M	6000	14000	5	1450	1450	0.02	QS	n	
s6138-II	M	7000	14000	11	769	769	0.04	QF	n	
s6184	L	0	1000		24.8	27.8	0.02	DF	y	
s6184	L	0	1800	5	370	370	0.01	QF	n	
s6184	L	1000	2300		26	26	0.014	DF	y	
s6184	L	1800	2300		83	83	0.01	DL	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s6184	L	2300	3000		26	26	0.026	DF	y	
s6184	L	2300	3800	3.5	418	418	0.024	QF	n	
s6184	L	3000	3500		25	25	0.008	DF	y	
s6184	L	3500	3800		26	26	0.014	DF	y	
s6184	L	3800	5000		27	27	0.017	DF	y	
s6184	L	3800	5000		74	74	0.016	DL	y	
s6184	L	5000	5500		26	26	0.013	DF	y	
s6184	L	5000	6000	2.5	390	390	0.01	QF	n	
s6184	L	5500	6000		27	27	0.011	DF	y	
s6316	L	0	2500		20	20	0.009	D1	y	
s6316	L	0	2000	6	294	294	0.001	Q1	y	
s6316	L	2500	4000		20	20	0.011	D1	y	
s6316	L	4000	6000		20.5	20.5	0.006	D1	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s6316	L	4000	6000		77	83	0.008	DL	n	
s6316	L	4000	5000	5	227	227	0.015	Q1	y	
s6316	L	6000	10000		21.8	21.8	0.01	D1	y	
s6316	L	10000	18000		21	23	0.016	D1	y	
s6316	L	10000	15000		87	87	0.012	DL	n	
s6316	L	11000	18000	21	375	375	0.009	Q1	y	
s6316	L	15000	18000		88	88	0.011	DL	n	
s6528-I	L	0	6000		18	23	0.015	D2	y	
s6528-I	L	0	6000	13	666	666	0.024	QF	n	Implied fundamental
s6528-I	L	6000	12000		21	21	0.01	D2	y	
s6528-I	L	6000	12000	5	1299	1299	0.024	QS	n	Implied fundamental
s6528-I	L	12000	18000		23	23	0.011	D2	y	
s6528-I	L	12000	18000	6	885	885	0.013	QF	n	Implied fundamental

Continued on next page

Table 9 -- continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments	
s6528-I	L	18000	20800		23	23	0.009	D2	y	Implied fundamental	
s6528-I	L	18000	20800		32	32	0.008	D1	y		
s6528-I	L	18000	20800	2	1052	1052	0.017	QF	n		
s6528-II	L	0	6000		24	24	0.007	D2	y		
s6528-II	L	0	6000	21	370	370	0.015	Q2	n		
s6528-II	L	6000	10000		24	24	0.008	D2	y		
s6528-II	L	6000	10000		37	37	0.005	D1	y		
s7222-I	Super	0	2000		27	27	0.01	D2	y		
s7222-I	Super	0	10000	8	1282	1282	0.05	QF	n		Implied fundamental
s7222-I	Super	2000	4000		26	26	0.005	D2	y		
s7222-I	Super	2000	4000	5	625	476	0.02	Q1	n		
s7222-I	Super	3000	6000		27	29	0.011	D2	y		
s7222-I	Super	6000	7500		31	31	0.011	D2	y		

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s7222-I	Super	6000	8000		38	38	0.013	D1	y	
s7222-I	Super	6000	7500	6	270	270	0.01	Q2	y	
s7222-I	Super	7000	9500	5	588	588	0.04	Q1	n	
s7222-I	Super	7500	9500		29	29	0.015	D2	y	
s7222-II	Super	0	2500		34	34	0.013	D2	y	
s7222-II	Super	0	2500	3	549	549	0.045	Q1	n	
s7222-II	Super	2500	4000		29	29	0.015	D2	y	
s7222-II	Super	4500	8000	3	1136	1136	0.038	QF	n	
s7222-II	Super	5000	9000	8	333	333	0.035	Q2	y	
s7222-II	Super	5000	10000	5	555	555	0.037	Q1	n	
s7222-II	Super	5500	7500		28	28	0.022	D2	y	
s7222-II	Super	7500	8500		30	30	0.016	D2	y	
s7222-II	Super	8500	9000		32	32	0.009	D2	y	

Continued on next page

Table 9 - continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s7222-II	Super	9500	1000		34	34	0.01	D2	y	
s7301	M	0	2000		33	33	0.009	D2	y	
s7301	M	2000	4000		32	30	0.015	D2	y	
s7301	M	4000	6000		32	34	0.015	D2	y	
s7301	M	4000	7000	8	454	454	0.036	Q2	n	
s7301	M	6000	7000		29	29	0.014	D2	y	
s7301	M	7000	9000		32	32	0.012	D2	y	
s7301	M	7000	9000	5	456	456	0.05	Q2	n	
s7301	M	9000	0		29	31.8	0.014	D2	y	
s7311	M	0	0		90	90	0.014	DL	n	
s7311	M	0	1000	3	312	312	0.014	Q1	y	
s7311	M	1800	3500	8	294	294	0.017	Q1	y	
s7311	M	2000	3500		19	19	0.008	D1	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s7311	M	3500	6000		20	20	0.006	D1	y	Implied fundamental
s7311	M	6000	15000		14	14	0.015	D2	y	
s7311	M	6000	7400		20	20	0.007	D1	y	
s7311	M	6000	10000	12	316	316	0.008	Q1	y	
s7311	M	7000	9000	4	476	476	0.028	QF	n	
s7311	M	7400	10000		20	21	0.006	D1	y	
s7311	M	10000	13000	9	333	333	0.028	Q1	y	
s7311	M	10500	12500		21	21	0.01	D1	y	
s7311	M	12500	0		21	22	0.009	D1	y	
s7311	M	13000	16000	6	345	345	0.015	Q1	y	
s7311	M	13000	16000	4	666	666	0.016	QF	y	
s7342	M	0	0		20	20	0.006	D1	y	
s7342	M	0	2000	7	236	236	0.009	Q2	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s7342	M	1800	6000	5	714	714	0.006	QF	n	
s7342	M	2000	6000	10	321	321	0.01	Q1	y	
s7342	M	6000	0		14	14	0.015	D2	y	
s7342	M	6000	8000		38	38	0.009	DF	y	
s7342	M	6000	10000	14	136	136	0.005	Q2	n	
s7342	M	8000	9200	5	256	256	0.015	Q1	y	
s7342	M	8000	10000	4	500	500	0.009	QF	y	
s7621	L	0	600		28	28	0.007	DF	y	
s7621	L	600	2000		27	30	0.011	DF	y	
s7621	L	4000	6000	5	454	454	0.012	QF	y	
s7621	L	6000	11000	11	454	454	0.027	QF	n	
s7621	L	8000	10000		31	33	0.008	DF	y	
s7621	L	11000	14000		18	18	0.01	D1	y	

Continued on next page

Table 9 – continued from previous page

Name	Type	Osc start	Osc end	Cycles	P start	P end	Mag	Class	Old?	Comments
s7621	L	11000	14000		38	38	0.013	DF	y	

8 Further Analysis of VW Hydri Data

8.1 Introduction

In the previous chapter, we discussed the many new QPOs and lpDNOs which wavelet analysis enabled us to detect in VW Hyi. In this chapter we use these results to extend the investigation of the evolution of DNOs and QPOs during outburst and quiescence discussed in Woudt & Warner (2002a) and Warner & Woudt (2006a).

Unlike the majority of dwarf novae, which show DNOs and QPOs on the rise to outburst maximum, DNOs and QPOs are usually observed in VW Hyi toward the end of normal and superoutbursts, as the system is returning to minimum light (regions I, II and III in figure 162). The evolution of DNO and QPO behaviour for $T = -0.2$ to 1.15 d, described in Warner & Woudt (2006a), is summarized as follows:

From $-0.2 < T(d) < 0.2$, P_{DNO} shows the monotonic single-valued increase in from ~ 14 to 30s (region I in figure 162) characteristic of the standard LIMA model for dwarf novae early in the decline from maximum. From $T = 0.2$ to 1.15 d, propellering causes P_{DNO} to increase linearly from ~ 30 to 104s (region II in figure 162), after which the period decreases slightly, stabilizing at ~ 90 s at $T = 1.4$ d (region III in figure 162). After $T = 1.4$ d no further DNOs or QPOs are observed. Between $T = 0.2$ to 0.5 d the period increase shows little scatter, but between $T = 0.5$ and 1.15 d there is considerable scatter, largely due to superoutburst observations.

The first appearance of the DNO first harmonic (D1), and the QPO first




Figure 162: Averaged outburst profiles for the four different types of outburst seen in VW Hya. Solid line indicates superoutburst, dashed line indicates long normal outburst, dash-dot indicates medium normal and dotted indicates short normal outburst. Regions I, II and III show different DNO and QPO behaviour.

harmonic, is at about 0.2 d, and the D1 amplitude first dominates fundamental DNO (DF) amplitude at $T = 0.25$ d. The DNO second harmonic (D2) appears at about $T = 0.36$ d. Between 0.5 and 1.0 d, the D2 amplitude increases to dominate the D1 amplitude, and at 0.8 d the QPO second harmonic, (Q2), makes its first appearance. Either or both D1 and D2 are generally present in most runs after $T = 0.3$ d, and DF makes rare appearances up to $T = 1.1$ d. Occasionally the QPO subharmonic (at twice P_{QPO}) is seen. For DNO-related QPOs, the ratio $R = P_{QPO}/P_{DNO}$ ranges from 15.2 at $T = 0.2$ d to 11.0 at $T = 1.15$ d (Warner & Woudt, 2006a).

8.2 DNOs

Since the identification of QPOs as DNO-related or not relies on the detection of DNOs in the run, we have perforce included DNOs in our analysis. As we have only found 2 DNOs not reported by Woudt & Warner (2002a) and Warner & Woudt (2006a), and all those previously detected were found in our analysis, our analysis here parallels that of Woudt & Warner (2002a) or Warner & Woudt (2006a), providing a useful check of the accuracy of our results. Figure 163 shows all the DNOs identified in our analysis, plotted by outburst type (cf Warner & Woudt (2006a), figure 1). As expected, the DNOs form three groups, showing the same phenomena discussed in Warner & Woudt (2006a), and summarized in the previous section. The highest group, containing fundamental period DNOs, begins with the non-linear increase in period from $T = -0.2$ to 0.2 d predicted by the LIMA model (region I in figure 163), followed by the linear phase due to propellering (region II in figure 163). Between $T = 0.2$ and 1.15 d both the middle group (first harmonic DNOs) and the lowest group (second harmonic DNOs) follow the linear trend due to propellering (region II in figure 163). From $T = 1.14$ d the D2 period decreases slightly (region III in figure 163).

The dotted line in figure 163 is the least squares fit to the second harmonic, for $T = 0.2$ to 1.15 d, and is also shown multiplied by three to give the implied fundamental period. The dashed line is the least squares fit to the first harmonics, for $T = 0.2$ to 1.15 d, and is also shown multiplied by two to give the implied fundamental period. The solid line is the least squares fit to the observed or implied fundamental for $T = 0.2$ to 1.15 d. Figure 164 shows the observed or implied fundamental DNO period, and the least

squares fit for the range $T = 0.2$ to 1.15 d, $P_{DNO} = 64.9T(\text{d}) + 20.1$, which is comfortably close to the $P_{DNO} = 65.06T(\text{dd}) + 19.0$ found by Warner & Woudt (2006a) (see Figure 3 of Warner & Woudt (2006a))

Figure 165 illustrates the changes in number of DF, D1 and D2 for $T = -0.4$ to 1.4 d. The number of DFs reaches a maximum between $T = 0$ and 0.1 days, and then falls off rapidly to zero at $T = -0.4$ d. D1s are first observed at $T = 0.1$ d, and rapidly reach maximum numbers between $T = 0.2$ and 0.3 d. The fall off in D1 numbers is less rapid than that of DFs; the last D1s are seen at $T = 1$ d, 0.8 d after their first appearance. D2s are initially seen at $T = 0.3$ d, but only show significant numbers at $T = 0.7$ d. Numbers of D2s decrease to zero at $T = 1.4$ d, but appear to show a spike at $T = 1.1$ d.

To investigate the amplitude evolution of DNOs during outburst, we converted our magnitudes (m_{DNO}) to intensity (I_{DNO}), using the outburst templates to provide the correct system magnitude (m_{sys}) for a given phase and outburst type:

$$I_{DNO}(\text{Jy}) = 3880(\text{Jy})10^{\frac{m_{sys}-m_{DNO}}{-2.5}} \quad (44)$$

where $1 \text{ Jy} = 10^{-26} \text{ W/m}^2/\text{Hz}$. Figure 166 plots the DF, D1 and D2 amplitudes over $T = -0.4$ to 1.4 d. Solid circles show normal outburst DNOs, while superoutburst DNOs are shown as open triangles. There is a sharp decline in DNO amplitude between $T = -0.2$ and 0.2 d (region I), coinciding with the standard LIMA period increase phase, which has not been previously detected. It also appears, although there are too few points to be sure, that this phase occurs somewhat later in superoutbursts. From $T = 0.2$ d onwards there is no discernable evolution of DNO intensity, either in the




Figure 163: DNOs observed in VW Hyi. Supercroutburst data are indicated by filled triangles, long normal outburst data by squares, medium normal outburst data by circles and short normal outburst data by stars. The solid line is the least squares fit to the observed or implied fundamental for all normal outbursts; the dashed and dotted lines are the least squares fits to the first and second harmonic respectively (also shown multiplied by 2 and 3 respectively).

fundamental or harmonics, although there might be a decrease in D2 amplitude from $T = 1.2$ d onwards, coinciding with the phase of period decrease (region III in figure 166).

The very different amplitude behaviour seen in regions I and II add weight to the interpretation that the period increase in these regions is due to different physical mechanisms - the LIMA model in region I, and propellering in region II.




Figure 164: Evolution of the observed or implied fundamental DNO period. Superoutburst data are indicated by filled triangles, long normal outburst data by squares, medium normal outburst data by circles and short normal outburst data by stars.




Figure 165: The evolution of the number of fundamental, first and second harmonic DNOs observed during late outburst in VW Hyi.




Figure 166: Evolution of the amplitudes of DNOs during late outburst. Solid circles indicate normal outburst data, open triangles superoutburst data.

8.3 lpDNOs

In previous work, DNOs with periods ≥ 70 s have been classified as DFs or lpDNOs almost arbitrarily; we have chosen to classify every DNO with period ≥ 70 s as a lpDNO. Using this classification, there are 8 previously observed lpDNOs, and we have discovered a further 7.

The period and amplitude evolution of the lpDNOs observed during outburst is shown in figure 167. For $T = -0.4$ to 1.5 d, the period of the lpDNOs increases from 70 to 100 s; this is the first time that evolution of the lpDNO period during outburst has been seen with certainty; Warner, Woudt & Pretorius (2003) suggested it as a possibility, but had to few observations to be sure. A linear fit of P_{lpDNO} as a function of phase T yields $P_{lpDNO} = 78.7 + 8.3T(\text{d})$, and is shown as a dotted line in figure 167, top panel. The evolution of the DNO fundamental (observed or implied, short dashed line) and the lpDNOs is shown in figure 168, where it can be seen that the implied fundamental DNO period tends to the lpDNO at the end of outburst (long dashed line), merging at $T = -1.15$ d, after which there is no further evolution of either DNO type.

We calculated the lpDNO amplitude in the same way as DNO amplitudes (figure 167 lower panel). The lpDNO amplitude appears to evolve during late outburst; when plotted together with DNO amplitudes, as in figure 169, it is clear the the lpDNOs are following a similar decline in amplitude until $T = 0.2$ d, followed by constant amplitude until quiescence is reached. The lpDNO amplitudes are comparable to DNO amplitudes throughout the outburst, with a maximum of about 0.5 mJy, although the two lpDNOs observed during early superoutburst, at $T = -10.75$ d and $T = -5.25$ d, both

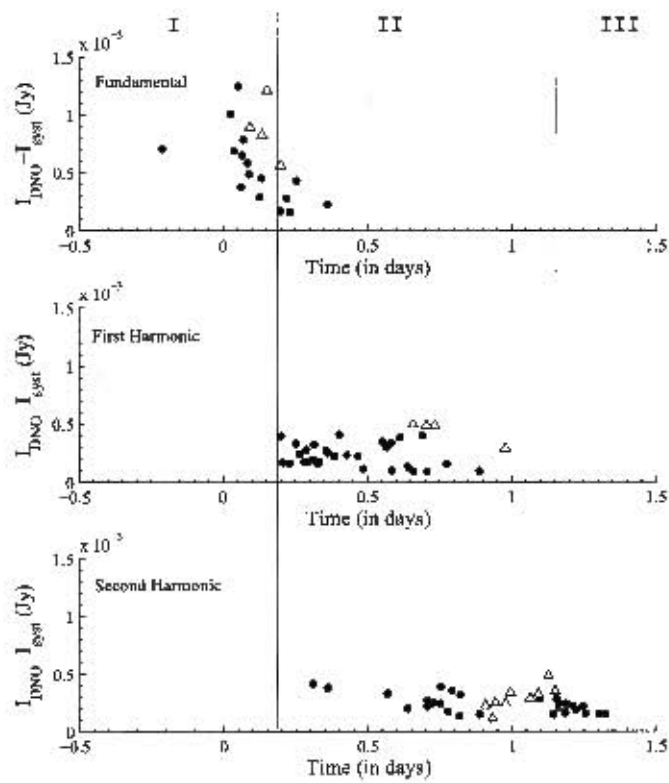


Figure 166: Evolution of the amplitudes of DNOs during late outburst. Solid circles indicate normal outburst data, open triangles superoutburst data.

show far higher amplitudes than any of the later lpDNOs.

In section 1.4.3, we discussed two possible models for lpDNOs. That the amplitude behaviour of the lpDNOs parallels that of the DNOs suggests that they are due to similar physical mechanisms, favoring the two-pole magnetic accretion model over the shock-induced oscillations just outside the co-rotation radius. However, further observations are needed to determine if the lpDNO period also shows different behaviour in the LIMA phase and the propeller phase.

We have included the number of lpDNOs at each phase in figure 165; while the DNOs DF, D1 and D2 show a distinct evolution over the course of the outburst decline, the number of lpDNOs observed appears to remain relatively constant, although this may be due to the small sample size.




Figure 167: a) lpDNO period plotted by outburst type. Circles indicate long normal outburst data, triangles indicate medium normal data, squares indicate short normal outburst data and filled diamonds indicate superoutburst data. The dotted line indicates the evolution of the fundamental DNO period for $T=0.2-1.15d$. b) lpDNO amplitude evolution during outburst.




Figure 168: Period evolution of DNOs and lpDNOs. Circles indicate fundamental period DNOs, squares indicate first harmonic DNOs, triangles indicate second harmonic DNOs and diamonds indicate lpDNOs. Normal outburst oscillations are indicated by open shapes, superoutburst oscillations by filled shapes.




Figure 169: Evolution of the amplitudes of both (fundamental or implied fundamental) DNOs (circles) and lpDNOs (triangles) during late outburst. Open shapes indicate normal outburst data, filled shapes indicate superoutburst data.

8.4 Outburst QPOs

Figure 170 shows all the QPOs detected in our analysis (for ease of viewing, DNO-related QPO harmonics are shown at the implied fundamental). Previously studied QPOs (Woudt & Warner (2002a) or Warner & Woudt (2006a)) are indicated by filled shapes; open shapes indicate QPOs that have been detected using wavelet analysis. As mentioned in the previous chapter we have added some 62 QPOs to the existing 44, enabling us to draw more accurate, detailed conclusions about the evolution of QPOs than those presented previously. Circles represent DNO-related QPOs, triangles those QPOs which have no DNO counterpart and squares indicate quiescent QPOs. We divide the investigation of QPOs into three sections: first, DNO-related QPOs (QF, Q1, Q2, QS) are discussed, then QPOs observed during outburst that do not appear to be related to DNOs (QU), and finally QPOs observed in quiescence (QQ). Note that the large numbers of QPOs observed around $0 < T(d) < 1.15$ is at least partly due to the observing campaign of Brian Warner and Patrick Woudt, who selectively observed VW Hyi during late outburst.

8.4.1 DNO-related QPOs

Out of 24 outburst runs showing DNOs, including all those covered in Warner & Woudt (2006a), we found DNO-related QPOs in each run. In addition, after finding a triplet of QPOs in s1616, close inspection of the FT and wavelet showed evidence of DNOs, although of very low coherence.

Figure 171 shows two plots of the observed or implied QPO fundamental

frequency, the first (a) plotted by outburst type, the second (b) by QPO harmonic type. It is clear from (a) that while all QPOs show similar evolution over time, superoutburst QPOs (filled diamonds) exhibit greater scatter than normal outburst QPOs (open shapes). For this reason, the least squares fit for the range $0.2 < T(\text{d}) < 1.15$ was calculated using only normal outburst data. Our fit for this range is $P_{QPO}(s) = 679T(\text{d}) + 368$, which has a slightly steeper slope than the $P_{QPO}(s) = 575T(\text{d}) + 373$ found by Warner & Woudt (2006a), and results in the ratio $R = P_{QPO}/P_{DNO}$ changing from 15.3 at $T = 0.2$ d to 12.1 at $T = 1.15$ d.

Figure 172 shows all DNO-related QPOs; circles indicate fundamental, squares the first harmonic and triangles and diamonds the second and subharmonics respectively. The least squares fit to the first and second harmonics are indicated by dashed and dotted lines respectively, also shown multiplied by 2 or 3 to give the period of the implied fundamental period. The fitted implied fundamental line is indicated by a solid line, also shown multiplied by 2 to give the subharmonic period evolution.

Figure 173 shows the evolution of the QPO fundamental, plotted by outburst type. Filled diamonds indicate QPOs occurring during superoutburst. Superoutburst QPOs, especially those observed after $T = 0.6$ d, show a much greater scatter than that observed in normal outburst QPOs. The superoutburst QPOs for $T > 0.6$ d are from 3 different runs, and each run shows a similar level of scatter. The overall trend for these runs is a decrease in period, similar to that seen in D2s near the end of outburst, but starting somewhat earlier.

The number of QPOs present over $-0.4 < T(\text{d}) < 1.4$ is shown in figure 174.

The evolution in QPO numbers is markedly different from that of the DNOs. The first fundamental period QPOs appear at $T = -0.3$ d, increasing to reach maximum numbers at $T = 0.1$ d, at the same time as the maximum number of DFs, but then fluctuates considerably, with the last QFs appearing at $T = 1.1$ d, 0.7 days after the last DF is seen. The first Q1s are seen at $T = 0.1$ d, at the same time as the first D1s, with the maximum number of Q1s visible at between $T = 0.2$ and 0.3 d, again following the D1 evolution. However, after $T = 0.3$ d, the numbers of Q1s fluctuates in the same manner as the QFs; the last Q1s are also visible at $T = 1.1$ d. The Q2s first appear at $T = 0.2$ d, a little earlier than the first D2s. The maximum number of Q2s is reached at $T = 0.7$ d. While the numbers of DNOs of each type are similar, and show a similar evolution, there are considerably fewer Q2s than QFs and Q1s, although they all show similar fluctuations in number. There are too few subharmonic QPOs observed to draw any useful conclusions; it is not even clear whether they are subharmonics of the fundamental frequency QPOs, or if they should be ranked as unclassified QPOs.

Figure 175 shows the amplitude of QF, Q1, Q2 and QS for $T = -0.5$ to 1.5 d, in intensity. QF, Q1 and QS have a maximum amplitude of ~ 0.0015 Jy, while Q2s show a lower maximum of ~ 0.001 Jy. There is no clear evolution of amplitude over the course of the outburst, even during $T = -0.2$ to 0.2 d when the DNO amplitude is decreasing. There is no obvious distinction between superoutburst QPO amplitude and normal outburst QPO amplitude.

An alternative model for the behaviour of DNOs, based on QPOs, has been suggested by Warner & Woudt (2006a). If the mechanism exciting the QPOs causes 1st or 2nd harmonics in the travelling wave, then the QPO profile will be either double- or triple-humped, causing the rotating DNO beam




Figure 170: All QPOs analysed. Circles indicate DNO-related QPOs, triangles indicate unclassified outburst QPOs and squares indicate quiescent QPOs. Filled shapes are QPOs previously analysed in Warner I, II or IV; open shapes are QPOs added by this study.

to generate DNOs at twice or three times the fundamental DNO frequency. They noted that the observed simultaneous appearance of D1s and Q1s, and D2s and Q2s, is well-explained by this model. That our increased pantheon of QPOs continue to show this pattern, and that we always observe QPOs when DNOs are present, gives further credence to this model. In addition, the fewer numbers of Q2s that we observe are readily explained: second harmonics would be harder to excite than first harmonics, and so would be rarer. Warner & Woudt (2006a) also point out that this model ensures the $P_{QPO}/P_{DNO} = 15$ relationship. However, they note that this model does not facilitate a simple explanation for simultaneous D1s and D2s, or the observed abrupt jumps from direct to synodic DNOs.




Figure 171: a) shows the observed or implied QPO fundamental period, plotted by outburst type. Circles indicate long normal outburst data, triangles indicate medium normal data, squares indicate short normal outburst data and filled diamonds indicate superoutburst data. b) show the same data, but plotted by QPO harmonic type. Fundamental QPOs are indicated by circles, first harmonics by triangles, second harmonics by squares and subharmonics by diamonds.




Figure 172: QPOs by harmonic, showing linear fits to the second harmonic (dotted line), the first harmonic (dashed line) and the observed or implied fundamental (solid line). Fundamental QPOs are indicated by circles, first harmonics by triangles, second harmonics by squares and subharmonics by diamonds.




Figure 174: The evolution of the number of fundamental, and first, second and sub-harmonic QPOs observed during late outburst in VW Hyi.

8.4.2 Unclassified Outburst QPOs

Figure 176 shows the unclassified outburst QPOs, plotted by outburst type as in figure 171a), with filled shapes indicating unclassified QPOs, which are also labeled by run number and period. The 307 s QPO in run s0111 is the only previously undetected QU; the others have been reported in Woudt & Warner (2002a).

Figure 177 shows where each of the unclassified QPOs appears on the outburst profile. The wavelet spectra of all the unclassified outburst QPOs have been discussed in detail in the previous chapter.

There are only two normal outburst QPOs that are unclassified - s0026 and s0030 - both from long outbursts. Since VW Hyi takes longer to return to minimum light after a long normal outburst, the 1042 s QPO from s0026

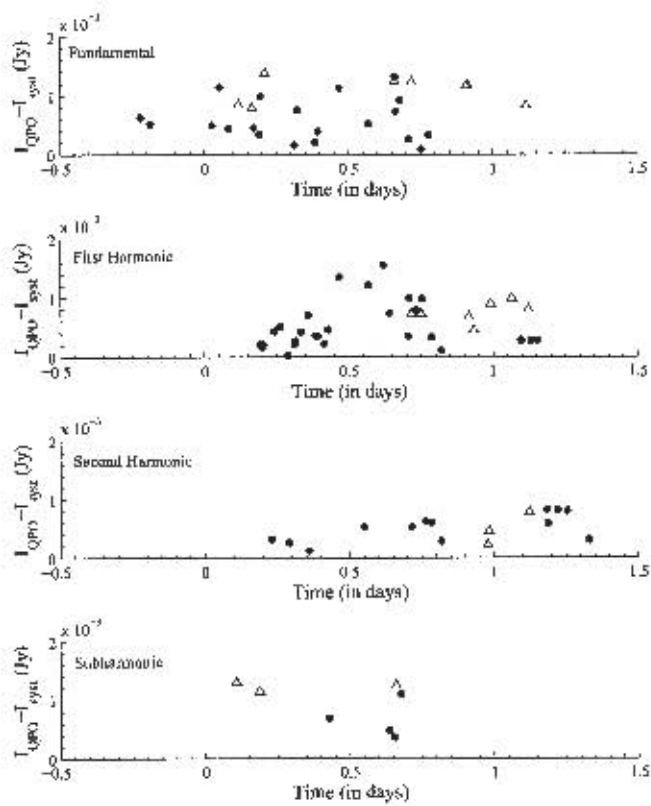


Figure 175: Evolution of the amplitudes of QPOs during late outburst. Solid circles indicate normal outburst data, open triangles superoutburst data.

occurs at a magnitude usually reached 1 day earlier, at which time the implied QPO fundamental is ~ 1000 s. The 1124 s QPO seen in s0030 should probably be classed as a quiescent QPO, since by this time VW Hyi has reached minimum light.

The RASNZ outburst profile for the s0122 and s1571 superoutbursts (figures 178 and 179) show that these QPOs occur at brighter magnitudes than those predicted by the average outburst profile. These magnitudes usually occur at $T_1 = -1$ d and $T_2 = -0.5$ d respectively (see figure 177). However, even correcting for this (T_1 and T_2 on figure 177), these QPOs still occur considerably earlier than any DNO-related QPOs. The s2241 714 s (4 cycles) and 740 s QPOs (5 cycles) occur earlier still, at a lower magnitude than any previously observed (8 mag) - reached only during superoutburst. We note that the 307 s QPO of s0111 occurs when the system is at the same magnitude, on the rise to supermaximum, at which ~ 300 s QPOs are observed on the fall to minimum (T_3 on figure 177).




Figure 176: All outburst QPOs. Fundamental QPOs are indicated by circles, first harmonics by triangles, second harmonics by squares and subharmonics by diamonds. Open shapes are DNO-related QPOs, filled shapes occur in runs with no DNOs.




Figure 177: Occurrence of unclassified QPOs in the outburst profile.




Figure 178: s0122 outburst




Figure 179: s1571 outburst




Figure 180: Quiescent QPOs

8.5 Quiescent QPOs

Figures 180a) and 180b) show the period and amplitude evolution of the few QPOs observed in quiescence. We have not converted the amplitudes to intensity, as no evolution in magnitude is observed, and there is no clear baseline to use for calibration, as VW Hyi fluctuates in magnitude during quiescence. While the QPOs in s0102 and s0105 are only visible during the orbital hump, as has been observed in the unclassified superoutburst QPOs, the other quiescent QPOs are not limited in duration. Clearly further examples of quiescent QPOs are needed.

8.6 Discussion

In the previous chapter, using wavelet analysis and EMD we elicited detailed information about the amplitude and period behaviour of individual DNOs, lpDNOs and QPOs in quiescent and outburst lightcurves of VW Hyi. In this chapter, we combined the information from all our observations to extend the evolutionary description of these oscillations given in Woudt & Warner (2002a) and Warner & Woudt (2006a). We have both confirmed their description and discovered several new phenomena, including:

- The amplitude of both DNOs and lpDNOs shows a significant decrease during the non-linear period increase phase, but a similar change is not seen in QPOs during this phase.
- The period of lpDNOs increases over the course of the outburst, in a similar manner to the period increase observed in DNOs.
- Whenever DNOs are observed in VW Hyi, related QPOs are present.

This chapter explores just a few of the ways in which information from wavelet spectra and EMD can be used to gain a more detailed view of DNOs and QPOs. We have not, for example, investigated the way in which amplitudes and phases vary; measurements for both of these are possible with EMD and the MHAT wavelet, and will form the basis of future study. However, even using average amplitudes, we have already learned a great deal about the behaviour of DNOs and QPOs, which will be of use in the development and testing of physical models of these phenomena.

9 Mathematical Models for QPOs

As discussed in section 2, Robinson & Nather (1979) modeled QPOs as second-order autoregressive processes. In this chapter, we investigate both stationary and time-varying autoregressive processes as suitable ways of modeling QPOs. We begin with an overview of stationary auto-regressive processes, followed by a discussion of our results in modeling QPOs in RU Peg and VW Hyi using this method. We then discuss the decomposition of lightcurves using a time-varying auto-regressive model, and our results in applying this model to our data.

9.1 Stationary Auto-regressive Modeling

A time-series x_t can be described by an autoregressive moving-average process (ARMA(p,q)) if

$$x_t = \underbrace{\phi_1 x_{t-1} + \phi_2 x_{t-2} + \dots + \phi_p x_{t-p} + \varepsilon_t}_{\text{AR}(p)} + \underbrace{\theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2} + \dots + \theta_q \varepsilon_{t-q}}_{\text{MA}(q)}, \quad (45)$$

where ε_k are independent and identically distributed $N(0, \sigma^2)$ variables (Harvey, 1993). The process is weakly stationary if the roots of $1 - \phi_1 z - \phi_2 z^2 - \dots - \phi_p z^p = 0$ lie outside the the unit circle (Hamilton, 1994). The power spectrum of an AR(p) process (i.e. a process with no moving-average

component) is given by

$$f(\omega; \psi) = \left(\frac{\sigma^2}{2\pi} \right) \left(\frac{1}{|1 - \phi_1 e^{-i\omega} + \dots + \phi_q e^{-iq\omega}|^2} \right) \quad (46)$$

where $\psi \equiv (\sigma^2, \phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q)$. While the power spectrum characterizes the time-dependent properties of a process in the frequency domain, the autocorrelation function (ACF) is an alternative way of characterizing the time-dependent properties, using the time domain. The ACF is found by dividing the autocovariances (defined in section 4.2) by the variance ($\rho_j = \gamma_j/\sigma^2$), and is in fact the Fourier transform of the power spectrum. For an AR(p) process, the ACF has the form

$$\rho_j = \phi_1 \rho_{j-1} + \phi_2 \rho_{j-2} + \dots + \phi_p \rho_{j-p} \quad (47)$$

for $j = 1, 2, \dots$

If the process is stationary, then the minimum number of AR coefficients needed to model hump or peak in a spectrum is 2. As mentioned in section 2, an AR(2) model describes a randomly damped harmonic oscillator, with resonant frequency $\omega_{AR(2)}$, mean amplitude $\sigma_{AR(2)}$ and damping time $\gamma_{AR(2)}$. These parameters can be estimated from the periodogram or the ACF, or calculated from the AR parameters found by fitting an AR(2) model directly to the time series. While Robinson & Nather (1979) used the periodogram and correlogram to estimate $\omega_{AR(2)}$, $\sigma_{AR(2)}$ and $\gamma_{AR(2)}$, we have chosen to fit an AR model directly to the data, using the statistics package SPSS.

9.2 AR Modelling of RU Peg and VW Hyi

We used the prewhitened RU Peg and VW Hyi data the previous chapter. We analysed each lightcurve in overlapping sections, with length at least 5 times the longest QPO period, and overlap chosen to be shorter than QPO period. To decide how many AR coefficients to use initially, we looked at the periodogram of the data, and estimated 2 coefficients for every significant hump. We fitted an initial model in SPSS, and continued fitting with additional AR coefficients (if the ACF appeared to contain periodicities) until the ACF of the residuals had no significant coefficients, and the periodogram of the residuals showed no peaks or humps. In a few cases it was necessary to add a moving-average coefficient to successfully model the noise, but this was rare. For some data it was not possible to remove all significant coefficients in the ACF by adding AR or MA coefficients. We then focused on ensuring that the periodogram of the residuals showed no significant peaks or humps.

Figure 181 shows the first 4 sections of s1674, which have length 640 s and an overlap of 480 s. The estimated AR coefficients are as follows (standard errors in parentheses):

	I 1 - 640s	II 160 - 800s	III 320-960s	IV 480-1120
ϕ_1	1.1(0.2)	0.5(0.4)	0.7(1.6)	1.01(0.01)
ϕ_2	-0.1(0.2)	0.2(0.5)	-0.0(0.5)	0.03(0.12)
ϕ_3	0.3(0.1)	0.1(0.4)	-0.0(0.5)	0.09(0.04)
ϕ_4	-1.1(0.1)	-0.2(0.4)	-0.5(0.2)	-1.26(0.05)
ϕ_5	0.6(0.2)	0.1(0.4)	-0.6(0.9)	0.61(0.06)
ϕ_6		-0.3(0.4)	0.0(0.5)	0.07(0.10)
ϕ_7		-0.6(0.5)	-0.5(0.4)	
ϕ_8		0.6(0.3)	0.2(0.5)	
ϕ_9			-0.2(0.2)	
ϕ_{10}			0.1(0.4)	

We note that the number of coefficients needed to model each section changed significantly from section to section, and in all cases more than two pairs of coefficients were needed. Particularly in the third section, the standard errors on the fitted coefficients suggest that the model is ill-defined. Most importantly, the values for each coefficient change significantly from section to section.

In figure 182 we have plotted the periodograms of each section, and the power spectrum resulting from the estimated parameters. In all four sections, the coefficients of the fitted AR process result in a peak in the periodogram that roughly corresponds to the 12 s DNO, which is increasing in amplitude and coherence during this section. The 50 s QPOs, however, is not modeled in the third section, where the periodogram shows that it is at its largest amplitude.




Figure 181: The first 4 sections of s1674, which each have length 640 s and an overlap of 480 s.

Increasing the overlap did result in overlapping sections showing closer coefficients, but over an entire run coefficients still changed significantly. Varying the length of the sections used resulted in best fit models with completely different coefficients, and seldom improved the fit if the initial fit was ill-defined.

Further problems were encountered when modeling VW Hyi. Figure 183 shows the periodograms and fitted power spectra for the first 2800 s and 4800 s of VW Hyi s0484. In the shorter section model, the 33 s DNO appears to be modeled, but none of the longer period QPOs are. In the longer section the DNO is no longer modeled, and the QPOs are included in a red continuum. In fact, only when the QPOs are of particularly large amplitude and coherence are they modeled at all. DNOs, however, were frequently well modeled by an AR process.

9.3 Time-varying Autoregressive Decomposition

The time-varying nature of the AR coefficients found in the previous section led us to consider the time-varying autoregressive decomposition (TVAR) used by West, Prado & Krystal (1999), who successfully analysed EEG data from patients undergoing electroconvulsive therapy. West, Prado and Krys-




Figure 182: Periodograms of each section of s1674, and the power spectrum resulting from the estimated parameters.




Figure 183: The periodograms and fitted power spectra for the first 2800 s and 4800 s of VW Hyi s0484

as $\mathbf{G}_t = \mathbf{E}_t \mathbf{\Lambda}_t \mathbf{E}_t^{-1}$, where $\mathbf{\Lambda}_t$ is the diagonal matrix of eigenvalues and \mathbf{E}_t is the matrix with columns formed by the corresponding eigenvectors. Defining $\mathbf{H}_t = \text{diag}(\mathbf{E}_t' \mathbf{F}) \mathbf{E}_t^{-1}$, and linearly transforming \mathbf{x}_t to $\gamma_t = \mathbf{H}_t \mathbf{x}_t$, equation 49 is reparameterized as

$$\mathbf{x}_t = \mathbf{1}' \gamma_t, \quad \gamma_t = \mathbf{\Lambda}_t \mathbf{K}_t \gamma_{t-1} + \delta_t \quad (51)$$

where $\mathbf{1} = (1, 1, \dots, 1)'$, $\delta_t = \mathbf{H}_t \omega_t$ is a zero-mean normal innovation with a structured and singular variance matrix, $\mathbf{K}_t = \mathbf{H}_t \mathbf{H}_t^{-1}$, and $\gamma_t = (\gamma_{t,1}, \gamma_{t,2}, \dots, \gamma_{t,p})'$. \mathbf{x}_t is thus the sum of the individual $\gamma_{t,j}$ processes.

If we assume that number of complex and real eigenvalues of \mathbf{G}_t stays the same over time, and denote the complex conjugate eigenvalue pairs $r_{t,j} \exp(\pi \omega_{t,j})$ for $j = 1, \dots, c$ and the real eigenvalues by $r_{t,j}$ for $j = 2c + 1, \dots, p$, then by construction

$$\mathbf{x}_t = \sum \gamma_{t,j} = \sum_{j=1}^c z_{t,j} + \sum_{j=2c+1}^d y_{t,j} \quad (52)$$

where

$$y_{t,j} = r_{t,j} y_{t-1,j} + \nu_{t,2c+j} \text{ for } j = 1, \dots, r \quad (53)$$

is a TVAR(1) process and

$$z_{t,j} = 2r_{t,j} \cos(\omega_{t-1,j}) z_{t-1,j} - r_{t,j}^2 z_{t-2,j} + \psi_{t,j} \text{ for } j = 1, \dots, c \quad (54)$$

is a TVARMA(2,1) quasi-periodic process with characteristic frequency $\omega_{t,j}$ and slowly varying modulus $r_{t,j}$. Thus x_t has been decomposed into a sum of slowly varying trends (the TVAR(1) processes) and quasi-periodic components (the TVARMA(2,1) processes), which reflect the latent frequency structure of the data.

There are several points to note about the above decomposition:

1. It is assumed \mathbf{G}_t has p distinct eigenvalues; this is the case for TVAR models of practical interest (West, Prado & Krystal, 1999)
2. If the number of real and complex eigenvalues changes, or the frequency content of one of the latent processes changes, then frequency bands may appear to "switch" between components.
3. If $\mathbf{K}_t = \mathbf{H}_t\mathbf{H}_{t-1}$ differs significantly from the identity (i.e the TVAR coefficients are changing quickly), then the interpretation of the frequency components is unclear; it is only if the TVAR coefficients are slowly changing, so that \mathbf{G}_t and \mathbf{G}_{t-1} are similar, that the decomposition can be interpreted as identifying the latent components in the signal

Estimation of the model begins with the fitting of an initial prior, followed by sequential updating and retrospective filtering to compute the eigenvectors and eigenvalues of \mathbf{G}_t ; the details are discussed in depth in the Appendix of West, Prado & Krystal (1999).

9.4 TVAR Decompositon of VW Hyi

We used the matlab software² written by Raquel Prado and Mike West, to decompose the prewhitened VW Hyi data discussed in the previous chapter. For each decomposition we had to pick an upper and lower bound for p , from which the best initial prior was calculated. We chose to start with $p \in [4, 8]$, and then if a model at either extreme was produced, to extend the range appropriately.

Table 9.4 give the results of the decompositions. ‘DNOs’ and ‘QPOs’ indicate whether DNOs or QPOs have been found in the data using wavelet analysis. The order of the initial TVAR model used is listed in column ‘p’, while the (most common) number of complex pairs is listed under ‘nc’. The number of times the decomposition switches the number of complex pairs is included under ‘Switching’, while the type of decomposition which models the QPOs (if present) is given under ‘QPOs detected in’.

For the majority of runs analysed, the QPO(s) where found to be modeled in the TVAR(1) components of the decomposition; only 3 of 36 runs showed QPOs modeled in a TVARMA(2,1) component.

Figure 184 shows the eigenstructure decomposition of VW Hyi run s1616, which is known to contain an lpDNO at 87s, a fundamental QPO at 1282s, a first harmonic QPO at 694s and a second harmonic QPO at 333s. Next to each component we have added a periodogram of the component, to check what frequency range it models. The series is modeled with a TVAR(8) process, which is decomposed into 3 TVARMA(2,1) processes, and 2 TVAR(1)

²Available at <http://xpress.isds.duke.edu:8080/softwarelinks/tvar.html>

Run	Phase	DNOs	QPOs	p	nc	Switching	QPOs detected in
s0110	-14.85	n	y	5	2	0	TVAR(1)
s0111	-14.77	n	n	6	3	0	TVAR(1)
s0112	-14.67	n	n	10	4	5	TVAR(1)
s3434	-11.7	n	n	9	3	3	TVAR(1)
s2241	-11.34	n	y	9	4	3	TVAR(1)
s2243	-10.36	n	y	16	7	0	TVAR(1)?
s3078	-5.32	y	y	7	3	1	TVAR(1)
s0480	-5.19	n	n	13	6	0	
s2911	-2.71	n	n	8	3	0	TVAR(1)?
s0122	-1.82	n	y	8	3	0	TVAR(1)
s1277	-1.77	n	n	7	2	0	TVARMA(2,1)
s1571	-1.04	n	y	14	6	0	TVAR(1)
s0124	-0.79	n	n	9	4	0	TVARMA(2,1)
s1307	-0.3	y	y	16	9	5	TVAR(1) & TVARMA(2,1)
s0018	-0.11	y	y	9	4	2	TVAR(1)
s6184	0.06	y	y	9	4	0	TVAR(1)
s0127	0.17	y	y	14	6	1	TVAR(1)
s6138	0.54	y	y	7	3	1	TVAR(1)
s3416	0.56	y	y	9	4	1	TVAR(1)
s1322	0.7	y	y	10	4	4	TVAR(1)
s5248	0.76	y	y	7	2	1	TVAR(1)
s2623	0.78	y	y	9	4	4	TVAR(1)
s0484	0.79	y	y	13	5	3	TVAR(1)
s0019	0.94	y	y	9	3	4	TVAR(1)
s1616	0.96	n	y	8	3	1	TVAR(1)
s0129	1.04	y	y	6	2	0	TVAR(1)
s0026	1.92	y	y	13	5	6	TVAR(1) & TVARMA(2,1)
s2917	2.34	n	n	8	3	0	
s0030	2.96	n	y	6	2	1	TVAR(1)
s0073	q	n	y	9	3	3	TVAR(1)
s0077	q	n	y	9	3	2	TVAR(1)
s0085	q	n	y	5	2	1	TVAR(1)
s0093	q	n	y	10	4	5	TVAR(1)
s0102	q	n	y	14	5	2	TVAR(1)
s0105	q	n	y	9	4	2	TVAR(1)
s1414	q	n	n	8	2	1	

processes, except for a few points before $t=6000s$, when the the number of TVARMA(2,1) components falls to 2, and the number of TVAR(1) components rises to 4. While it can be seen from the periodogram that the lpDNO is modeled in the second TVARMA(2,1) component (Comp 2 in figure 184), the QPOs are only seen in the TVAR(1) processes (Comp 4-6 in figure 184).

The top panel of figure 185 shows the number of complex eigenpairs at each time, while the panel is a measure of the divergence of \mathbf{K}_t from the identity. The number of eigenpairs is 3 except for the points mentioned previously just before 6000s. Where this change occurs, \mathbf{K}_t clearly differs significantly from the identity matrix. There is also a point near the 500s where \mathbf{K}_t diverges significantly from the identity. For the rest of the decomposition, \mathbf{K}_t is close enough to the identity that we may assume that the decomposition is faithful to the true latent processes in the data.

The bottom two panels of figure 185 show the period and amplitude of the complex components. The period for the second component, in which the lpDNO is modeled, hovers around 10s, which suggests that the lpDNO is modeled in the amplitude change of this component, rather than the frequency component. However, no such fluctuation is visible in the amplitude of the second component (seen in the bottom panel of figure 185).

Figure 186 shows the decomposition of run s1307, which is known to contain a fundamental DNO at 31s, a first harmonic DNO at 18 seconds, a fundamental QPO at 500s, a QPO first harmonic at 250s and a QPO second harmonic at 175s. The run is modeled as a TVAR(16) process, with 8 or 9 eigenpairs. From the periodograms it appears that the DF and D1 are modeled in the third and fourth components respectively, which is confirmed

by period decomposition (third panel of 187), which shows that the characteristic period for the fundamental DNO components is about 25s and 15s for the first harmonic. The first and second QPO harmonics appear to be modeled by the first two components of the decomposition, but the characteristic periods of these components are much less than the periods of the QPOs. Again, it might be that the QPOs are modeled in the amplitudes of the components, but this is not clear from the amplitude plot. The 500s QPO is modeled in the 9th and 10th components, which are TVAR(1).

The number of complex pairs changes 5 times over the course of the run, which results in switching between components which complicates the analysis. It is clear that K_t differs considerably from the identity for much of the decomposition, resulting in us being unable to clearly interpret the components in a meaningful way.

The model order required to model our VW Hyi data ranged from 5 to 16; generally longer data sets required more TVAR coefficients, but this was not always the case. In practise, the TVARMA(2,1) components usually modeled the higher frequency aspects of the data - the DNOs, while the QPOs were modeled by the TVAR(1) processes. We found that K_t was generally close to the identity for runs in which the QPOs were modeled in TVAR(1) components, with a few large spikes, but for the 3 runs in which the QPOs were modeled in TVARMA(2,1) components, K_t often differed substantially from the identity. Most decompositions also had a least one change in the number of complex pairs, resulting in switching between components.




Figure 184: Eigen-decomposition of s1616




Figure 185: Number of complex eigenvalues, K-identity measure, period decomposition and amplitude decomposition for the eigen-decomposition of s1616.




Figure 186: Eigen-decomposition of s1307

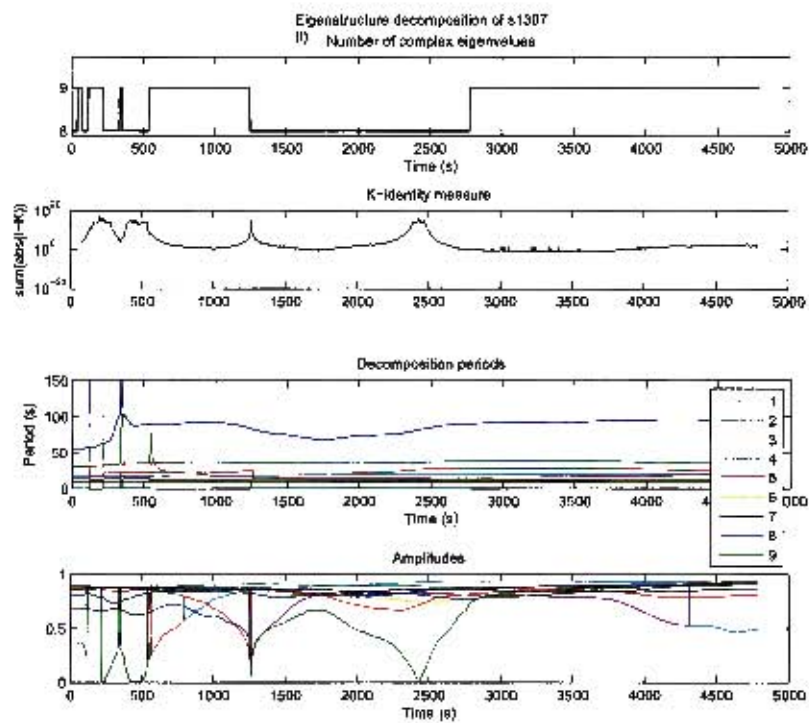


Figure 187: Number of complex eigenvalues, K-identity measure, period decomposition and amplitude decomposition for the eigen-decomposition of s1307.

9.5 Discussion

That DNOs could be modeled with an AR(2) process indicates that while the method might be helpful for investigating the time-varying nature of a series, it cannot be assumed that the model accurately depicts the physics of the situation, since DNOs do not, in fact, vary stochastically, as an AR(2) model assumes. We thus felt that even if we had managed to model QPOs successfully with AR(2) processes, perhaps by smoothing the data first to remove the DNOs, we would not gain anything from the results, as they could not be used to deduce information about the physics of QPOs.

While the TVAR decomposition did model DNOs, the period/amplitude results were not meaningful; again, it did not seem worthwhile to pursue this type of analysis for QPOs.

10 Conclusion

In this thesis, we developed and tested methods of analysing the low signal-to-noise, non-stationary time series with quasi-period components characteristic of cataclysmic variable data. We began the thesis with an overview of cataclysmic variables, and the oscillatory behaviour they display. In chapter 2, we reviewed the existing literature on QPOs in cataclysmic variables, focusing on the methods that have been used to investigate them, and the characteristics that they exhibit. We ended the section with a discussion of the characteristics which a tool capable of detecting QPOs would have to encompass, concluding that in order to detect QPOs, we needed a joint time-frequency technique which could detect relatively small numbers of cycles of oscillations with changing amplitudes in relatively high noise levels, and allow us to see the evolution of the frequency and amplitude of oscillations in the time-series over time.

In chapter 3 we gave an overview of time-frequency methods used in time series analysis in astronomy, focussing on the type of data analysed and the type of analysis undertaken. While the wavelet transform is perhaps the best known, we also discussed several other methods which have been known to be superior for some data types. The majority of the data analysed in the literature have periodicities large enough to be seen in the periodogram, but this is not the case for our data. Our analysis also requires a broader frequency range than that needed for any of the analyses reviewed. Significance testing is also of greater importance for our analysis than in any of the papers reviewed, since our signals are often buried in high noise levels. These difference between our data and those previously analysed in the literature

led us to conduct an exploration of the the various time-frequency techniques, using synthetic data mimicking the features of our data, in chapter 4.

One of the main features of time-frequency representations other than the wavelet scalogram and the Gabor spectrogram is the presence of interference terms, which can be mistaken for signal terms if there is no additional means of differentiation (as is the case for our data). We investigated a selection of time-frequency representations which allow variable time-frequency attenuation, but found that if, as is required for our data, we ensure that the interference terms are completely attenuated, then the time and frequency resolution of these representations are insufficient for our analysis. We also discussed the noise properties of each of the time-frequency distributions analysed; while the majority of the time-frequency representation allow calculation of confidence contours based on the power spectrum of the signal, the cone-kernel representation, which has perhaps the best time and frequency resolution and interference term attenuation, does not, making it an unsatisfactory choice for our analysis. The Gabor spectrogram does not suffer from interference terms, but, over the broad frequency range we require, either has poor time resolution at short periods or poor frequency resolution, neither of which is acceptable. We concluded that the wavelet spectrum, with confidence contours and wavelet ridges, was an appropriate tool for analysing DNOs and QPOs, since it has good time resolution at high frequencies, good frequency resolution at low frequencies, and a well-defined stochastic distribution.

In Chapter 5 we gave a more detailed introduction to wavelet theory, focussing on a discussion of factors affecting choice of analysing wavelet and

instantaneous frequency estimation. We also included an number of tests of the wavelet spectrum with synthetic data mimicking features of DNO and QPO behaviour. We showed that even in very low signal-to-noise conditions (< 0.6) we were still often able to identify signal periodicities. We also showed that the Morlet wavelet is ideal for frequency estimation and signal detection, while the MHAT wavelet, which has good time resolution, can be used to detect discontinuities, and to find the extrema of an oscillation.

While wavelet analysis had been shown to be useful for detection of quasi-periodicities, in chapter 6 we introduced another time-frequency method, Empirical Mode Decomposition, as a means of analysing the amplitudes of detected oscillations in more detail. We used synthetic data to probe the behaviour of EMD for signals such as ours, and conducted Monte Carlo tests to investigate the effect of noise on the amplitude of the extracted signal, concluding that higher noise levels result in decreased estimated amplitudes. We also showed that EMD can be an effective tool for prewhitening data.

We used wavelet analysis and EMD to investigate data from RU Peg, U Gem, TY PsA and VW Hyi, and discussed our results in some detail in chapter 7. We chose data sets which have already been extensively analysed, in order to compare our results with those of standard techniques such as periodogram analysis. In the case of VW Hyi, this enabled us to build up a set of criteria for detection of QPOs, which we used to identify a further 62 new QPOs. We also detected 7 new lpDNOs, doubling the existing number.

We used the results of this detailed analysis to extend the existing model of DNO and QPO evolution in chapter 8. In addition, we found that the period of lpDNOs increases linearly during late outburst, although at a slower rate

than that seen in normal DNOs. We also showed that during the early non-linear period increase phase seen in VW Hyi's DNOs, the amplitude of the DNOs decreases, but during the linear phase the amplitude remains roughly constant. No amplitude evolution during outburst was found in QPOs. We discussed the implications for our new findings for the various models of DNOs and QPOs currently hypothesized, although we could not use our results to exclude any models.

We attempted to model QPOs using AR and TVAR processes; section 9 covered the theory of the models we used, and the results of applying these models to RU Peg and VW Hyi data. We concluded that neither stationary AR models nor time-varying autoregressive decomposition are useful in the analysis of QPOs.

References

- Adami, C., Slezak, E., Durret, F., Conselice, C. J., Cuillandre, J. C., Gallagher, J. S., Mazure, A., Pelló, R., Picat, J. P. & Ulmer, M. P., 2005, *A&A*, **429**, 39.
- Africano, J., Quigley, R., 1978, *PASP*, **90**, 191.
- Albano, A. M., Muench, J., Schwartz, C., Mees, A. I., Rapp, P. E., 1988, *Phys. Rev. A*, **38**, 3017.
- Albano, A., Passamante, A., and Farrell, M., *Nucl. Physica D*, **54**, 85.
- Anderson, N., 1988, *ApJ*, **325**, 266.
- Andronov, I.L., 1997, *A&A*, **125**, 207.
- Andronov, I.L., 1999, *IAUS*, **191**, 201.
- Andronov, I.L., 2000, *Baltic Astronomy*, **9**, 532.
- Andronov, I.L., 2003, *ASPC*, **292**, 391.
- Andronov, I.L., 2005, *ASPC*, **335**, 37.
- Andronov, I., Chinarova, L., 2003, *ASPC*, **292**, 211.
- Andronov, I.L., Arai, K., Chinarova, L.L., Dorokhov, N.I., Dorokhova, T.N., Dumitrescu, A., Nogami, D., S. V. Kolesnikov, A. V., Lepardo, A., Mason, P.A., Matsumoto, K., Oprescu, G., Pajdosz, G., Passuelo, R., Patkós, L., Senio, D.S., Sostero, G., Suleimanov, V.F., Tremko, J., Zhukov, G.V., and Zola, S., 1999, *ApJ*, **117**, 574.
- Aruajo-Betancor, S. et al., *A&A*, **430**, 629.
- Arentoft, T., & Sterken, C., 2000, *J. Ast. Data* **6**, **4C**, 1.

- Astley, V., Komm, R., Howe, R., 2004, *AAS*, **204**, 5311.
- Auger, F., Flandrin, P., Gonçalves, P., & Lemoine, O., 1996, *Time-Frequency Toolbox: For Use with MATLAB*, CNRS(France) & Rice University (USA), available for download from <http://tftb.nongnu.org/>.
- Bailey, J., 1979, *MNRAS*, **187**, 645.
- Barbera, M., Bocchino, F., Damiani, F., Micela, G., Sciortino, S., Favata, F., & Harnden Jr., F., 2002, *A&A*, **387**, 463.
- Barret, D., Olive, J. & Kluzniak, W., 2005, *MNRAS*, **357**, 1288.
- Barthès, D., Mattei, J., 1997, *AJ*, **113**, 373.
- Bateson, F. M., 1990, *R. A. S. N. Z., Pub. Var. Star Sec.*, **17**, 1.
- Bedding, T., Zijlstra, A., Jones, A., Foster, G., 1998, *MNRAS*, **301**, 1073.
- Belloni, T., Parolin, I., Casella, P., 2004, *A&A*, **423**, 969.
- Ben-Mizrachi, A., Procaccia, I., 1984, *Phys. Rev. A*, **29**, 975.
- Bergeron, P., Wesemael, P., Lamontagne, R., Fontaine, G., Saffer, R., Allard, N., 1995, *ApJ*, **449**, 258.
- Bertrand, J. & Bertrand, P., 1992, *J. Math. Phys.*, **33**, 2515.
- Blinchikoff, H. J., & Zverev, A. I., *Filtering in the Time and Frequency Domains*(John Wiley & Sons)
- Boashash, B., 1992, *Proc. IEEE*, **80**, NO. 4, 520.
- Boberg, F. & Lundstedt, H., 2002, *ESASP*, **477**, 79.
- Bond, H., Grauer, A., Burstein, D. and Marzke, O., 1987, *PASP*, **99**, 1097.

- Borisov, G. V., 1992, *A&A*, **261**, 154.
- Boyd, P., & Smale, A., 2004, *ApJ*, **612**, 1006.
- Boyd, P., Carter, P., Gilmore, R. & Dolan, J., 1995, *ApJ*, **445**, 861.
- Box, G. E. P., Jenkins, G. M., *Time Series Analysis: forecasting and control*
Revised Edition, (Holden-Day) Oakland, California.
- Brockwell, P. J., Davis, R. A., *Time Series: Theory and Methods* Second
Edition, (Springer-Verlag) New York
- Brosch, N., Hoffman, Y., 1999, *MNRAS*, **305**, 241.
- Bruch, A., 1992, *A&A*, **266**, 237.
- Bruch, A., 1996, *A&A*, **312**, 97.
- Bruch, A., Steiner, J., Gneiding, C., 2002, *PASP*, **112**, 237.
- Buchler, J. & Kolláth, Z., 2001, *ASSL*, **257**, 185.
- Buchler, J., Kolláth, Z., & Cadmus, R., 2004, *ApJ*, **613**, 532.
- Buchler, J., Kolláth, Z., & Serre, T., 1996, *ApJ*, **462**, 489.
- Buchler, J., Kolláth, Z., & Serre, T., 1996, *ASPC*, **135**, 231.
- Burg, J.P., 1972, *Geophysics*, **37**, 375.
- Cannizzo, J. K., Goodings, D. A., Mattei, J., 1990, *ApJ*, **357**, 235.
- Carothers, J., & Swenson, E., 1992, *SBIR Phase 1 Final Report: Unconventional Signal Processing Using the Cone Kernel Time-Frequency Representation*, Space Application Corporation, Univ. Washington.

- Chatfield, C., *The analysis of time series: theory and practice*, Chapman and Hall Ltd, London.
- Chen, A.-L., 1994, 'A Survey of cataclysmic variables from the Edinburgh-Cape Blue Object Survey', PHD Thesis, University of Cape Town.
- Chinarova, L., Andronov, I., 2000, *OAP*, **13**, 116.
- Choi, H., Williams, W., 1989, *IEEE Trans ASSP*, **37**, 862.
- Chui, C. K., 1997, *Wavelets: A Mathematical Tool for Signal Processing* (Philadelphia: Society for Industrial and Applied Mathematics)
- Claasen, T., & Mecklenbräuker, W., 1984, *IEEE Int. Conf. ASSP*, 41B7.1,
- Cohen, L., 1989, *Proc. IEEE*, **77**, 941.
- Cohen, L., 1995, *Time-frequency analysis* (Englewood Cliffs, NJ: Prentice-Hall)
- Cohen, A., Kovačević, J., 1996, *Proceedings of the IEEE*, **84**, 514.
- Cook, M. C., Warner, B., 1981, *MNRAS*, **196**, 55P.
- Cook, M. C., Warner, B., 1984, *MNRAS*, **207**, 705.
- Cordova, F., Chester, T., Mason, K., Kahn, S., Garmier, G., 1984, *ApJ*, **278**, 739.
- Coughlin, K., & Tung, K., 2002, *AdSpR*, **34**, 323.
- Da Rocha, C., Mendes de Oliveira, C., 2005, *MNRAS*, **364**, 1069.
- Daminai, F., Maggio, A., Micela, G., & Sciortino, A., 1997, *ApJ*, **483**, 370.
- Daubechies, I., 1992, *Ten Lectures on Wavelets*, Capital City Press, Montpelier, Vermont.

- Deeming, T.J., 1970, *A. J.*, **75**, 1027.
- Deeming, T.J., 1975, *Astrophysics and Space Science*, **36**, 137.
- Delprat, N., Escudié, B., Guillemain, P., Kronland-Martinet, R., Tchamitchian, P., Torrèsani, B., 1992, *IEEE Trans. Inf. Theory*, **38**, 644.
- Dessart, L., & Osocki, S., 2002, *A&A*, **393**, 991.
- Djurić, P. M., Kay, S. M., 1998, in *The Digital Signal Processing Handbook*, eds V. K. Madisetti & D. B. Williams, CRC Press in Cooperation with IEEE Press, USA, Chapter 14.
- Dolan, J., Hill, R., Boyd, P., Silvis, J., Percival, J., and van Citters, G., 1998, *A&A*, **331**, 1037.
- Donahue, M., et al, 2002, *ApJ*, **569**, 689.
- Eason, E. L. E., Worden, W. P., Klimke E., Africano J. L., 1984, *PASP*, **96**, 372.
- Elsworth, Y., & James, J., 1982, *MNRAS*, **198**, 889.
- Elsworth, Y., & James, J., 1986, *MNRAS*, **220**, 895.
- Engle, R., *Econometrica*, **50**, 987.
- Figueiredo, A., Nave, M., & EFDA-JET contributors, 2004, *Rev. Sci. Instr.*, **75**, 4268.
- Flandrin, P., 1999, *Time-Frequency/Time-Scale Analysis*, San Diego, CA: Academic.
- Flandrin, P., 2000, in *Nonlinear and Nonstationary Signal Processing*, Ed. Fitzgerald, W., Smith, R., Walden, A. & Young, P., Cambridge University Press, Cambridge, UK, p275.

- Flandrin, P., & Martin, W., 1984, *Lecture Notes in Control & Information Science*, Springer Berlin, Heidelberg, p17.
- Flandrin, P., Gonçalves, P., & Rilling, G., 2004, *XII. European Signal Processing Conference*, Vienna, Austria.
- Flandrin, P., Rilling, G., & Gonçalves, P., 2004, *IEEE Sig. Proc. Let.*, **11**, 112.
- Foster, G., 1996, *Astron. J.*, **112**, 1709.
- Frescura, F., Engelbrecht, C., Frank, B., 2007, *MNRAS*, **temp**, 767.
- Frick, P., Baliunas, S. L., Galyagin, D., Sokoloff, D., Soon, W., 1997, *ApJ*, textbf483, 426.
- Fritz, T., Bruch, A., 1998, *Astron. Astrophys.*, **332**, 586.
- Gabor, D., 1946, *IEEE: Proceedings of the Radio Section*, **93 III**, 429.
- Gänsicke, B., et al., 2006, *MNRAS*, **365**, 969.
- Garnavich, P. M., Szkody, P., 1992, *J. AAVSO*, **21**, 81.
- Gill, A., & Henriksen, R., 1990, *ApJ*, textbf365, L27.
- Gonçalves, P., Flandrin, P., & Chassande-Mottin, E., 1997, *Time-frequency methods in time-series data analysis*, Second Workshop on Gravitational Wave Data Analysis, (M. Davier, P. Hello eds.), pp. 35–46, Frontiers 1997.
- Gonçalves, P., & Baraniuk, R., 1998, *IEEE Trans. Sig. Proc.*, **46**, 1505.
- González Pérez, J., Solheim, J., Dorokhova, T., & Dorokhov, N., 2003, *Baltic Astr.*, **12**, 125.

- Goswami, J. C., Chan, A. K., 1999 *Fundamentals of Wavelets: Theory, Algorithms, and Applications* (New York: John Wiley & Sons)
- Goupil, M., Auvergne, M., & Baglin, A., 1990, *ASPC*, **11**, 578.
- Goupil, M., Auvergne, M., & Baglin, A., 1991, *A&A*, **250**, 89.
- Grassberger, P., Procaccia, I., 1983, *Phys. Rev. Lett.*, **50**, 346.
- Gnedin, Yu. N., Nagovitsyn, Yu. A., Natsvlshvili, T. M., 1999, *Astron. Rep.*, **43**, 462.
- Haefner, R., Schoembs, R., and Vogt, N., 1977, *Astr. Ap.*, **61**, L37.
- Halevin, A., Andronov, I., Shakhovskoy, N., Pavlenko, E., Kolesnikov, S., Ostrova, N., 2002, *ASPC*, **261**, 155.
- Hamilton, J., 1994, *Time Series Analysis*, Princeton University Press, Princeton, New Jersey.
- Harris, F., 1978, *Proc. IEEE*, **66**, 51.
- Harvey, A., 1993, *Time Series Models*, MPG Books Ltd, Bodmin, Cornwall.
- Hearon, A., & Amin, M., 1995, *IEEE Trans. Sig. Proc.*, **43**, 1258.
- Hellier, C., 2001, *Cataclysmic Variable Stars: How and why they vary* Praxis Publishing Ltd, Chichester, UK.
- Henry, B., Lovell, N., and Camacho, F., *Nonlinear Dynamics Time Series Analysis*, Draft.
- Hesser, J. E., Lasker, B. M., Osmer, P. S., 1972, *ApJ*, **176**, L31.
- Hildebrand, R., Spillar, E., Stiening, R., 1981, *ApJ*, **248**, 268.

- Hiriart, D., Salas, L., & Cruz-González, I., 2004, *ApJ*, **128**, 291.
- Hlawatsch, F., 1991, Technical Report 1291-0001, Dept. Elec. Eng., University of Rhode Island.
- Hlawatsch, F. & Boudreaux-Bartels, G., 1992, *IEEE Sig. Proc. Mag*, **April**, 21.
- Hlawatsch, F., & T., Urbanke, R., 1994, *IEEE Trans. Sig. Proc*, **42**, 357.
- Hlawatsch, F., Manickam, T., Urbanke, R. & Jones, W., 1995, *Signal Processing*, **43**, 149.
- Hollander, A., van Paradijs, J., 1992, *A&A*, **265**, 77.
- Horne, J. H., Baliunas, S. L., 1986, *ApJ*, **302**, 757.
- Horne, J., Gomer, R., 1980, *ApJ*, **237**, 845.
- Howarth, J., 1991, *J. Br. Astron. Assoc.*, **101**, 101.
- Howarth, J., & Greaves, J., 2001, *MNRAS*, **325**, 1383.
- Huang, N. E., Shen, Z., Long, S. R., Wu, M. C., Shih, H. H., Zheng, Q., Yen, N., Tung, C. C., Liu, H. H., 1998, *Proc R. Soc. Lond. A*, **454**, 903-995.
- Huang, N. E., Shen, Z., Long, S. R., 1999, *Ann. Rev. Fluid Mech.*, **31**, 417-457.
- Husser, T.-O., Dreizler, S., Solanki, S., Thomas, R., 2005, *Astron. Nachr.*, **326**, 628.
- Isham, V., 1993, in *Networks and Chaos Statistical and Probabilistic Aspects* (London: Chapman and Hall), Chapter 3.

- Ivanović, V., Daković M., & Stanković, L., 2003, *IEEE Trans. Sig. Proc.*, **51**, 77.
- Jenkins, J., 2002, *ApJ*, **575**, 493.
- Jones, M., & Watson, M., 1992, *MNRAS*, **257**, 633.
- Kalman, R., 1960, *J. Basic Eng.*, **82**, 35.
- Kammler, D. W., 2000, *A First Course in Fourier Analysis* (Upper Saddle River, New Jersey: Prentice Hall).
- Kato, T., 2002b, *PASJ*, **54**, 87.
- Kato, T., Hirata, R., 1990, *IBVS*, **3489**.
- Kato, T., & Kunjaya, C., 1995, *PASJ*, **47**, 163.
- Kato, T., Starkey, D. R., 2002, *IBVS*, **5358**.
- Kato, T., Hirata, R., Mineshige, S., 1992, *PASJ*, **44**, L215.
- Kato, T., et al., 2003, *PASJ*, **55**, 489.
- Kato, T., et al., 2004, *MNRAS*, **347**, 861.
- Katsiyannis, A., Williams, D., McAteer, R., Gallagher, P., Mathioudakis, M., Keenan, F., 2002, *ESASP*, **505**, 441.
- Kendall, M., Stuart, A., & Ord, J. K., 1983, *The Advanced Theory of Statistics, Volume 3, 'Design and analysis, and time-series'*, The University Press (Belfast), Northern Ireland.
- Kiss, L., & Szatmáry, K., 1999, *IAUS*, **191**, 133.
- Kiss, L., Szatmáry, K., 2002, *A&A*, **390**, 585.

- Kiss, L., Szabó, G., Bedding, T., 2000, *MNRAS*, **372**, 1721.
- Kiss, L., Szatmáry, K., Szabó, B., Mattei, J., 2000, *A&ASS*, **145**, 283.
- Kiss, L., Szatmáry, K., Cadmus, R., Mattei, J., 1999, *AA*, **346**, 542.
- Kodera, K., Villedary, C., & Gendrin, R., 1976, *Phys. Earth & Plan. Int.*, **12**, 142.
- Koen, C., 1990, *ApJ*, **348**, 700.
- Kolláth, Z., Buchler, J., 1997, *Nonlinear Signal and Image Analysis. Edited by J. Robert Buchler and Henry Kandrup*, **808**, 116.
- Kolláth, Z., Buchler, J., Serre, T., & Mattei, J., 1998, *A&A*, **329**, 147.
- Kolláth, Z. & Csubry, Z., 2006, *Memorie della Societa Astronomica Italiana*, **77**, 109.
- Komm, R., Hill, F., & Howe, R., 2001, *ApJ*, **558**, 428.
- Kozhevnikov, V., 2003, *A&A*, **398**, 267.
- Kraicheva, Z., Stanishev, V., Genkov, V., Iliev, L., 1999, *A&A*, **351**, 607.
- Kuulkers, E., Hollander, A., Oosterbroek, T., and van Paradijs, J., 1991, *A&A*, **242**, 401.
- Kurtz, D., 1985, *Mon. Not. R. astr. Soc*, **213**, 773.
- Lachowicz, P. & Czerny, B., 2005, *MNRAS*, **361**, 645.
- Lachowicz, P., Czerny, B. & Abramowicz, M., 2006, *MNRAS*, **368**, 1025.
- Lamb F., Al J., Cook, M., Lamb D., 1983, *ApJ*, **274**, 71L.
- Lazzati, D., Campana, S., Rosati, P. et al., 1999, *ApJ*, **524**, 414.

- Lefèvre, L., et al., 2005, *MNRAS*, **360**, 141.
- Lépine, S., & Moffat, A., 1996, *ApJ*, **466**, 392.
- Loughlin, P., Pitton, J., Atlas, L., 1993, *IEEE Trans. Sig. Proc.*, **41**, 750.
- Loughlin, P., & Cohen, L., 2004, *IEEE Trans. Sig. Proc.*, **52**, 1218.
- Li, M., Jia, W., & Zhao, W., 1999, *Proc. of IEEE icics99*, Dec. 7-10, Singapore.
- Li, Q., Wu, J., Xu, Z., Wu, J., 2007, *ChA&A*, **31**, 261.
- Liang, Hl, Lin, Q., & Chen, J., 2005, *IEEE Trans. Bio. Eng.*, **52**, 1692.
- Madgwick, D., Hewett, P., Mortlock, D., Lahav, O., 2002, *MNRAS*, **334**, 209.
- Mallat, S., 1989, *A Wavelet Tour of Signal Processing*, Academic Press, California, USA.
- Marchenko, S. & Moffat A., 1998, *ApJ*, **499**, 195.
- Marino, B. F., Walker, W.S.G., 1984, *South. Stars*, **30**, 389.
- Marple, S., 1999, *IEEE Trans. Signal Proc.*, **47**, 2600.
- Marsh, M., Walsh, R., & Bromage, B., 2002, *ESASP*, **508**, 319.
- Masciadri, E., & Raga, E., 2004, *ApJ*, **611**, L137.
- Mathioudakis, M., Seiradakis, J. H., Williams, D. R., Avgoloupis, S., Bloomfield, D. S., McAteer, R. T. J., 2003, *A&A*, **409**, 1101.
- Signal Processing Toolbox User's Guide, The MathWorks, Inc.
- Mauche, C.W., 2002, *ApJ*, **580**, 423.

- Mazeh, T., Tal, Y., Shaviv, G., Bruch, A., Budell, R., 1985, *A&A*, **149**, 470.
- McAteer, R., Gallagher, P., Bloomfield, D., Williams, D., Mathioudakis, M., Keenan, F., 2004, *ApJ*, **602**, 436.
- McSaveney, J., Cottrell, P., Pollard, K., Mattei, J., 2002, *ASPCS*, **259**, 576.
- Mecklenbräuker, W., & Hlawatsch, F., 1997, *The Wigner Distribution: Theory and Applications in Signal Processing*, Elsevier, The Netherlands.
- Meiksin, A., 2000, *MNRAS*, **314**, 566.
- Meyers, S., Kelly, B., & O'Brien, J., 1993, *Mon. Weather Rev.*, **121**, 2858.
- Middleditch, J., Córdova, F. A., 1982, *ApJ*, **255**, 585.
- Misra, R., Harikrishnan, K., Ambika, G. & Kembhavi, A., 2006, *ApJ*, **643**, 1114.
- Mitra-Kraev, U., Harra, L., Williams, D., Kraev, E., 2005, *A&A*, **436**, 1041.
- Moffat, A., Lépine, Henriksen, R., & Robert, C., 1994, *A&SS*, **216**, 55.
- Mumford, G. S., 1969, *Sky Telesc.*, **37**, 288.
- Nather, E., 1973, *Vistas in Astronomy*, **15**, 91.
- Nather, R., Robinson, E., 1974, *ApJ*, **190**, 637.
- Nilsson, R., Uthas, H., Ytre-Eide, M., Solheim, J-E., & Warner, B., 2006, *MNRAS*, **370**, L56.
- Nogami, D., Baba, H., Kato, T., Novák, R., 1998, *PASJ*, **50**, 297.
- Nogami, D., Kato, T., Baba, H., 2002, *PASJ*, **54**, 987.
- Nyblom, J., 1989, *JASA*, **84**, 223.

- O'Brien, K., Horne, K., Boroson, B., Still, M., Gomer, R., Oke, Boyd, P., & Vrtilek, S., 2001, *MNRAS*, **326**, 1067.
- O'Donoghue, D., 1986, *MNRAS*, **220**, 23P.
- O'Donoghue, D., 1995, *Baltic Astr.*, **4**, 517.
- O'Neill, J., & Flandrin, P., 2000, *IEEE Trans. Sig. Proc.* **48**, 2641.
- O'Neill, J., & Williams, W., 1999, *IEEE Trans. Sig. Proc.*, **47**, 133.
- Oh, S., & Marks, R., 1992, *IEEE Trans. Sig. Proc.*, **40**, 1735.
- Orr, W., & Hoffman, H., 1974, *IEEE Trans. Bio. Eng.*, **21**, 130.
- Osaki, Y., 1974, *PASJ*, **26**, 429.
- Ostriker, J.P., Hesser, J.E., 1968, *ApJ*, **153**, L151.
- Otazu, X., Ribó, M., Peracaula, J. & Núñez, J., 2002, *MNRAS*, **333**, 365.
- Otazu, X., Ribó, M., Paredes, J., Peracaula, J. & Núñez, J., 2004, *MNRAS*, **351**, 215.
- Oung, H., & Forsberg, F., 1998, *IEEE Trans. Sig. Proc.*, **46**, 2616.
- Packard, N., Crutchfield, J., Farmer, J., and Shaw, R., *Phys. Rev. Let.*, **277**, 250.
- Paczynski, B., 1978, in Zytkow A., ed., *Nonstationary Evolution of Close Binaries*. Polish Scientific Publ., Warsaw, p. 89.
- Page, C., 1952, *J. Appl. Phys*, **23**, 103.

- Papadaki, C., Boffin, H., Cuypers, J., Stanishev, V., Kraicheva, Z., Genkov, V., 2004. In: Hilditch, R.W., Hensberge, H., Pavlovski, K. (Eds.), *Spectroscopically and Spatially Resolving the Components of Close Binary Stars*, ASP Conference Series, **318**, Astronomical Society of the Pacific, 399.
- Papadaki, C., Boffin, H., Sterken, C., Stanishev, V., Cuypers, J., Boumis, P., Akras, S., Alikakos, J., 2006, *A&A*, **456**, 599.
- Papandreou-Suppappola, A., 1999, *IEEE Signal Processing/Communication Chapter Meeting, 15 November 1999, Abstract*.
- Patterson, J., 1979, *ApJ*, **234**, 978.
- Patterson, J., 1981, *ApJS*, **45**, 517.
- Patterson, J., 1991, *PASP*, **103**, 1149.
- Patterson, J., 1995, *PASP*, **107**, 657.
- Patterson, J., Skillman, D. R., 1994, *PASP*, **106**, 1141.
- Patterson, J., Robinson, E. L., Nather, R. E., 1977, *ApJ*, **214**, 144.
- Patterson, J., Branch, D., Chincarini, G., Robinson, E. L., 1980, *ApJ*, **240**, L133.
- Patterson, J., Thorstensen, J. R., Fried, R., Skillman, D.R., Cook, L. M., Jensen, L., 2001, *PASP*, **113**, 72.
- Patterson, J., Fenton, W. H., Thorstensen, J. R., Harvey, D. A., Skillman, D. R., Fried, R. E., Monard, B., O'Donoghue, D., Beshore, E., Martin, R., Niarchos, P., Vanmunster, T., Foote, J., Bolt, G., Rea, R., Cook, L. M., Butterworth, N., Wood, M., 2002, *PASP*, **114**, 1364.
- Patterson, J., et al., 2002, *PASP*, **114**, 721.

- Patterson, J., et al., 2002, *PASP*, **114**, 65.
- Pavlenko, E.P., 1998, *Astr. Rep.*, **42**, 200.
- Pavlenko, E.P., Shugarov, S.Yu., 1999, *A&A*, **343**, 909.
- Percival, D., 1995, *Biometrika*, **82**, 619.
- Percival, D., & Walden, A., 2000, *Wavelet Methods for Time Series Analysis*,
Cambridge: Cambridge University Press.
- Percy, J. & Kastrukoff, R., 2001, *JAAVSO*, **30**, 16.
- Pezzuto, S., Bernacca, P. L., Stagni, R., 1992, *A&A*, **257**, 523.
- Pizzolato, N., Maggio, A., & Sciortino, S., *A&A*, **361**, 614.
- Posch, T., 1990, *IEEE Trans. ASSP*, **38**, 725.
- Press, W., Flannery, B., Teukolsky, S. & Vetterling, W., 1988, *Numerical Recipes in C*, Cambridge University Press, USA.
- Pretorius, M. L., Unpublished.
- Pretorius, M. L., Warner, B., Woudt, P.A., 2006, *MNRAS*, **368**, 361.
- Priestley, M., 1981, *Spectral Analysis and Time Series* (London, UK: Elsevier Academic Press)
- Provencal, J., et al., 1995, *ApJ*, **445**, 927.
- Qiu, L., 1993, *Proc. IEEE Singapore Int. Conf. Inf. Eng.*, **1**, 388.
- Rabiner, L. R., 1975, *Theory and Application of Digital Signal Processing*
(Englewood Cliffs, New Jersey: Prentice Hall)

- Régulo, C., Almenara, M., Alonso, R., Deeg, H., & Roca Cortés, T., 2007, *A&A*, **467**, 1345.
- <http://reylab.bidmc.harvard.edu/DynaDx/case-study/seizure/def/hilbert>
accessed 12 Jan 2002.
- Rihaczek, W., 1968, *IEEE Trans. Informat. Theory*, **14**, 369.
- Riera, A., Raga, C. A., Reipurth, B., Amram, P., Boulesteix, J., Cantó, J., & Toledano, O., 2003, *ApJ*, **126**, 327.
- Rilling, G., Flandrin, P., & Gonçalves, P., 2003, *IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing, NSIP-03*, Grado (I).
- Rioul, O. & Flandrin, P., 1992, *IEEE Trans. Sig. Proc*, **40**, 1746.
- Roberts, P., Yu, Z., Russell, C., 2007, *GApFD*, **101**, 11.
- Robertson, S., Leiter, D., 2003, *ApJ*, **596**, 203.
- Robinson, E. L., Nather, R.E., 1979, *ApJ*, **39**, 461.
- Robinson, E. L., Warner, B., 1972, *MNRAS*, **157**, 85.
- Robinson, E. L., Warner, B., 1984, *ApJ*, **277**, 250.
- Robinson, E. L., Shafter, A. W., Allen Hill, J., Wood, M. A., 1987, *ApJ*, **313**, 772.
- Rodríguez-Gil, P., Casares, J., Martínez-Pais, I.G., Hakala, P., Steeghs, D., 2001, *ApJ*, **548**, L49.
- Rodríguez-Gil, P., Casares, J., Martínez-Pais, I.G., Hakala, P., 2002, *The Physics of Cataclysmic Variables and Related Objects, ASP Conf. Proc*, **261**, 533.

- Rosati, P. et al., 1995, *ApJ*, **445**, L11.
- Rosen, S. R., Watson, T. K., Robinson, E. L., Prinja, R. K., Misselt, K.,
Shafter, A. W., 1995, *A&A*, **300**, 392.
- Rossi, V., & Salusti, E., 1995, *An. Geo.*, **13**, 437.
- Ruzmaikin, A., Santee, M., Schwartz, M., Froidevaux, L., Pickett, H., 2007,
Geophys. Res. Let., **34**, L02819.
- Sabin, L., & Zijlstra, A., 2006, *Mem. S. A. It.*, **77**, 148.
- Sabin, L., & Zijlstra, A., 2006, *Mem. S. A. It.*, **77**, 933.
- Salon, S., Crise, A., Picco, P., de Marinis, E., Gasparini, O., 2003, *An. Geo.*,
21, 833.
- Scargle, J. D., 1981, *ApJS*, **45**, 1.
- Scargle, J. D., 1982, *ApJ*, **263**, 835.
- Scargle, J. D., 1992, *Chaotic Processes in Astronomical Data*,
Penn State Conference 1992. Available for download at
<http://trotsky.arc.nasa.gov/jeffrey/>.
- Scargle, J. D., 1997, *Wavelet and Other Multi-resolution Methods for Time
Series Analysis*, Penn State Conference 1997. Available for download at
<http://trotsky.arc.nasa.gov/jeffrey/>.
- Scargle, J., Steiman-Cameron, T., Young, K., Donoho, D., Crutchfield, J.,
Imamura, J., 1993, *ApJ*, **411L**, 91.
- Schoembs, R., Stolz, B., 1981, *IBVS*, **1986**.
- Semeniuk, I., Schwarzenberg-Czerny, A., Duerbeck, H., Hoffman, M., Smak,
J., Stepień, K., Tremko, J., 1987, *Act. Astr.*, **37**, 197.

- Serre, T., Kolláth, Z., & Buchler, J., 1996, *A&A*, **311**, 833.
- Shafter, A., Wheeler, J., & Cannizzo, J., 1986, *ApJ*, **305**, 261.
- Singh, J., Rao, V., Agrawal, P. C., Apparao, K. M. V., Manchanda, R. K.,
Sanwal, B. B. and Sarma, B. K., 1993, *ApJ*, **419**, 337.
- Skillman, D. R., Patterson, J., Thorstensen, J. R., 1995, *PASP*, **107**, 545.
- Smak, J., 1971, *Acta. Astr.*, **21**, 15.
- Smak, J., 1979, *Acta. Astr.*, **29**, 309.
- Smak, J., 1984, *Acta. Astr.*, **34**, 93.
- Smith, D., Dhillon, V., Marsh, T., 1998, *MNRAS*, **296**, 465.
- Smith, E., Zhou, X., Ruzmaikin, A., 2006, *American Geophysical Union*,
Fall Meeting 2006, abstract SH53B-1506.
- Stanković, L., 2001, *Signal Processing*, **81**, 621.
- Stanković, L., 2002, *IEEE Sig. Proc. Let.*, **9**, 286.
- Stanković, L. & Katkovnik, V., 1999, *IEEE Trans. Sig. Proc.*, **47**, 1099.
- Starck, J., Murtagh, F., 1994, *A&A*, **288**, 342.
- Starck, J-L., Siebenmorgen, R., & Gredel, R., 1997, *ApJ*, **482**, 1011.
- Starck, J-L., Bijaoui, A., Lopez, B., & Perrier, C., *A&A*, **283**, 349.
- Stellingwerf, R.F., 1978, *ApJ*, **224**, 953.
- Suleimanov, V. F., Zhukov, G. V., Senio, D. S., 1996, *Odessa Astr. Pub.*, **9**,
57.
- Szatmary, K., Gál, J., & Kiss, L., 1996, *A&A*, **308**, 791.

- Szatmary, K., Kiss, L., & Bebesi, Z., 2003, *A&A*, **398**, 277.
- Szatmary, K., Vinko, J., 1992, *MNRAS*, **256**, 321.
- Szatmary, K., Vinko, J., Gál, J., 1994, *A&ASS*, **108**, 377.
- Szkody, P., 1976, *ApJ*, **207**, 190.
- Szkody, P., 1990, *ApJ*, **100**, 546.
- Takens, R., 1981, in *Dynamical Systems and Turbulence* (Warick)
- Taylor, C. J., Thorstensen, J. R., Patterson, J., 1999, *PASP*, **111**, 184.
- Templeton, M., Mattei, J., & Willson, L., 2005, *ApJ*, **130**, 776.
- Terradas, J., Oliver, R., & Ballester, J., 2004, *ApJ*, **614**, 435.
- Theuns, T., & Zaroubi, S., 2000, *MNRAS*, **317**, 989.
- Thayaparan, T., & Kennedy, S., 2003, *Application of Joint Time-Frequency Representations to a Maneuvering Air Target in Sea-Clutter: Analysis Beyond FFT*, Defence R & D Canada, Ottawa.
- Timmer, J., Koenig, M., 1995, *A&A*, **300**, 707.
- Torrence, C., and Compo, G., 1998, *Bulletin of the American Meteorological Society*, **79**, 61.
- Tovmassian, G. G., 1988, *Adv. Space Res.*, **8**, (2)329.
- Townsend, R., 1999, *MNRAS*, **310**, 851.
- Townsley, D., & Bildsten, L., 2003, 2003, *ApJ*, **596**, 227.
- Tremko, J., Andronov, I. L., Chinarova, L. L., Kumsiashvili, M. I., Luthardt, R., Pajdosz, G., Patkós, Rössiger and Zola, S., 1996, *A&A*, **312**, 121.

- Udalski, A., 1988, *IBVS*, **3239**.
- Uemura, M., Kato, T., Ishioka, R., Novak, R., Pietz, J., 2002, *PASJ*, **54**, 299.
- Unterberger, A., 1993, *Commun. Part. Diff. Equations*, **9**, 1179.
- van der Klis, M., 2002, *ARA&A*, **38**, 717.
- van Paradijs, J., 1983, *A&A*, **125**, L16.
- Vaughan, S., 2005, *A&A*, **431**, 391.
- Ville, J., 1948, *Cable Transmissions*, **2A(1)**, pp.61-74. Translation by I. Selin, "Theory and applications of the notion of complex signal," Report T-92, RAND corporation, Santa Monica, CA.
- Warner, B., 1974, *Mon. Not. R. astr. Soc.*, **168**, 235.
- Warner, B., 1975, *Mon. Not. R. astr. Soc.*, **173**, 37P.
- Warner, B., 1981, *Mon. Not. R. astr. Soc.*, **195**, 101.
- Warner, B., 1985, *Astr. Sp. Sc.*, **118**, 271.
- Warner, B., 1995, *Cataclysmic Variable Stars* (Cambridge: Cambridge University Press).
- Warner, B., in Buckley, D, Warner, B., eds, ASP Conf. Ser. Vol 85, Cape Workshop on Magnetic Variables, Astron. Soc. Pac., San Fransisco, p. 343.
- Warner, B., 2004, *PASP*, **116**, 115.
- Warner, B., Brickhill, A. J., 1978, *MNRAS*, **182**, 777.

- Warner, B., Nather, R. E., 1971, *MNRAS*, **152**, 209.
- Warner, B., Nather, R. E., 1971, *MNRAS*, **152**, 219.
- Warner, B., Nather, R., 1988, *IBVS*, 3140.
- Warner, B., O'Donoghue, D., 1987, *MNRAS*, **224**, 733.
- Warner, B., O'Donoghue, D., Wargau, W., 1989, *MNRAS*, **238**, 73.
- Warner, B., Pretorius, R., 2008, *MNRAS*, **383**, 1469.
- Warner, B., Robinson, E.L., 1972, *MNRAS*, **159**, 101.
- Warner, B., Woudt, P.A., 2002, *MNRAS*, **335**, 84.
- Warner, B., Woudt, P.A., 2003, *Magnetic Cataclysmic Variables ASP Conference Series*, Eds M. Cropper & S. Vrielman, **315**, 39.
- Warner, B., Woudt, P.A., 2004, *Variable Stars in the Local Group ASP Conference Series*, Eds D. Kurtz & K. Pollard, **310**, 382.
- Warner, B., Woudt, P.A., 2005, *PASP*, **334**, 453.
- Warner, B., Woudt, P.A., 2006, *MNRAS*, **367**, 1562.
- Warner, B., Woudt, P.A., 2008, *arXiv*, **0806:1317v1**, astro-ph.
- Warner, B., Woudt, P., Pretorius, M., 2003, *MNRAS*, **344**, 1193.
- Watson, M. G., King, A. R., Osborne, J., 1985, *MNRAS*, **212**, 917.
- West, M., Prado, R., & Krystal, D., 1999, *J.A.S.A.*, **94**, 1.
- Wood, J. H., Irwin, M. J., Pringle, J. E., 1985, *MNRAS*, **214**, 475.
- Williams, G., Hiltner, W., 1984, *MNRAS*, **211**, 629.

- Williams, W., 1998, *Recent Advances in Time-Frequency Representations: Some Theoretical Foundations*, in *Time Frequency and Wavelets in Biomedical Signal Processing*, Edited by M. Akay, IEEE Press, New York, USA.
- Wood, J. H., Horne, K., Berriman, G., Wade R., O'Donoghue, D., Warner, B., 1986, *MNRAS*, **219**, 629.
- Woudt, P. A., Warner, B., 2002, *MNRAS*, **333**, 411.
- Woudt, P. A., Warner, B., 2002, *MNRAS*, **334**, 44.
- Woudt, P. A., Warner, B., 2003, *MNRAS*, **340**, 1011.
- Woudt, P. A., Warner, B., 2003, *MNRAS*, **348**, 599.
- Woudt, P., Warner, B., Pretorius, M & Dale, D., 2005, *PASP*, **330**, 325.
- Wright, D., Lynas-Gray, A., Kilkenny, D., Cottrell, P., Shobbrook, R., Koen, C., van Wyk, F., Kilmartin, P., Martinez, P., Gilmore, A., 2006, *MNRAS*, **369**, 2049.
- Xu, T., Wu, J., Wu, Z., Li, Q., 2008, *ChJAA*, **8**, 337.
- Zhang, E., Robinson, E. L., Ramseyer, T. F., Shetrone, M., Stiening, R. F., 1991, *ApJ*, **381**, 534.
- Zhao, Y., Atlas, L. & Marks, R., 1990, *IEEE Trans. ASSP*, **38**, 1084.
- Zhou, A., 2004, *Chin. J. Astron. Astrophys.*, **4**, 335.
- Zhou, A., & Qiu, Y., 2002, *Ap&SS*, **288**, 327.
- Zijlstra, A., Bedding, T., Mattei, J., 2002, *MNRAS*, **334**, 498.

Zijlstra, A., Bedding, T., Markwick, A., Loidl-Gautschy, R., Tabur, V.,
Alexander, K., Jacob, A., Kiss, L., Price, A., Matsuura, M. & Mattei, J.,
2004, *MNRAS*, **352**, 325.

Zola, S., 1988, *Acta. Astr.*, **39**, 45.

Zwitter, T., Cadez, A., Dintinjana, B., Mikuz, H., 1994, *IBVS*, **3975**.